UNIVERSITY OF SOUTHAMPTON

SCHOOL OF MATHEMATICS

Aspects of the Curve Complex and the Mapping Class Group

by

Kenneth J. Shackleton

Thesis for the degree of Doctor of Philosophy

October 2005

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

SCHOOL OF MATHEMATICS

Doctor of Philosophy

ASPECTS OF THE CURVE COMPLEX AND THE MAPPING CLASS GROUP

by Kenneth J. Shackleton

In this thesis we discuss two related and important objects in the study of geometric group theory and Teichmüller theory, namely the *curve complex* and *mapping class group*. The original material is entirely contained in Chapter 2 and Chapter 3. The two chapters are self-contained and can be read independently, and independent of the introduction in Chapter 1. Both chapters form part of my long-term interest in the computational and large-scale geometry of the curve complex and in the structure of the mapping class group.

In Chapter 2, we discuss some of the computability aspects of the path-metric on the 1-skeleton of the curve complex, called the *curve graph*. Specifically, we develop an algorithm for computing distances in the curve graph by constructing (all) tight geodesics between any two of its vertices. Our work produces bounds on intersection numbers associated to tight geodesics without the need for taking geometric limits. Thus, we complement the work of Masur-Minsky and of Bowditch in this area. We then discuss some of the implications for the action of the *mapping class group* on the curve graph. We discover a computability version of the acylindricity theorem of Bowditch and we recover all of the weak proper discontinuity of Bestvina-Fujiwara. Our methods are entirely combinatorial.

In Chapter 3, we study combinatorial rigidity questions regarding the curve complex and the mapping class group. We see that every embedding of one curve complex to itself or another, whose dimensions do not increase from domain to codomain, is induced by a surface homeomorphism. While previous approaches due to Ivanov (for automorphisms) and Irmak (for superinjective maps) for the same surface require a deep result proven by Harer and by Bowditch-Epstein, among others, regarding the existence of a triangulation of *Teichmüller space* and the approach due to Luo makes essential use of a modular structure, we shall require little more than the connectivity of links in the curve complex. The techniques we develop only require the local injectivity of an embedding on curve complexes, giving us Theorem 3. We then deduce that every "local" injection from any finite index subgroup of a mapping class group to a mapping class group whose complexity is no greater is typically the restriction of a unique inner automorphism, stated as Theorem 4. Thus, mapping class groups do not admit faithful actions on another curve complex of the same dimension. Given the simplicity of our approach, it would seem applicable to very many surface-related complexes and groups.

Contents

List of Figures

ACKNOWLEDGEMENTS

First, it is a great pleasure to offer my deepest thanks to my supervisor, Brian H. Bowditch, for, without your support, guidance and encouragement, this thesis could not exist. I also thank you and your alter ego, Yumimizo-sensei, for the many fun and exciting real-world travels we have had during the last three years: I will never forget all those good times we spent in Dijon, Bonn and Tokyo.

I thank both Ian J. Leary and Graham H. Niblo for your support and advice, and I thank both my examiners, Graham H. Niblo and Caroline M. Series, for so patiently and thoroughly reading this thesis when you had so many pressing things to do. I will always be grateful to my collaborators, James W. Anderson and Javier Aramayona, for all those fun times we spent playing with mapping class groups. I offer my sincerest thanks to Sadayoshi Kojima for making me so welcome in Tokyo, to Makoto Sakuma for inviting me to Osaka, to Koji Fujiwara for inviting me to Sendai and to John S. Crisp for inviting me to Dijon, and I thank Tokyo Institute of Technology, Max-Planck-Institut für Mathematik and Université de Bourgogne for your warm and generous hospitality. I offer my sincerest thanks to Ruth Kellerhals for inviting me to EPFL and to Makoto Sakuma for inviting me to Osaka as a JSPS fellow early next year.

I will always be indebted to the financial support from the EPSRC and the School of Mathematics, University of Southampton, and to the clerical staff. Without your tireless work, so many of the good experiences I have been lucky enough to enjoy would never have come to be.

On a more personal level, I wish to thank all those dear friends who stood by me and supported me when I needed them most: Without you I could not have come this far. I offer my sincerest thanks to Mike, Yoshika, Akika and Emani for all the wonderful times you showed me on Rokko Island and in Osaka. Finally, my sincerest thanks to my parents, Dave and Marilyn, my brothers, Kevin and Paul, and to my grandparents, John, Betty, Alan and Barbara. I thank you all for everything.

Chapter 1

Introduction

ABSTRACT: We define the curve complex and discuss its geometry and the role it has played in the last twenty to thirty years, paying particular attention to its connections with mapping class groups, Teichmüller spaces and the ending lamination theorem. We put the work of Chapter 2 and Chapter 3 into context.

1.1 Curve complex

Let Σ be a closed orientable surface and let $\Pi \subseteq \Sigma$ be a finite (possibly empty) subset. The curve complex $\mathcal{C}(\Sigma, \Pi)$ associated to the pair (Σ, Π), or the punctured surface $\Sigma - \Pi$, was first considered by Harvey [Harv] in the early 1980s and is defined as follows. We shall say that an embedded loop in $\Sigma - \Pi$ is *trivial* if it bounds a disc and *peripheral* if it bounds a once-punctured disc. Let $X = X(\Sigma, \Pi)$ denote the set of all free homotopy classes of non-trivial and non-peripheral embedded loops in $\Sigma - \Pi$. The elements of X will be referred to as *curves*, and we shall always use the Greek alphabet to denote these. We take X as the vertex set of $\mathcal{C}$ and deem a family of distinct curves $\{\gamma_0, \gamma_1, \ldots, \gamma_k\}$ to span a k-simplex if and only if any two of these curves can be realised disjointly in $\Sigma - \Pi$, that is to say any two curves contain disjoint representatives. Indeed, it can be shown that any family of pairwise disjoint curves can be simultaneously realised disjointly in $\Sigma - \Pi$. The resulting simplicial complex is what we shall call the *curve complex.*

The simplicial dimension $C(\Sigma, \Pi)$ of $\mathcal{C}(\Sigma, \Pi)$, when this is defined, is seen to be equal to $3genus(\Sigma)+|\Pi|-4$ where, we note, $3genus(\Sigma)+|\Pi|-3$ is precisely the size of a maximal collection of distinct and disjoint curves. Examples of natural subcomplexes are furnished by the curve complexes associated to subsurfaces of $\Sigma - \Pi$. In particular, the curve complex associated to the complement of a curve α corresponds to the codimension one subcomplex whose vertices are all the neighbours of α; such a subcomplex is referred to as a *vertex link.* Lastly, we may equally well speak of bordered surfaces and this distinction goes undetected at the level of the curve complex. This is relevant in Chapter 2 where we cut surfaces along curves and, to argue inductively as we will, we must speak of bordered surfaces rather than punctured surfaces. We shall clarify this later.

With the exception of only a few cases, namely Σ is a 2-sphere and $|\Pi| \leq 4$ (so that $\Sigma - \Pi$ is a sphere, once-punctured sphere, twice-punctured sphere, thrice-punctured sphere or four-times punctured sphere) and Σ is a torus and $|\Pi| \leq 1$ (so that $\Sigma - \Pi$ is a torus or a once-punctured torus), X is non-empty and the curve complex is connected. The latter is a very elementary but very useful fact easily verified once we introduce the intersection number of any two curves α and β, defined to be equal to the minimal integer $|a \cap b|$ arising among all representative loops $a \in \alpha$ and $b \in \beta$. Given this the path-connectivity of $\mathcal{C}$ in the non-exceptional cases is established by an induction on the intersection number of the two curves concerned that proceeds by surgering α relative to β, yielding a third curve α' of smaller intersection number with β.

We note that when $C(\Sigma, \Pi)$ is zero, so that $\Sigma - \Pi$ is either a four-times punctured sphere or a once-punctured torus, the curve complex is a countable set of points as no two distinct curves can be realised disjointly. For these two cases alone we modify the standard definition of the curve complex by declaring two vertices connected by an edge if and only if they intersect minimally: For the four-times punctured sphere this means they intersect twice and for the once-punctured torus this means they intersect once. This variation gives a very special complex as both are naturally isomorphic to the *Farey graph*. (See Figure 1.1.)

We give all of Luo's curve complex classification up to isomorphism by recalling that two apparently distinct and non-exceptional curve complexes are isomorphic if and only if one surface is the five-holed sphere and the other is the twice-punctured torus or one surface is the six-times punctured sphere and the other is the closed surface of genus two. This classification makes use of the homotopy type of the curve complex, determined by Harer [Hare].

Theorem *(Harer) The curve complex $\mathcal{C}(\Sigma, \Pi)$ is homotopic to a bouquet of spheres all of dimension*
i). $2genus(\Sigma) + |\Pi| - 3$, if $genus(\Sigma) > 0$ and $|\Pi| > 0$,
ii). $2genus(\Sigma) - 2$, if $\Pi = \emptyset$, and
iii) $|\Pi| - 4$, if $genus(\Sigma) = 0$.

(For an implied classification by a combinatorial argument not using Harer's theorem, please see Chapter 3.) In particular every curve complex of simplicial dimension at least two is simply-connected but non-contractible, from which it follows that the curve complex does not support a CAT(0)-metric. It would be of much interest to find an elementary proof of simple-connectivity.

We can regard the curve complex as a *path-metric* space, by agreeing that each edge should have length one (and that the simplices should be euclidean) and then declaring the distance between two of its vertices to be the minimal

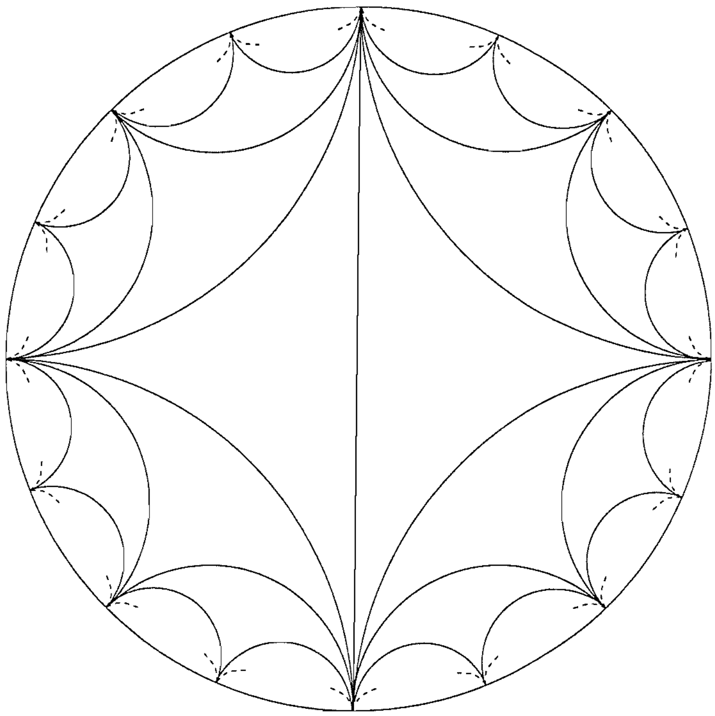

Figure 1.1: An artistic representation of the Farey graph, drawn on the Poincaré disc.

length among all paths that end there. As far as understanding this metric is concerned, it is better to pass to the 1-skeleton $\mathcal{G}(\Sigma, \Pi)$ of the curve complex, called the *curve graph*, where all distances are integers. We do not lose much by doing this, for the curve complex and the curve graph share the same large-scale geometry (that is, they are *quasi-isometric*).

It is easy enough to provide examples of two curves a short distance apart in this graph: Any curve is distance zero from itself, two distinct but disjoint curves are distance one apart and two curves that intersect and are disjoint from a third curve are distance two apart. Thereafter, it is not so easy to recognise distances by eye and generically we can only say that any two intersecting curves that, between them, cut the surface into a disjoint union of discs and once-punctured discs are at least distance three apart. However, Schleimer [Schl] has found an example of a pair of curves in the five times punctured sphere at distance four. In future we shall say that a pair of such curves *fills* the surface.

We can however see upper bounds for distances in terms of intersection numbers, for given two curves α, β their distance in the curve graph is at most $\iota(\alpha, \beta) + 1$. Even this upper bound is not entirely satisfactory, for $\iota(\alpha, \beta)$ can be made arbitrarily large even when $d(\alpha, \beta)$ is held fixed and at least two. An elegant argument attributed to Luo and recorded in [MaMi1], making use of the same geodesic laminations we describe later, does however assure us that this metric on the curve graph, and on the curve complex, is unbounded. It is such difficulties that motivate the question: Given two vertices of the curve graph, can we compute their distance? An affirmative answer is given in Chapter 2, where we deduce, among other things, the following theorem.

Theorem *(Shackleton) There is an explicit algorithm which takes as input a closed surface Σ, a finite subset Π and two curves and returns the distance between these two curves in the curve graph $\mathcal{G}(\Sigma, \Pi)$.*

The curve complex has enduring appeal for not only will we see that it has a very special type of geometry, making it interesting in its own right, but it also has strong interactions with other important surface related groups and spaces and is a key tool in understanding 3-manifolds. The simplicity of its construction correlates with the many challenging problems it continues to pose. It serves as a natural model space for the extended mapping class group and its subgroups and combinatorially describes otherwise analytic objects such as Teichmüller space, the space of all marked constant curvature metrics supported by $\Sigma - \Pi$. It is intimately related to certain splittings of 3-manifolds, called *Heegaard splittings*, and fundamental properties of such splittings are naturally expressed in terms of curve graph distances (see [AbrSch], [Hemp]). Finally, the last few years have seen remarkable advances made in low-dimensional topology and, of most relevance here, in the classification of hyperbolic 3-manifolds where the curve complex assumes a central role.

1.2 Large-scale geometry

In the early 1980s Gromov [Gro] introduced a notion of large-scale equivalence and negative curvature for metric spaces and groups. This brought to life geometric group theory as a subject in its own right, for now many problems that were once viewed as group theoretic became amenable to geometric approaches. There are many subsequent accounts of this and related ideas we can recommend here: See [BriHae], [Bow4], [CDP] and [GhD], and references therein.

One natural class of maps to consider in comparing two metric spaces on a large scale is the class of quasi-isometries. We shall say that two metric spaces are quasi-isometric if one space is a linearly bounded distorsion, say by stretching, tearing or thickening, of the other. More precisely, we say that a function $f : X \longrightarrow Y$ between two metric spaces is a (k, C)*-quasi-isometry* if for each pair of points x_1 and x_2 in X, we have:

$$d_X(x_1, x_2)/k - C \leq d_Y(f(x_1), f(x_2)) \leq k d_X(x_1, x_2) + C$$

with $k \geq 1$ and $C \geq 0$. We say that f is a quasi-isometric embedding if there exist k, C such that f is a (k, C)-quasi-isometry. We say that f is a *quasi-isometry* if it is both a quasi-isometric embedding and is almost surjective, that is its image is ϵ-dense in Y for some non-negative ϵ.

Given a quasi-isometry $f : X \longrightarrow Y$ a coarse inverse is readily constructed. For each point $y \in Y$, there exists $z_y \in f(X)$ uniformly close to y. Now choose $x_y \in f^{-1}(z_y)$, and define $g : Y \longrightarrow X$ by $g(y) = x_y$ for each $y \in Y$. It can be verified that g is a quasi-isometry and that both $g \circ f$ and $f \circ g$ are uniformly close to the respective identity maps. It is now easily seen that the notion of quasi-isometry is reflexive, symmetric and transitive. We stress that the idea of a quasi-isometry is a geometrical and a not topological one, for there is no insistance that a quasi-isometry be continuous. All bounded metric spaces are quasi-isometric, and the euclidean space $\mathbb{R}^n$ and the lattice $\mathbb{Z}^n$ with standard metric are quasi-isometric. Closer to our discussion, we can generalise the curve complex by instead declaring a family of distinct curves to span a simplex if and only if any two have intersection number at most a uniform constant. With similarly defined path-metrics, these complexes are all quasi-isometric.

A key instance of quasi-isometries at work is provided by the family of "word-metrics" on a finitely generated group. Given any group G and a symmetric finite generating set S, each element of G may be written as a word in S. For $g \in G$, we define the norm $|g|$ to be the length of a shortest word in S representing g. The *S-word metric* on G is defined by $d(g, h) = |g^{-1}h|$ for all $g, h \in G$. Any two such metrics are in fact quasi-isometric, for changing one generating set by another uniformly stretches or shrinks word-length. We can therefore speak of *the* word-metric on G. By way of example, any two finite groups are quasi-isometric, $\mathbb{Z}^m$ and $\mathbb{Z}^n$ are quasi-isometric only if m and n are equal, and any two finitely generated free groups of rank at least two are quasi-isometric.

An important invariant of a quasi-isometry is a certain large-scale negative curvature. For a metric space X with basepoint x_0 we define, following Gromov, the inner product of two points $x, y \in X$ relative to this basepoint as

$$(x.y)_{x_0} = (d(x_0, x) + d(x_0, y) - d(x, y))/2.$$

When X is a metric tree the Gromov product is precisely the overlap of the two geodesics, that is distance realising paths, connecting x_0 to x and to y. We say that X is k-hyperbolic, in the sense of Gromov, relative to x_0 if for any three points $x, y, z \in X$ we have the following inequality:

$$(x.y)_{x_0} \geq min\{(x.z)_{x_0}, (y.z)_{x_0}\} - k.$$

If X is k-hyperbolic for one basepoint, then X is $2k$-hyperbolic for any basepoint and if $k_1 \leq k_2$ and X is k_1-hyperbolic then X is also k_2-hyperbolic. For such reasons, we may say that a metric space is *hyperbolic* if it is k-hyperbolic, for some non-negative k and relative to some basepoint. The property of being hyperbolic is thus basepoint invariant. By way of example, we note that any tree, or any $\mathbb{R}$-tree, is 0-hyperbolic and that each hyperbolic space $\mathbb{H}^n$ is also Gromov hyperbolic. As free groups are quasi-isometric to trees, so any finite rank free group is hyperbolic. Fundamental groups of compact hyperbolic manifolds are quasi-isometric to the hyperbolic space of the right dimension, and so are Gromov hyperbolic.

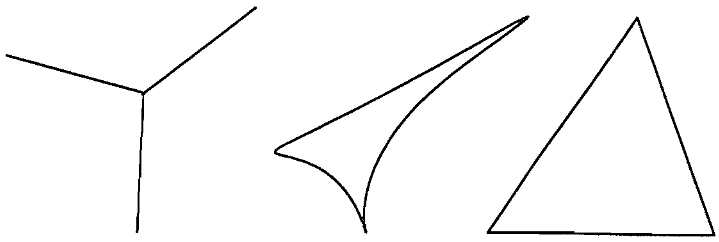

Figure 1.2: Three geodesic triangles, the leftmost is 0-thin.

Farey graphs, being quasi-isometric to a tree, are hyperbolic. It follows that the curve complex associated to the four times punctured sphere and the once-punctured torus are hyperbolic. It is a deep and startling theorem of Masur-Minsky [MaMi1] that *every* curve complex is hyperbolic in the sense of Gromov.

Theorem *(Masur-Minsky) Every curve complex is unbounded and hyperbolic.*

(See Bowditch [Bow1] for a more succinct proof of the hyperbolicity of the curve complex, where the author also computes hyperbolicity constants approximately logarithmic in surface complexity.)

Among those spaces in which any two points are connected by a geodesic there are various characterisations of hyperbolicity, and these are most relevant in studying the curve complex. We shall say that a geodesic triangle is k-thin if there is at least one point at distance at most k from each side (see Figure 1.2). It can be shown that a geodesic metric space is k-hyperbolic if and only if every geodesic triangle is k-thin. Equivalently, a geodesic metric space is k-hyperbolic if and only if given any geodesic triangle, each side is contained in the $2k$-neighbourhood of the union of the other two. When X is hyperbolic, non-parallel geodesics diverge exponentially. This is characteristic of hyperbolic spaces. For other important characterisations, making use of linear isoperimetric inequalities, see [Bow4].

There are more exotic characterisations of hyperbolicity. A metric space X is hyperbolic if and only if every asymptotic cone of X is an $\mathbb{R}$-tree, that is X is an $\mathbb{R}$-tree when viewed from very far away. There are characterisations for the hyperbolicity of word-metrics on groups in terms of the bounded cohomology of the group (see Mineyev [Mine] and references therein).

Any hyperbolic metric space admits a canonical boundary, and the curve complex is no exception. This boundary is constructed as follows. A quasi-geodesic in X is a quasi-isometric embedding from a subinterval of $\mathbb{R}$ into X. For instance we note that any geodesic in X is also a quasi-geodesic and point out that images of quasi-geodesics need not be connected sets (we are allowed to make bounded jumps). We say that a quasi-geodesic is a quasi-ray if the interval is semi-infinite and bi-infinite if the interval is $\mathbb{R}$. When X is hyperbolic any two quasi-geodesics with common ends remain a uniformly bounded distance apart, this bound depending only on the choice of hyperbolicity constant and the two pairs of parameters. Now choose a basepoint x_0 for X and let $\mathcal{R}$ denote the set of all quasi-rays issuing from x_0. We declare two quasi-rays f and g equivalent if $(f(i).g(j))_{x_0} \longrightarrow \infty$ as $i, j \longrightarrow \infty$. Loosely speaking, two quasi-rays issuing from this basepoint are equivalent if they remain close for all time. Since X is hyperbolic, it is not hard to see that changing the basepoint sets up a natural correspondence between the two sets of quasi-rays. We have an equivalence relation on $\mathcal{R}$, and we denote the quotient by this relation by ∂X. The correspondence due to a change of basepoint respects this relation, and so we can make reference to ∂X as *the* boundary of X. We refer to the points of ∂X as *ideal* points.

The Gromov product, with reference to the same basepoint, extends naturally to ∂X and we can use this to construct a metric on the boundary in which any two points are close if their representatives remain close to one another over large distances. Changing the basepoint sets up a uniformly continuous boundary homeomorphism, so at the very least ∂X has a well-defined topology to which quasi-isometries extend as boundary homeomorphisms. By way of example, the boundary of $\mathbb{H}^{n+1}$ is the n-sphere and the boundary of a tree is totally disconnected. We may always connect two ideal points by a quasi-geodesic, that goes via the basepoint if need be, and when X is proper this quasi-geodesic can be taken to be a geodesic and quasi-rays are equivalent to geodesic rays. This visual connectivity by geodesics can even be found among some hyperbolic metric spaces that fail to be proper; perhaps a striking example of this is the curve graph, which is locally infinite (see Bowditch [Bow2]).

The boundary of most interest to us is of course the boundary of the curve complex: That this has an explicit description is a remarkable result of Klarreich [Kla]. Before we give her theorem we must briefly explain some of the terminology and ideas of Thurston, and we shall be making use of these again. As we will see in §1.5, the ideal points of the curve complex correspond to the non-conformal end invariants of Thurston's ending lamination conjecture.

Fix a complete hyperbolic metric on Σ of finite volume. A *geodesic lamination* in Σ is a compact set foliated by simple geodesics. The most straightforward examples are provided by collections of distinct and disjoint curves when we realise each curve by its unique simple closed geodesic. The idea of a lamination was first considered by Thurston and gives a well-defined limit for a sequence of curves.

We may topologise the set $\mathcal{GL}(\Sigma)$ of all geodesic laminations in Σ by regarding $\mathcal{GL}(\Sigma)$ as a subspace of the Hausdorff topology on compact subsets of Σ. This can be metrised by defining $d_{\mathcal{H}}(A, B) = \inf\{r : A \subseteq N_r(B), B \subseteq N_r(A)\}$ for all compact subsets A and B of Σ, and where $N_r(A)$ denotes the set of all points in Σ at most distance r from A. Convergence in this topology can be seen in the surface Σ, and we give an example to illustrate this: Take two simple but intersecting closed geodesics α, β and take whole twists of α around

β. The resulting sequence converges in the Hausdorff topology to a geodesic lamination that contains β and spirals into β from both sides. We shall say that a lamination is *minimal* if it does not properly contain a non-empty lamination and we shall say that it is *filling*, or *fills* Σ, if it has non-empty intersection with every simple closed geodesic in Σ.

Any ideal point of the curve complex is by definition represented by a sequence of curves, and it turns out that this sequence also converges on a well-defined geodesic lamination that fills the surface. For more or less the same reason, the limiting lamination is also minimal, and we see that ideal points correspond naturally to minimal filling laminations. This is essentially Klarreich's point, and it remains to topologise the set of minimal filling laminations, which we denote as $\mathcal{MF}(\Sigma)$, in the right way.

Every geodesic lamination admits a transverse measure whose support, where it assigns non-zero measure, is the union of all its minimal sublaminations. In some sense, transverse measures detect where and the extent to which a lamination runs self-parallel. The simplest examples of these are provided by assigning non-negative weights to each component of a multicurve. The set of measured laminations, which we denote by $\mathcal{ML}(\Sigma)$, is to carry the weak*-topology so that convergence is on measures. If we allow the positive real numbers to act on $\mathcal{ML}(\Sigma) - \{0\}$ by scaling measures, then we have the orbit space $\mathcal{PML}(\Sigma)$ called the *projective measured lamination* space. Toplogically, $\mathcal{ML}(\Sigma)$ is a ball of dimension $6genus(\Sigma) - 6$ and $\mathcal{PML}(\Sigma)$ a sphere of dimension $6genus(\Sigma) - 7$.

In passing from a projective measured minimal and filling lamination to its support, we have the so-called "measure-forgetting" topology on $\mathcal{MF}(\Sigma)$. Klarreich's theorem may now be stated as follows.

Theorem *(Klarreich) The natural map from the Gromov boundary of the curve complex to the space of minimal filling laminations $\mathcal{MF}(\Sigma)$ given the measure-forgetting topology is a homeomorphism.*

(For a second proof, making use of a particular train track complex, see Hamenstädt [Ham1].) Besides this topological description, to the best of my knowledge not a great deal is known about this boundary. It is known to be uncountable, to carry a complete metric and to be second countable but nothing appears to be known about its remaining topological properties. In particular, it is not known whether the boundary is connected.

1.3 Mapping class groups

The term mapping class group has come to refer to many slightly different flavours of group associated to a surface. Here, the *extended mapping class group* $Map^*(\Sigma)$ of a connected orientable surface Σ is the group of all self-homeomorphisms of Σ up to homotopy, with group multiplication defined by composing two representative homeomorphisms before taking homotopy classes. The elements of $Map^*(\Sigma)$ shall be referred to as *mapping classes*. The subgroup $Map(\Sigma)$ of index two comprising all orientable mapping classes is what we shall refer to as the *mapping class group*. It is an observation of Nielsen that the mapping class group is the outer automorphism group of $\pi_1(\Sigma)$. The *pure mapping class group* $PMap(\Sigma)$ is the subgroup of $Map^*(\Sigma)$ whose elements fix each boundary component setwise or puncture of Σ.

While these families of groups play a central role in the study of surface topology, they continue to be of intense interest in their own right for they straddle many aspects of group theory and geometric group theory by simultaneously exhibiting many traits and differences found in other classes of group. For example, mapping class groups and arithmetic groups all have finite virtual cohomological dimension, verify a Tits alternative and are residually finite and yet mapping class groups are not arithmetic (they have finite index in their abstract commensurator groups, see [Iva2]). Furthermore, mapping class groups are hyperbolic in the sense of Gromov only in a few low complexity cases (where they do not have high rank free abelian subgroups) and they are *weakly* hyperbolic (in the sense of Farb) relative to a finite family of curve stabilisers and yet they are not *strongly* hyperbolic (in the sense of Bowditch or, equivalenty, Osin) relative to any finite family of proper subgroups (see Anderson-Aramayona-Shackleton [AAS1] for one proof). For a proof making use of convergence group actions, see [Bow5].

The study of mapping class groups related to surfaces dates back to the 1920s and was started by Dehn and Nielsen, who took very different perspectives. In particular, Dehn first made use of the natural action of the mapping class group on the set of all curves associated to a surface, which he refers to as the arithmetic field. This action is defined by representing a curve by a loop and a mapping class by a homeomorphism before passing to the free homotopy class of the image loop, and extends naturally to what is now the standard simplicial action on the curve complex. Their interaction is fundamental in the study of the mapping class group.

These ideas were then taken on by Ivanov, who further studied the mapping class group and its interaction with the curve complex. In particular, Ivanov [Iva1] succeeded in showing that the automorphism group of the curve complex of a surface of genus at least two is a quotient of the extended mapping class group. This is the analogue of Royden's celebrated theorem [Roy] for the Teichmüller metric, later generalised by Earle-Kra [EKra]. The remaining genus zero and genus one cases, with the exception of the two-holed torus, were settled by Korkmaz [Kor1] in his thesis. We summarise their combined result.

Theorem *(Ivanov-Korkmaz) The automorphism group of the curve complex of a finite type surface that is not the two-holed torus is isomorphic to a quotient of the mapping class group.*

A new proof, an induction on the dimension of the curve complex making use of a modular structure on Farey graphs, was later put forward by Luo [Luo1]. In the same paper, Luo first identifies an automorphism of the curve complex associated to the two-holed torus not induced by a mapping class. Making use of Ivanov's arguments, Irmak [Irm1] generalises Ivanov's theorem to show that the so-called superinjective maps, that by definition are simplicial and preserve the non-disjointness of curves, of the curve complex associated to a surface of genus two with at least two punctures or of genus at least three are induced by a mapping class. The purpose of this work is to establish a virtual co-Hopfian result for mapping class groups.

In Chapter 3, we bring this process of generalising Ivanov's work [Iva1] to what is, perhaps, a natural conclusion by showing that in fact every embedding, and moreover every local embedding, on any two curve complexes whose dimensions do not increase from domain to codomain is induced by a mapping class. From this we deduce a strong local co-Hopfian result for mapping class groups, and that no mapping class group can admit a faithful action on another curve complex of the same dimension. Recall that a *star* in the curve graph is the union of all edges incident on a common vertex.

Theorem *(Shackleton) Suppose that Σ_1 and Σ_2 are two orientable surfaces of finite type such that the complexity of Σ_1 is at least that of Σ_2 and that whenever they have equal complexity at most three they are homeomorphic or one is the three-holed torus. Then, every simplicial map from the curve complex $\mathcal{C}(\Sigma_1)$ to the curve complex $\mathcal{C}(\Sigma_2)$ injective on each star (that preserves the separating type of a curve when both surfaces are the two-holed torus) is induced by a mapping class.*

Meanwhile, the work of Nielsen largely concerns individual mapping classes and was later developed by Thurston in his classification of surface diffeomorphisms (see [Thur3]); the elements of the mapping class group either are periodic, fix a multicurve or are *pseudo-Anosov*. By definition, periodic mapping classes are the torsion elements and have finite order, reducible elements fix a multicurve (that is, fix each vertex of a simplex in $\mathcal{C}(\Sigma)$) and pseudo-Anosovs stabilise a pair of projective measured laminations, one attracting and the other repelling under forward iteration. In the language of §1.2, these laminations represent ideal points of the curve complex for they are both minimal and together fill the surface (see [CasBl] for more details). According to [FLP], any mapping class h is pseudo-Anosov if and only if for every curve α we have $\iota(\alpha, h^n\alpha) \longrightarrow \infty$ as $n \longrightarrow \infty$; in particular, they are characterised in leaving no curve invariant. Centralisers of pseudo-Anosovs are virtually cyclic. While pseudo-Anosovs enjoy these special properties, they are in fact typical elements of the mapping class group and can be generated by a composition of two suit-

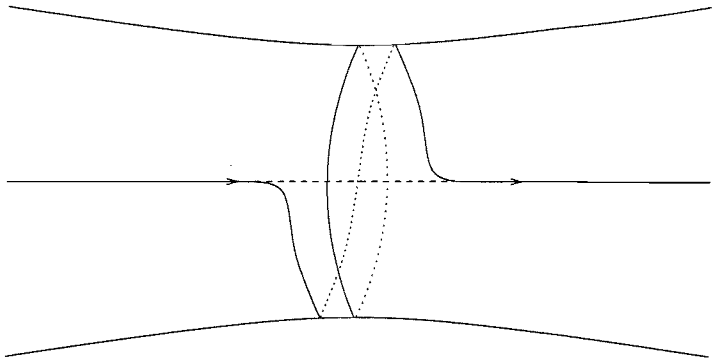

Figure 1.3: A right Dehn twist of one curve around a second.

ably high powers of (Dehn) twists about curves that together fill the surface (see [Fathi] for one proof, and [Long] or [Pen] for proofs when two or more curves are involved).

Problems regarding the generation and presentation of the mapping class groups were first approached by Dehn [Dehn2]. A key discovery in his work is the family of elements that have since become known as *Dehn twists* and are, informally, constructed by cutting the surface along a curve and gluing the surface back together with an integer's worth of twisting. (See Figure 1.3.) These give a very intuitive feel for many of the elements of mapping class groups. It was known to Dehn that the pure mapping class groups are generated by finite sets of Dehn twists. This idea was later developed by Lickorish [Lic1] (see [Lic2] for his corrigendum). In the influential work of Hatcher-Thurston [HatThur], the authors establish the following.

Theorem *(Hatcher-Thurston) The mapping class group of a finite type surface is finitely-presented.*

As they say, the presentations found using their techniques are rather complicated and in need of some simplification. Building on their work, Wajnryb [Waj] identifies a simpler finite presentation for the mapping class group of a closed surface or a one-holed surface comprising $2genus(\Sigma) + 1$ Dehn twists. This culminates in the work of Gervais [Ger], who finds a clean finite presentation for the mapping class group of an arbitrary surface, and Luo [Luo2], who finds a second simple but infinite presentation exploiting relations that were known to Dehn.

The growth of any non-trivial class, in any finite generating set, is known to be linear (see [MaMi2] and [FLM]). The mapping class group is both automatic (see Mosher [Mos]) and biautomatic (see [Ham2]), and the conjugacy problem for pseudo-Anosovs can be solved in linear time [MaMi2]. The mapping class groups of punctured spheres and of the closed surface of genus two are known to be linear (see Bigelow-Budney [BigeBud] and Korkmaz [Kor2]) for they have strong connections with braid groups (that these are linear is a remarkable result of Bigelow [Bige] and then Krammer [Kra]). The following important question remains open.

Question *Does the mapping class group of a finite type surface $\Sigma - \Pi$ always admit a faithful and finite dimensional linear representation?*

This is a challenging problem, and a complete solution seems some way off. In a recent note of Storm [Sto], resting on the work of Hamenstädt [Ham2] and Kato [Kat], it is shown that the mapping class group verifies the Novikov conjecture. Many other properties of linear groups may be attributed to the mapping class groups, for example they have uniformly exponential growth (see Anderson-Aramayona-Shackleton [AAS2]). There are still many more properties to enquire of.

There are many important subgroups of the mapping class group. The *Johnson kernel* $\mathcal{K}(\Sigma)$ associated to the surface Σ is the subgroup of the mapping class group $Map(\Sigma)$ generated by all Dehn twists around separating curves. For this reason it is intimately related to the flag subcomplex of the curve complex whose vertices are all separating curves, namely the *separating curve complex*. In any case, this subgroup was first studied by D. Johnson who shows in [J1] that, when Σ is a closed surface of genus at least three, $\mathcal{K}(\Sigma)$ has infinite index in the *Torelli group* $\mathcal{T}(\Sigma)$, the kernel of the natural action of the extended mapping class group on the first homology group $H_1(\Sigma, \mathbb{Z})$, by considering the short exact sequence

$$1 \longrightarrow \mathcal{K}(\Sigma) \longrightarrow \mathcal{T}(\Sigma) \longrightarrow \Lambda^3 H_1(\Sigma,\mathbb{Z})/\Lambda H_1(\Sigma,\mathbb{Z}) \longrightarrow 1.$$

The quotient homomorphism has since become known as the *Johnson homomorphism*. It is a result of Johnson's [J2] that the Torelli group associated to a closed surface of genus at least three is generated by a finite subset of Dehn twists around separating curves and around bounding pairs. However, it is subsequently shown by Biss-Farb [BisF] that no Johnson kernel associated to a closed surface of genus at least three is finitely generated. This builds on the work of McCullough-Miller [McCM] who prove this for the closed surface of genus two. The mapping class group, and its various subgroups, have many connections with the topology of 3-manifolds. Morita [Mor] establishes connections between the Torelli group of a closed surface of genus at least three and Casson's invariant of homology 3-spheres. We may reasonably ask whether the methods developed in Chapter 3 apply to the separating curve complex, and hence the Johnson kernel and the Torelli group.

On the face of it, we seem to be extremely well-placed: We are studying an interesting group with a slew of interesting subgroups, all of which just so happen to act very naturally on a hyperbolic metric space, and there is an established theory for just such a situation, as we described in outline in §1.2. However, this theory applies best to proper metric spaces, such as hyperbolic groups in their word-metrics, and the curve complex is far from being proper. We can now see this for, if we take two intersecting curves α and β in Σ both disjoint from a third curve γ, then by Dehn twisting α around β we generate an infinite family of distinct curves each disjoint from γ. In particular, these all belong to the link of γ.

Nonetheless, much can be said about the nature of the action of the mapping class group on the curve complex. Bestvina-Fujiwara [BesFu] introduced the notion of weak proper discontinuity, hence forth referred to as WPD, for a group action by isometries and showed that the action of the extended mapping class group on the curve complex satisfies this. Roughly speaking, we say that the action of a group on a metric space satisfies WPD if the number of mapping classes moving both any point and its translate by a large power of any loxodromic group element a small distance is finite. They do so by the use of a geometric limiting argument, going from an infinite sequence of curves to geodesic laminations. From this they deduce, among other things, that the dimension of the real second bounded cohomology of the mapping class group, and any subgroup that is not virtually abelian, is infinite.

Inspired by their line of argument Bowditch [Bow2] generalises their result. Bowditch introduces the notion of acylindricity for an isometric action, and shows that the action of the mapping class group on the curve complex also satisfies this. Acylindricity has its roots in the study of 3-manifolds and Sela's work on splittings and group actions on trees. Roughly speaking once more, we say that an isometric action by a group is *acylindrical* if the number of group elements moving a long geodesic a short distance is uniformly finite. While the work of Minsky et al towards the ending lamination conjecture uses the curve complex to understand hyperbolic 3-manifolds, the argument of Bowditch makes use of hyperbolic 3-manifolds to study the action of the mapping class group on the curve complex.

Both the approach to acylindricity and the approach to WPD make use of geometric limit arguments. An important feature to bear in mind is that limiting arguments are, by nature, non-constructive and the information they yield non-computable. For this reason, the WPD theorem of Bestvina-Fujiwara is re-established by computational methods and a computable version of Bowditch's acylindricity theorem (see Chapter 2, Theorem 23) is established in this thesis. A weaker form of acylindricity satisfied by the action of the mapping class group is also shown by Hamenstädt [Ham3] using a certain train track complex. Our computable version of acylindricity is stronger than this and stronger than WPD, though slightly weaker than that of Bowditch for the bounds we provide are non-uniform. However, these bounds are computable and explicit for they can be expressed in terms of intersection numbers and the topology of $\Sigma - \Pi$. This is our reward.

1.4 Teichmüller space

Given a topological manifold Σ, it is natural to ask what constant negative curvature riemannian metrics, if any, can be endowed on Σ. It is a consequence of Mostow rigidity that if Σ is a topological n-manifold, with $n \geq 3$, that admits a hyperbolic metric of finite volume then this is the only such metric up to isometry. When Σ is a surface of negative Euler characteristic, rigidity no longer applies and indeed we have a whole family of such metrics referred to as *Teichmüller space*. This space carries the topology of an open ball of finite dimension, obtained by any one of several parameterisations. We shall, after defining Teichmüller space, construct this topology via the Frenchel-Nielsen coordinates. For more details, we refer to Abikoff [Abi] and Imayoshi-Taniguchi [ImaTan].

For ease of exposition, we shall only deal with a fixed closed, connected and orientable topological surface Σ such that $genus(\Sigma) \geq 2$, though much of what we say either still holds when Σ is allowed to have punctures and boundary components or can be translated to this setting. A *marking* of Σ is a pair (f, S) where S is a riemannian surface and $f : \Sigma \longrightarrow S$ is a homeomorphism. We declare two such markings equivalent if, up to isotopy, they differ by an isometry: The two markings (f_i, S_i) $(i \in \{1, 2\})$ are equivalent if there exists a riemannian isometry $h : S_1 \longrightarrow S_2$ such that $hf_1 \simeq f_2$. It can be verified that this does define an equivalence relation on the set of all markings. We define the *Teichmüller space* $Teich(\Sigma)$ associated to Σ as the set of equivalence classes of markings. Each marking class on Σ contains a marking (f, S) such that the pull-back of S via f is a hyperbolic metric on Σ. It is therefore appropriate to view each marking class as a hyperbolic metric on Σ and the Teichmüller space as the set of all marked hyperbolic metrics.

The group of all self-homeomorphisms has a natural action on the set of all markings by precomposition and, on taking homotopy classes, this descends to an action of the extended mapping class group on Teichmüller space. The moduli space of Σ is the orbit space of this action, and it is a result of Mumford that the "thick" part of moduli space, as will be defined soon, is compact.

A *pants decomposition* P of Σ is a maximal collection of distinct and disjoint curves in Σ, and corresponds to a top simplex in the curve complex. Each pants decomposition corresponds to a top multicurve and to a top dimensional simplex in the curve complex. By a pair of pants, we shall mean topologically a sphere with three open discs removed. The complement of a pants decomposition of Σ is a disjoint collection of interiors of pairs of pants. We shall use a pants decomposition P to parameterise Teichmüller space. To each marked hyperbolic surface (f, S) we can associate the $2N$-tuple $(l_S(f(\alpha_i)), \theta_S(f(\alpha_i))) \in (0, \infty)^N \times \mathbb{R}^N$, where $l_S(f(\alpha_i)$ denotes the S-length of the closed geodesic homotopic to $f(\alpha_i)$ and $\theta_S(f(\alpha_i)$ the twisting parameters. Each tuple encodes the assembling data of S via hyperbolic pairs of pants, with their boundary lengths specified by the $l_S(f(\alpha_i))$ and the twisting round boundary components by the $\theta_S(f(\alpha_i))$.

There is a natural correspondence between the elements of Teichmüller space and such tuples, and we give $Teich(\Sigma)$ the topology in which this correspondence defines a homeomorphism. Although this parameterisation is dependent on the choice of pants decomposition, the topologies on $Teich(\Sigma)$ resulting from different choices of P all differ by a homeomorphism. Since $|P| = 3genus(\Sigma) - 3$, the Teichmüller space associated to Σ is topologically an open ball of dimension $6genus(\Sigma) - 6$.

This topology on $Teich(\Sigma)$ is meterizable in at least two important and somewhat contrasting ways, namely by the Teichmüller metric (equal to Kobayashi's metric) and by the Weil-Petersson metric. We shall say something on both, but it is the Teichmüller metric that is most closely related to the curve complex and therefore is most relevant here. Given two marked hyperbolic surfaces $\sigma_i = [f_i, S_i]$ ($i \in \{1, 2\}$) we define their Teichmüller distance $d_T(\sigma_1, \sigma_2)$ to be equal to

$$\inf\{(1/2)\log K(f_1 f_2^{-1}) : g_i \in [f_i, S_i], i = 1, 2\},$$

where $K(f)$ denotes a certain "dilation" of a homeomorphism f between two riemannian surfaces. This defines a metric on $Teich(\Sigma)$ compatible with the topology we just described. (The triangle inequality follows from a certain multiplicative property of the dilation.) Teichmüller's metric is non-riemannian but it is complete, uniquely geodesic and geodesically complete. Moreover, geodesics in d_T are determined by *quadratic differentials*, or *singular euclidean structures*, via the *Teichmüller map* and quadratic differentials in turn determine, and are determined by, metrics on Σ flat away from a finite number of singularities (where all the negative curvature is to be thought of as being concentrated). See Strebel [Str] for a thorough account of quadratic differentials. It is a celebrated result of Royden's [Roy] that the isometry group of the Teichmüller metric is isomorphic to the extended mapping class group.

We remark that the Teichmüller metric for the torus is isometric to the hyperbolic plane. It is largely for this reason that it was briefly suspected that every Teichmüller metric is hyperbolic in the sense of Gromov. However, the failure of this is largely due to the presence of *Margulis thin regions* in $Teich(\Sigma)$ which coarsely resemble product spaces with a sup-product metric [Mi1]. Thus, $Teich(\Sigma)$ exhibits properties of positive curvature. This is, in some sense, the only obstruction to hyperbolicity.

There are other good reasons why d_T is not hyperbolic. For instance, d_T admits more than one natural boundary. There are two compactifications of Teichmüller space, due to Bers and due to Thurston, but they are non-homeomorphic for Kerckhoff [Ker] (see also [KerT]) explains that the action of the mapping class group on $Teich(\Sigma)$ fails to extend continuously to Bers's boundary (indeed, this boundary is base point dependent and the mapping class group moves the base point) whereas the action does extend continuously to Thurston's boundary. A short proof of the non-hyperbolicity of d_T is provided by Bowditch [Bow5], by considering convergence group actions of the mapping class group. An alternative argument that exploits the exponential divergence of geodesics in a hyperbolic space, which fails in d_T, is given by Ivanov [Iva3].

Although not hyperbolic, the Teichmüller metric is coarsely related to the curve graph. This construction requires a version of the *collar lemma* (details and proof can be found in [Abi]), from which we find a positive constant ϵ such that, for any surface Σ of genus at least two, any hyperbolic metric $\sigma \in Teich(\Sigma)$ and any two intersecting curves α and β in Σ, if $l_\sigma(\alpha)$ is at most ϵ then $l_\sigma(\beta)$ is greater than ϵ. That is to say, any curve intersecting a "short" curve must be "long". It is standard to refer to such a choice of ϵ as a *Margulis constant.* For any curve α, let $T_\alpha \subseteq Teich(\Sigma)$ be the set of all those marked hyperbolic metrics in which the length of α is strictly less than ϵ. Each T_α is an open and non-empty subset of $Teich(\Sigma)$ and the complement in $Teich(\Sigma)$ of their union is often referred to as the *thick part* of Teichmüller space. We see that $T_\alpha \cap T_\beta$ is empty if and only if α and β can be realised disjointly. We construct a nerve of $\{T_\alpha\}_\alpha$ by introducing to each T_α a new point and connecting this to each point of T_α by an edge of length one. The resulting space $Teich_{el}(\Sigma)$ is a *length space* and hence carries a natural path-metric, in much the same way as we saw for the curve complex. We call this metric the *electric (Teichmüller) metric*, and record that the usual path-metric on the curve complex and the electric metric are quasi-isometric. Thus, the electric metric is hyperbolic in the sense of Gromov. In particular, it is the thin regions that prevent d_T from being hyperbolic.

While the Teichmüller metric and the curve complex fail to be quasi-isometric, geodesics in the first determine quasi-geodesics in the second by tracking the thin regions they pass through. These are important observations for both the approach of Masur-Minsky and the approach of Bowditch to the hyperbolicity of the curve complex.

For each point $\sigma \in Teich(\Sigma)$, the space of *holomorphic quadratic differentials* $\mathcal{Q}(\sigma)$ may be identified with the cotangent space $T^*_\sigma(\Sigma)$. Now $\mathcal{Q}(\sigma)$ carries a natural L^2-inner product, defined by the weighted integral

$$< \phi, \eta >= \int_\Sigma \phi\eta/\sigma^2.$$

The riemannian part of the dual of $<,>$ is referred to as the *Weil-Petersson* metric, denoted d_{WP}. It is a result of Wolpert's [Wolp] that the Weil-Petersson metric is not complete since curves collapse to cusps in finite time, and that the completion of d_{WP} may be realised by adjoining *noded surfaces* to $Teich(\Sigma)$. Further, d_{WP} has strict negative-sectional curvatures neither bounded away from 0 nor $-\infty$ (so that d_{WP} is CAT(0)), has isometry group isomorphic to the extended mapping class group of $\Sigma - \Pi$ (see Masur-Wolf [MaWolf]), is geodesically convex and the length function associated to any curve is strictly convex.

There is a combinatorial description of the Weil-Petersson metric in terms of the *pants complex* due to Hatcher-Thurston, whose vertices are the pants decomposition of Σ we have just discussed. Whereas the Teichmüller metric and the curve complex fail to be quasi-isometric, the Weil-Petersson and the path-metric on the pants complex are quasi-isometric. This is due to Brock-Farb [BroF]. The construction of the quasi-isometry makes use of the Bers constant (for every hyperbolic metric on Σ there is a pants decomposition of Σ in which the length of each curve is at most this constant). In [Bro], Brock describes a coarse Lipschitz relationship between the volume of the convex core of a quasi-Fuchsian group and the Weil-Petersson distance between the two conformal end invariants (they are roughly proportional). In [BroF], it is shown that the Weil-Petersson metric associated to Σ is hyperbolic if and only if Σ is a four-times punctured sphere, a five-times puncture sphere or a twice punctured torus. The case of the five-times punctured sphere (and the two-holed torus) were shown independently by Aramayona [Ara] and by Behrstock [Behr] via very different means.

1.5 The theory of ends

The study of 3-manifolds has seen remarkable progress in the last few years alone, invoking many ideas developed over the past century. This subject underwent a revolution during the 1970s with the influential work of Thurston, who put forward three significant conjectures. Of most relevance to us are the tameness conjecture and the ending lamination conjecture, and we shall give careful statements of these. The third, the geometrisation conjecture, roughly states that any 3-manifold can be given a geometric structure by first giving a geometric structure on each piece of its JSJ-decomposition. Many of the ideas involved inspired the work of Dunwoody and Stallings, among others, on group splittings. In 2003, Perelman [Per1], [Per2] announced a proof of the geometrization conjecture making use of Ricci curvature. His result implies the well-known Poincaré conjecture.

Tameness conjecture *Let M be a complete hyperbolic 3-manifold with finitely generated fundamental group. Then M is the interior of a compact 3-manifold.*

A manifold which is the interior of a compact manifold is said to be *topologically tame*, or just *tame*. The first example of a 3-manifold which is not tame was found by Whitehead, and McMillen proved that there are uncountably many non-tame 3-manifolds. We recall that a hyperbolic metric is a riemannian metric with constant sectional curvature -1. The universal cover of a complete hyperbolic 3-manifold is $\mathbb{H}^3$ with boundary S^2_∞ the Riemann sphere. (See Kapovich [Kap] for an excellent survey.) We actually gave Marden's formulation of the tameness conjecture, posed after he verified tameness for "geometrically tame" complete hyperbolic 3-manifolds (see [Mard]). As we shall see, the question of tameness reduces to looking at the peripheral pieces known as ends. The theory of geometrically finite ends was largely developed by, among others, Ahlfors, Bers, Kra, Marden, Maskit, Sullivan and Thurston between 1960 and 1980. Later, the non-geometrically finite ends were studied by Bonahon, Brock, Canary, Minsky and Souto, among others.

Particular cases of the tameness conjecture were treated by Ohshika [Oh], Canary-Minsky [CanMi] and Souto [Sou] before independent and simultaneous proofs of the tameness conjecture were announced by Agol [Ag] and Calegari-Gabai [CalGa] in 2004. Later, Soma [Som] found new arguments simplifying those of Calegari-Gabai. This verifies, among other things, the Ahlfors measure conjecture, which correctly asserts that the limit set of a finitely generated Kleinian group has full or zero measure.

Ending lamination conjecture *Every complete and orientable hyperbolic* 3*-manifold M with finitely generated fundamental group is determined by its topological type, by the conformal boundaries of its geometrically finite ends and by the ending laminations of its geometrically infinite ends.*

This is posed assuming the tameness conjecture to give us the lamination invariant, and represents a classification of hyperbolic 3-manifolds. It was first explicitly stated by Thurston in [Thur2]. (For a survey of the standard Thurston machinery, see [CanEG] and [Thur1].) The two end invariants naturally correspond to points in Thurston's compactified Teichmüller spaces: The conformal boundaries belong in the interior and the ending laminations correspond to boundary points, modulo measures. The classification for finite volume hyperbolic 3-manifolds is the Mostow rigidity theorem we recalled in §4: If two 3-manifolds with isomorphic fundamental groups both admit a hyperbolic metric of finite volume, then they are isometric. Very recently, a complete proof of the ending lamination conjecture was announced by Brock-Canary-Minksy (see [Mi3] and [BroCM] for the indecomposable case, the general case is to follow), building on a significant volume of work due in large part to Minsky. A second and subsequent proof of the indecomposable case is given by Bowditch [Bow6].

In 1910, Dehn put forward a key lemma on 3-manifolds that was subsequently named after him. As Kneser points out though, the argument Dehn gives has serious gaps first addressed by Kneser himself. It was not until nearly fifty years after Dehn's work was published that the first correct proof, due to Papakyriakopoulos, was given. Subsequent simplifications and improvements on Dehn's original claim were given by Shapiro-Whitehead and then Stallings. The following statement of Dehn's lemma was given by Stallings.

Lemma *Let M be a* 3*-manifold and B a component of ∂M and N a normal subgroup of $\pi_1(B)$. Suppose that there are elements of $ker(\pi_1(B) \longrightarrow \pi_1(M))$ not contained in N. Then, there is a simple loop in B bounding an embedded disc in M and not contained in N.*

Each argument is constructive, making essential use of Papakyriapoulos's tower of covers with base M, and in parallel proves the analogue of Dehn's lemma for spheres, subsequently known as the sphere theorem. Both have important applications in knot theory, answering some conjectures of Hopf, but Dehn's lemma features prominently in Scott's proof (see [Sco]) of his "core" theorem. This is a key result in formulating both the tameness and the ending lamination conjectures, and is the starting point for Bonahon's work. Scott proved that every finitely generated group is the fundamental group of a compact 3-manifold. His core theorem, stated as follows, implies this compact manifold can be chosen to be a submanifold when considering fundamental groups.

Theorem *Let M be a 3-manifold with $\pi_1(M)$-finitely generated. Then there is a compact submanifold V of M such that the inclusion of V in M induces an isomorphism on fundamental groups.*

The argument is delicate and makes use of Dehn's lemma to find discs at which to surger a candidate submanifold, reducing a complexity each time. It is due to McCullough-Miller-Swarup [McCMS] that this core is unique up to homeomorphism when M is orientable, a result generalised to non-orientable 3-manifolds by Harris-Scott [HarrSco] (they show that there are only finitely many up to homeomorphism). There is a relative version of Scott's core theorem due to McCullough [McC], producing a compact core meeting each cusp of M in an annulus or a torus.

It is due to Bonahon [Bon] that the ends of M, as defined by Freudenthal, are in one-to-one correspondence with the complementary components of the core, and so we define an end of M in this way. It follows from Tucker [Tu] that a 3-manifold is tame if and only if each of its ends is tame. The tameness conjecture therefore boils down to the study of the complementary components of the core in M.

Of further importance in understanding the structure of a complete hyperbolic 3-manifold M of finitely generated fundamental group is a second core, to which M retracts, and is defined as follows. We may represent M as the quotient of $\mathbb{H}^3/\Gamma$ where Γ is both a discrete and torsion-free subgroup of $PSL(2,\mathbb{C})$. When Γ is not torsion free, the quotient manifold is said to be an *orbifold*. The limit set of Γ, denoted Λ_Γ, is the set of all accumulation points in $\mathbb{H}^3 \cup S^2_\infty$. As Γ is discrete, $\Lambda_\Gamma \subseteq S^2_\infty$. The complement of $\Lambda\Gamma$ in S^2_∞ is denoted by Ω_Γ and is referred to as the *domain of discontinuity of* Γ. Notice that Γ has a natural action on each of S^2_∞, Ω_Γ and the convex hull of Λ_Γ in $\mathbb{H}^3$. The *convex core* of Γ, sometimes referred to as the *Nielsen core*, is the quotient of this convex hull by Γ and is to be denoted $C(M) \subset \mathbb{H}^3/\Gamma$. Whenever $C(M)$ has finite volume, we say that M is *geometrically finite*.

Three mutually exclusive types of end of M are described as follows. We say that an end is a *cusp* if its fundamental group corresponds to a parabolic subgroup of Γ, so that it fixes an ideal point of $\mathbb{H}^3$ and is isomorphic either to $\mathbb{Z}$ or $\mathbb{Z} \oplus \mathbb{Z}$. This means a cuspidal end is homeomorphic to the solid open torus or torus cross $\mathbb{R}$, respectively. Among other things, cusps of M are contained in the (Margulis) thin part of M. While there may well be infinitely many non-cuspidal components of the Margulis thin region (each referred to as Margulis tubes), it is Sullivan's finiteness theorem that there are only finitely many $\mathbb{Z} \oplus \mathbb{Z}$-cusps.

We say that an end is *geometrically finite* if it has an open neighbourhood disjoint from the convex core $C(M)$. These "diverge exponentially" and can be compactified by gluing a conformal surface, found as one of the components of Ω_Γ/Γ. It follows from the Ahlfors finiteness theorem that this orbit space is a finite type orbifold. When M is geometrically finite it is determined, up to isometry, by its fundamental group and by these conformal boundaries at infinity. A non-cuspidal and non-geometrically finite end is sometimes referred to as geometrically infinite.

An end E is said to be *simply degenerate* if it has closure homeomorphic to the product $S \times [0, \infty)$, where S is a component of the boundary of Scott's core in M, and there exists a sequence of pinched negative curvature surfaces that exit down E, that is are contained in E and eventually leave every compact subset in the closure $\overline{E}$, and are homotopic to $S \times \{0\}$ in M. Bonahon [Bon] showed that the simply degenerate ends are exactly those down which a sequence of simple closed geodesics exit (a geometrically finite end therefore can not be simply degenerate, as such a sequence stays in the convex core). We can project each of these simple closed geodesics to curves in S and these converge, in the Hausdorff topology on S, to a geodesic lamination. Thurston showed that this lamination is an invariant of the simply degenerate end, that is it does not depend on the chosen sequence of simple closed geodesics disappearing down E, and there is always an approximating sequence of simple closed geodesics of uniformly bounded length, depending only on S. Furthermore, Thurston proved that these laminations both are minimal and fill the corresponding surface S. These laminations, and the conformal boundaries of geometrically finite ends, are the end invariants appearing in the statement of the ending lamination conjecture.

Since parabolics in Γ often give only technical complications, for the remainder of this section let us assume that Γ is parabolic free. We can always reduce to a close approximation of this by cutting out horoballs around parabolic points to which sequences of simple closed geodesics collapse. We shall say that M is *geometrically tame* if every end of M is either geometrically finite or simply degenerate. Bonahon [Bon] proved that if the Scott core has incompressible boundary in M, then M is geometrically tame. Any geometrically tame manifold is topologically tame, and the converse is due to Canary [Can1], [Can2]. It follows that Marden's topological version of the tameness conjecture and Thurston's geometric formulation are equivalent. To complete a proof of the tameness conjecture, it is now enough to verify that a non-geometrically finite end of M is simply degenerate (or cuspidal).

This is the point of view adopted by Calegari-Gabai. Indeed, they construct the required subsurfaces explicitly by "shrinkwrapping" certain incompressible subsurfaces relative to a finite collection of simple closed geodesics. Combining this with the main result of Souto [Sou], that every M a nested union of compact cores is topologically tame, the tameness conjecture follows. Meanwhile, the approach of Agol is somewhat different: Rather than prove that every geometrically infinite end is simply degenerate (or cuspidal), he shows that a particular limit of tame manifolds is again tame.

In parallel, much of the work towards a proof of the ending lamination conjecture was driven by Minsky. During the early 1990s, Minsky [Mi2] verified this among the bounded geometry 3-manifolds (where simple closed geodesics do not degenerate to points). This is considerably easier to understand than the full ending lamination statement and, while the arguments will not generalise, the ideas it encompasses do prove instructive. The main strategy from here is to construct, combinatorially and piecewise, a model Riemannian manifold M^* for M with a map $f : M^* \longrightarrow M$ that is bi-Lipschitz away from the solid tori in M^* corresponding to Margulis tubes in M. Sullivan's rigidity theorem [Sul] implies, among other things, that a bi-Lipschitz map between hyperbolic 3-manifolds of finitely generated fundamental groups is in fact an isometry. From this the ending lamination conjecture will follow.

The construction of the model for a simply degenerate end makes fundamental use of the curve complex and its geometric properties. An important component in this is the hyperbolicity of the curve complex, due to Masur-Minsky [MaMi2] (see Bowditch [Bow1] for a more succinct proof) as we recall. Given this, the curve complex has a canonical boundary whose points are, as we saw in §4, the invariants of simply degenerate ends.

Before this can be exploited, the local finiteness issues of the curve complex must be overcome. This is a significant obstacle, and takes some heavy machinery introduced by Masur-Minsky to make this work. The two main constructions are the tight geodesic and the hierarchy. The former is defined thoroughly in Chapter 2, but for now we remark that tight geodesics have the highly desirable property that neighbouring curves drag one another around the surface. In any case, these exist and there are only finitely many between any two vertices. We recover many of the properties of a locally finite hyperbolic graph, in particular we can use tight geodesics and a diagonal sequence argument to extract a bi-infinite (tight) geodesic between any two ideal points.

Meanwhile, hierarchies represent a controlled system of tight geodesics and are the main device Minsky and his collaborators use to construct a model for a simply degenerate end. This construction is a little technical, but is simplified by restricting to the case of the five-times punctured sphere where every geodesic in the curve graph is tight (see [Mi4]). For a simply degenerate end E meeting the Scott core in a surface S, there is a sequence of simple closed geodesics contained in and exiting down E that projects to a sequence of curves in S that lie on a tight geodesic ray in the curve graph of S ending on the lamination invariant of E. We can "thicken" this ray to a hierarchy of tight geodesics. The model is then obtained by gluing together standard blocks organised by the hierarchy. It should be noted that the resulting model does not depend on the original choice of geodesic ray, up to bi-Lischitz equivalence.

Since we can compute tight geodesics and hierarchies, it is worth investigating whether we can compute finite pieces of the model manifold.

Chapter 2

Distances in the curve graph

ABSTRACT: We give explicit bounds on the intersection number between any curve on a tight multigeodesic and the two ending curves. We use this to construct all tight multigeodesics and so conclude that distances are computable. The algorithm applies to all surfaces. We recover the finiteness result of Masur-Minsky for tight geodesics. The central argument makes no use of the geometric limit arguments seen in the recent work of Masur-Minsky (2000) and Bowditch (2003). We apply these methods to study the computability of stable lengths of all mapping classes and to give a computable version of both the WPD and acylindricity theorems, and remark they can be used to compute hierarchies. Our methods are entirely combinatorial.

KEYWORDS: Curve complex, multigeodesic, train track.

2.1 Introduction

Let Σ be a closed, connected and orientable surface and let $\Pi \subseteq \Sigma$ be a finite subset. In [Harv], Harvey associates to the pair (Σ, Π) a simplicial complex $\mathcal{C}(\Sigma, \Pi)$ called the curve complex. This is defined as follows. We shall say that an embedded loop in $\Sigma - \Pi$ is trivial if it bounds a disc and peripheral if it bounds a once punctured disc. Let $X = X(\Sigma, \Pi)$ be the set of all free homotopy classes of non-trivial and non-peripheral embedded loops in $\Sigma - \Pi$. The elements of X will be referred to as curves. We take X to be the set of vertices and deem a family of distinct curves $\{\gamma_0, \gamma_1, \ldots, \gamma_k\}$ to span a k-simplex if any two curves can be disjointly realised in $\Sigma - \Pi$. The mapping class group has a cocompact simplicial action on the curve complex and this action has been exploited by various authors, see for example [BesFu], [Hare] and [Iva2].

With the exception of only a few cases, namely Σ is a 2-sphere and $|\Pi| \leq 4$ and Σ is a torus and $|\Pi| \leq 1$, X is non-empty and the curve complex is connected. For these non-exceptional cases, it can be verified that the simplicial dimension C of $\mathcal{C}$ is equal to $3genus(\Sigma) + |\Pi| - 4$. We see that (Σ, Π) is in fact non-exceptional if and only if $C(\Sigma, \Pi) > 0$. In each subsequent section of this chapter, it is to be assumed that (Σ, Π) is such that $C(\Sigma, \Pi) > 0$.

When $C(\Sigma, \Pi) > 0$ the curve complex can be endowed with a path-metric by declaring all edge lengths to be equal to 1. All that is important here is the 1-skeleton $\mathcal{G}$ of $\mathcal{C}$, with which $\mathcal{C}$ is quasi-isometric and where vertex distances are integers. The induced metric on $\mathcal{G}$ is known to be unbounded and hyperbolic in the sense of Gromov [MaMi1], [Bow1]. The boundary of $\mathcal{G}$ is homeomorphic to the space of minimal geodesic laminations filling $\Sigma - \Pi$, given any hyperbolic metric on $\Sigma - \Pi$, endowed with the "measure forgetting" topology [Kla], [Ham1]. The curve complex plays a key role in Minsky et al's approach to the indecomposable case of Thurston's ending lamination conjecture, further studied by Bowditch [Bow6]. A third approach has been proposed by Rees [Re].

All this at first sight suggests that we may apply the methods of hyperbolic groups and spaces to study various groups acting on $\mathcal{G}$, in particular the mapping class group and its subgroups. The curve graph, though, is not locally finite or even fine: As early as the 2-ball around any vertex of $\mathcal{G}$ these problems are manifest.[1] Tight multigeodesics, introduced in [MaMi2] and further studied in [Bow2] and [Bow3], address this problem. Their introduction has been fruitful: Masur-Minsky used these to study the conjugacy problem in the mapping class group and Bowditch used these to describe the action of the mapping class group on the curve complex.

Masur-Minsky [MaMi2] showed that there are only finitely many tight multigeodesics between any two vertices of the curve graph and Bowditch [Bow2] improved on this, showing that there are only uniformly boundedly many curves in any given slice. We go some way to re-establishing these results, though our bounds depend on the intersection number of the two ending vertices. Since the arguments given here do not rely on passing to geometric limits, our work can be viewed as addressing the local finiteness problems as well as offering computability. Furthermore, our construction is local and therefore not sensitive to the topological type of $\Sigma - \Pi$. We see how to construct geodesics, all tight multigeodesics and compute the distance between any two vertices. In fact, their notions of tightness are slightly stronger than we shall need but finding a succinct weakening proves problematic.

We introduce the key idea of chords and pulses to measure the interleaving in the surface of curves lying on a geodesic in $\mathcal{G}$, turning an inherent complexity to our advantage. This measure is preserved by the action of the mapping class group. We construct a canonical train track relative to the ends of the geodesic and carrying the first curve, in so doing pulses determine a measure on this train track that verifies a coarse version of the switch condition. We use this to establish bounds on pulses that apply to all geodesics and, combining these with an appropriate tightness criterion, we establish the finiteness of tight multigeodesics. These methods are readily applicable to related complexes.

[1] We say that a graph is *fine* if any edge is contained in only finitely many circuits of any given length

In §2.8 we use our results to compute stable lengths of all mapping classes, modulo a non-computed uniform power. We then study the usual action of the mapping class group on the curve complex, finding a computable version of Bowditch's acylindricity theorem in §2.9 and a new proof of its weak proper discontinuity in §2.10. It is worth pointing out that we only make essential use of the hyperbolicity of the curve complex in §2.8.

2.2 Tight geodesics and our main results

Let us remind ourselves of a few definitions. Associated to any two curves α and β is their geometric intersection number $\iota(\alpha, \beta)$, namely the minimal cardinality of the set $a \cap b$ among all $a \in \alpha$ and $b \in \beta$. Note $\iota(\alpha, \alpha) = 0$ for all α, $\iota(\alpha, \beta) = \iota(\beta, \alpha)$ for all α and β, and $d(\alpha, \beta) \leq 1$ if and only if $\iota(\alpha, \beta) = 0$. For any two curves α and β, we have $d(\alpha, \beta) \leq \iota(\alpha, \beta) + 1$ (see [Bow1] for a logarithmic bound).

For us, paths in the curve graph shall be sequences of vertices $\gamma_0, \gamma_1, \ldots, \gamma_n$ such that $\gamma_i \neq \gamma_{i+1}$ and $\iota(\gamma_i, \gamma_{i+1}) = 0$, that is γ_i and γ_{i+1} are adjacent, for each i, or on occasion viewed as sets of vertices by a convenient abuse of notation. A geodesic in $\mathcal{G}$ is a distance realising path.

We shall recall the notion of tight multigeodesic due to Bowditch [Bow2], but that of Masur-Minsky [MaMi2] works equally well here. Recall that a multicurve is a collection of pairwise distinct curves of pairwise zero intersection number, corresponding to a simplex in the curve complex. Intersection number on multicurves is defined additively. Recall that a multipath is a sequence of multicurves $(v_i)_0^n$ such that $(\gamma_i)_0^n$ is a path for any $\gamma_i \in v_i$, each i. We shall refer to each v_i as a vertex of the multipath. We say that a multipath $(v_i)_0^n$ is *tight at* v_j $(1 \leq j \leq n-1)$ if for all curves δ, whenever $\iota(\delta, v_j) > 0$ we have $\iota(\delta, v_{j-1}) + \iota(\delta, v_{j+1}) > 0$. In other words, a multipath is tight at v_j if any curve intersecting v_j also intersects at least one of the neighbouring multicurves v_{j-1} and v_{j+1}. We say that $(v_i)_0^n$ is *tight* if tight at each v_j $(1 \leq j \leq n-1)$. A multipath $(v_i)_0^n$ is a *multigeodesic* if $d(\gamma_i, \gamma_j) = j - i$ for each $i < j$ and each $\gamma_i \in v_i$ and $\gamma_j \in v_j$. In which case, we speak of v_i and v_j as being aligned, for each $i < j$. Finally, for any non-negative integer k the multipath $(v_i)_0^n$ is *locally k-geodesic* if v_i and v_j are aligned for $1 \leq j - i \leq k$, so that $d(\gamma_i, \gamma_j) = j - i$ for $1 \leq j - i \leq k$.

We remark that every path in the curve graph associated to the five-times punctured sphere and the twice-punctured torus is tight since the dimension of the corresponding curve complexes leaves no alternative. The existence of tight multigeodesics in general was established by Masur-Minsky ([MaMi2], Lemma 4.5), and since it will be important later in this chapter we shall now give their construction. Suppose that v_0, v_1, v_2 is a multigeodesic. Realise v_0 and v_2 in general position, and take a regular open neighbourhood N of their union. Attach to N all the disc and once-punctured disc components in its complement. We denote the resulting subsurface of $\Sigma - \Pi$ by N'. Now N' is well-defined up to isotopy, and is what we shall call the subsurface filled by v_0 and v_2. The

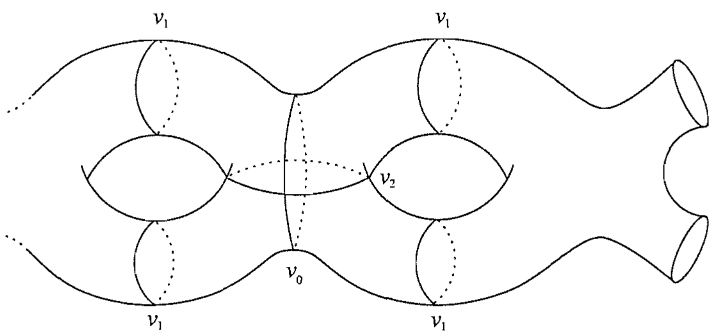

Figure 2.1: The multicurve v_1 is the relative boundary of v_0 and v_2, making v_0, v_1, v_2 a tight geodesic. The subsurface bordered by the components of v_1 is that filled by v_0 and v_2.

boundary of N', denoted $\partial(v_0, v_2)$ and called the relative boundary of v_0 and v_2, is a non-empty multicurve such that $v_0, \partial(v_0, v_2), v_2$ is a tight multigeodesic. (See Figure 2.1.)

Let us now suppose that v_0, v_1, v_2, v_3 is a multigeodesic tight at v_2. We replace v_1 with the relative boundary of v_0 and v_2. It can be shown that this does not affect tightness at v_2, that is $\partial(v_0, v_2), v_2, v_3$ is a tight multigeodesic. This tightening procedure is therefore robust, and can be applied in any order to any geodesic in the curve graph to find a tight multigeodesic with the same ends. Moreover, any tight multipath is necessarily locally 2-geodesic. Whether we can always connect two vertices of $\mathcal{G}$ by a tight geodesic, rather than having to use multicurves, remains open. Perhaps weaker forms of tightness can be developed to address this.

The central result may be stated as follows.

Lemma 1 *There is an explicit function $F : \mathbb{N}^3 \longrightarrow \mathbb{N}$ such that the following holds. Let Σ be any closed and orientable surface and let Π be a finite subset such that $\Sigma - \Pi$ is non-exceptional. For $(v_i)_0^n$ any multigeodesic in $\mathcal{G}(\Sigma, \Pi)$ tight at v_1, we have $\iota(v_1, v_n) \leq F(\iota(v_0, v_n), genus(\Sigma), |\Pi|)$.*

Note that $F(s, genus(\Sigma), |\Pi|)$ grows exponentially with s, perhaps leaving some room for improvement. In particular, the loss of the uniformity of the bounds of [Bow2] appears to be the current price of computability. Even so, these bounds are enough to deduce that any two ideal boundary points of $\mathcal{G}$ are connected by a bi-infinite tight multigeodesic. A proof of this visual connectivity, using a diagonal sequence argument, is given by Bowditch [Bow2]. For distances at most 5, there are geodesics satisfying the inequality given in Lemma 1 with $F(\iota(\alpha, \beta), genus(\Sigma), |\Pi|)$ cubic in $\iota(\alpha, \beta)$.

Consequences of Lemma 1 include the following.

Theorem 2 *(Masur-Minsky) Suppose $\Sigma - \Pi$ is non-exceptional. Then, between any two vertices of $\mathcal{G}(\Sigma, \Pi)$ there are only finitely many tight multigeodesics. In addition, (S) their number is uniformly bounded in terms of the topology of $\Sigma - \Pi$ and the intersection number of the two ending vertices.*

From this it follows that between any two vertices in the curve graph of the five times punctured sphere, or the twice-punctured torus, there are only finitely many geodesics. Using Theorem 2, we can algorithmically construct all finite length tight multigeodesics.

Theorem 3 *There exists an explicit algorithm which takes as input Σ, Π and any two curves α and β in $\Sigma - \Pi$ and returns all tight multigeodesics connecting α to β.*

From this we can deduce an algorithm for constructing all *hierarchies*, as introduced by Masur-Minsky [MaMi2], and deduce that distances are computable.

Theorem 4 *There exists an explicit algorithm which takes as input Σ, Π and any two curves α and β in $\Sigma - \Pi$ and returns the distance between α and β in $\mathcal{G}(\Sigma, \Pi)$.*

(A version of Theorem 4 for closed surfaces of genus at least two is given in the unpublished thesis of Jason Leasure [Lea]. We learned of this after submission for publication, and we are grateful to Richard P. Kent IV for alerting us.)

We remark that the arguments we shall give in this paper enable the limiting argument of Bowditch [Bow2] for tight multigeodesics to apply in bounded time, that is without having to go all the way to the limit and moreover we know when we can stop. Fix a hyperbolic metric ρ on $\Sigma - \Pi$. Let α and β be any two curves and suppose that $(v_0^i, v_1^i, \ldots, v_n^i)_i$ is an infinite sequence of distinct tight multigeodesics with $v_0^i = \alpha$ and $v_n^i = \beta$ for all i. Suppose further that this sequence is ordered lexicographically according to ρ-length. Then, there exists an explicit natural number N, depending only on $\iota(\alpha, \beta)$ and the topology of $\Sigma - \Pi$, such that for all $i \geq N$ we have to conclude that $v_0^i, v_1^i, \ldots, v_n^i$ is not a geodesic.

2.3 An overview of the main argument

In this section, we give the heart of our main argument before giving bounds on intersection number for the easy low distance cases of Lemma 1. All subsequent sections are devoted to formalising this outline so as to compute the function F of Lemma 1 for larger distances.

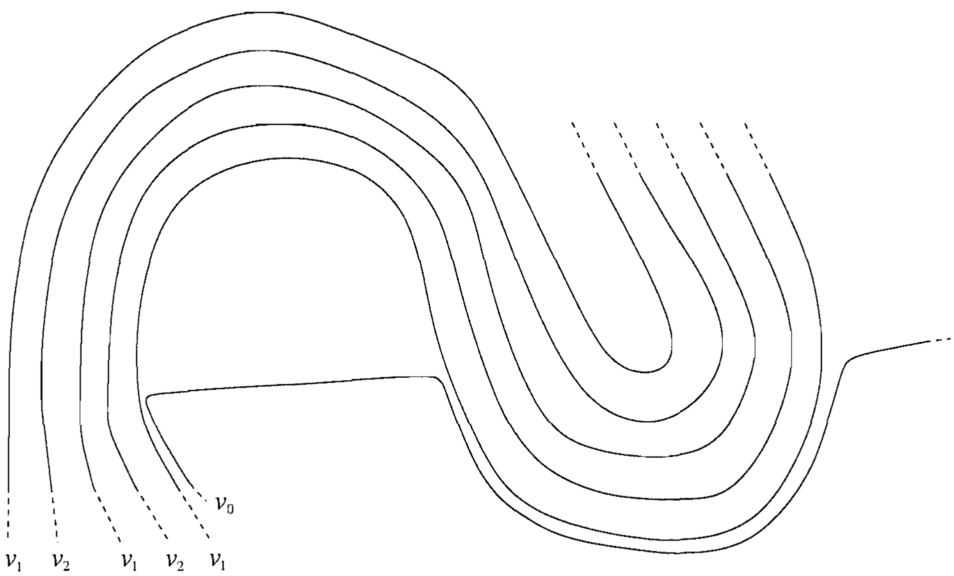

Figure 2.2: v_0, v_1, v_2 is a tight geodesic. v_1 is shielding v_2 from v_0.

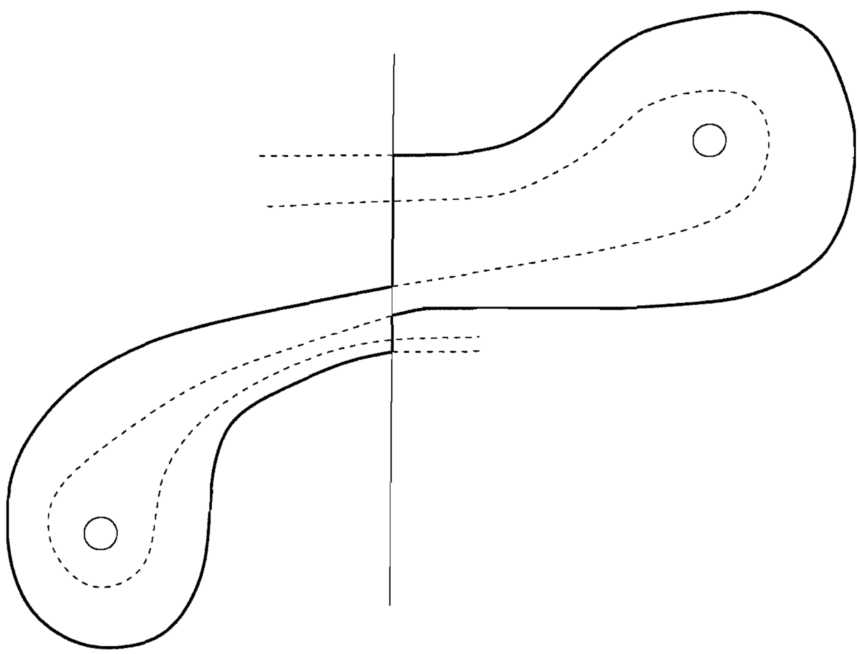

Figure 2.3: A surgery of c_1 along a component of $c_n - c_0$ to form a new curve. In this case, the curve bounds two punctures. When Π is empty we can form a new curve from a single subarc ending on $c_n - c_0$.

Neighbouring multicurves on a tight multigeodesic tend to drag one another round the surface and shield each other from other curves. (See Figure 2.2 and Proposition 5.) Consider any multigeodesic $(v_i)_0^n$ and any simple realisation c_i for v_i, each i, such that $c_i \cap c_{i+1} = \emptyset$ for each i, $|c_i \cap c_n| = \iota(v_i, v_n)$ each $i \leq n-2$ and $c_i \cap c_j \cap c_n = \emptyset$ for each $i < j \leq n-2$. Suppose that two components J_1 and J_2 of $c_n - c_0$ are connected by three subarcs of c_1, denoted g_1, g_2 and g_3, and are homotopic relative to $c_n - c_0$. Tightness at v_1 implies that the ends of at least two of these subarcs, say g_1 and g_2, are separated by a point from $c_2 \cap c_n$. To see this, we can otherwise readily construct a curve intersecting v_1 but disjoint from both v_0 and v_2, violating tightness at v_1. (See Figure 2.3 above for a depiction of such a surgery.)

Now c_1 and c_2 are disjoint so we conclude that there must be a subarc h of c_2 connecting J_1 to itself and sandwiched between g_1 and g_2. This subarc of c_2 is shielded from c_0 by c_1. Suppose that $J_1 = J_2$. If the ends of h are not separated on J_1 by a point from $c_3 \cap c_n$ then we may carry out a surgery of h along c_n to find a new simple loop c_2' disjoint from both c_0 and c_3. In particular, when Π is empty c_2' represents a curve, denoted γ_2', and we have succeeded in finding a new multipath v_0, γ_2', v_3, contradicting $d(v_0, v_3) = 3$. We conclude that the ends of h must be separated by c_3. The case when Π is non-empty is similarly treated, except we ask for the second return to J_1.

This analysis continues along $(v_i)_0^n$ to higher indices. Suppose that a long subarc of c_{n-2} is fellow travelled by two long subarcs of c_1, one on either side, and kept apart from c_1 by long subarcs from $c_2, c_3, \ldots, c_{n-3}$ in turn. Then we may perform a surgery on c_{n-2} along $c_n - c_0$ to arrive at a new curve γ_{n-2}' having zero intersection with v_0. Furthermore, $c_{n-1} \cap c_n = \emptyset$ and so this time we are guaranteed $\iota(\gamma_{n-2}', v_{n-1}) = 0$. We have succeeded in finding a multipath connecting v_0 and v_{n-1} of length less than $n-1$ and hence we have a contradiction. (See Figure 2.4 overleaf for an illustration of this argument.)

We shall see that if we start with a multigeodesic $(v_i)_0^n$ which is tight at v_1 and is such that $\iota(v_1, v_n)$ is large relative to $\iota(v_0, v_n)$ (we shall quantify this in terms of F in §2.6) then this is exactly the situation we find ourselves in. Furthermore, the argument only requires tightness at v_1.

In the absence of an analogue for the tight geodesic these arguments would not seem to translate to the 1-skeleton of the pants complex. With the probable exception of punctured spheres, they would also seem to not translate to the 1-skeleton of the arc complex. Although there is an analogue of the tight geodesic, we find that performing a surgery on an arc running self-parallel over much of its length does not address the possibility of a second arc intersecting near the ends. It would be of much interest to overcome these issues, and find distance computing algorithms for these graphs.

Returning to the curve graph, we now deal with the cases $n = 2$ and $n = 3$ separately and hereafter assume $n \geq 4$.

Proposition 5 *For each multigeodesic v_0, v_1, v_2 we have $\iota(v_1, v_2) = 0$. For each multigeodesic v_0, v_1, v_2, v_3 tight at v_1 we have $\iota(v_1, v_3) \leq 2\iota(v_0, v_3)$.*

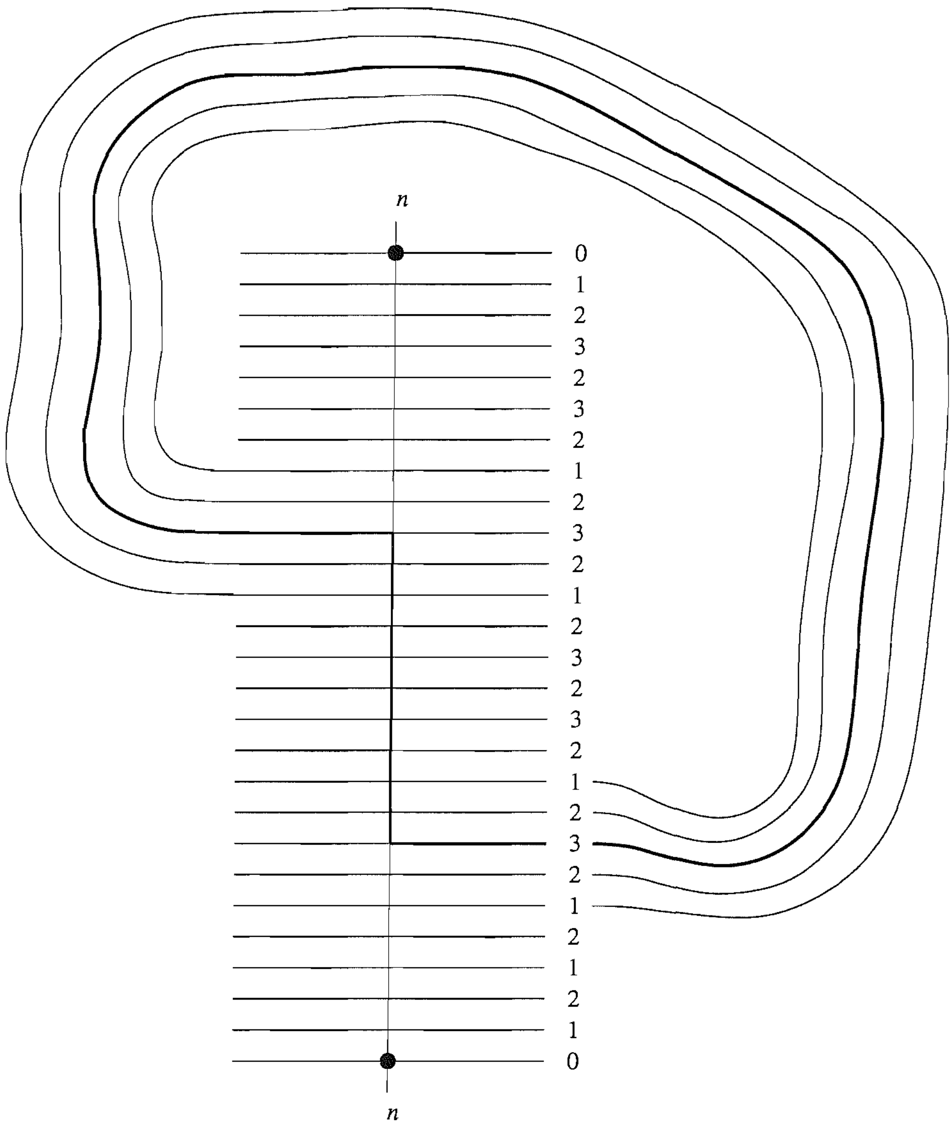

Figure 2.4: A section of consecutive multicurves on some tight geodesic. In general, the first multicurve traps long subarcs of the second, the second of the third and so on. The highlighted curve, resulting from a surgery on the third multicurve, is disjoint from the zeroth and fourth multicurves. When $n = 5$, the highlighted curve is always disjoint from the fourth multicurve.

Proof The first statement is obvious. For the second, we note that for any multigeodesic u_0, u_1, u_2 tight at u_1 and any multicurve z, we have $\iota(z, u_1) \leq 2(\iota(z, u_0)+\iota(z, u_2))$. Before proving this inequality in the next paragraph, we see that when $z = v_3$ we have $\iota(v_1, v_3) \leq 2(\iota(v_0, v_3) + \iota(v_2, v_3)) = 2(\iota(v_0, v_3) + 0) = 2\iota(v_0, v_3)$ and we are done.

Denote by N a subsurface of $\Sigma - \Pi$ representing that filled by u_0 and u_2, and choose a simple loop $c \in z$. Let a denote an essential component of $N \cap c$. Now there are many ways to extend a to a non-trivial and non-peripheral simple loop a^* meeting ∂N essentially and such that $N \cap a^*$ is equal either to a or has precisely two components, each homotopically parallel to a relative to ∂N. Denote the free homotopy class of a^* by $[a^*]$. Since u_0, u_1, u_2 is tight, it follows that at least one of $\iota([a^*], u_0)$ and $\iota([a^*], u_2)$ is non-zero. That is to say, each component of $N \cap c$ makes a positive contribution to the intersection sum $\iota(z, u_0) + \iota(z, u_2)$ and the required inequality follows. $\diamondsuit$

The same argument fails for $n \geq 4$ since v_2 and v_n are no longer adjacent.

2.4 The idea of pulse

We introduce a measure of the interleaving in the surface $\Sigma - \Pi$ of curves lying on a tight multigeodesic in $\mathcal{G}(\Sigma, \Pi)$, allowing us to formalise the outline argument of §2.3 and compute the stated upper bounds on intersection numbers. The same ideas can be applied to geodesics and multigeodesics tight at a given vertex.

For any given positive integer n, let $\mathcal{F}_n$ denote the free monoid of rank n generated by the set $\{e_1, e_2, \ldots, e_n\}$. We shall refer to the elements of $\mathcal{F}_n$ as words. For any word $w \in \mathcal{F}_n$, denote by $w(i)$ the ith letter appearing in w. By a subword u of $w \in \mathcal{F}_n$, we shall mean a finite word $u_1 u_2 \ldots u_k$ such that, for each i, there is an index $j(i)$ with $u_i = w_{j(i)}$ and, for each $i \leq k-1$, we have $j(i) < j(i+1)$. That is, u_i and u_{i+1} need not be consecutive in w. For each i, j define $|e_i - e_j| = |i - j|$. Let $\mathcal{R}$ denote the relation set $\{e_i e_j = e_j e_i : i < j - 1\}$ and let $\mathcal{N}$ denote the congruence on $\mathcal{F}_n$ generated by $\mathcal{R}$. We form the quotient $\mathcal{F}_n/\mathcal{N}$ and refer to the elements of this monoid as *chords*, denoting the $\mathcal{N}$-congruence class of the word w by $\overline{w}$.

Chords naturally arise in the context of paths and multipaths in the curve graph. Consider a path or multipath $(v_i)_0^n$ and choose simple representatives c_i for v_i once more, so that $c_i \cap c_{i+1} = \emptyset$ each i, $|c_i \cap c_n| = \iota(v_i, v_n)$ for each $i \leq n-2$ and $c_i \cap c_j \cap c_n = \emptyset$ for each $i < j \leq n-2$. Let J be any component of $c_n - c_0$. Orient J and use this orientation to enumerate the points of $J \cap \bigcup_0^{n-2} c_i$. This enumeration spells out an element w of $\mathcal{F}_n$ by identifying a point from c_i with the ith generator e_i of $\mathcal{F}_n$, each i. Tightness at v_i implies that w cannot be of the form $w = w_1 e_i^3 w_2$, for some $w_1, w_2 \in \mathcal{F}_n$ and each i. Later, we will consider various subsets of $J \cap \bigcup_0^{n-2} c_i$, and enumerate their elements with the orientation on J to determine the m-pulse.

Now suppose that $|w(i) - w(i+1)| > 1$. Then $\gamma_{w(i)}$ and $\gamma_{w(i+1)}$ have non-zero intersection number and we may homotop both $c_{w(i)}$ and $c_{w(i+1)}$ near J so as to transpose the two points of intersection corresponding to i and $i+1$. If we re-enumerate, we arrive at a second word $w' \in \mathcal{F}_n$ with $\overline{w} = \overline{w'}$. In this way, paths may be viewed as defining chords and tight multigeodesics "pinched" chords. Note also that each word in a chord induced by a path or multipath can be induced by the same path or multipath, just by considering transpositions.

Let us set about defining the m-pulse of a given word and then of a given chord, for each $2 \leq m \leq n$. For each word $w \in \mathcal{F}_n$ we consider subwords u satisfying the following three conditions. First, both the initial and the final letters in u are equal to e_1. Second, for each i we have $|u(i) - u(i+1)| \leq 1$. Third, between any two successive e_1's in u there is exactly one e_m. We define the m-pulse of such a subword u to be equal to the number of times e_m appears in u. We define the m-pulse of w to be the maximal m-pulse arising among all such subwords u of w and denote it by $p_m(w)$. Even when u satisfies these criteria and is maximal with respect to inclusion among all such subwords, it need not realise the m-pulse of w.

Lemma 6 *Suppose that $v, w \in \mathcal{F}_n$ represent the same chord, that is $\overline{v} = \overline{w}$. Then, the m-pulse of v is equal to that of w for each $m \geq 2$.*

Proof Any two elements of a chord are related by a finite sequence of transpositions using $\mathcal{R}$. The result follows by an induction on the length of such sequences, noting that each transposition fixes every subword satisfying our three criteria and so preserves m-pulse for each $m \geq 2$. $\diamondsuit$

For $2 \leq m \leq n$, we define the *m-pulse* of a given chord to be equal to the m-pulse of one (hence any) representative word, and denote this by $p_m(\overline{w})$. We have just seen that this is well-defined.

Let us complete this section with a few examples and remarks. Chords may be represented by, and are determined by, words from $\{e_1, e_2, \ldots, e_n\}$. For instance, $e_1e_2e_3e_4e_2e_3e_4$ and $e_1e_2e_3e_2e_4e_3e_4$ represent the same chord since the first e_4 and the second e_2 may be transposed. The words $e_1e_2e_1e_2e_1$ and $e_1e_1 \ldots e_1e_2e_1e_1 \ldots e_1e_2e_1$ represent different chords although their 2-pulses both equal 2. The words $e_1e_1e_1e_2e_1e_1e_1e_2e_1$ and $e_1e_2e_2e_2e_1e_2e_2e_2e_1$ define different chords although both their 2-pulses and their lengths are equal. The 2-pulse of either word associated to the component of $c_n - c_0$ depicted in Figure 5 is equal to 6. When chords are induced by a given tight geodesic, the tightness property prevents consecutive repetition of single letters e_i. If we bound the m-pulse on such a chord, for each m, then we bound its length.

Notice that pulse is symmetric and almost additive: For each m, the m-pulse of the concatenation of two words is either the sum of each m-pulse or one more than this sum. The 1-pulse of a word or chord should always be regarded as zero and, for each $2 \leq m \leq n-1$, the $m+1$-pulse of a chord is at most the m-pulse.

In certain circumstances, summing the 2-pulses over each component of $c_n - c_0$ closely approximates $\iota(v_1, v_n)$ as follows.

Lemma 7 *Suppose that $(v_i)_0^n$ is a multigeodesic tight at v_1. Then $\sum p_2(\overline{w}) \leq \iota(v_1, v_n) \leq 2(\iota(v_0, v_n) + \sum p_2(\overline{w}))$, where the summations are taken over the components of $c_n - c_0$. In particular, if K denotes the maximal 2-pulse over each component of $c_n - c_0$, we have $\iota(v_1, v_n) \leq 2(1+K)\iota(v_0, v_n)$.*

Proof Let J be any component of $c_n - c_0$. Orient J and let w be the word associated to this oriented component. Any point of $c_2 \cap J$ contributing to $p_2(\overline{w})$ is, by definition, trapped between two points from $c_1 \cap c_n$. It follows that $p_2(\overline{w}) \leq |c_1 \cap J|$. As $|c_1 \cap c_n|$ is minimal, summing over all components of $c_n - c_0$ gives the first inequality, $\sum p_2(\overline{w}) \leq \iota(v_1, v_n)$.

For the second inequality, we note that, at the very least, for any three consecutive points on J of $c_1 \cap J$ there are two separated on J by a point from $c_2 \cap J$. Thus, $|c_1 \cap c_n| \leq 2p_2(\overline{w}) + 2$, and summing over each component of $c_n - c_0$ gives $\iota(v_1, v_n) \leq 2\iota(v_0, v_n) + 2\sum p_2(\overline{w})$.

The final statement is an immediate consequence of the second inequality. $\diamondsuit$

The following section introduces all the machinery we need to control K.

2.5 Train tracks relative to the ends of a geodesic

Let us fix a smooth structure on Σ. Recall that the smooth double of a connected subsurface E of Σ is formed by taking a copy of E, $\widetilde{E}$, and identifying all corresponding pairs of non-singular points of the boundaries ∂E and $\partial \widetilde{E}$. A *train track*, τ, in $\Sigma - \Pi$ is a smooth branched 1-submanifold such that the Euler characteristic of the smooth double of each component of $\Sigma - (\Pi \cup \tau)$ is negative. This rules out discs, once-punctured discs and discs with one or two boundary singularities as complementary regions. Train tracks were introduced by Thurston to study geodesic laminations and they shall prove useful here in handling long curves.

It is standard to refer to the branch points of a train track as *switches* and the edges between switches as *branches.* We say that τ is *generic* if each switch has valence three. By sliding branches along branches, if need be, we can take a train track and return a generic train track. This is a convenient option since it greatly simplifies our counting arguments. A *train subpath* $p : I \longrightarrow \tau$ is a continuous map on a closed interval $I \subseteq \mathbb{R}$ such that $p(n)$ is a switch for each $n \in \mathbb{Z} \cap I, p^{-1}(v) \in \mathbb{Z}$ for each switch v and $\partial I \subseteq \mathbb{Z} \cup \{\pm\infty\}$.

A smooth simple closed loop c is *carried* by τ if there exists a smooth map $\phi : \Sigma - \Pi \longrightarrow \Sigma - \Pi$ homotopic to the identity map on $\Sigma - \Pi$ such that the restriction $\phi|_c$ is an immersion and $\phi(c) \subseteq \tau$. We refer to ϕ as a carrying map or supporting map. If τ carries c then we have a measure on the branch set of τ by counting the number of times the subpath $\phi(c)$ traverses any given branch. This measure satisfies a switch condition: At each switch, the total inward measure is equal to the total outward measure.

We recall a useful combinatorial lemma relating the number of switches and the number of branches of a train track to the Euler characteristic of $\Sigma - \Pi$. This is Corollary 1.1.3 from [PenH].

Lemma 8 *Let τ be any train track in $\Sigma - \Pi$, let s denote the number of switches and e the number of branches. Then:*
i). $s \leq -6\chi(\Sigma - \Pi) - 2|\Pi|$;
ii). $e \leq -9\chi(\Sigma - \Pi) - 3|\Pi|$.

Let $(v_i)_0^n$ be any multipath in $\mathcal{G}(\Sigma, \Pi)$ and choose smooth and simple realisations c_i for v_i, each i, such that $c_i \cap c_{i+1} = \emptyset$ for each i, $c_i \cap c_n = \iota(v_i, v_n)$ for each $i \leq n-2$ and $c_i \cap c_j \cap c_n = \emptyset$ for each $i < j \leq n-2$. We construct a train track τ which will carry all of c_1 and all those subarcs of each c_i $(2 \leq i \leq n-2)$ which end on c_n and which are trapped between subarcs of c_1 over large distances.

There exists a smooth surjection $\phi : \Sigma - \Pi \longrightarrow \Sigma - \Pi$ homotopic to the identity map such that the restriction of ϕ to c_1 is an immersion onto a smooth branched 1-submanifold τ of $\Sigma - \Pi$ with the characterising properties that any two components of $c_1 - c_n$ homotopic relative to $c_n - c_0$ are carried into the same edge of τ, each component of $c_n - c_0$ intersects τ at most once and τ is to be disjoint from c_0.

Another way to explain the construction of τ is to first form a homotopy $H : \Sigma - \Pi \times [0,1] \longrightarrow \Sigma - \Pi$, such that $H(\_, t) : \Sigma - \Pi \longrightarrow \Sigma - \Pi$ is injective for all $t \in [0,1)$ and $H(\_, 1)$ is an immersion, by sliding all parallel components of $c_1 - c_n$ along $c_n - c_0$ and onto a single arc. Any two such arcs incident on a common component of $c_n - c_0$ are to be coincident. Last, we insist that H smooths out c_1 so that $H(c_1, 1)$ is a smooth 1-submanifold of $\Sigma - \Pi$. This homotopy is depicted in Figure 2.5, and we can take our carrying map ϕ to be precisely $H(\_, 1) : \Sigma - \Pi \longrightarrow \Sigma - \Pi$.

Now each branch point of τ necessarily belongs to one component of $c_n - c_0$ and each component of $c_n - c_0$ contains at most one branch point. We next check that τ defines a train track, with each branch point viewed as a switch and each edge thought of as a branch, and that τ is unique up to isotopy.

Lemma 9 *τ is a train track.*

Proof Note that no region complementary to τ can be diffeomorphic to a disc with smooth boundary or a monogon (disc with one outward pointing singularity) by the minimality of $|c_1 \cap c_n|$.

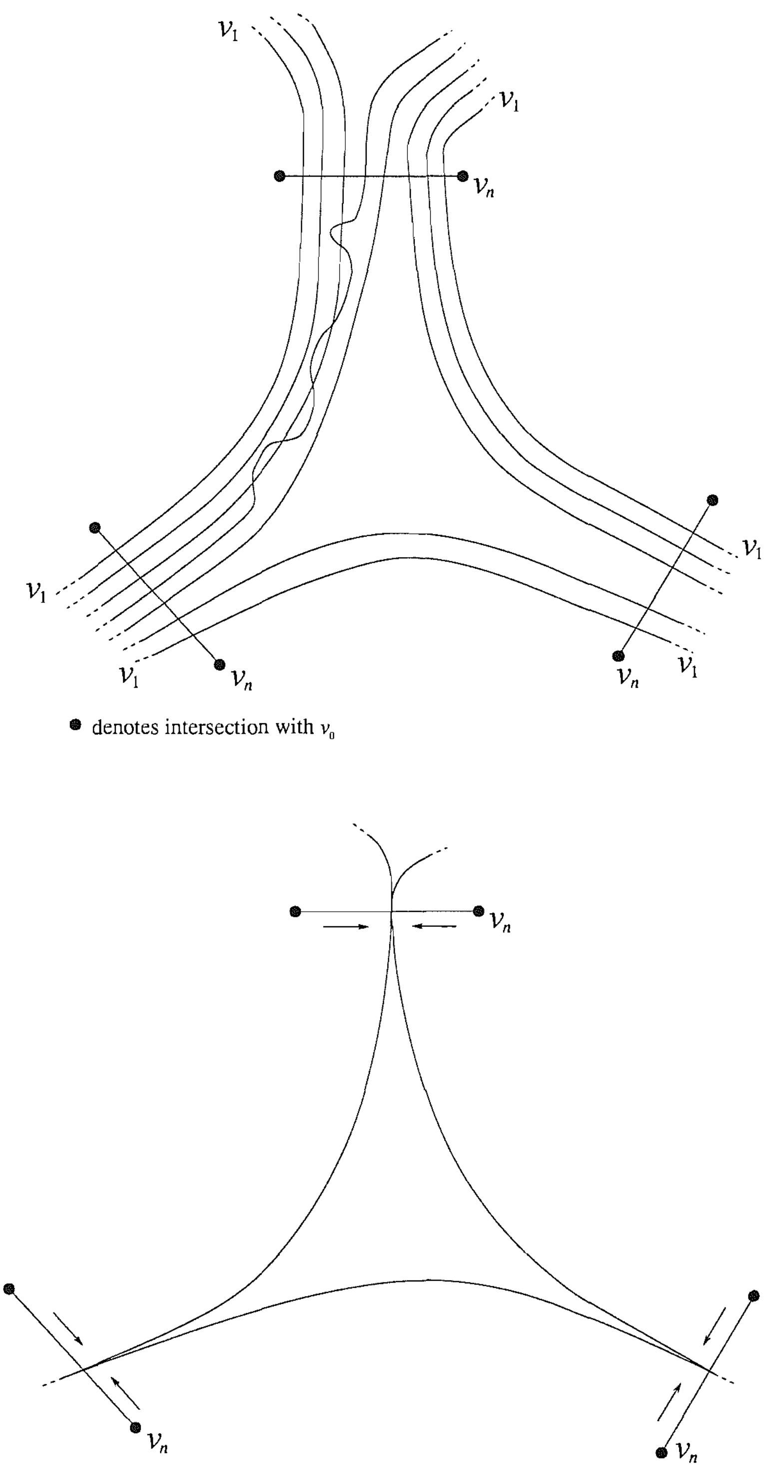

Figure 2.5: Subarcs of v_1 above collapse to branches of τ below, taking nearby and parallel subarcs of other multicurves with them.

Suppose for contradiction that E is a bigon component, that is a disc with two outward pointing singularities, of $\Sigma - (\Pi \cup \tau)$. The two subarcs of ∂E connecting the two singularities of ∂E are homotopic to one another relative to $c_n - c_1$. Hence E must intersect c_0, for otherwise these two subarcs of ∂E would have been collapsed into a single branch of τ. Since τ and c_0 are disjoint, so ∂E and c_0 are disjoint. Hence E contains a component of c_0 which is therefore homotopically trivial. This is absurd, and we conclude that τ is a train track. $\diamondsuit$

Lemma 10 *Suppose that $c_0^i, c_1^i, \ldots, c_n^i$ $(i = 1, 2)$ are two such realisations for $v_0, v_1, \ldots, v_n$ and that τ_1 and τ_2 are the resulting train tracks, respectively. Then τ_1 and τ_2 are isotopic.*

Proof This follows since $c_0^1 \cup c_1^1 \cup c_n^1$ and $c_0^2 \cup c_1^2 \cup c_n^2$ are isotopic. $\diamondsuit$

It is worth pointing out, though we shall not be making use of it, that the same construction for c_i $(i = 2, 3, \ldots, n-2)$ relative to c_0 and c_n will not necessarily yield a train track but instead a bigon train track, where we allow the complementary regions to be bigons. Each complementary bigon will contain at least one point from $c_0 \cap c_n$ so there would be at most $\iota(v_0, v_n)$ bigons. If we wished, we could also proceed from here and deduce different bounds to those described in Lemma 1.

Now to each switch z of τ we can associate the finite set $\phi^{-1}(z) \cap \bigcup_0^{n-2} c_i$, which we henceforth denote by $D(z)$. If need be, we are free to homotop $c_1, c_2, \ldots, c_{n-1}$ so that $|D(z)|$ is henceforth minimal for each switch z. Orient c_n and use this orientation to enumerate the points of $D(z)$. This gives us a word w in F_{n-2} which, by the minimality of $|D(z)|$, begins and ends with e_1. Thus, for each integer $2 \leq m \leq n-2$, we may associate to the switch z the m-pulse of the chord $\overline{w}$.

We now use pulse on switches to define a certain measure on the branch set of τ. Suppose that z_1 and z_2 are adjacent switches of τ connected by a branch b. We need to verify that it will make no difference whether we use either z_1 or z_2 to define the pulse on b. In what follows, the topological closure of b in $\Sigma - \Pi$ is denoted $cl(b)$.

Lemma 11 *The chords on $c_n - c_0$ defined by $\phi^{-1}(cl(b)) \cap D(z_1)$ and by $\phi^{-1}(cl(b)) \cap D(z_2)$ are equal.*

Proof For each $x \in \phi^{-1}(cl(b)) \cap D(z_1)$, define $q(x)$ to be the end on $\phi^{-1}(cl(b)) \cap D(z_2)$ of the subarc of the c_i containing x, beginning at x and collapsing to b. We denote this subarc by $[x, q(x)]$. The subarc $[x, q(x)]$ crosses a second subarc $[y, q(y)]$ only if the corresponding multicurves c_i and c_j satisfy $|i - j| \geq 2$, and the lemma follows. $\diamondsuit$

We define the m-pulse of the branch b to be equal to the m-pulse associated to $\phi^{-1}(cl(b)) \cap D(z_1)$, or indeed the m-pulse associated to $\phi^{-1}(cl(b)) \cap D(z_2)$, and denote it by $p_m(b)$. This measure on the branch set of τ is invariant under the action of the mapping class group and satisfies a certain coarse switch condition. For any switch z we may choose one of two directions at z and partition the branches incident on z as outgoing and incoming. Denote the corresponding branches by $b_1, b_2, \ldots, b_s$ and $b_{s+1}, b_{s+2}, \ldots, b_{s+t}$, respectively.

When τ is generic, we have the following.

Lemma 12 *Suppose that τ is generic and that b_1 and b_2 are outgoing. For each $m \geq 2$ we have:*
i). $0 \leq p_m(b_3) - p_m(b_1) - p_m(b_2) \leq 1$;
ii). $p_m(z) - 1 \leq p_m(b_1) + p_m(b_2) \leq p_m(z)$;
iii). $p_m(b_3) = p_m(z)$.

We shall only require Lemma 12i) and we prove it directly. The remaining two points follow by similar considerations.

Proof Suppose that v is a trivalent switch with outgoing branches b_1 and b_2 and incoming branch b_3. Each component of $c_1 \cap \phi^{-1}(b_3)$ goes on to be carried either by b_1 or by b_2. All those components of $c_m \cap \phi^{-1}(b_3)$ that are trapped between components of $\phi^{-1}(b_1 \cup \{z\} \cup b_i)$ go on to be supported by b_i $(i = 2, 3)$. However, where these components of c_1 diverge subarcs of c_m may escape. This reduces the total outgoing m-pulse by at most one, if at all. That is, we have part i) of Lemma 12. $\diamondsuit$

An induction argument on the valency at any given switch, that proceeds by sliding branches along branches to reduce valency and uses Lemma 12 in the base case, yields the following.

Corollary 13 *"Coarse switch condition." For each integer $2 \leq m \leq n-2$, we have:*
i). $|\sum_1^s p_m(b_i) - \sum_{s+1}^{s+t} p_m(b_j)| \leq s+t-2$;
ii). $p_m(z) - s + 1 \leq \sum_1^s p_m(b_i) \leq p_m(z)$;
ii′). $p_m(z) - t + 1 \leq \sum_{s+1}^{s+t} p_m(b_j) \leq p_m(z)$.

In passing, we remark that the sum of all 2-pulses over each branch of τ resembles a reduced intersection number for c_1 and c_n relative to c_0. That is, suppose that $\iota(v_0, v_n)$ is large. Then most of the regions complementary to $c_0 \cup c_n$ are squares. Whenever c_1 meets an edge of one of these squares, necessarily from c_n, it must go on to meet the edge opposite. Consider a sequence $S_1, S_2, \ldots, S_k$ of closed squares whose interiors are complementary to $c_0 \cup c_n$ and such that

$S_i \cap S_{i+1} \subseteq c_n$, each i. Whenever c_1 meets one of the outer edges of $\bigcup_1^k S_i$ then c_1 remains trapped in $\bigcup_1^k S_i$ and goes on to meet each edge $S_i \cap S_{i+1}$ before exiting at the other outer edge. However, the components of $c_1 \cap \bigcup_1^k S_i$ all collapse into a single branch of τ. If any two such components of c_1 are separated by a subarc of c_2 then we conclude that the 2-pulse on this branch is precisely one less than the number of components of $c_1 \cap \bigcup_1^k S_i$.

Returning to our main thread, each arc g carried by τ and whose ends are carried to switches by ϕ defines a train subpath g_ϕ of τ. We are primarily interested in those train subpaths that occur in this way. For a train path q in τ we define $Supp(q)$ to be the set of all subarcs g of each c_i $(i = 1, 2, \dots, n-2)$ with $q = g_\phi$. Suppose that $q : I \longrightarrow \tau$ is a train path with $0 \in I \subseteq [0, \infty)$. We call the branch of τ containing $q(i + 1/2)$ the ith branch of q. For each integer $m \geq 2$ and each $i \geq 0$ we define the m-pulse of the ith branch of q to be equal to the m-pulse of the subset $\phi^{-1}(q(0)) \cap \bigcup Supp(q|_{[0,i+1]})$ of $c_n - c_0$. We denote this by $p_{m,q}(i)$. Note the argument used in the proof Lemma 11 but applied to train subpaths tells us that this quantity is well-defined. Note also that $p_{m,q}(0)$ is precisely the m-pulse of the 0th branch traversed by q, for each $m \geq 2$.

To better understand the behaviour of $p_{m,q}$, we record two properties.

Lemma 14 *"Trains run out of fuel." Let $q : [0, \infty) \longrightarrow \tau$ be any train path in τ. Then:*

i). $p_{m,q}(i+1) \leq p_{m,q}(i)$ for each $m \geq 2$ and for each i;

ii). $p_{m,q}(i) \longrightarrow 0$ as $i \longrightarrow \infty$ for each $m \geq 2$.

Proof Each time q arrives at a switch of τ, subarcs of c_1 that have so far induced q may diverge and there can be no gain in pulse. Hence i) holds. If $p_{m,q}(i) \geq 1$ for all i then, since Σ is orientable, we conclude that c_1 has two freely homotopic components and this is absurd. Hence ii) holds. $\diamondsuit$

2.6 Proof of Lemma 1

Let $(v_i)_0^n$ be any multigeodesic in $\mathcal{G}(\Sigma, \Pi)$ tight at v_1 with $n \geq 4$. Let c_i be any realisation for v_i, each i, such that $c_i \cap c_{i+1} = \emptyset$ for each $i \leq n-1$, $|c_i \cap c_n| = \iota(v_i, v_n)$ each $i \leq n-2$ and $c_i \cap c_j \cap c_n = \emptyset$ each $i < j \leq n-2$. Recall the construction of the generic train track τ relative to v_0 and v_n and carrying all of v_1, as described in §2.5. We can endow each branch of τ with a family of measures and each of these satisfies a coarse switch condition. For any train subpath q of τ we defined a "time" measure associated to q by considering those subarcs of each c_i that induce q via ϕ.

For reasons we will soon make clear, let us define a function $K_n : \{2, 3, \dots, n-1\} \times \mathbb{N} \times \mathbb{N} \longrightarrow \mathbb{N}$ by the recurrence relation

$$K_n(j, s, t) = 2^{12(3s+t-3)}(1 + K_n(j+1, s, t)),$$

for all $j = 2, 3, \ldots, n-2$, with the boundary condition $K_n(n-1, s, t) = 0$. Note when $s = genus(\Sigma)$ and $t = |\Pi|$ the exponent is precisely twice the upper bound on the number of branches of a train track in $\Sigma - \Pi$ given in Lemma 8ii), where the scaling by two will allow for the losses in pulse described by Lemma 12i).

The following is an important and versatile result we shall be using time and again. In particular, it makes no specific reference to geodesics.

Lemma 15 *Let $m \geq 2$. Suppose that all the branches of τ have $m+1$-pulse at most $K_n(m+1)$ and that at least one branch of τ has m-pulse greater than $K_n(m)$. Then there exists a subarc h of c_m disjoint from c_0 and whose ends lie on, but are otherwise disjoint from, one component of $c_n - (c_0 \cup c_{m+1})$.*

Proof The choice of K is entirely motivated by our needs, thus our proof really amounts to a derivation and explanation of the defining recurrence relation. We begin with $K_n(n-2, genus(\Sigma), |\Pi|)$. Whenever parallel subarcs of c_{n-2} carried by the branch b, and thus trapped by subarcs of $c_1, c_2, \ldots, c_{n-3}$ in turn, approach either end of b they may either both continue along the same branch or diverge. This is reflected in Lemma 12i), where the $n-2$-pulse on the two outgoing branches sums to either $p_{n-2}(b) - 1$ or $p_{n-2}(b)$.

Let us travel along τ from b by always following the branch that receives most $n-2$-pulse from the previous branches. Combining Lemma 8ii) and Lemma 12i), if we are given that $p_{n-2}(b)$ is at least $2^{-2(9\chi(\Sigma-\Pi)+3|\Pi|)}$ then we are guaranteed to return to b. To be more precise, in the language of §2.5 we can find a circuit $q : [0, k+1] \longrightarrow \tau$ such that $q(1/2) = q(k+1/2) = b$ and $p_{n-2,q}(k) \geq 1$. We deduce that q is the ϕ-image of a subarc h of c_{n-2} beginning and ending on the same component of $c_n - c_0$. Further, since h is disjoint from c_1 and homotopically parallel to two subarcs of c_1 relative to $c_n - c_0$, one subarc on either side of h, the subarc h is also disjoint from c_0. Finally, $c_{n-1} \cap c_n$ is empty so h certainly begins and ends on the same component of $c_n - (c_0 \cup c_{n-1})$.

The general case follows the same principal. We again need the m-pulse on some component J of $c_n - (c_0 \cup c_{m+1})$ to be at least $2^{-2(9\chi(\Sigma-\Pi)+3|\Pi|)}$. A sufficient condition for this to happen is $p_m(b) > (2^{-2(9\chi(\Sigma-\Pi)+3|\Pi|)} + 2)p_{m+1}(b)$, which in turn holds whenever $p_m > K_n(m, genus(\Sigma), |\Pi|)$ and $p_{m+1}(b) \leq K_n(m+1, genus(\Sigma), |\Pi|)$.

In which case, we may again travel around τ and back to b by always following those branches that receive most m-pulse. It might be that on the first return to b we arrive at the wrong component of $c_n - (c_0 \cup c_m)$. We therefore continue our journey until we first return to the same component of $c_n - (c_0 \cup c_m)$. Notice that the resulting train subpath can not lose m-pulse more than once at any switch it visits more than twice. $\diamondsuit$

Corollary 16 *The m-pulse on each branch of τ is at most $K_n(m, genus(\Sigma), \Pi)$, each $m \geq 2$.*

Proof Suppose, for contradiction, that there exists m and a branch of τ whose m-pulse is at least $K_n(m, genus(\Sigma), \Pi)$. We take m to be maximal subject to this property. Now by Lemma 15 we deduce that there exists a subarc h of c_m beginning and ending on the same component of $c_n - (c_0 \cup c_{m+1})$ and disjoint from c_0. When Π is empty we know that the union of h and the subinterval of $c_n - (c_0 \cup c_m)$ connecting its ends defines a curve, denoted δ, and we know that this curve has zero intersection with both v_0 and v_{m+1}. We have found a multipath v_0, δ, v_{m+1} of length two. Since $d(v_0, v_{m+1}) = m + 1 \geq 2 + 1 = 3$, we have a contradiction.

When Π is non-empty we only have to be slightly more careful since δ may otherwise be peripheral. Instead, we can ask for the second return to the same component of $c_n - (c_0 \cup c_{m+1})$. By considering the boundary components of a regular neighbourhood of the union of h and the subinterval of $c_n - c_0$ connecting the ends of h, we again find a curve δ which again has zero intersection with both v_0 and v_{m+1}. This type of surgery is depicted in Figure 2.3. $\diamondsuit$

Corollary 17 *Suppose that $(v_i)_0^n$ is a multigeodesic tight at v_1. Then $\iota(v_1, v_n) \leq 2(1 + K_n(2, genus(\Sigma), |\Pi|))\iota(v_0, v_n)$.*

Proof We have seen that the 2-pulse on each branch of τ is bounded above by $K_n(2, genus(\Sigma), |\Pi|)$. Since τ is generic, we can, if need be, perturb each component of $c_n - c_0$ so that it crosses only one branch rather than one switch. As $(v_i)_0^n$ is tight at v_1 we therefore have, by Lemma 7, $\iota(v_1, v_n) \leq 2(1 + max\{p_2(b) : b$ is a branch of $\tau\})\iota(v_0, v_n) \leq 2(1 + K_n(2, genus(\Sigma), |\Pi|))\iota(v_0, v_n)$. $\diamondsuit$

We conclude the proof of Lemma 1.

Proof (of Lemma 1) Since $d(v_0, v_n) \leq \iota(v_0, v_n) + 1$ and $d(v_0, v_n) = n$ we have $K_n(2, genus(\Sigma), |\Pi|) \leq K_{\iota(v_0,v_n)+1}(2, genus(\Sigma), \Pi)$. Hence $F(j, s, t) = 2(1 + K_{j+1}(2, s, t))j$ suffices. $\diamondsuit$

In summary, we have established the following bounds on pulse applicable to all geodesics in the curve graph.

Theorem 18 *Suppose that $\Sigma - \Pi$ is non-exceptional. For any geodesic $\gamma_0, \gamma_1, \ldots, \gamma_n$ in the curve graph $\mathcal{G}(\Sigma, \Pi)$ and for each $m \geq 2$, the m-pulse on any branch of the train track associated to γ_1 and relative to γ_0 and γ_n is at most $K_n(m, genus(\Sigma), \Pi)$.*

2.7 Proof of Theorems 2, 3 and 4

In this section, we prove the three main implications of Lemma 1: We establish a finiteness result for tight multigeodesics and we establish the computability of tight multigeodesics and hence the computability of distances in the curve graph.

Lemma 19 *There exists an explicit increasing function* $F_1 : \mathbb{N}^3 \longrightarrow \mathbb{N}$ *such that the following holds. Let* $(v_i)_0^n$ *be any tight multigeodesic. Then,* $\iota(v_j, v_n) \leq F_1(\iota(v_0, v_n), genus(\Sigma), |\Pi|)$ *for all* j.

Proof For $1 \leq k \leq n-1$ and non-negative integers s and t, put $G_{s,t}(k) = F(k, s, t)$. Lemma 19 follows by an induction on j, using Lemma 1 and considering $F_1(k, s, t)$ defined equal to the composition $G_{s,t}^k(k)$. $\diamondsuit$

Being an iterated composition using F, this choice of F_1 has superexponential growth in its first factor. While this is highly undesirable, our choice does have at least one interesting advantage. If, in future work, we succeed in taking the function F of Lemma 1 to be polynomial in its first factor, this immediately improves F_1. For reasons we shall not go into, at the time of writing such an improvement in F does not seem entirely unrealistic.

Now replacing $(v_{n-i})_0^n$ for $(v_i)_0^n$ in the statement of Lemma 19, we deduce the following.

Corollary 20 *There exists an explicit increasing function* $F_2 : \mathbb{N}^3 \longrightarrow \mathbb{N}$ *such that the following holds. Let* $(v_i)_0^n$ *be any tight multigeodesic. Then,* $\iota(v_0, v_j)$ *and* $\iota(v_j, v_n)$ *are both at most* $F_2(\iota(v_0, v_n), genus(\Sigma), |\Pi|)$ *for all* j.

Let α and β be any two vertices of the curve graph. The set of multipaths of length at most $\iota(\alpha, \beta) + 1$ connecting α to β for which each multicurve verifies the bounds in Corollary 20 has size uniformly and explicitly bounded in terms of $\iota(\alpha, \beta)$, $genus(\Sigma)$ and $|\Pi|$. Since this set includes all the tight multigeodesics connecting α to β, we conclude the proof of Theorem 2. The algorithm of Theorem 3 is then given by a search for all the tight multigeodesics contained in this bounded set. Since tight multigeodesics are distance realising multipaths, we can then read off the distance between α and β and so deduce Theorem 4.

2.8 The computability of stable lengths

Given a metric space X and an isometry $h : X \longrightarrow X$ we define the stable length, $||h||$, of h to be equal to $lim_{n \longrightarrow \infty} d(x, h^n x)/n$. See [BaGS] for more details. It is easily verified that $||h||$ does not depend on the choice of x and is always finite. We say that a mapping class h is pseudo-Anosov if for any two curves α and β we have $\iota(\alpha, h^n(\beta)) \longrightarrow \infty$ as $n \longrightarrow \infty$ (see [FLP]). Equivalently, no non-zero power of h fixes a curve. We consider $\mathcal{G}$ endowed with usual path-metric, and prove the computability of the stable length of any given pseudo-Anosov mapping class.

Let us first recall a few results. In [Bow1], not only is the hyperbolicity of the curve complex re-established but also hyperbolicity constants, logarithmic in surface complexity, are computed. In [Bow2], it is established that there exists a positive integer $N = N(genus(\Sigma), |\Pi|)$ such that for each pseudo-Anosov mapping class h, h^N has a geodesic axis in $\mathcal{G}$. This implies the stable lengths of pseudo-Anosov mapping classes are both positive and uniformly rational. It is not known how to compute N or, indeed, whether we may take N equal to one. Last, in a k-hyperbolic geodesic metric space each geodesic rectangle is $8k$-narrow (so that any point on any one side of the rectangle is within $8k$ of the union of the other three). We now state the result in full:

Theorem 21 *There is an explicit algorithm which takes as input Σ, Π, N and a pseudo-Anosov mapping class h and returns $||h||$.*

Proof Fix a choice of k such that $\mathcal{G}$ is k-hyperbolic. Choose an integer $M \geq 18k$. Let us suppose that h is the Nth power of a pseudo-Anosov. Then, h is again pseudo-Anosov and has a geodesic axis denoted L. Choose any curve α and construct a geodesic $[\alpha, h^M\alpha]$ from α to $h^M\alpha$ in $\mathcal{G}$. A central vertex β of $[\alpha, h^M\alpha]$ must lie within $8k$ of L. Now construct a geodesic $[\beta, h^M\beta]$ from β to $h^M\beta$. We have $M||h|| = ||h^M|| \leq d(\beta, h^M\beta) + 16k \leq Md(\beta, h\beta) + 16k$. Hence $||h|| \leq d(\beta, h\beta) + 16k/M < d(\beta, h\beta) + 1$. As $||h||$ is an integer, so $||h|| \leq d(\beta, h\beta)$.

On the other hand, $d(\beta, h^M\beta) \leq ||h^M|| + 16k = M||h|| + 16k$ and so $d(\beta, h^M\beta)/M \leq ||h|| + 16k/M < ||h|| + 1$.

Combining the two inequalities, we have $||h|| \leq d(\beta, h^M\beta)/M < ||h|| + 1$. Hence $\lfloor d(\beta, h^M\beta)/M \rfloor = ||h||$. $\diamondsuit$

Notice that in the above we do not find an axis L, nor do we see how to compute an appropriate value for N. It would be interesting to find ways of doing so.

2.9 Acylindricity

We prove a computable version of the acylindricity theorem due to Bowditch. In [Bow2], he formulated acylindricity as follows.

Definition 22 *We say that an isometric action by a group Γ on a metric space X is acylindrical if for all $r \geq 0$, there exist constants $R, N \geq 0$ such that for all $x, y \in X$ with $d(x, y) \geq R$ there are at most N distinct elements h of Γ such that $d(x, hx) \leq r$ and $d(y, hy) \leq r$.*

(Such an action is said to be *weakly acylindrical* if, instead of asking for uniform bounds on the number of such mapping classes, we ask for their number to be finite.) In our study of the action of the mapping class group on the curve graph, we shall be replacing N with a computable function of r, the intersection number of the two curves concerned and the topology of $\Sigma - \Pi$. The main result of this section is thus:

Theorem 23 *There exists an explicit function $N : \mathbb{N}^4 \longrightarrow \mathbb{N}$ such that the following holds. Suppose that $\Sigma - \Pi$ is non-exceptional and let r be any non-negative integer. There exists a non-negative integer R, depending only on r, such that for any two curves α and β with $d(\alpha, \beta) \geq R$, the number of mapping classes $h \in Map(\Sigma - \Pi)$ such that $d(\alpha, h\alpha) \leq r$ and $d(\beta, h\beta) \leq r$ is bounded above by $N(\iota(\alpha, \beta), r, genus(\Sigma), |\Pi|)$.*

There are already a number of studies on this action so we note that our version is stronger than the acylindricity of Hamenstädt [Ham3], weak acylindricity and the weak proper discontinuity of Bestvina-Fujiwara [BesFu] though weaker than, but not implied by, the acylindricity theorem of Bowditch [Bow2] for the bounds we provide are non-uniform. However, the bounds we give are both explicit and computable and this is our reward.

Once we prove the next supporting lemma, which may well be of independent interest, Thereom 23 will quickly follow. For a non-negative integer r and two curves α and β, we define $T_r(\alpha, \beta)$ as the collection of all tight multigeodesics starting in $B_r(\alpha)$ and ending in $B_r(\beta)$, and whose ends are single curves.

Lemma 24 *There is an explicit function $F_5 : \mathbb{N}^4 \longrightarrow \mathbb{N}$ such that the following holds. Let r be any non-negative integer and let α and β be any two filling curves such that $d(\alpha, \beta) \geq 14r + 5$. Then, $|\bigcup T_r(\alpha, \beta) - (B_{7r}(\alpha) \cup B_{7r}(\beta))|$ is at most $F_5(\iota(\alpha, \beta), r, genus(\Sigma), |\Pi|)$.*

Proof of Lemma 24. We make use of the idea of m-pulse, and we remark that this makes just as much sense in the context of multipaths. We can similarly construct a train track τ relative to the ending multicurves and the branch set of τ inherits a family of measures satisfying a coarse switch condition. Thus, Lemma 15 holds in this setting. Having gone to great pains in writing down explicit bounds in the previous sections, we are justified in being a little more relaxed about this (losing nothing in rigour while being more transparent). Instead of chasing explicit bounds, we make statements like "can be uniformly and explicitly bounded". In particular, any bounds we need will derive directly from Lemma 15.

For any geodesic $z \in T_r(\alpha, \beta)$, suppose there exists a tight multipath connecting α to β, containing $z - (B_{7r}(\alpha) \cup B_{7r}(\beta))$, and such that no two vertices spanning a subpath of length at least three contain two curves, one in each, of distance at most two. Denote a shortest such multipath by q. Then, Lemma 15 implies that each vertex of q has intersection number with α and intersection number with β uniformly and explicitly bounded in terms of $\iota(\alpha,\beta), genus(\Sigma)$ and $|\Pi|$. Each vertex of $z - (B_{7r}(\alpha) \cup B_{7r}(\beta))$ has intersection number with α and intersection number with β similarly bounded.

All that remains is to establish the existence of such multipaths. For any two multipaths p and p' with common ends, we shall write $p \longrightarrow p'$ if there exist consecutive vertices $u_0, u_1, \ldots, u_i$ of p, with $i \geq 2$, and consecutive vertices $w_0, \ldots, w_j$ of p', with $j \leq min\{2, i-1\}$, such that $p - \{u_0, u_1, \ldots, u_i\} = q - \{w_0, \ldots, w_j\}$, $w_0 \subseteq u_0$ and $w_j \subseteq u_i$. For any locally 3-geodesic multipath q and a multipath q' with the same ends but distinct from q, we write $q \Longrightarrow q'$ if q' results from tightening q at a single vertex. Notice if q is locally 3-geodesic and $q \Longrightarrow q'$, we do not presume q' to be locally 3-geodesic. However, if $q \Longrightarrow q'$ then q and q' agree on all but one vertex.

Choose any geodesic $z \in T_r(\alpha, \beta)$ and form a new multipath p by concatenating a tight multigeodesic connecting α to the end of z in $B_r(\alpha)$, z and a tight multigeodesic connecting the end of z in $B_r(\beta)$ to β. Now take any sequence $p = p_0^1 \longrightarrow p_1^1 \longrightarrow \ldots \longrightarrow p_{k_1}^1 \Longrightarrow p_0^2 \longrightarrow p_1^2 \longrightarrow \ldots \longrightarrow p_{k_{t-1}}^{t-1} \Longrightarrow p_0^t$, or $\ldots \longrightarrow p_{k_{t-1}}^{t-1} = p_0^t$ if $p_{k_{t-1}}^{t-1}$ is already tight, that is maximal with respect to inclusion, noting that the length of this sequence must be at most $4r$. Moreover, p_0^t is tight and contains no subpaths of length at least three whose ends contain two curves, one in each, of distance at most two. By construction, $z - (B_{7r}(\alpha) \cup B_{7r}(\beta)) \subseteq p_0^t - (B_{7r}(\alpha) \cup B_{7r}(\beta))$. $\diamondsuit$

By improving on this argument, or appealing to the hyperbolicity of the curve compex, we can probably do better than $7r$. Nonetheless, this is already more than adequate for our needs. In the statement of Theorem 23, we now take R to be $14r+5$ and we connect α and β by a tight multigeodesic z. For any such mapping class h, the h-translate of z belongs to $T_r(\alpha, \beta)$. Now, $|\bigcup T_r(\alpha,\beta) - (B_{7r}(\alpha) \cup B_{7r}(\beta))|$ is uniformly bounded in terms of $r, \iota(\alpha, \beta)$ and topology. Moreover, the length of $z - (B_{7r}(\alpha) \cup B_{7r}(\beta))$ is at least three. Combining these two facts with with a well known result, stating that the stabiliser in $Map(\Sigma - \Pi)$ of a pair of curves at least distance three in the curve graph is uniformly bounded in terms of $genus(\Sigma)$ and $|\Pi|$, we complete the proof of Theorem 23. $\diamondsuit$

2.10 Weak proper discontinuity

We recall the notion of weak proper discontinuity (WPD) due to Bestvina and Fujiwara [BesFu], and then show how to give a direct computability proof that the action of the extended mapping class group on the curve graph satisfies WPD for $\Sigma - \Pi$ non-exceptional.

Definition 25 *We say that the action of a group Γ on a metric space X satisfies WPD if Γ is not virtually cyclic, if Γ contains at least one element that acts on X as a loxodromic and if, for every loxodromic $h \in \Gamma$, every $x \in X$ and every*

$C > 0$, there exists a positive integer N such that the set

$$\{g \in \Gamma : d(x, g(x)) \leq C, d(h^N x, gh^N x) \leq C\}$$

is finite.

We denote this set by $\Gamma(x, h, C, N)$. Comparing the definitions of acylindricity and WPD, we note that the constant R and exponent N play similar roles (they ensure that we are shifting long geodesics). Notice that N depends on the choice of pseudo-Anosov whereas R does not. For this reason, we see that the weak acylindricity statement is stronger than WPD.

Theorem 26 *i). (Bestvina-Fujiwara) Suppose that $\Sigma - \Pi$ is non-exceptional. Then, the action of the mapping class group on the curve graph satisfies WPD.*
ii). (S) In addition, the cardinality of the set $\Gamma(\alpha, h, C)$ is bounded in terms of the topology of $\Sigma - \Pi, C, N$ and $\iota(\alpha, h^N \alpha)$.

Proof The first two points in the definition of WPD are straightforward. For instance, every infinite cyclic subgroup of $Map^*(\Sigma - \Pi)$ is contained in a free subgroup of rank 2 or a free abelian subgroup of rank 2 as an infinite index subgroup and so the mapping class group is not virtually cyclic, and this action certainly has loxodromics for, according to Masur-Minsky [MaMi1], all mapping classes containing a pseudo-Anosov homeomorphism act as loxodromics.

To establish the third point we use a uniform version of a covering distance result due to Hempel (see Theorem 2.5 in [Hemp]). For this, we must recall his idea of a covering distance. A covering space $p : S \longrightarrow \Sigma - \Pi$ separates two curves α and β in $\Sigma - \Pi$ if there is a component in the p-preimage of any simple realisation $a \in \alpha$ disjoint from a component in the p-preimage of a simple realisation $b \in \beta$. A finite covering p is called sub-solvable if p can be factored as a composition of cyclic coverings. The *covering distance $cd(\alpha, \beta)$* between two curves α and β is precisely

$$1 + min\{k : \Sigma - \Pi \text{ has a degree } 2^k \text{ subsolvable covering separating } \alpha, \beta\}.$$

Among other things, Hempel proves that the covering distance is always a lower bound for distances in the curve graph.

Lemma 27 *(Hempel) Suppose that $\Sigma - \Pi$ is non-exceptional. For any pseudo-Anosov mapping class $h \in Map(\Sigma - \Pi)$ and any non-negative integer m, there exists a non-negative integer N such that, for each curve α, we have $cd(\alpha, h^N \alpha) \geq m$.*

Proof of Lemma 27. This a version of the argument given by Hempel. Let H be the intersection of all 2^m index subgroups of $\pi_1(\Sigma - \Pi)$. Then H is normal in $\pi_1(\Sigma - \Pi)$ and we let $p : Z \longrightarrow \Sigma - \Pi$ be the cover associated with H. Now H is characteristic and in particular is invariant under the homomorphism h_*. Any two curves separated by a degree 2^m covering are therefore separated by p as well.

Note h lifts to a pseudo-Anosov homeomorphism $\widetilde{h} : Z \longrightarrow Z$. Let $\widetilde{\alpha}$ denote any component of $p^{-1}(\alpha)$. As h is not periodic so there exists a positive integer N, at most $\kappa(Z)+1$, such that $\widetilde{h}^N\widetilde{\alpha}$ intersects every curve component of $p^{-1}(\alpha)$. Since our cover p is regular, we see that p fails to separate α and $h^N\alpha$. Hence every 2^m sub-solvable cover fails to separate α and $h^N\alpha$. We conclude that the covering distance $cd(\alpha, h^N\alpha)$ is at least m. $\diamondsuit$

Returning to the proof of Theorem 26, we fix the value of m from Lemma 27 to be $14C + 5$. For every curve α we have $d(\alpha, h^N\alpha) \geq cd(\alpha, h^N\alpha) \geq 14C + 5$. The finiteness of the given set $\Gamma(\alpha, h^N, C)$ now follows from Lemma 24. Indeed, we can go on to add uniform and explicit bounds in terms of the topology of $\Sigma - \Pi, C, N$ and the intersection number $\iota(\alpha, h^N\alpha)$. $\diamondsuit$

Chapter 3

Combinatorial rigidity

ABSTRACT: The significant advances made by Ivanov, McCarthy, Korkmaz, Luo and Irmak, among others, on the relationship between the curve complex and mapping class group lead us to the natural though, until now, unaddressed question of whether a local embedding on curve complexes is induced by a surface homeomorphism. In this chapter, using different arguments we prove that local embeddings betwen two curve complexes whose complexities do not increase from domain to codomain are induced by surface homeomorphism. This is our first main result (given as Theorem 3). From this we deduce our second (given as Theorem 4), a strong local co-Hopfian result for almost all mapping class groups. It follows that such mapping class groups do not admit a faithful action on another curve complex of the same or lower dimension.

KEYWORDS: Curve complex, mapping class group.

3.1 Introduction

The curve complex $\mathcal{C}(\Sigma)$ associated to a surface Σ was introduced by Harvey [Harv] to encode the large scale geometry of Teichmüller space, and serves as a good model space for the mapping class group as well playing a central role in the proof of Brock-Canary-Minsky [BroCM] of Thurston's ending lamination conjecture.

We start by defining the curve complex, and throughout our surfaces will be compact, connected and orientable. In contrast to Chapter 2 our arguments oblige us to look at holed rather than punctured surfaces. We say that a simple loop in Σ is trivial if it bounds a disc and peripheral if it bounds an annulus whose other boundary component belongs to $\partial\Sigma$. A curve in Σ is a free homotopy class of a non-trivial and non-peripheral simple loop and we denote the set of these by $X(\Sigma)$; we say that two curves intersect minimally if they intersect once or they intersect twice with zero algebraic intersection and refer to either as the type of minimal intersection. We will define the *complexity* of Σ, denoted $\kappa(\Sigma)$, as equal to the maximal number of distinct and disjoint curves that can be realised simultaneously. When $\kappa(\Sigma) \geq 2$, the curve graph is the graph whose vertex set is $X(\Sigma)$ and we deem two distinct curves to span an edge if and only if they can be realised disjointly in Σ. When $\kappa(\Sigma) = 1$, we say that two distinct curves are joined by an edge if and only if they intersect minimally. The *curve complex* associated to Σ is the curve graph when $\kappa(\Sigma) = 1$, making it isomorphic to a Farey graph, and the flag simplicial complex whose 1-skeleton is the curve graph when $\kappa(\Sigma) \geq 2$ and has dimension precisely $\kappa(\Sigma) - 1$.

For each curve α we denote by $X(\alpha)$ the set of all curves distinct and disjoint from α in Σ, that is the vertex set of the link of α. This link is always connected whenever $\kappa(\Sigma)$ is at least three, and whenever $\kappa(\Sigma)$ is two any two elements of $X(\alpha)$ may be "connected" by a finite sequence of curves in which consecutive curves have minimal intersection.

In this chapter, we shall be discussing embeddings between two curve complexes whose complexities do not increase from domain to codomain and we shall find that these are all induced by surface homeomorphism, so long as we place a necessary but consistent hypothesis in one sporadic case. The argument we give is by an induction on complexity and requires little more than the connectivity of links in the curve complex over and above this. As such, our approach does not discriminate topologically. Moreover, we actually only require the local injectivity of an embedding and we shall say more on this towards the end of this section. All told this generalises the seminal work of Ivanov [Iva1] for automorphisms of curve complexes associated to surfaces of genus at least two, further treated by Korkmaz [Kor1] for holed spheres and holed tori and later by Luo [Luo1] in all cases. Making use of their combined result, Margalit [Marg] establishes the analogue for automorphisms of a second surface complex called the pants complex. There are several analogues for other surface complexes using Ivanov's argument, see for example Schmutz Schaller's [Sch].

Our first result is stated as follows.

Theorem 1 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces of positive complexity such that the complexity of Σ_1 is at least that of Σ_2, and that when they have equal complexity at most three they are homeomorphic or one is the three-holed torus. Then, any simplicial embedding from $\mathcal{C}(\Sigma_1)$ to $\mathcal{C}(\Sigma_2)$ (preserving the separating type of each curve when the two surfaces are homeomorphic to the two-holed torus) is induced by a surface homeomorphism.*

This covers all possibilities. We remind ourselves that there exist isomorphisms between the curve complex of the closed surface of genus two and the six-holed sphere, the two-holed torus and the five-holed sphere and finally the one-holed torus and the four-holed sphere and that there exists an automorphism of the curve complex associated to the two-holed torus that sends a non-separating curve to an outer curve (see [Luo1] for more details). These are examples of embeddings not induced by a surface homeomorphism. Finally, we point out that there exist embeddings on curve complexes with complexity increasing from domain to codomain not induced by a surface embedding: Easy examples are provided by taking Σ_1 to be a subsurface of Σ_2 of cocomplexity at least one and changing the induced embedding on curve complexes by instead taking just one curve of Σ_1 to the curve lying outside of Σ_1.

The work on Theorem 1 was in part motivated by the work of Irmak on simplicial self-maps of a curve complex that send intersecting curves to intersecting curves, namely superinjective maps. In [Irm1], the author shows that a superinjective self-map is induced by surface homeomorphism provided that the surface has genus at least two and at least two boundary components or genus at least three. From this, Irmak deduces a "virtual co-Hopfian" result for the corresponding mapping class group. The *mapping class group* $Map(\Sigma)$ is the group of all self-homeomorphisms of the surface Σ, up to homotopy. This is sometimes known as the extended mapping class group, for it contains the group of orientation preserving mapping classes as an index two subgroup. Some of its other subgroups, namely the Johnson kernel and the Torelli group, are of wide interest (see, respectively, Brendle-Margalit [BreMarg] and Farb-Ivanov [FarIva], and references therein).

The mapping class group has a natural simplicial action on the curve complex determined by first lifting a curve to a representative loop and then taking the free homotopy class of the image under a representative homeomorphism. The kernel of this action is almost always trivial; the only exceptions lie in low complexity, where this kernel is isomorphic to $\mathbb{Z}_2$ and generated by the hyperelliptic involution when Σ is the one-holed torus, the two-holed torus or the closed surface of genus two or isomorphic to $\mathbb{Z}_2 \oplus \mathbb{Z}_2$ and generated by two hyperelliptic involutions when the four-holed sphere (this is due to Birman [Bir] and Viro [V]). For a detailed account of the mapping class group and its subgroups, see also Ivanov [Iva2], [Iva4].

Theorem 1 implies the following strong co-Hopfian result for mapping class groups. Among other things, it follows that the commensurator group of any such mapping class group is isomorphic to the same mapping class group, that the outer automorphism group of any such mapping class group is trivial and that such mapping class groups do not admit a faithful action on another curve complex of the same or lower dimension. Moreover, there is only one faithful action by any such mapping class group on its curve complex, up to conjugation.

Theorem 2 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2 and at least two, and that whenever they both have complexity equal to three they are homeomorphic or one is the three-holed torus and when they both have complexity two they are both homeomorphic to the five-holed sphere. Then, a finite index subgroup H of $Map(\Sigma_1)$ injects into $Map(\Sigma_2)$ only if Σ_1 and Σ_2 are homeomorphic. In which case, unless both surfaces are homeomorphic to the closed surface of genus two, every such injection is the restriction of a unique inner automorphism of $Map(\Sigma_1)$. In this exceptional case, every such injection is the restriction of an inner automorphism of $Map(\Sigma_1)$ possibly composed with a hyperelliptic involu-*

tion.

This is a generalisation of a result of Ivanov-McCarthy (Theorem 4 from [IvaMcCar]) where the two authors consider injections defined on mapping class groups associated to surfaces of positive genus. The superinjectivity theorem of Irmak implies Theorem 2 when the two surfaces under consideration are homeomorphic and have genus at least two with two boundary components or genus at least three. An extension of this result to cover the surfaces of genus two with at most one hole can be found in [Irm2], where the author also finds non-inner injections for the closed surface of genus two. Behrstock-Margalit prove an extension of Irmak's theorem to homeomorphic surfaces of complexity at least three. Their arguments can be found in [BehrMarg], where they also find a commensurator of the mapping class group of the two-holed torus not induced by an inner automorphism. The approach we need for Theorem 2 follows that given by Ivanov, translating an injection on a finite index subgroup to an embedding on curve complexes. This is now a well-established strategy on which we have very little to add, and a thorough account can be found in the work of Bell-Margalit [BellMarg]. For completeness, and for its application to Theorem 4, we recall the argument in §3.3.

Though our arguments are phrased in terms of embeddings, they only ever need the simplicial and local injectivity properties of such maps. We can therefore record the following generalisation of Theorem 1, the first of two main results. Recall that a *star* in the curve graph is the union of all edges incident on a common vertex.

Theorem 3 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces of positive complexity such that the complexity of Σ_1 is at least that of Σ_2, and that when they have equal complexity at most three they are homeomorphic or one is the three-holed torus. Then, any simplicial map from $\mathcal{C}(\Sigma_1)$ to $\mathcal{C}(\Sigma_2)$ injective on every star (and preserving the separating type of each curve when both surfaces are homeomorphic to the two-holed torus) is induced by a surface homeomorphism.*

Again, this covers all possibilities. We remark that proving a local embedding is induced by a surface homeomorphism would appear the most direct way of seeing that it must also be a global embedding. From Theorem 3 we can deduce, using the same strategy of Ivanov, the following local version of Theorem 2. Among other things, it follows that a self-homomorphism of a mapping class group that is injective on every curve stabiliser is the restriction of a unique inner automorphism.

Theorem 4 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2 and at least two, and that* *whenever they both have complexity equal to three they are homeomorphic or one is the three-holed torus and when they both have complexity two they are both homeomorphic to the five-holed sphere. Suppose that H is a finite index subgroup of the mapping class group $Map(\Sigma_1)$. Then, unless both surfaces are homeomorphic to the closed surface of genus two, every homomorphism of H into $Map(\Sigma_2)$ that is injective on every curve stabiliser in H is the restriction of a unique inner automorphism of $Map(\Sigma_1)$. For the one exceptional case, any such homomorphism is the restriction of an inner automorphism possibly composed with a hyperelliptic involution.*

Note that the existence of such a homomorphism therefore implies the two surfaces are homeomorphic. Investigations into homomorphisms from a mapping class group associated to a closed surface of genus at least one to another mapping class group associated to a closed surface of smaller genus have been made by Harvey-Korkmaz [HarvKor], the authors finding that every such homomorphism has finite image. Their approach seems to make essential use of the existence of torsion in mapping class groups and, as mapping class groups are virtually torsion free, it would be of much interest to find a way around this so as to consider finite index subgroups. Somewhat different rigidity questions have been addressed by Farb-Masur [FarMa], the authors finding that homomorphisms from lattices in semi-simple Lie groups with the mapping class group as target also have finite image.

A complete proof Theorem 4, adapting some of the work of Ivanov and Ivanov-McCarthy, is given in §3.3, and a couple of new observations are made.

3.2 Curve complex embeddings

For any compact, connected and orientable surface Σ the *complexity* $\kappa(\Sigma)$ of Σ is defined to be equal to $3genus(\Sigma) + |\partial\Sigma| - 3$. This is slightly non-standard, since usually complexity is taken to be equal to the simplicial dimension of the curve complex, but the additivity of κ best suits our induction argument. By way of example, the one-holed torus and the four-holed sphere are the only surfaces of complexity one, the two-holed torus and the five-holed sphere are the only surfaces of complexity two and the closed surface of genus two, the three-holed torus and the six-holed sphere are the only surfaces of complexity three. On occasion we refer to these as the *low complexity* surfaces.

In what follows, we shall continue to abuse notation slightly by viewing each curve as a vertex, as a class of loops and as a simple loop already realised in Σ. Our interpretation will be apparent from the context. We say that a curve is *separating* if its complement is not connected, or equivalently if it is null homologous, and say it is *non-separating* otherwise. We say that a curve is an *outer curve* if it is separating and if it bounds a two-holed disc (equivalently, a three-holed sphere). These are usually known as boundary curves in the literature, but here we need to avoid confusing these with the components of $\partial\Sigma$.

A pair of distinct and disjoint non-separating curves is a *bounding pair* if their complement is not connected. We say that a separating curve or a bounding pair has a given complexity if one of its two complementary components has this complexity. Thus, they may simultaneously hold two different complexities. A *pants decomposition* of Σ is a maximal collection of disjoint curves, and a *pair of pants* in Σ is an essential subsurface homeomorphic to a three-holed sphere. We say that two curves are *adjacent* in a pants decomposition P if they appear in the boundary of a three-holed sphere complementary to P. This is slightly unfortunate terminology that seems to be a long way to becoming standard; we hope that any confusion between adjacency in the curve complex and adjacency in a pants decomposition will be obviated by the context. Lastly, a collection of distinct and disjoint curves shall be referred to as a *multicurve*.

The structure of our argument is broadly as follows. We establish a short list of topological properties verified by any embedding on curve complexes from which we easily deduce, among other things, that the existence of such an embedding implies the two surfaces have equal complexity and then, with more work, almost always means the two surfaces under consideration are homeomorphic. For the time being, we refer to embeddings between two apparently distinct curve complexes as *cross-embeddings*. Dealing with embeddings in low complexity typically requires individual arguments and it therefore streamlines our work if we do this separately, as we do in Lemma 14. The proof of Theorem 1 is then completed by an induction on complexity, where we cut the surface along a curve. As embeddings preserve the topological type of a curve, the resulting surfaces are again homeomorphic. For the induction argument to pass through complexity one (sub)surfaces, we will need to show that embeddings preserve minimal intersection.

We start by showing, in turn, that embeddings send pants decompositions to pants decompositions, they preserve a form of small intersection and they preserve adjacency *and* non-adjacency in a pants decomposition. Notice that in the subsequent lemmata, an appropriate local injectivity hypothesis will suffice.

Lemma 5 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2. Then, any simplicial embedding ϕ from $\mathcal{C}(\Sigma_1)$ to $\mathcal{C}(\Sigma_2)$ sends pants decompositions to pants decompositions.*

Proof This follows for complexity reasons and because ϕ is injective and simplicial. $\diamondsuit$

To make sense of the following lemma, we need to define what we mean by the *subsurface of Σ filled by two curves α and β*. Letting $N(\alpha \cup \beta)$ denote a regular neighbourhood of $\alpha \cup \beta$ in Σ, we augment $N(\alpha \cup \beta)$ by taking its union with all the complementary discs whose boundary is contained in $N(\alpha \cup \beta)$ and all the complementary annuli with one boundary component in $\partial\Sigma$ and the other in $N(\alpha \cup \beta)$. The resulting subsurface of Σ is well-defined up to homotopy,

π_1-injects and is what we mean by the subsurface filled by α and β. Whenever a third curve enters the subsurface filled by two curves, it must intersect at least one of these two curves.

Lemma 6 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2. Let ϕ be any simplicial embedding from $C(\Sigma_1)$ to $C(\Sigma_2)$ and let α, β be any two curves in Σ_1 that fill either a four-holed sphere or a one-holed torus. Then, $\phi(\alpha)$ and $\phi(\beta)$ fill either a four-holed sphere or a one-holed torus in Σ_2.*

Proof Let Q be any maximal multicurve in Σ_1 such that each curve is disjoint from both α and β. For complexity reasons, $\phi(Q)$ is a maximal multicurve disjoint from both $\phi(\alpha)$ and $\phi(\beta)$. In particular, as ϕ is injective and simplicial so $\phi(\alpha)$ and $\phi(\beta)$ must together fill either a four-holed sphere or a one-holed torus. $\diamondsuit$

We shall say that two curves have *small intersection* if they together fill either a four-holed sphere or a one-holed torus, and refer to either as the *type* of the small intersection. Any two curves that intersect minimally have small intersection, but the converse does not hold.

Lemma 7 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2. Let P be any pants decomposition of Σ_1 and let ϕ be any simplicial embedding from $C(\Sigma_1)$ to $C(\Sigma_2)$. Then, any two curves adjacent in P are sent by ϕ to two adjacent curves in $\phi(P)$ and two curves in P that are not adjacent in P are sent by ϕ to two curves not adjacent in $\phi(P)$.*

Proof The first part follows from Lemma 6: For any two curves α_1 and α_2 adjacent in P, there exists a curve δ having small intersection with both and disjoint from every other curve in P. This is preserved under ϕ and so $\phi(\alpha_1)$ and $\phi(\alpha_2)$ are adjacent in $\phi(P)$.

Similarly, if two curves α_1, α_2 are not adjacent in P we can find two disjoint curves δ_1, δ_2 such that δ_1 has small intersection with α_1 but is disjoint from α_2 and δ_2 has small intersection with α_2 but is disjoint from α_1 and both δ_1, δ_2 are disjoint from every other curve in P. If $\phi(\alpha_1)$ and $\phi(\alpha_2)$ are adjacent in $\phi(P)$ then $\phi(\delta_1)$ and $\phi(\delta_2)$ must intersect. As ϕ is simplicial, this is a contradiction. $\diamondsuit$

The import of Lemma 5, Lemma 6 and Lemma 7 is perhaps best understood by associating to a pants decomposition P a certain graph. The vertices of this graph are the curves in P, and any two distinct vertices span an edge if and only if they correspond to adjacent curves in P. Lemma 7 not only tells us that any embedding ϕ induces a map between adjacency graphs, but that this

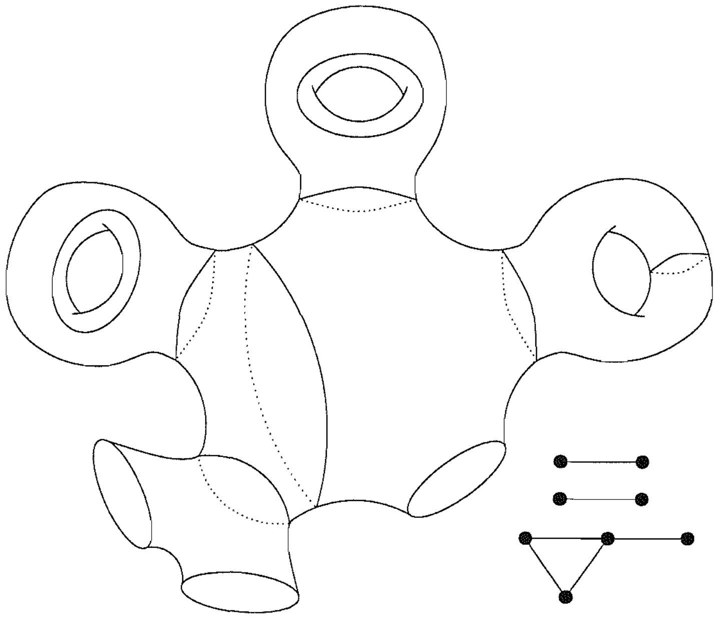

Figure 3.1: A codimension 1 multicurve, with its adjacency graph.

map is actually an isomorphism. Cut points in the graph correspond to non-outer separating curves, and a complexity of a non-outer separating curve in P corresponds to the number of vertices on a side of the corresponding cut point.

This graph, and the ideas bound by Lemma 7, were independently and simultaneously discovered by Behrstock-Margalit. Their approach can be found in [BehrMarg] and the arguments they give will deal with all superinjective maps for two homeomorphic surfaces of complexity at least three. From this they also deduce that the commensurator group of a mapping class group is isomorphic to the same mapping class group. We both refer to such a graph as an *adjacency graph.*

We can just as well speak of an adjacency graph associated to a multicurve Q, in which the vertices again correspond to the curves in Q and any two vertices are declared adjacent if their corresponding curves border a common pair of pants in the surface complement of Q. There is a subtle point to be made here, namely that the complementary graph of a vertex in a pants adjacency graph will *not* in general be the adjacency graph of the multicurve that results by removing the corresponding curve from the pants decomposition. It will however be the adjacency graph that results from cutting the surface along this curve. By way of example, on removing a curve α from a pants decomposition P the curves that together bound the complementary four-holed sphere will not necessarily be adjacent in the adjacency graph of $P - \{\alpha\}$. (See Figure 3.1 for an illustrated example.) This observation will be important later when we come to look at outer curves. It does however hold that a curve complex embedding induces an isomorphism between multicurve adjacency graphs.

Lemma 8 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2. Let Q be any multicurve of Σ_1 and let ϕ be any simplicial embedding from $C(\Sigma_1)$ to $C(\Sigma_2)$. Then, ϕ induces an isomorphism from the adjacency graph of Q to the adjacency graph of $\phi(Q)$.*

Proof We make use of Lemma 7. Extend Q to a pants decomposition P of Σ_1. If two curves are adjacent in Q then they either border a pair of pants with a third curve from Q or they border a pair of pants meeting $\partial\Sigma$. This remains so in P, and is preserved on applying ϕ. To show non-adjacency is preserved, consider any two curves not adjacent in Q and arrange for them to be non-adjacent in P. This is preserved under ϕ. $\diamondsuit$

As embeddings between curve complexes induce isomorphisms on adjacency graphs and graph isomorphisms send cut points to cut points, so embeddings must send non-outer separating curves to non-outer separating curves.

Lemma 9 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2. Then, any simplicial embedding ϕ from $C(\Sigma_1)$ to $C(\Sigma_2)$ sends every non-outer separating curve to a non-outer separating curve of the same complexity.*

A similar argument gives an analogue for bounding pairs of positive complexities when we note that such a pair in a pants decomposition corresponds to a pair of non-cut points in the adjacency graph whose complementary graph is not connected.

Lemma 10 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2. Then, any simplicial embedding ϕ from $C(\Sigma_1)$ to $C(\Sigma_2)$ sends bounding pairs to bounding pairs of the same complexity.*

We use Lemma 7 and the adjacency graph to distinguish between non-separating and outer curves.

Lemma 11 *Suppose Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2 and that whenever they have equal complexity at most three they are homeomorphic and not the two-holed torus. Let $\phi : C(\Sigma_1) \longrightarrow C(\Sigma_2)$ be any simplicial embedding. Then, ϕ takes non-separating curves to non-separating curves.*

Proof We note that the ϕ-image of a non-separating curve can never be a non-outer separating curve, for otherwise we see a non-cut point sent to a cut point in some pants adjacency graph. Suppose that α is a non-separating curve in Σ_1. When $\kappa(\Sigma_1)$ is at least four we can find a pants decomposition P extending α in which α corresponds to a vertex in the adjacency graph of P of valence three or four. As ϕ induces an isomorphism on the adjacency graph, so $\phi(\alpha)$ must have the same valence. As outer curves only ever correspond to vertices of valence at most two, so $\phi(\alpha)$ can only be non-separating.

With the exception of the two-holed torus, all cases in which Σ_1 has complexity at most two hold since there is only ever one type of curve. In complexity three, when Σ_1 is the six-holed sphere our claim holds vacuously and when Σ_1 is the closed surface of genus two our claim follows from Lemma 9 by noting that every pants decomposition contains at most one separating curve.

Really, the only non-trivial case in low complexity is that of Σ_1 and Σ_2 both homeomorphic to the three-holed torus. In which case, there are only two pants adjacency graphs, up to isomorphism, but three different pants decompositions, up to the action of the mapping class group. For this reason, we need to argue differently. If there is a non-separating curve sent by ϕ to an outer curve, then there is an outer curve α sent by ϕ to a non-separating curve. To see this, extend this non-separating curve to a pants decomposition containing a non-outer separating curve. By appealing to Lemma 9, we see that the third curve in this pants decomposition will suffice. Now extend α to a second pants decomposition containing two non-separating curves δ_1 and δ_2. The ϕ-image of at least one of these, say δ_1, is again a non-separating curve. Choose any two disjoint curves γ_1, γ_2 in Σ_1 that have small intersection with δ_1 and α but disjoint from α and δ_1, respectively. Now $\phi(\delta_1)$ and $\phi(\alpha)$ border a common pair of pants in Σ_2 invaded by $\phi(\gamma_1)$ and $\phi(\gamma_2)$. We see that the ϕ-images of both γ_1 and γ_2 are forced to intersect, and this is a contradiction. $\diamondsuit$

Lemma 12 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2, and that whenever they have equal complexity at most three they are homeomorphic and not the two-holed torus. Then, any simplicial embedding ϕ from $\mathcal{C}(\Sigma_1)$ to $\mathcal{C}(\Sigma_2)$ sends outer curves to outer curves.*

Proof We note that this holds vacuously when $|\partial\Sigma_1|$ is at most one. We also remark, again paranthetically, that when $|\partial\Sigma_1|$ is at least three any outer curve represents an extreme point in the adjacency graph of some pants decomposition. Graph isomorphisms send extreme points to extreme points and, as ϕ induces an isomorphism on the adjacency graph, so ϕ must send outer curves to outer curves.

Suppose, for contradiction, that α is an outer curve in Σ_1 sent by ϕ to a non-outer curve in Σ_2. We note that $\phi(\alpha)$ can not be a separating curve, for α can never correspond to a cut point in a pants adjacency graph, and so $\phi(\alpha)$

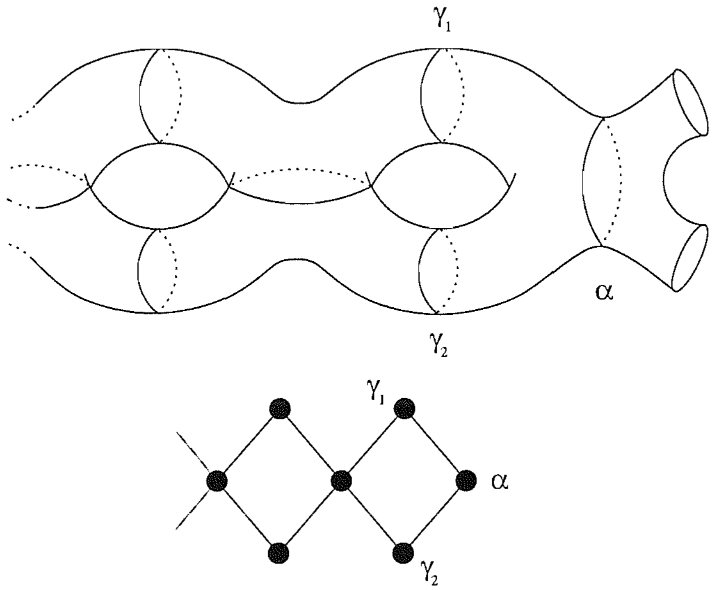

Figure 3.2: A convenient extension of α to a pants decomposition.

must be a non-separating curve. If $\kappa(\Sigma_1)$ at least four then we can extend α to a pants decomposition P in which the two curves adjacent to α, denoted γ_1 and γ_2, are not adjacent in the adjacency graph of $P - \{\alpha\}$. Using Lemma 10, $\phi(\gamma_1)$ and $\phi(\gamma_2)$ either both separate or they form a bounding pair. As α is an outer curve, we note that γ_1 and γ_2 are adjacent in P. However, this does not remain so on applying ϕ and we contradict Lemma 7. (See Figure 3.2.)

Once more, the only remaining non-trivial case in low complexity is that of both surfaces homeomorphic to the three-holed torus. Suppose that α is an outer curve sent to a non-separating curve by ϕ. Extend α to a pants decomposition P containing a separating curve. Then the non-separating curve in P is sent to an outer curve by ϕ, and this is contrary to Lemma 11. $\diamondsuit$

We can now safely state that any such embedding will send separating curves to separating curves. It follows that small and peripheral subsurfaces can not change topological type under embeddings.

Lemma 13 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2 and that when both have equal complexity at most three they are homeomorphic and not the two-holed torus. Let Z be any essential π_1-injective subsurface of Σ_1 of complexity one and bordered by a single curve β. Then, for any simplicial embedding ϕ from $C(\Sigma_1)$ to $C(\Sigma_2)$, the π_1-injective minimal subsurface $\phi(Z)$ of Σ_2 filled by $\phi(X(Z))$ is homeomorphic to Z.*

Proof Note Z may be represented as either a four-holed sphere meeting $\partial\Sigma$ in three components or a one-holed torus whose boundary is a curve in Σ. Such a change in topology would otherwise force ϕ to send a non-separating curve to an outer curve or an outer curve to a non-separating curve, contrary to Lemma 11 and Lemma 12 respectively. $\diamondsuit$

We can finally rule out cross-embeddings, and here after we regard the two surfaces as being homeomorphic and denote them by Σ.

Lemma 14 *Suppose that Σ_1 and Σ_2 are two compact and orientable surfaces such that the complexity of Σ_1 is at least that of Σ_2, and that whenever they have complexity at most two they are homeomorphic and whenever they have complexity equal to three they are either homeomorphic or one is the three-holed torus. Then, there is a simplicial embedding $\phi : C(\Sigma_1) \longrightarrow C(\Sigma_2)$ only if Σ_1 and Σ_2 are homeomorphic.*

Proof The existence of such an embedding implies the complexities $\kappa(\Sigma_1)$ and $\kappa(\Sigma_2)$ are equal. When the two surfaces have complexity at least four, we know that any such embedding must send separating curves to separating curves and non-separating curves to non-separating curves. We recall that the size of a maximal collection of distinct and disjoint separating curves in Σ_1 is precisely $2genus(\Sigma_1)+|\partial\Sigma_1|-3$. By our earlier work, this is precisely $2genus(\Sigma_2)+|\partial\Sigma_2|-3$. The only possible solution is $genus(\Sigma_1) = genus(\Sigma_2)$ and $|\partial\Sigma_1| = |\partial\Sigma_2|$. That is, Σ_1 and Σ_2 are homeomorphic.

Among the low complexity surfaces, there are no embeddings from the curve complex associated to the six-holed sphere or the closed surface of genus two to the three-holed torus. To see this, extend an outer or non-separating curve α in Σ_1 to a pants decomposition P consisting only of outer or non-separating curves, respectively, and choose a separating curve β disjoint from both curves in $P - \{\alpha\}$ and having small intersection with α. We assume that if any curve in $\phi(P)$ is outer then it is $\phi(\alpha)$. Now $\phi(\beta)$ is a non-outer separating curve intersecting $\phi(\alpha)$ and it follows that $\phi(\beta)$ must intersect another curve in $\phi(P)$. This is a contradiction.

The remaining cases, namely from the curve complex of the three-holed torus to the curve complex of the six-holed sphere or the closed surface of genus two, are covered as follows. For any pants decomposition P in Σ_1 choose a non-outer separating curve β meeting only two curves in P. When Σ_2 is the six-holed sphere, each curve in P goes to an outer curve. By Lemma 6, small intersection is preserved. Now any non-outer separating curve in the six-holed sphere meets either only one curve or all three curves in a pants decomposition made up entirely of outer curves. It follows that $\phi(\beta)$ meets every curve in $\phi(P)$ and this is a contradiction. Given the isomorphism classification of curve complexes, this simultaneously deals with Σ_2 the closed surface of genus two. $\diamondsuit$

To allow the induction argument to pass through complexity one surfaces unhindered, we need the following lemma on minimal intersection in those subsurfaces bordered by a single curve. This relies on what is a well-established argument, first given by Ivanov [Iva1] for the intersection one property. Guided by this, Luo [Luo1] proves the analogue for intersection two with zero algebraic intersection.

Lemma 15 *Suppose that Σ is a compact and orientable surface of positive complexity and not homeomorphic to the two-holed torus. Suppose that Z is an essential subsurface of Σ of complexity one and bordered by a single curve β. Then, any simplicial embedding $\phi : \mathcal{C}(\Sigma) \longrightarrow \mathcal{C}(\Sigma)$ preserves minimal intersection and its type on $X(Z)$.*

Proof Note Z may be represented as either a four-holed sphere meeting $\partial\Sigma$ in three components or a one-holed torus whose boundary is a curve in Σ. In either case, let α_1 and α_2 be the two curves in Z intersecting minimally. Choose any two disjoint curves γ_1 and γ_2 such that γ_1 (respectively γ_2) intersects α_1 (respectively α_2) minimally and γ_1 (respectively γ_2) is disjoint from α_2 (respectively α_1) and both have small intersection with β. We now note that if $\phi(\alpha_1)$ and $\phi(\alpha_2)$ fail to intersect minimally then $\phi(\gamma_1)$ will intersect $\phi(\gamma_2)$ which is absurd. Since for any two of the five curves we consider there is a third disjoint from both, this conclusion holds even when ϕ is only assumed to be locally injective.

As β can not see any change in topology under ϕ, by Lemma 13, so the type of minimal intersection is also preserved. $\diamondsuit$

This closes our study of the topological properties of curve complex embeddings, and the promised induction argument now starts with a look at the Farey graph.

Lemma 16 *Every simplicial embedding on a Farey graph $\mathcal{F}$ to itself is an automorphism.*

Proof We note that each edge in $\mathcal{F}$ separates and belongs to exactly two 3-cycles and that such a map sends 3-cycles to 3-cycles. Thus any embedding ϕ on $\mathcal{F}$ induces an embedding ϕ^* on the dual graph. This graph is a tree in which every vertex has the same valence, hence the induced map is a surjection. It follows that every 3-cycle of $\mathcal{F}$ is contained in the image of ϕ. That is to say, ϕ is also a surjection. $\diamondsuit$

It is a well-known fact (indeed, it was known to Dehn [Dehn2]) that the automorphisms of $\mathcal{C}(\Sigma)$ are all induced by surface homeomorphisms when Σ is either a four-holed sphere or a one-holed torus. This completes the base case of the induction.

We now furnish the inductive step. Let $\phi : \mathcal{C}(\Sigma) \longrightarrow \mathcal{C}(\Sigma)$ be any embedding satisfying the hypotheses of Theorem 1. Let α be any curve in Σ. Our previous work on the topological properties of ϕ tells us that the complement of α and the complement of $\phi(\alpha)$ are homeomorphic. Therefore, after conjugating by a suitable mapping class, ϕ restricts to a self-embedding on the curve complex associated to each component of $\Sigma - \alpha$. The embeddings arising in this way are very natural for they inherit many of the properties verified by ϕ, for instance they also preserve the separating type of a curve. This is of particular relevance when cutting the surface Σ along a curve and finding a two-holed torus complementary component. In [Luo1], the author explains how to find automorphisms of the curve complex associated to the two-holed torus not induced by a surface homeomorphism. No such automorphism can arise as a restriction, nor can any embedding, as outer curves in this two-holed torus correspond to separating curves in Σ and we would otherwise contradict either Lemma 9 or Lemma 12, as appropriate.

Our inductive hypothesis therefore applies and it tells us that each restriction of ϕ associated to a positive complexity component of $\Sigma - \alpha$ is induced by a surface homeomorphism. In gluing back together by identifying the boundary components of $\Sigma - \alpha$ corresponding to α, we have a countable family of surface homeomorphisms distinguished up to homotopy where any two differ by a precomposed power of a Dehn twist around α. In particular, these all agree on the complement of α. We must somehow decide which one of these, if any, is appropriate.

What saves us is that this construction applies equally well for every curve in Σ, in particular any curve β adjacent to α. Between them, the families of homeomorphisms associated to α and associated to β determine a coset in $Map(\Sigma)$ of the free abelian rank two subgroup A generated by a single Dehn twist τ_α around α and a single Dehn twist τ_β around β. This coset is represented by a mapping class determined by any one of the homeomorphisms associated to α or to β. The family of homeomorphisms associated to α and the family of homeomorphisms associated to β therefore correspond to cosets in A of the cyclic subgroups generated by τ_α and τ_β, respectively. The images in the automorphism group of $\mathcal{C}(\Sigma)$ of these two cosets intersect in a single point.

We need to verify that for any three curves α, β_1 and β_2 such that α is adjacent to both β_1 and β_2, the action of a surface homeomorphism f_1 so-determined by α and β_1 agrees with that of a surface homeomorphism f_2 so-determined by α and β_2. To do this, we exploit the fact that the link of α is either connected or chain-connected (in the sense that between any two vertices there is a finite sequence of curves with consecutive curves intersecting minimally). First, we remark that α has the same image under f_1 and f_2, so $f_2^{-1}f_1$ fixes α. We can therefore express the mapping class $[f_2^{-1}f_1]$ as the commuting product of a Dehn twist power τ_α^n with a mapping class $[g]$ of Σ fixing α. By construction, f_1 and f_2 agree on the complement of α and so, by changing n appropriately, we can

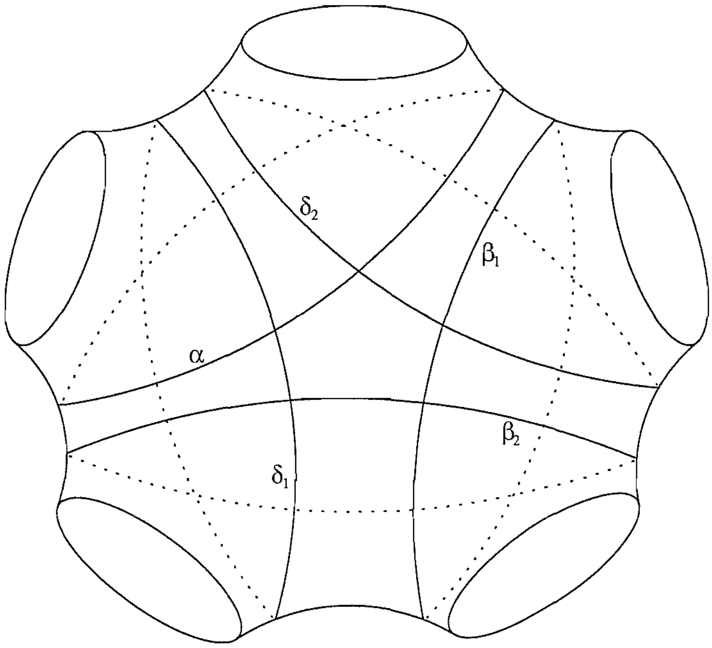

Figure 3.3: The case Σ a five-holed sphere.

and do take $[g]$ to be the trivial mapping class. Now suppose, for contradiction, that this Dehn twist power is non-trivial and take two disjoint curves δ_1, δ_2 in Σ where each has small intersection with α and is disjoint from β_1 and β_2, respectively. (See Figure 3.3, the five-holed sphere.) Then, $\iota([f_1]\delta_1, [f_2]\delta_2) = \iota([f_2][g]\tau_\alpha^n\delta_1, [f_2]\delta_2) = \iota([g]\tau_\alpha^n\delta_1, \delta_2) = \iota(\tau_\alpha^n\delta_1, \delta_2) \geq |n| > 0$. However, $\phi(\delta_i)$ and $[f_i]\delta_i$ are equal, for each i, and so $\iota([f_1]\delta_1, [f_2]\delta_2)$ is zero as ϕ is simplicial. This is a contradiction, and we deduce that f_1 and f_2 determine the same automorphism of the curve complex.

We have seen that any edge in $\mathcal{C}(\Sigma)$ is uniquely prescribed a surface homeomorphism induced automorphism agreeing with ϕ on the link of either vertex and that any two edges with a vertex in common are prescribed the same such automorphism. Since $\mathcal{C}(\Sigma)$ is connected, it follows that every edge is allocated the *same* automorphism Φ. All we need do now is verify that this automorphism is the correct one. To do this, we only need to remark that any curve α spans an edge with a second curve β. The two are prescribed the automorphism Φ which agrees with ϕ on both $X(\alpha)$ and $X(\beta)$. In particular, Φ agrees with ϕ on $X(\beta)$ which contains α. This completes the proof of Theorem 1.

3.3 Global versus local injections

This section is devoted entirely to giving a complete proof of Theorem 2 and of Theorem 4. There are perhaps several different ways of going about this, but we shall adapt Bell-Margalit's proof from [BellMar] of Theorem 2 for homeomorphic surfaces homeomorphic to a holed sphere (their argument does not discriminate). Their argument combines several results from the work of Ivanov and of Ivanov-McCarthy. Consequently, we contribute no further ideas in giving a proof to our Theorem 2. Any proof of Theorem 4 along these lines requires a careful consideration of just where the injectivity of ϕ is used and, where it is, whether our local alternative will suffice. All the originality of this section lies in tackling this issue.

Beginning with Theorem 2 then, the main aim is to show that ϕ induces a simplicial injection between the corresponding curve complexes. To this end, we need to verify that ϕ sends powers of Dehn twists to powers of Dehn twists. At this point, we invoke a few important results of Ivanov [Iva4], [Iva5] and Ivanov-McCarthy [IvaMcC], translated to suit our needs. The first result is Ivanov's algebraic characterisation of non-trivial Dehn twist mapping classes, and allows us to control the ϕ-image of a Dehn twist. The version we will give combines Theorem 2.1, Theorem 2.2 and Theorem 2.3 of [Iva5] and the subsequent advances made in §11 of [IvaMcC], in particular Theorem 11.6.

Before we state this result, we must explain some of the terminology it uses. As might be standard, for a group G we denote by $Z(G)$ its centre and, for any element $g \in G$, we denote by $C_G(g)$ the centraliser of g in G. By the rank of a group G, denoted rank(G), we mean the maximal size of any subset of G freely generating a free abelian subgroup in G. We shall say that a mapping class $g \in Map(\Sigma)$ is *pure* if it satisfies the following condition, taken from page 3 of [Iva4]: There is a self-homeomorphism f representing g and a 1-submanifold, C, of Σ such that the components of C each define distinct curves, f fixes C pointwise and $\Sigma - C$ componentwise, and on the closure in Σ of any component of $\Sigma - C$ we either have f restricts to the identity or it restricts to a pseudo-Anosov (so that for any two curves α and β in the closure of this component, $\iota(\alpha, g^n(\beta)) \longrightarrow \infty$ as $n \longrightarrow \infty$). We shall say that a subgroup of $Map(\Sigma)$ is *pure* if it comprises only of pure mapping classes. Finally, the *canonical reduction system* for a mapping class g is the maximal multicurve v fixed by g such that g fixes no curve β with $\iota(v, \beta) > 0$.

There are plenty of finite index pure subgroups of $Map(\Sigma)$; it is Theorem 3 of [Iva4] that the kernel of the natural action of $Map(\Sigma)$ on the finite group $H_1(\Sigma, \mathbb{Z}/m\mathbb{Z})$ is pure whenever $m \geq 3$. Note that pure mapping classes and pure subgroups are not to be confused with elements and subgroups, respectively, of $PMap(\Sigma)$. Although working with finite index pure subgroups of mapping class groups does lead to more notation, pure subgroups are easier to understand and more in accord with our intuition. While Theorem 17 is implicit in the literature, it has yet, to the best of my knowledge, to be stated in this form and so we record a proof.

Theorem 17 *(Ivanov, Ivanov-McCarthy) Let Γ be a finite index pure subgroup of $Map(\Sigma)$, where Σ is a connected and orientable surface of finite type with complexity $\kappa(\Sigma)$ at least two. Then, any element $g \in \Gamma$ is a non-trivial power of a Dehn twist if and only if $Z(C_\Gamma(g)) \cong \mathbb{Z}$ and rank$(C_\Gamma(g)) \geq 2$.*

Proof Suppose that g is a non-trivial power of a Dehn twist. Then, g extends to a rank $\kappa(\Sigma)$ free abelian subgroup of Γ by taking high power Dehn twists about the components of any multicurve fixed by g. It follows that the rank of $C_\Gamma(g)$ is at least two. The remaining condition is a special case of Theorem 11.6 from [IvaMcC], for the canonical reduction system of g is just a single curve.

Now suppose that $g \in \Gamma$ satisfies our two criteria. First, pseudo-Anosov mapping classes have cyclic centralisers in Γ and, as we assumed $\text{rank}(C_\Gamma(g)) \geq 2$, so g can not be pseudo-Anosov. It follows that g is either reducible or periodic by the Nielsen-Thurston classification of mapping classes. Since g is pure it is certainly not periodic, and must therefore be reducible. We now know that g fixes at least one curve. Second, if c the denotes the number of curves in a reduction system for g and p the number of pseudo-Anosov components of g, then $\text{rank}(Z(C_\Gamma(g))) = c + p$. Since $c \geq 1$, we deduce $p = 0$ and $c = 1$. That is, g is a Dehn twist power. $\diamondsuit$

In words, $Z(C_\Gamma(g))$ is the centre of the centraliser in Γ of g, or the set of all elements of Γ that commute with every element that commutes with g. The second result we require tells us that injections between groups of equal and finite rank behave well on centres of centralisers. It is an adapted version of Lemma 12.2 from [IvaMcC].

Lemma 18 *(Ivanov-McCarthy) Let A and B be two groups of equal and finite rank. Let G be a maximal rank free abelian subgroup of A. Then, for any injection $\phi : A \longrightarrow B$ and any $g \in G$, we have rank$(Z(C_B(\phi(g)))) \leq$ rank$(Z(C_A(g)))$.*

We now set about showing that ϕ must send powers of Dehn twists to powers of Dehn twists. Let Γ_1 and Γ_2 be pure finite index subgroups of H and $Map(\Sigma_2)$, respectively, such that $\phi(\Gamma_1) \subseteq \Gamma_2$. We could, for instance, take Γ_2 to be the kernel of the natural action of $Map(\Sigma_2)$ on $H_1(\Sigma_2, \mathbb{Z}/3\mathbb{Z})$ and then take Γ_1 to be $\phi^{-1}(\Gamma_2) \cap P \cap H$, where P is the kernel of the action of $Map(\Sigma_1)$ on $H_1(\Sigma_1, \mathbb{Z}/3\mathbb{Z})$. Notice $\phi^{-1}(\Gamma_2) \cap P \cap H$ has finite index in H so, to avoid excessive notation, let us assume we took H to be pure, with $\phi(H) \subset \Gamma_2$, from the outset. That is, H is a finite index subgroup of $Map(\Sigma_1)$ whose elements are all pure mapping classes whose ϕ-images are also pure mapping classes.

Choose any curve α in Σ_1 and any non-trivial Dehn twist τ in $Map(\Sigma_1)$ that fixes α. Denote by $Stab_H(\alpha)$ the set of all elements of H that fix α. Since H has finite index in $Map(\Sigma_1)$ and τ has infinite order, there is a positive integer k such that $\tau^k \in H$. Let g denote τ^k. By Theorem 17 we have $Z(C_H(g)) \cong \mathbb{Z}$, for g is also a Dehn twist. Using Lemma 18 and the injectivity of ϕ, we have $\text{rank}(C_{\Gamma_2}(\phi(g))) = \kappa(\Sigma_1)$ and, furthermore, $\kappa(\Sigma_1) = \kappa(\Sigma_2)$. Appealing to Theorem 17 once more, we see that $\phi(g)$ is a Dehn twist. As it happens, $\phi(\tau)$ is also a Dehn twist but we shall not be making use of this.

The promised map on curve complexes, $\Phi : \mathcal{C}(\Sigma_1) \longrightarrow \mathcal{C}(\Sigma_2)$, is defined by taking any curve $\alpha \in X(\Sigma_1)$ to the underlying curve of the ϕ-image of a suitably high power Dehn twist around α. This map is well-defined. As ϕ is injective on all rank 2 free abelian subgroups of H, and two non-trivial Dehn twists about distinct curves commute if and only if the two curves have zero intersection (see Theorem 4.2 of [IvaMcC]), so Φ is injective. Finally, as ϕ is a homomorphism, Φ is simplicial.

We now invoke Theorem 1 to find a mapping class $h \in Map(\Sigma_1)$ inducing Φ. In particular, we now know that Σ_1 and Σ_2 are homeomorphic. Since the action of the mapping class group on the curve complex of the closed surface of genus two is not faithful, there is more than one choice for h. For all other cases, the action of the mapping class group on the curve complex is faithful and this gives the uniqueness of h. Anyway, post-composing ϕ with any isomorphism $Map(\Sigma_2) \longrightarrow Map(\Sigma_1)$, if need be, we may regard ϕ as a self-injection of $Map(\Sigma_1)$ and Φ as an automorphism of $\mathcal{C}(\Sigma_1)$. To deduce that $\phi(g) = hgh^{-1}$ for all $g \in H$, it is now enough to verify that $\phi(g)h(\alpha) = hg(\alpha)$ for all $\alpha \in X(\Sigma_1)$. This follows by a brief calculation involving Dehn twists, and so we conclude the proof of Theorem 2.

Turning to Theorem 4 now, any deduction from Theorem 3 along the same lines requires a careful examination of exactly where global injectivity was used and, where it was, whether our version of local injectivity will suffice. First, note that for any non-trivial Dehn twist power $g \in H$, we have $\mathrm{rank}(C_{\Gamma_2}(\phi(g))) \geq \kappa(\Sigma_1)$ as g extends to a free abelian subgroup of H, of rank $\kappa(\Sigma_1)$, on which ϕ is injective. Since $C_{\Gamma_2}(\phi(g))$ is a subgroup of $Map(\Sigma_2)$ and $\mathrm{rank}(Map(\Sigma_2)) = \kappa(\Sigma_2)$, we deduce $\kappa(\Sigma_1) = \kappa(\Sigma_2)$ and $\mathrm{rank}(C_{\Gamma_2}(\phi(g))) = \kappa(\Sigma_2)$. As it stands, the only potential obstruction to applying Theorem 17 is in verifying that $Z(C_{\Gamma_2}(\phi(g))) \cong \mathbb{Z}$.

While we can not apply Lemma 18 to help establish this, we are able to modify its proof from [IvaMcC] in the context of mapping class groups. This time, we take G to be a free abelian group extending $g = \tau^k$, generated by $\kappa(\Sigma_1)$ Dehn twists and fixing α. Let K be the subgroup of Γ_2 generated by $\phi(G)$ and $Z(C_{\Gamma_2}(\phi(G)))$, and let L be equal to $\phi(G) \cap Z(C_{\Gamma_2}(\phi(G)))$. As $g \in G$ and G is abelian, we have $\phi(G) \subseteq C_{\Gamma_2}(\phi(g))$ and so K is abelian. A standard result in group theory tells us $\mathrm{rank}(\phi(G))+\mathrm{rank}(Z(C_{\Gamma_2}(\phi(g)))) = \mathrm{rank}(K)+\mathrm{rank}(L)$.

At this point, we invoke the injectivity of ϕ on curve stabilisers to deduce $\mathrm{rank}(\phi(G)) = \kappa(\Sigma_1)$. Further, as $K \leq Map(\Sigma_2)$, so $\mathrm{rank}(K) \leq \kappa(\Sigma_2) = \kappa(\Sigma_1)$. It follows that $\mathrm{rank}(Z(C_{\Gamma_2}(\phi(g)))) \leq \mathrm{rank}(L)$.

We now appeal to a second group-theoretical result, given as Lemma 12.1 in [IvaMcC]. Our application is somewhat different. The lemma states that for any subgroup A_1 of a group A_2 and a subset B of A_2, we have $Z(C_{A_2}(B)) \cap A_1 \subseteq Z(C_{A_1}(B))$. In our case, we consider $\{\phi(g)\} \subseteq \phi(Stab_H(\alpha)) \leq \Gamma_2$ to deduce $Z(C_{\Gamma_2}(\phi(g))) \cap \phi(Stab_H(\alpha)) \subseteq Z(C_{\phi(Stab_H(\alpha))}(\phi(g)))$. However, ϕ is injective on $Stab_H(\alpha)$ and so $Z(C_{\phi(Stab_H(\alpha))}(\phi(g))) \cong Z(C_{Stab_H(\alpha)}(g))$. It follows that $\mathrm{rank}(Z(C_{\phi(Stab_H(\alpha)}(g)))) = 1$ and so $\mathrm{rank}(L) \leq 1$. At last, we can correctly state $rank(Z(C_{\Gamma_2}(\phi(g)))) \leq 1$. As g has infinite order, belongs to $Stab_H(\alpha)$ and ϕ is injective on $Stab_H(\alpha)$, so $\phi(g)$ has infinite order. It follows that $rank(Z(C_{\Gamma_2}(\phi(g)))) = 1$. We have satisfied both the criteria of Theorem 17, and we deduce that $\phi(g)$ is a power of a Dehn twist.

The map Φ between curve complexes is defined exactly as before. Since ϕ is a homomorphism, Φ is simplicial, and, since ϕ is injective on all curve stabilisers in H, Φ is injective on all stars in $\mathcal{C}(\Sigma_1)$. Appealing to Theorem 3, we see that Σ_1 and Σ_2 are homeomorphic, regarded equal, and Φ is induced by a mapping class $h \in Map(\Sigma_1)$. To finally verify that ϕ is everywhere equal to the inner automorphism of $Map(\Sigma_1)$ corresponding to h, we argue exactly as for Theorem 2. This concludes the proof of Theorem 4.

In closing, our work begs us to further weaken the injectivity hypothesis on ϕ. For example, if ϕ is only assumed to be injective on all free abelian subgroups of $Map(\Sigma_1)$, is it still the restriction of an inner automorphism of $Map(\Sigma_1)$?

Bibliography

[Abi] W.Abikoff, *The real analytic theory of Teichmüller space* : Lecture Notes in Mathematics **820**, Springer, Berlin (1980).

[AbrSch] A. Abrams, S. Schleimer, *Distances of Heegaard splittings* : Geometry & Topology **9** (2005) 95-119.

[Ag] I. Agol, *Tameness of hyperbolic 3-manifolds* : preprint arXiv:math.GT/0405 568.

[AAS1] J. W. Anderson, J. Aramayona, K. J. Shackleton, *A simple criterion for non-relative hyperbolicity and one-endedness of groups* : preprint, arXiv:math. GT/0504271.

[AAS2] J. W. Anderson, J. Aramayona, K. J. Shackleton, *Uniformly exponential growth and mapping class groups* : preprint, arXiv:math.GT/0508321.

[Ara] J. Aramayona, *The Weil-Petersson geometry of the five-times punctured sphere* : to appear in "Spaces of Kleinian Groups", LMS Lecture Notes, editors Y. N. Minsky, M. Sakuma and C. M. Series.

[BaGS] W. Ballmann, M. Gromov, V. Schröder, *Manifolds of non-positive curvature* : Progress in Mathematics **61**, Birkhäuser (1985).

[Behr] J. Behrstock, *Asymptotic geometry of the mapping class group and Teichmülller space* : PhD thesis (2004).

[BehrMarg] J. Behrstock, D. Margalit, *Curve complexes and finite index subgroups of mapping class groups* : preprint arXiv:math.GT/0504328.

[BellMar] R. Bell, D. Margalit, *Injections of Artin groups* : preprint, arXiv: math.GR/0501051.

[BesFu] M. Bestvina, K. Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : Geometry & Topology **6** (2002) 69 - 89.

[Bige] S. Bigelow. *Braid groups are linear* : Journal of the American Mathematical Society **14** (2001) 471-486.

[BigeBud] S. Bigelow, R. Budney, *The mapping class group of a genus 2 surface is linear* : Algebraic and geomtric topology, No. 1 (2001) 699-708.

[Bir] J. S. Birman, *Braids, links and mapping class groups* : Princeton University Press, Annals of Mathematical Studies **82** (1974).

[BisF] D. Biss, B. Farb, K_g *is not finitely generated* : preprint (2004).

[Bon] F. Bonahon, *Bouts des variétés de dimension 3* : Annals of Mathematics **124** (1986) 71-158.

[Bowl] B. H. Bowditch, *Intersection numbers and the hyperbolicity of the curve complex* : preprint, Southampton (2002).

[Bow2] B. H. Bowditch, *Tight geodesics in the curve complex* : preprint, Southampton (2003).

[Bow3] B. H. Bowditch, *Length bounds on curves arising from tight geodesics* : preprint, Southampton (2003).

[Bow4] B. H. Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces* : in "Group theory from a geometrical viewpoint" (ed. E. Ghys, A. Häfliger, A. Verjovsky), World Scientific (1991) 64-167.

[Bow5] B. H. Bowditch, *Hyperbolic 3-manifolds and the geometry of the curve complex* : preprint, Southampton (2004).

[Bow6] B. H. Bowditch, *Geometric models for hyperbolic* 3-*manifolds* : preprint, Southampton (2005).

[BreMarg] T. E. Brendle, D. Margalit, *Commensurations of the Johnson kernel* : Geometry & Topology **8** (2004) 1361-1384.

[BridHae] M. R. Bridson, A. Haefliger, *Metric spaces of non-positive curvature* : Springer-Verlag (1999).

[Bro] J. F. Brock, *The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores* : Journal of the American Mathematical Society **16** No. 3 (2003) 495-535.

[BroCM] J. F. Brock, R. D. Canary, Y. N. Minsky, *The classification of Kleinian surface groups, II: The ending lamination conjecture* : preprint, arXiv:math.GT/ 0412006.

[BroF] J. F. Brock, B. Farb, *Curvature and rank of Teichmuller space* : to appear in American Journal of Mathematics.

[CalGa] D. Calegari, D. Gabai, *Shrinkwrapping and the taming of hyperbolic 3-manifolds* : preprint, arXiv:math.GT/0407161.

[Can1] R. D. Canary, *Hyperbolic structures on 3-manifolds with compressible boundary* : PhD thesis, Princeton University (1989).

[Can2] R. D. Canary, *Ends of hyperbolic 3-manifolds* : Journal of the American Mathematical Society **6** (1993) 1-35.

[CanEG] R.D.Canary, D. B. A. Epstein, P. Green, Notes on notes of Thurston : in Analytic and geometric aspects of hyperbolic space, London Math. Society Lecture Notes Series No. 111, (ed. D. B. A. Epstein) Cambridge University Press (1987) 392.

[CanMi] R. D. Canary, Y. N. Minsky, *On limits of tame hyperbolic 3-manifolds* : Journal of Differential Geometry **43** No. 1 (1996) 141.

[CasBl] A. J. Casson, S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston* : London Mathematical Society, Student Texts **9** (1982).

[CDP] M. Coornaert, T. Delzant, A. Papadopoulos, *Géométrie de groupes: les groupes hyperboliques de Gromov*, Springer-Verlag (1990).

[Dehn1] M. Dehn, *Über die Topologie des dreidimensionalen Raumes* : Mathematical Annals **69** (1910) 137-168.

[Dehn2] M. Dehn, *Papers on group theory and topology* : ed. J. Stilwell, Springer (1987).

[EKra] C. J. Earle, I. Kra, *On isometries between Teichmüller spaces* : Duke Mathematics Journal **41** (1974) 583-591.

[FarIva] B. Farb, N. V. Ivanov, *The Torelli geometry and its applications* : Mathematical Research Letters **12** (2005) 293-301.

[FLM] B. Farb, A. Lubotzky, Y. Minsky, *Rank one phenomena for mapping class groups* : Duke Mathematics Journal **106** No. 3 (2001) 581-597.

[FarMa] B. Farb, H. A. Masur, *Superrigidity and mapping class groups* : Topology **37** No. 6 (1998) 1169-1176.

[Fathi] A. Fathi, *Dehn twists and pseudo-Anosov diffeomorphisms* : Inventiones Mathematicae **87** No. 1 (1987) 129-151.

[FLP] A. Fathi, F. Laudenbach, V. Poénaru, *Travaux de Thurston sur les surfaces, seconde édition* : Société Mathématique de France **66-67** (1991).

[Ger] S. Gervais, *A finite presentation of the mapping class group of an oriented surface* : Topology **40** No. 4 (2001) 703-725.

[GhD] E. Ghys, P. de la Harpe, *Sur les groupes hyperboliques d'aprés Mikhael Gromov* : Birkhäuser (1990).

[Gro] M. Gromov, *Hyperbolic groups* : Essays in group theory (ed. S. M. Gersten), MSRI Publications **8**, Springer-Verlag (1987).

[Ham1] U. Hamenstädt, *Train tracks and the Gromov boundary of the complex of curves* : to appear in "Spaces of Kleinian groups" (Y. Minsky, M. Sakuma, C. Series, eds.), London Math. Soc. Lec. Notes 329 (2005), 187-207.

[Ham2] U. Hamenstädt, *Train tracks and mapping class groups I* : Preliminary version, preprint (2004).

[Ham3] U. Hamenstädt, *Geometric properties of the mapping class group* : preprint (2005).

[Hare] J. L. Harer, *The virtual cohomological dimension of the mapping class groups of an orientable surface* : Inventiones Mathematicae **84** (1986) 157-176.

[HarrSco] L. Harris, G. P. Scott, *The uniqueness of compact cores for 3-manifolds* : Pacific Journal of Mathematics **172** No. 1 (1996) 139-150.

[Harv] W. J. Harvey, *Boundary structure of the modular group* : in "Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference" (ed. I. Kra, B. Maskit), Annals of Mathematical Studies No. 97, Princeton University Press (1981) 245-251.

[HarvKor] W. J. Harvey, M. Korkmaz, *Homomorphisms from mapping class groups* : Bulletin of the London Mathematical Society **37** No. 2 (2005) 275-284.

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

[Hemp] J. Hempel, *3-manifolds as viewed from the curve complex* : Topology **40** No. 3 (2001) 631-657.

[ImaTan] Y. Imayoshi, M. Taniguchi, *An introduction to Teichmüller spaces* : Springer-Verlag (1991).

[Irm1] E. Irmak, *Superinjective simplicial maps of complexes of curves and injective homomorphisms of subgroups of mapping class groups II* : preprint, arXiv:math.GT/0311407.

[Irm2] E. Irmak, *Complexes of non-separating curves and mapping class groups* : preprint, arXiv:math.GT/0407285.

[Iva1] N. V. Ivanov, *Automorphisms of complexes of curves and of Teichmüller spaces* : International Mathematics Research Notices (1997) 651-666.

[Iva2] N. V. Ivanov, *Mapping class groups* : in "Handbook of geometric topology" (ed. R. Daverman and R. Sher), Elsevier (2001) 523-633.

[Iva3] N. V. Ivanov, *A short proof of non-Gromov hyperbolicity of Teichmüller spaces* : Annales Academie Scientiarum Fennice, Mathematica, **27** F. 1 (2002) 3-5.

[Iva4] N. V. Ivanov, *Subgroups of Teichmüller modular groups* : Translations of Mathematical Monographs **115**, American Mathematical Society (1992).

[Iva5] N. V. Ivanov, *Automorphisms of Teichmüller modular groups* : Topology and geometry - Rohlin seminar, Lecture Notes in Mathematics **1346** 199-270 (1988).

[IvaMcCar] N. V. Ivanov, J. D. McCarthy, *On injective homomorphisms between Teichmüller modular groups I* : Inventiones Mathematicae **135** (1999) 425-486.

[J1] D. L. Johnson, *An abelian quotient of the mapping class group $\mathcal{T}_g$* : Mathematische Annalen **249** (1980) 225-242.

[J2] D. L. Johnson, *The structure of the Torelli group I: A finite set of generators for $\mathcal{T}$* : Annals of Mathematics **118** (1983) 423-442.

[Kap] M. Kapovich, *Hyperbolic manifolds and discrete groups* : Progress in Mathematics **183**, Birkhäuser Boston (2001)

[Kat] T. Kato, *Asymptotic Lipschitz maps, combable groups and higher signatures* : GAFA **10** (2000) 51-110.

[Ker] S. P. Kerckhoff, *The asymptotic geometry of Teichmüller space* : Topology **19** No. 1 (1980) 23-41.

[KerT] S. P. Kerckhoff, W. P. Thurston, *Noncontinuity of the action of the modular group at Bers's boundary of Teichmüller space* : Inventione Mathematicae **100** (1990) 25-47.

[Kla] E. Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space* : preprint (1999).

[Kor1] M. Korkmaz, *Automorphisms of complexes of curves on punctured spheres and on punctured tori* : Topology and its Applications **95** (1999) 85-111.

[Kor2] M. Korkmaz, *On the linearity of certain mapping class groups* : Turkish Journal of Mathematics **24** No. 4 (2000) 367-371.

[Kra] D. Krammer, *Braid groups are linear* : Annals of Mathematics (2) No. 1 **155** (2002) 131-156.

[Lea] J. Leasure, *Geodesics in the complex of curves of a surface* : PhD thesis, University of Texas (2002).

[Lic1] W. B. R. Lickorish, *A finite set of generators for the homeotopy group of a 2-manifold* : Proceedings of the Cambridge Philosophical Society **60** (1964) 769-778.

[Lic2] W. B. R. Lickorish, *On the homeotopy group of a 2-manifold (corrigendum)* : Proceedings of the Cambridge Philosophical Society **62** (1966) 679-681.

[Long] D. D. Long, *Constructing pseudo-Anosov maps* : in "Knot theory and manifolds (Vancouver, B.C., 1983)", Lecture Notes in Mathematics **1144** (1985) 108-114.

[Luo1] F. Luo, *Automorphisms of the complex of curves* : Topology **39** (2000) 283-298.

[Luo2] F. Luo, *A presentation of the mapping class group* : Mathematical Research Letters **4** (1997) 735-739.

[Mard] A. Marden, *The geometry of finitely generated Kleinian groups* : Annals of Mathematical Studies **99** (1974) 383-462.

[Marg] D. Margalit, *Automorphisms of the pants complex* : Duke Mathematical Journal **121** No. 3 (2004).

[MaMi1] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves I: Hyperbolicity* : Inventiones Mathematicae **138** (1999) 103-149.

[MaMi2] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves II: Hierarchical structure* : Geometry & Functional Analysis **10** (2000) 902-974.

[MaWolf] H. A. Masur, M. Wolf, *The Weil-Petersson isometry group* : Geometriae Dedicata **93** (2002) 177-190.

[McC] D. McCullough, *Compact submanifolds of 3-manifolds with boundary* : Quart. J. Math. Oxford **37** 299-306.

[McCM] D. McCullough, A. Miller, *The genus 2 Torelli group is not finitely generated*, Topology and its Applications **22** No. 1 (1986) 4349.

[McCMS] D. McCullough, A. Miller, G. A. Swarup, *Uniqueness of cores of noncompact 3-manifolds* : Journal of the London Mathematical Society **52** (1985) 548-556.

[Mine] I. Mineyev, *Bounded cohomology characterizes hyperbolic groups* : Quarterly Journal of Mathematics **53** (2002) 59-73.

[Mi1] Y. N. Minsky, *Extremal length estimates and product regions in Teichmüller space* : Duke Mathematical Journal **83** (1996) 249-286.

[Mi2] Y. N. Minsky, *Teichmüller geodesics and ends of hyperbolic manifolds* : Topology **32** (1993) 625-647.

[Mi3] Y. N. Minsky, *The classification of Kleinian surface groups I: Models and bounds* : preprint, arXiv:math.GT/0302208.

[Mi4] Y. N. Minsky, *Combinatorial and geometrical aspects of hyperbolic* 3-*manifolds* : in Kleinian Groups and Hyperbolic 3-Manifolds, London Mathematical Society Lecture Notes **299** (2003) 3-40.

[Mor] S. Morita, *Cassons invariant for homology 3-spheres and characteristic classes of surface bundles* : Topology **28** (1989) 305-323.

[Mos] L. Mosher, *Mapping class groups are automatic* : Annals of Mathematics **142** No. 2 (1995) 303-384.

[Oh] K. Ohshika, *Kleinian groups which are limits of geometrically finite groups* : preprint.

[Pen] R. C. Penner, *A construction of pseudo-Anosov homeomorphisms* : Transactions of the American Mathematical Society **310** No. 1 (1998).

[PenH] R. C. Penner with J. L. Harer, *Combinatorics of train tracks* : Annals of Mathematical Studies, Princeton University Press (1992).

[Per1] G. Perelman, *The entropy formula for Ricci flow and its geometric applications* : preprint, arXiv:math.DG/0211159.

[Per2] G. Perelman, *Ricci flow with surgery on 3-manifolds* : preprint, math.DG/ 0303109.

[Re] M. Rees, *The geometric model and large Lipschitz equivalence direct from Teichmüller geodesics* : arXiv:math.GT/0404007.

[Roy] H. L. Royden, *Automorphisms and isometries of Teichmüller space* : in L. V. Ahlfors, et al., *Advances in the Theory of Riemann Surfaces, 1969 Stony Brook Conference*, Annals of Mathematical Studies **66**, Princeton University Press (1971) 369-383.

[Sch] P. Schmutz Schaller, *Mapping class groups of hyperbolic surfaces and automorphism groups of graphs* : Compositio Mathematica **122** No. 3 (2000) 243-260.

[Schl] S. Schleimer, *Notes on the complex of curves* : preprint, Rutgers (2005).

[Sco] G. P. Scott, *Compact submanifolds of 3-manifolds* : Journal of the London Mathematical Society **7** (1973) 246-250.

[Som] T. Soma, *Existence of polygonal wrappings in hyperbolic* 3-*manifolds* : preprint, Tokyo Denki (2005).

[Sou] J. Souto, *A note on the tameness of hyperbolic 3-manifolds* : Topology **44** (2005) 459-474.

[Sto] P. Storm, *The Novikov conjecture for the mapping class groups as a corollary of Hamenstädt's theorem* : preprint, arXiv:math.GT/0504248.

[Str] K. Strebel, *Quadratic differentials* : Springer-Verlag (1984).

[Sul] D. Sullivan, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions* : in "Riemannian surfaces and related topics: Proceedings of the 1978 Stony Brook Conference", Annals of Mathematical Studies **97**, Princeton (1981).

[Thur1] W. P. Thurston, *The geometry and topology of hyperbolic 3-manifolds* : Princeton University Lecture Notes, 1982.

[Thur2] W. P. Thurston, *Three-dimensional manifolds, Kleinian groups, and hyperbolic geometry* : Bulletin of the American Mathematical Society **6** (1982) 357-381.

[Thur3] W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces I* : Bulletin of the American Mathematical Society **19** No. 2 (1988).

[Tu] T. W. Tucker, *Non-compact 3-manifolds and the missing boundary problem* : Topology **13** (1973) 267-273.

[V] O. Viro, *Links, two-sheeted branched coverings and braids* : Soviet Math, Sbornik **87** (2) (1972) 216-228.

[Waj] B. Wajnryb, *A simple presentation for the mapping class group of an orientable surface* : Israel Journal of Mathematics **45** No. 2-3 (1983) 157-174.

[Wolp] S. A. Wolpert, *Geometry of the Weil-Petersson completion of Teichmüller space* : Surveys in Differential Geometry VIII: Papers in honor of Calabi, Lawson, Siu and Uhlenbeck, editor S. T. Yau, International Press (2003).