Asymptotic Geometry of the Mapping Class Group and Teichmüller Space

A Dissertation, Presented

by

Jason Alan Behrstock

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Mathematics

Stony Brook University

May 2004

Stony Brook University

The Graduate School

Jason Alan Behrstock

We, the dissertation committee for the above candidate for the Doctor of Philosophy degree, hereby recommend acceptance of this dissertation.

Yair Minsky
Professor of Mathematics, Yale University
Dissertation Director

Anthony Phillips
Professor of Mathematics
Chairman of Dissertation

Dennis Sullivan
Distinguished Professor of Mathematics

Lee Mosher
Professor of Mathematics and Computer Science,
Rutgers University at Newark
Outside Member

This dissertation is accepted by the Graduate School.

Graduate School

Abstract of the Dissertation

Asymptotic Geometry of the Mapping Class Group and Teichmüller Space

by

Jason Alan Behrstock

Doctor of Philosophy

in

Mathematics

Stony Brook University

2004

In this thesis, we make progress towards understanding the asymptotic geometry of the mapping class group and Teichmüller Space. We begin by exhibiting some interesting properties of maps from the mapping class group of a surface S to the curve complex of a subsurface X—these maps are closely related to the map taking a curve in S to its intersection with X. Work of Masur and Minsky describes how these maps can be used to encode the coarse behavior of the word metric on the mapping class group and the Weil-Petersson metric on Teichmüller space. These "projection" maps lead to a description of the mapping class group as a subset of an infinite product of δ-hyperbolic spaces; we produce a description of the image restricted to finite subproducts. We extend this by discussing the relationship between these "projection" maps and the topology of the asymptotic cone of the mapping class group. In particular, we compute the dimension of certain subsets of the asymptotic cone and give a new proof that in the low complexity cases Teichmüller space with the Weil-Petersson metric is δ-hyperbolic.

To the Behrstock, Kaufman and Zaitlen families:
for showing me the joy to be found in spending a lifetime learning
in the company of friends.

Contents

List of Figures

Acknowledgements

I want to start by expressing my incredible gratitude to my advisor Yair Minsky. You have provided an inspiration to me mathematically over my years as a graduate student and I have benefitted immensely from your encouragement, enthusiasm, and mathematical insight. Also, I would like to thank Yair's young son Eyal for his suggestion at a crucial point in my writing process "if you say things more clearly, they will make more sense." Despite this advice, this thesis has taken a large number of iterations to get to the present state and I am deeply indebted to Yair for his careful reading and useful comments on many of these drafts.

I would also like to acknowledge the additional mentorship I have received from Lee Mosher. I have enjoyed and mathematically grown from the time we have spent at the various cafés around the city discussing mathematics. I also appreciate his detailed feedback on versions of this thesis, which helped create the clarity which exists (I hope) in this final version.

I want to thank the geometric topology group at Columbia University for many interesting discussions and for making me feel welcome there.

At various times in my graduate career I have been lucky to have had the opportunity for useful discussions with Mladen Bestvina, Benson Farb, Misha Gromov, and Dennis Sullivan who each introduced me to new ideas which have significantly changed my perspective on mathematics.

Various aspects of this work have benefitted from insightful discussions with Dan Margalit and Kevin Wortman, whom I look forward to working with more in the future.

I want to thank the many mathematicians at Berkeley who early in my career taught me the joys of topology and the fun of iterating maps, these include especially André de Carvalho, Moe Hirsch, Charles Pugh, and John Stallings. Also, I would be remiss to not mention the many graduate students I had the good fortune of befriending there and who all had an enormous impact on me as an impressionable undergraduate. I especially would like to thank Stephen Bigelow, Brandoch Calef, Danny Calegari, Ben Davis, Tom Insel, J.P. Lund, Diane Maclagan, Johanna Neaderhouser, Saul Schleimer, and Monica Vazirani.

At Stony Brook I was lucky to have the mathematical camaraderie of Joe Coffey, Hossein Namazi, Kasra Rafi, and Candida Silveira.

I want to thank the many friends whom I have had the good fortune of knowing over the years. I want to thank especially all the members of Metacomm, my many great friends in New York, and everyone I have had the pleasure of sharing a belay or playing SET with—you have kept me in touch with the world through many intellectual and (especially!) non-intellectual distractions.

I greatly appreciate the constant love and support I have received from my parents and sister.

Finally, I want to thank Kashi for her help mathematically—from midnight discussions of Margulis' Lemma while studying for my Quals to proofreading the final drafts of this work—and for her constant support and encouragement. And of course for much more...

And of course, warmest wishes to my many other friends and colleagues.

Chapter 1

Introduction

"For my own part I am pleased enough with surfaces—in fact they alone seem to me to be of much importance."

—Edward Abbey, Desert Solitaire

The question that motivates the work in this thesis is: to what extent is the algebraic structure of the mapping class group encoded by its large scale geometry? In this thesis we develop tools for analyzing the coarse geometry of the mapping class group and exhibit a new description of this space which helps elucidate questions concerning its dimension and geometric rank. Our techniques are also applied to obtain results about the asymptotic geometry of Teichmüller space.

1.1 Overview of Main Results

There are many interesting spaces associated to a surface S. For studying dynamical systems on S one considers the group of automorphisms of S to itself, known as the *mapping class group*. Geometers are interested in *Teichmüller space*, which is the space of hyperbolic structures on S. (Teichmüller space is also of interest to complex analysts as it happens that it can alternatively be viewed as the space of complex structures on S.) And then there is the *complex of curves*, a space which encodes the combinatorics of intersection patterns of simple closed curves on S.

The complex of curves was introduced by Harvey and has been used to prove many deep results about the mapping class group (see [Hare], [Harv], [Hat], [Iv]). More recently Masur and Minsky proved δ-hyperbolicity of the complex of curves [MM1] and then used this to establish new methods for studying the mapping class group [MM2]. The complex of curves has also found many applications to the study of 3-manifolds; in particular it has proved crucial in the recent proof by Brock, Canary, and Minsky of the Ending Lamination Conjecture [M, BrCM]. (For other applications see: [BaS], [BeFu], [Bo2], [Hem],[Ir], [Mar].)

The results contained in this dissertation are part of a program to understand the asymptotic geometry of the mapping class group and Teichmüller space; we proceed by proving several new results about these spaces via an analysis of the complex of curves. We now provide a brief summary of our main results.

For each subsurface $X \subset S$, we consider a "projection" map

$$\pi_X : \mathcal{MCG}(S) \to \mathcal{C}(X)$$

from the mapping class group of S into the curve complex of X, which is closely related to the map taking a curve in S to its intersection with X. The vertices of the curve complex of a surface, $\mathcal{C}(X)$, are homotopy classes of essential, non-peripheral closed curves on X. The distance between two distinct homotopy classes of curves, $\mu, \nu \in \mathcal{C}(X)$ is a measure of how complicated the intersection of representatives of these curves must be. For example as long as X has genus larger than one: $d_X(\mu, \nu) = 1$ if and only if μ and ν can be realized disjointly on X whereas $d_X(\mu, \nu) \geq 3$ if and only if μ and ν fill X (*fill* means that for any representatives of μ and ν the set $X \setminus (\mu \cup \nu)$ consists of disks and once-punctured disks).

The following tool provides the starting point for our analysis.

Theorem 4.2.1 (Projection estimates). *Let Y and Z be two overlapping subsurfaces of S. Then for any $\mu \in \mathcal{MCG}(S)$:*

$$d_{\mathcal{C}(Y)}(\partial Z, \mu) > M \implies d_{\mathcal{C}(Z)}(\partial Y, \mu) \leq M.$$

Where M depends only on the topological type of the surface S.

This theorem provides the base case of an inductive argument which yields a picture of certain aspects of the asymptotic geometry of the mapping class group. In particular, we have the following results which show that the geometric rank of certain approximations of $\mathcal{MCG}(S)$ are each at most $3g+p-3$. Given Θ, a finite set of subsurfaces of S, we can consider the image of the above projection map into $\prod_{Y\in\Theta} \mathcal{C}(Y)$. Collections of disjoint subsurfaces arise in the study of the mapping class group, since they can be used to build abelian subgroups. When U is a maximal collection of disjoint surfaces in Θ we define an object called a coarse tree-flat, denoted $\mathcal{F}(U) \subset \prod_{Y\in\Theta} \mathcal{C}(Y)$, which in the simplest cases is quasi-isometric to $\prod_{Y\in U} \mathcal{C}(Y)$ embedded in a nice way (in general it is obtained from such a product by an operation similar to a "blow-up"). One of our main results about these objects is the following.

Theorem 4.3.6 (Coarse tree-flats carry markings). *The map* $\pi = \prod_{X\in\Theta}\pi_X : \mathcal{M}CG(S) \to \prod_{X\in\Theta}\mathcal{C}(X)$ *satisfies:*

$$Image(\pi) \subset \bigcup_{U\in\mathcal{U}_\Theta} \mathcal{F}(U)$$

Where $\mathcal{U}_\Theta$ *is a finite set and each* $\mathcal{F}(U)$ *has rank* $\leq 3g+p-3$.

These theorems on projections have interesting consequences regarding the geometry of mapping class groups. In particular in Chapter 6 we study a map $\widehat{\psi} : \mathrm{Cone}_\infty\mathcal{M}\mathrm{CG(S)} \to \prod_{\alpha\in 2^{\mathbb{N}}}\mathrm{T}_\alpha$, from the asymptotic cone of the mapping class group to an uncountable product of $\mathbb{R}$-trees, where we use the above theorem to prove:

Theorem 6.3.4 *Let* K *be a subset of* $\mathrm{Cone}_\infty\mathcal{M}(\mathrm{S})$ *on which the restriction of* $\widehat{\psi}$ *is bi-Lipschitz, then* $\dim(\mathrm{K}) \leq 3\mathrm{g}+\mathrm{p}-3$.

An interesting special case of this is when K is a bi-Lipschitz flat.

One consequence of Thurston's work on classifying surface automorphisms is a structure theorem for abelian subgroups of the mapping class group. (See [BiLMc] for theorems on abelian subgroups. For the general classification see [BeH], [FaLaP], and [Th1].) In particular [BiLMc] compute that the torsion free rank of any abelian subgroup of the mapping class group is at most $3g+p-3$. For once-punctured surfaces [Mo] and then in full generality [FLM], it has been shown that these abelian subgroups are quasi-isometrically embedded in the mapping class group. Thus the *algebraic rank* provides a lower bound for the *geometric rank* of the mapping class group.

An important principle in studying groups of non-positive curvature is that, informally speaking, in directions orthogonal to maximal quasi-flats they behave hyperbolically. This principle motivated the approach taken by [EF] and [KlLe] in studying lattices in semi-simple Lie groups, one of the major successes in the program to geometrically classify groups. In this context, proving a classification theorem has been intertwined with developing an understanding of quasi-flats in symmetric spaces, culminating in a proof that these groups are geometrically rigid.

Corollary 6.3.4 provides evidence for the following Rank Conjecture for the mapping class group, a conjecture whose importance is indicated by the case of lattices in semi-simple Lie groups.

Conjecture 6.3.6 ([BrF] and [Ham][1]). $\mathcal{M}CG(S)$ *admits quasi-isometric*

[1] Hamenstädt has recently announced a proof of this conjecture using interesting techniques different from those contained in this work.

embeddings of $\mathbb{R}^n$ *if and only if* $n \leq 3g + p - 3$.

Closely related to the mapping class group is a simplicial complex called the *pants complex*, whose vertices are pants decompositions. Work of Brock proves that the pants complex is quasi-isometric to the Weil-Petersson metric on Teichmüller space [Br]. By techniques analogous to those used to prove Theorems 4.3.6 and 6.3.4 we obtain a similar theorem for the pants complex, which via Brock's result provides a theorem about the Weil-Petersson metric on Teichmüller space. This again allows us to prove certain cases of a Rank Conjecture for Teichmüller space. We also provide a new proof of the following result, which was first established by Brock and Farb in a paper which furthered our interest in this and related questions (see [BrF]).

Theorem 5.0.9 *With the Weil-Petersson metric, the Teichmüller spaces for the surfaces* $S_{1,2}$ *and* $S_{0,5}$ *are each* δ*-hyperbolic.*

1.2 Historical Context: Surfaces and Large Scale Geometry

The topological classification of surfaces was known since the 19th century in what may have been hailed at the time as the end of 2-manifold theory... but it turns out that a topological classification was only the beginning of the story! In the early 20th century Dehn and Nielsen each began detailed studies of the group of automorphisms of a surface taken up to homotopy, the *mapping class group*.

Nielsen proved the beginnings of a theorem classifying mapping class group elements, but it wasn't until the 1970's that Thurston's study of 3-manifolds which fiber over the circle led him to a complete classification theorem for elements of the mapping class group.

In studying the mapping class group Dehn became interested in the set of isotopy classes of simple closed curves on a surface, and eloquently called this the *arithmetic field of a surface*; though no longer called by this name this set is crucial in our work and shows up as the vertices of the *complex of curves.*

Another ingredient in the history of surfaces involves the space of geometric structures on a fixed topological surface, again considered up to homotopy. The first theorem on this space (a computation of its dimension) was announced by Riemann in 1859, but the first rigorous analysis on this space of complex structures occurred later and was due to Fricke and Klein. It was eventually named *Teichmüller space* after Teichmüller published two papers on the subject, in 1939 and 1943, which among other results introduced a metric on this space.

The mapping class group and Teichmüller space are closely related, indeed Thurston's proof of the classification of mapping classes involved a study of the action of the mapping class group on a compactification of Teichmüller space. A search for a satisfying description of the large scale geometry of these two objects motivates the work contained in this dissertation.

In the early 1980's Cannon, Gromov, and Rips each simultaneously (up to some additive constant) gave a definition of what it should mean for a metric space to be *hyperbolic*. Each had different motivations for their study, but further evidence for the importance of such studies was provided by the fact that their very different looking definitions all turn out to be equivalent. This occurred at around the same time that Gromov suggested (following a now classical result of Milnor and independently of Švarc) that a natural geometric invariant of a group is provided by considering (metrically) the Cayley graph of a group up to quasi-isometry (see [Tho] for a logician's interpretation of the naturality of this evolution). This created a new branch of geometric group theory: the study of large scale geometric invariants.

Perhaps the most general question to then ask is to classify all finitely generated groups up to quasi-isometry. This proposal of Gromov's is seen as a vastly ambitious program and a complete solution does not seem forthcoming. The current state of the art is to try to understand the set of groups quasi-isometric to one's favorite group or class of groups, which if one is lucky (as in the case of nilpotent groups or lattices in semi-simple Lie groups) includes only the groups to which it is obviously quasi-isometric, namely finite extensions of finite index subgroups.

We summarize with the following conjecture which serves as a motivating guide:

Conjecture 1.2.1. *Let S be a non-exceptional surface of finite type. For any finitely generated group G quasi-isometric to the mapping class group $\mathcal{MCG}(S)$, there exists a homomorphism $G \to \mathcal{MCG}(S)$ with finite kernel and finite index image.*

Informally, this question asks if the mapping class group is uniquely determined among finitely generated groups by its large scale geometry. (See [Mo2] for a further discussion and a proof of Conjecture 1.2.1 for mapping class groups of once punctured surfaces.)

1.3 Outline of Subsequent Chapters

In chapter 2 we discuss the tools of asymptotic geometry which are fundamental to the constructions and results in the rest of this work. All the results in this section are well known and thus we are as much setting up notation as we are reminding the reader of which tools to keep at their fingertips for the remainder of this work.

Surfaces are the main character throughout this thesis. In Chapter 3 we discuss the marking complex, which is our quasi-isometric model of choice for dealing with the mapping class group. We also remind the reader of various facts we will need concerning the mapping class group, the complex of curves, and Teichmüller space.

The original material in this thesis begins with the results in Chapter 4. Here we calculate estimates on the image of the "projection" map from the marking complex into the product of the curve complexes of the constituent subsurfaces. In this chapter we produce the main technical result, Theorem 4.2.1 (Projection estimates), which we use to analyze the mapping class group. In Theorem 4.2.2 (Projection estimates; geometric version) we provide a geometric interpretation of Theorem 4.2.1. We end this chapter with Theorem 4.3.6 (Coarse tree-flats carry markings) which provides a generalization of the picture established in Theorem 4.2.2.

We will end with two chapters on applications of the Projection estimates Theorem and the Coarse tree-flats carry markings Theorem. In the first of these chapters we will (re)prove hyperbolicity of the Weil-Petersson metric on Teichmüller space and of the mapping class group in some low complexity cases, Theorems 5.0.9 and 5.0.8. Finally, in Chapter 6 we will discuss how the Projection estimates Theorem helps produce a geometric picture of what the mapping class group looks like asymptotically and relate this to a discussion of the rank conjectures for the mapping class group and Teichmüller space. The main results in this chapter are in Theorems 6.2.1 and 6.3.4. The final section of Chapter 6 discusses a conjectural picture of the asymptotic cone of the mapping class group. Chapter 6 relies on some results from Dimension Theory, which we summarize in the Appendix.

Chapter 2

Large Scale Geometry

2.1 Geometric Group Theory

As was mentioned in the introduction, there are several equivalent definitions for hyperbolicity, throughout this paper we will have in mind Rips' *thin triangle* characterization of hyperbolicity, as defined below.

Recall that X is called a *geodesic metric space*, if for each $a, b \in X$ there exists a geodesic segment, which we denote $[a, b]$ whose length is equal to $d_X(a, b)$. The geodesic segment need not be unique.

Definition 2.1.1. A geodesic metric space X is called δ*-hyperbolic* if there exists a constant δ so that for each triple $a, b, c \in X$ and each choice of geodesics $[a, b], [b, c]$, and $[a, c]$ one has that $[a, b]$ is contained in a δ-neighborhood of the union of $[a, c]$ and $[b, c]$.

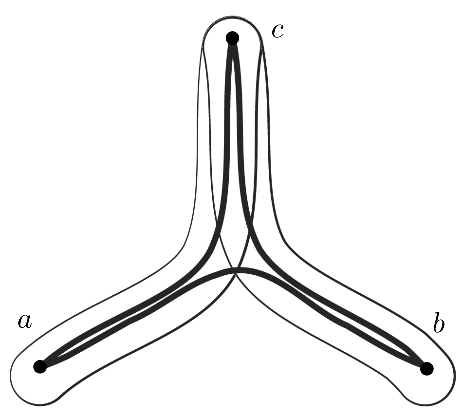

Figure 2.1: A thin triangle

If X is δ-hyperbolic for some constant δ then it is also δ'-hyperbolic for any $\delta' > \delta$. Often we will only care that there exists some δ, accordingly when the constant is unimportant we simply say X is *hyperbolic*. Of course our primary interest is in infinite diameter spaces, as it is easy to see that every finite diameter space is hyperbolic with *hyperbolicity constant* equal to the diameter of the space. An important infinite diameter example is provided by trees (both simplicial trees and their cousins $\mathbb{R}$-trees which we will define shortly).

Example 2.1.2. Trees are δ-hyperbolic with hyperbolicity constant 0.

The following justifies the name hyperbolic:

Example 2.1.3. Hyperbolic space $\mathbb{H}^n$, is δ-hyperbolic. In particular, a simple computation gives that $\mathbb{H}^n$ is hyperbolic with smallest hyperbolicity constant $\delta = \sinh^{-1}(1)$.

Definition 2.1.4. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $\phi : X \to Y$ is called a (K, C)-*quasi-isometric embedding* if there exist constants $K \geq 1$ and $C \geq 0$ such that for all $a, b \in X$

$$\frac{1}{K} d_X(a, b) - C \leq d_Y(\phi(a), \phi(b)) \leq K d_X(a, b) + C.$$

It is a (K, C)-*quasi-isometry* if additionally, ϕ has the property that every point of Y lies in the C-neighborhood of $\phi(X)$.

When there exists a (K, C)-quasi-isometry $\phi : X \to Y$ between two metric spaces X and Y, we say that they are (K, C)-*quasi-isometric*; When the choice of constants is not important (as will often be the case), we simply say X and Y are *quasi-isometric.*

Although the symmetry of the above definition is not immediately apparent, it is not difficult to show that quasi-isometry is an equivalence relation on the set of metric spaces.

For geodesic spaces it is not difficult to show that hyperbolicity is a quasi-isometry invariant; note that the hyperbolicity constant may change [GhHar].

In order to apply this definition to groups, we recall the natural left invariant metric on a group G with finite generating set $S = \{s_1, \ldots, s_n\}$. The *word metric on G relative to the generating set S* is given by $d_{G,S}(a, b) = |a^{-1}b|$ where $|a^{-1}b|$ is the smallest number of letters needed to represent the group element $a^{-1}b$ in terms of letters from S and their inverses. Although this metric depends on the choice of generating set, for finitely generated groups, its quasi-isometry type does not. Thus since hyperbolicity is a quasi-isometry invariant, the notion of a group being *hyperbolic* is well defined. Note that this metric is (K, C)-quasi-isometric with $K = 1 = C$ to the Cayley graph of G with respect to the generators in S. When we discuss the above metric, we usually are thinking of the Cayley graph, so that we may rely on the familiar notions of a path connected space while minimizing the need to refer to coarse notions like quasi-paths, etc.

The free group on two generators $F_2 =< a, b >$ is hyperbolic. This follows from Example 2.1.2 since with the word metric on the generating set $\{a, b\}$ the Cayley graph of F_2 is a tree with each vertex having valency 4.

A basic non-hyperbolic example where we use the notion of quasi-isometries is $\mathbb{Z}^n$ (with $n > 1$), which we observe is quasi-isometric to $\mathbb{R}^n$. This is easiest to see when we consider $\mathbb{Z}^n$ with the word metric generated by $S = \{(1, 0, 0, \ldots, 0), (0, 1, 0, \ldots, 0), \ldots, (0, 0, 0, \ldots, 1)\}$ which generates the l^1 metric on $\mathbb{Z}^n$, and $\mathbb{R}^n$ is also given the l^1 metric (which it is easy to see is quasi-isometric to the Euclidean metric). This easy fact is related to a classification theorem for finitely generated groups quasi-isometric to $\mathbb{R}^n$.

Theorem 2.1.5. [GhHar]. *Let G be a finitely generated group. G is quasi-isometric to $\mathbb{R}^n$ if and only if G contains a finite index subgroup isomorphic to $\mathbb{Z}^n$.*

It remains an open problem to provide an elementary proof of this theorem.

In any metric space X one can study quasi-isometric embeddings of $\mathbb{R}^n$ into X, these are called the *quasiflats of rank n*. Generalizing the above Euclidean example, one way to quantify non-hyperbolicity of a metric space X is by calculating its *geometric rank*, namely the largest rank of any quasi-flat in X. Having geometric rank greater than one is an obstruction to hyperbolicity, as it forces X to have arbitrarily fat triangles.

A useful tool for dealing with quasi-isometries is:

Theorem 2.1.6. (Milnor-Švarc). *Let X be a geodesic space. If Γ acts properly discontinuously and cocompactly by isometries on X, then Γ is finitely generated and for any choice of basepoint $b \in X$, the map $\gamma \mapsto \gamma \cdot b$ is a quasi-isometry.*

We end with a technical lemma which will be useful in Chapter 5 The proof is a standard exercise in thin triangles, which we provide for completeness. For a point $x \in X$ we will use the notation $B_\delta(x)$ to denote the closed ball of radius δ around x.

Lemma 2.1.7. *Let Z be a δ-hyperbolic space and γ a geodesic in Z. If $\phi : Z \to 2^\gamma$ sends each point of Z to the set of closest points on γ, then for each point $z \in Z$ we have* $\mathrm{diam}(\phi(\mathrm{z})) \leq 4\delta + 2$.

Proof: Fix a point $z \in Z$ and two points $x, y \in \phi(z)$. Then since Z is δ-hyperbolic, the triangle with edges $[x, y]$, $[x, z]$, and $[y, z]$ is δ-thin (for this argument we choose $[x, y]$ to refer to the geodesic contained in γ).

If the point $x' \in [x, z]$ a distance $\delta + 1$ from x is within δ of γ then there are points on γ closer to z and thus $x \notin \phi(z)$. Similarly for the point $y' \in [y, z]$ a distance $\delta + 1$ from y. Thus, x' must be within δ of a point of $[z, y'] = [z, y] \setminus B_\delta(y)$ and y' is in the δ neighborhood of $[x', z]$, so we have $[x', z]$ and $[y', z]$ are within δ of each other. In particular, if $d(x', y') > 2\delta$ this would violate either x or y being in $\phi(z)$. Thus $d(x', y') \leq 2\delta$ and thus $d(x, y) \leq 2\delta + 2\delta + 2 = 4\delta + 2$.

Now, $[x', z]$ must lie within δ of $[y, z]$. If $d(x', y') > 2\delta$, then the closest point to x' on $[y, z]$ is more than 2δ from y, and thus $d(x, z) < d(y, z)$ contradicting x and y being equidistant from z. Similarly with the roles of x and y reversed.

Thus we have $d(x', y') \leq 2\delta$ and thus $d(x, y) \leq 2\delta + 2\delta + 2 = 4\delta + 2$. □

2.2 Asymptotic Cones

Ultrafilters and ultraproducts have been used in logic since Gödel introduced them in the 1930's. Among other applications they can be used in a non-standard way to rigorize calculus, *a la* Robinson, and to give a slick proof of the compactness theorem of first order logic. van den Dries and Wilkie introduced these tools to topologists when they reformulated (and slightly strengthened) Gromov's famous Polynomial growth Theorem using this language (compare [G1] and [VW]). These techniques streamlined parts of the proof, replacing an iterated procedure of passing to subsequences with one ultralimit, and broadened the applicability of Gromov's construction, Gromov's original definition of asymptotic cone only applied to nilpotent groups whereas the ultrafied definition works for any finitely generated group.

A *non-principal ultrafilter* on the integers, denoted ω, is a nonempty collection of sets of integers with the following properties:

1. $\emptyset \notin \omega$.

2. If $S_1 \in \omega$ and $S_2 \in \omega$, then $S_1 \cap S_2 \in \omega$.

3. If $S_1 \subset S_2$ and $S_1 \in \omega$, then $S_2 \in \omega$.

4. It is *maximal* in the sense that for each $S \subset \mathbb{Z}$ exactly one of the following must occur: $S \in \omega$ or $\mathbb{Z} \setminus S \in \omega$.

5. ω does not contain any finite sets. (This is the *non-principal* aspect.)

Applying Zorn's Lemma it can be shown that these exist. Since we are only concerned with non-principal ultrafilter, we will abuse notation and simply refer to these as *ultrafilters*.

An equivalent way to view an ultrafilter is as a finitely additive probability measure, with no atoms, defined on all subsets of integers and which takes values in $\{0, 1\}$. This way of thinking is consistent with the intuition that the ultrafilter sees the behavior on only the large subsets of $\mathbb{N}$ and when this is done one says that ω-almost every integer has a given property when the set of integers with this property is in ω.

For an ultrafilter ω, a metric space (X, d), and a sequence of points $\langle y_i \rangle_{i\in\mathbb{N}}$. We define y to be the *ultralimit of* $\langle y_i \rangle_{i\in\mathbb{N}}$ *with respect to* ω, denoted $y = \omega\text{-}\lim_i y_i$, if and only if for all $\epsilon > 0$ one has $\{i \in \mathbb{N} : d(y_i, y) < \epsilon\} \in \omega$. Since metric spaces are Hausdorff, it follows that when the ultralimit of a sequence exists it is a unique point. It is easy to see that when the sequence $\langle y_i \rangle_{i\in\mathbb{N}}$ converges, then the ultralimit is the same as the ordinary limit; it is also useful to note that in a compact metric space ultralimits always exist.

Fix an ultrafilter ω and a family of based metric spaces (X_i, x_i, d_i). Using the ultrafilter, a pseudo-distance on $\prod_{i\in\mathbb{N}} X_i$ is provided by:

$$d_\omega(\langle a_i \rangle, \langle b_i \rangle) = \omega\text{-}\lim_i d_{X_i}(a_i, b_i) \in [0, \infty].$$

The *ultralimit of* (X_i, x_i) is then defined to be:[1]

$$\omega\text{-}\lim_i (X_i, x_i) = \{y \in \prod_{i\in\mathbb{N}} X_i : d_\omega(y, \langle x_i \rangle) < \infty\} / \sim,$$

where for two points $y, z \in \prod_{i\in\mathbb{N}} X_i$ we define $y \sim z$ if and only if $d_\omega(y, z) = 0$.

The pseudo-metric d_ω takes values in $[0, \infty]$, but requiring that points in $\omega\text{-}\lim(X_i, x_i)$ be a finite distance from the basepoint keeps the distance from obtaining the value ∞. The relation $\sim$ quotients out points whose distance from each other is zero; in an analogy to measure theory we think of this process as identifying sequences which agree almost everywhere. These two conditions combine to make $d_\omega(y, z)$ into a metric.

Just as ultralimits for sequences of points generalize the topological notion of limit, ultralimits of sequences of metric spaces generalize Hausdorff limits of metric spaces. In particular, for a Hausdorff precompact family of metric spaces the ultralimit of the sequence is a limit point with respect to the Hausdorff topology (see [KlLe] for a proof of this and other related facts).

Definition 2.2.1. The *asymptotic cone of* (X, x_0, d) *relative to the ultrafilter* ω is defined by:

$$\mathrm{Cone}_\omega(\mathrm{X}, \mathrm{x}_0) = \omega\text{-}\lim_{\mathrm{i}} (\mathrm{X}, \mathrm{x}_0, \frac{1}{\mathrm{i}}\mathrm{d})$$

[1]The term ultralimit is used by topologists and geometers; logicians tend to call this space the *ultraproduct of* (X_i, x_i)

When it is not a source of confusion we tend to suppress writing the basepoint. Also, when the choice of ultrafilter is unimportant we simply refer to the *asymptotic cone of* X and use the notation $\mathrm{Cone}_\infty \mathrm{X}$.

When taking such ultralimits, we often use the notation $\omega\text{-}\lim_i \frac{1}{i}(x_\alpha, 0_\alpha)$, when there is an implied ambient space T_α with basepoint 0_α, in order to emphasize that we are looking at an ultralimit in $\mathrm{Cone}_\infty(\mathrm{T}_\alpha, 0_\alpha)$.

An equivalent way to think of the asymptotic cone is as equivalence classes of sequences of points in X whose distance from the basepoint grow at a linear rate.

When considering the asymptotic cone, especially for the first time, the following quotation from Thomas Hobbes (1588–1679), will likely seem relevant.

> To understand this for sense it is not required that a man should be a geometrician or a logician, but that he should be mad.

Asymptotic cones provide a way to replace a metric space with one "limiting" space which carries information about sequences in the original space which leave every compact set. This process encodes the asymptotic geometry of a space into standard algebraic topology invariants of its asymptotic cone. As mentioned in the introduction, the first interesting example of this was due to Gromov [G1]—part of what he showed is summarized by:

Theorem 2.2.2. *Let G be a finitely generated group. If G has polynomial growth then every asymptotic cone of G is locally compact.*

A common use of asymptotic cones is in their relation to hyperbolicity as demonstrated by the next result and Theorem 2.2.5. First we give a preliminary definition.

Definition 2.2.3. An $\mathbb{R}$-tree is a metric space (X, d) such that between any two points $a, b \in X$ there exists a unique topological arc γ connecting them and γ is isometric to the interval $[0, d_X(a, b)] \subset \mathbb{R}$.

The simplest examples of $\mathbb{R}$-trees are simplicial trees, but in general $\mathbb{R}$-trees may be much more complicated. Indeed, $\mathbb{R}$-trees need not even be locally compact. An example of an $\mathbb{R}$-tree which is not locally compact, but nonetheless easy to visualize is $\mathbb{R}^2$ with the metric:

$$d_{tree}(a, b) = \begin{cases} |a-b| \text{ when } a = \alpha b \text{ for some constant } \alpha \in \mathbb{R} \\ |a| + |b| \text{ when } a \text{ and } b \text{ are linearly independent} \end{cases}$$

Sometimes called the *Paris metric*—distance is measured by the path distance along train tracks, where all trains run on straight lines through the center of the city. This metric has "uncountable branching" at the origin, but everywhere else locally looks just like $\mathbb{R}$.

Proposition 2.2.4. *For a sequence $\delta_i \to 0$, the ultralimit of δ_i-hyperbolic spaces is an $\mathbb{R}$-tree.*

In particular, if X is a δ-hyperbolic space, then $\mathrm{Cone}_\infty \mathrm{X}$ *is an $\mathbb{R}$-tree.*

If X is a hyperbolic space which is sufficiently complicated, then $\mathrm{Cone}_\infty \mathrm{X}$ is significantly more complicated than either of the examples considered above, indeed these $\mathbb{R}$-trees have uncountable branching at every point (the technical hypothesis needed is $|\partial X| > 2$ which rules out only the simplest cases).

A partial converse to Proposition 2.2.4 is provided by the following result which is well known. The statement of this theorem first appeared, without proof, in [G2]. A proof is given in [D].

Theorem 2.2.5. *If every asymptotic cone of a metric space X is an $\mathbb{R}$-tree, then X is δ-hyperbolic.*

Fix an ultrafilter ω and for a surface S define Seq to be the set of sequences of homotopy classes of essential, non-peripheral subsurfaces $X \subset S$ with $\xi(X) \neq 0$, considered up to the relation $\sim$, where two sequences $\alpha = \langle \alpha_i \rangle_{i \in \mathbb{N}}$ and $\beta = \langle \beta_i \rangle_{i \in \mathbb{N}}$ satisfy $\alpha \sim \beta$ if and only if $\alpha_i = \beta_i$ for each i in some set $K \in \omega$.

The following is an easy observation that simplifies the situation when one is dealing with finitely many equivalence classes of Seq.

Lemma 2.2.6. *For any finite set $\Gamma \subset$ Seq, the elements of Γ are pairwise distinct in* Seq *if and only if there exists a set $K \in \omega$ for which each $\gamma, \gamma' \in \Gamma$ has $\gamma_i \neq \gamma'_i$ for every $i \in K$*

Proof:

Fix a finite set $\Gamma \subset$ Seq.

($\Longrightarrow$).

The maximal clause in the definition of ultrafilter states that for each $K \subset \mathbb{N}$ either $K \in \omega$ or $\mathbb{N} \setminus K \in \omega$. We suppose that for each pair of elements $\gamma, \gamma' \in \Gamma$ we have $\gamma \nsim \gamma'$ and thus the set K of indices where $\gamma_i = \gamma'_i$must have $K \notin \omega$. Maximality of the ultrafilter then implies that $\mathbb{N} \setminus K \in \omega$. Thus for each $i \in \mathbb{N} \setminus K \in \omega$ we have $\gamma_i \neq \gamma'_i$, we define $\mathbb{N} \setminus K = K'_{\gamma,\gamma'}$.

Since ultrafilters are closed under finite intersections, the intersection over all pairs γ, γ' of $K_{\gamma,\gamma'}$ yields a set $J \in \omega$ where $\gamma_i \neq \gamma'_i$ for every $i \in J$ and every $\gamma, \gamma' \in \Gamma$.

($\Longleftarrow$). Let $K \in \omega$ be the set of indices for which $\gamma_i \neq \gamma_i'$ for each $i \in K$ and $\gamma, \gamma' \in \Gamma$.

Then for any γ, γ' we have $\{i : \gamma_i = \gamma_i'\} \in \mathbb{N} \setminus K$ and hence is not in ω. Thus $\gamma \nsim \gamma'$.

□

We end this section with a summary of some standard results about asymptotic cones that we will use in the sequel (see [KlLe] or [Ka]).

Proposition 2.2.7. *Fix a non-principal ultrafilter* ω

1. $\mathrm{Cone}_\omega(\mathrm{X})$ *is a complete metric space.*

2. $\mathrm{Cone}_\omega(\mathrm{X}_1 \times \mathrm{X}_2) = \mathrm{Cone}_\omega(\mathrm{X}_1) \times \mathrm{Cone}_\omega(\mathrm{X}_2)$

3. $\mathrm{Cone}_\omega \mathbb{R}^\mathrm{n} = \mathbb{R}^\mathrm{n}$

4. *The asymptotic cone of a geodesic space is a geodesic space.*

5. *A (K, C)-quasi-isometry between metric spaces induces a K-bi-Lipschitz map between their asymptotic cones.*

Chapter 3

Surfaces and their Friends

3.1 Complex of Curves

We use S to denote a connected, orientable surface of genus $g = g(S)$ with $p = p(S)$ punctures. (Note that we will use interchangeably the words puncture and boundary, as the distinction does not affect the results contained in this work.) We use the terms *subsurface* and *domain* to refer to a homotopy class of an essential, non-peripheral, connected subsurface of S (subsurfaces are not assumed to be proper unless explicitly stated). When we refer to the boundary of a domain this will mean the collection of non-peripheral closed curves which bound the domain as a subset of S.

Here we recall the construction of the complex of curves and relevant machinery developed by Masur and Minsky which we will use in our study; for further details consult [MM2]. We use $\xi(S) = 3g(S) + p(S) - 3$ to quantify the complexity of the surface S.[1] Recall that when positive, $3g(S) + p(S) - 3$ is the maximal number of disjoint homotopy classes of essential, non-peripheral simple closed curves which can be simultaneously realized on S. Naturality of $\xi(S)$ as a measure of complexity is justified by the property that it decreases when one passes from a surface to a proper subsurface (recall our convention of considering surfaces up to homotopy). Since our interest is in hyperbolic surfaces and their subsurfaces, we only consider subsurfaces with $\xi > -2$ (thus ignoring the disk and the sphere). Additionally, as it is not a hyperbolic surface nor does it appear as a subsurface of any hyperbolic surface we will usually ignore $S_{1,0}$ (although much of our discussion has analogues for this case); thus $\xi = 0$ is used only to denote the three punctured sphere.

Introduced by Harvey [Harv] to study the boundary of Teichmüller space,

[1] In [MM2] they use $\xi(S) = 3g(S) + p(S)$, but we use the current version because it has better additive properties when generalized to disconnected surfaces, as we need in Chapter 6

the complex of curves has proven to be a useful tool in the study of Teichmüller space, mapping class groups, and 3-manifolds. The complex of curves is a finite dimensional complex which encodes information about the surface via the combinatorics of simple closed curves. Analysis using the curve complex is necessarily delicate since the complex is locally infinite except for in a few low genus cases.

Definition 3.1.1. The *complex of curves for* S, denoted $\mathcal{C}(S)$, consists of a *vertex* for every homotopy class of a simple closed curve which is both non-trivial and non-peripheral. The N*-simplices* of $\mathcal{C}(S)$ are given by collections of $N+1$ vertices whose homotopy classes can all simultaneously be realized disjointly on S.

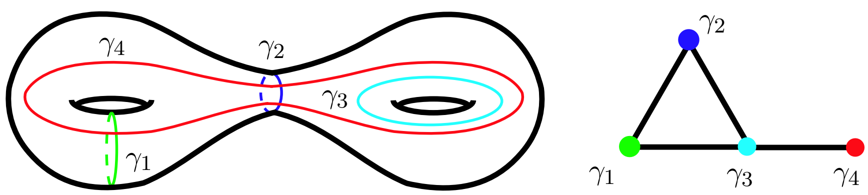

Figure 3.1: Example of vertices and edges of the curve complex

This definition works well when $\xi(S) > 1$, however it must be modified slightly for the surfaces of small complexity, which we refer to as the *sporadic cases*. The two cases where $\xi(S) = 1$ are $S_{1,1}$ and $S_{0,4}$ (the following discussion works for $S_{1,0}$ as well); here any two distinct homotopy classes of essential simple closed curves must intersect (at least once in $S_{1,1}$ and twice in $S_{0,4}$). In these cases the above definition of $\mathcal{C}(S)$ would have no edges, only vertices. Accordingly, we modify the definition so that an edge is added between two homotopy classes when they intersect the minimal amount possible on the surface (i.e., once for $S_{1,1}$ and twice for $S_{0,4}$). With this definition $\mathcal{C}(S_{1,1})$ and $\mathcal{C}(S_{0,4})$ are each connected, indeed they are each isometric to the classical Farey graph.

When $\xi(S) = 0, -2$, or -3 then $\mathcal{C}(S)$ is empty. The final modification we make is for the case where we have $A \subset S$ with $\xi(A) = -1$, the annulus. The annulus doesn't support a finite area hyperbolic metric, so our interest in it derives from the fact that it arises as a subsurface of hyperbolic surfaces. Indeed, annuli will play a crucial role as they will be used to capture information about Dehn twists. Given an annulus $A \subset S$, we define $\mathcal{C}(A)$ to be based homotopy classes of arcs connecting one boundary component of the annulus to the other. More precisely, denoting by $\widetilde{A}$ the annular cover of S to which A lifts homeomorphically, we use the compactification of $\mathbb{H}^2$ as the closed unit disk to obtain a closed annulus $\widehat{A}$. We define the vertices of $\mathcal{C}(A)$ to be homotopy classes of paths connecting one boundary component of $\widehat{A}$ to the other, where the homotopies are required to fix the endpoints. Edges of $\mathcal{C}(A)$ are pairs of vertices which have representatives with disjoint interiors. Giving edges a euclidean metric of length one, in [MM2] it is proven that $\mathcal{C}(A)$ is quasi-isometric to $\mathbb{Z}$.

The following foundational result, which is the main theorem of [MM1], will be useful in our later analysis. For another proof see [Bo1], which provides a constructive computation of a bound on the hyperbolicity constant.

Theorem 3.1.2. (Hyperbolicity of $\mathcal{C}(S)$, [MM1] and [Bo1]). *For any surface S, $\mathcal{C}(S)$ is an infinite diameter δ-hyperbolic space (as long as it is non-empty).*

Combined with Proposition 2.2.4, this yields the following consequence which we will use in Chapter 6.

Corollary 3.1.3. $\mathrm{Cone}_\infty \mathcal{C}(\mathrm{S})$ *is an* $\mathbb{R}$*-tree.*

Throughout this paper we use the convention that "intersection" refers to transverse intersection. Thus for example if we consider a subsurface $Y \subsetneq S$ and an element $\gamma \in \partial Y$ which is not homotopic to a puncture of S then we consider the annulus around γ to not intersect Y.

For a surface with punctures one can consider the *arc complex* $\mathcal{C}'(S)$, which is a close relative of the curve complex. When $\xi(S) > -1$ we define the vertices[2] of $\mathcal{C}'(S)$ (denoted $\mathcal{C}'_0(S)$) to consist of elements of $\mathcal{C}_0(S)$ as well as homotopy classes of simple arcs on S with endpoints lying on punctures of S, which don't bound a disk or a once punctured disk on either side. As done for the curve complex, we define N-simplices of $\mathcal{C}'(S)$ to be collections of $N+1$ vertices which can simultaneously be realized on the surface as disjoint arcs and curves. For annuli we define $\mathcal{C}(A) = \mathcal{C}'(A)$

The arc complex arises naturally when one tries to "project" an element $\gamma \in \mathcal{C}(S)$ into $\mathcal{C}(Y)$ where $Y \subset S$. When $\xi(Y) > 0$ we define

$$\pi'_Y : \mathcal{C}_0(S) \to 2^{\mathcal{C}'_0(Y)}$$

by the following:

- If $\gamma \cap Y = \emptyset$, then define $\pi'_Y(\gamma) = \emptyset$.

[2]This definition differs from that in [Hare] where an arc complex is considered consisting of only arcs and thus does not contain $\mathcal{C}(S)$ as a subcomplex. Our definition agrees with that in [MM2].

- If $\gamma \subset Y$, then define $\pi'_Y(\gamma) = \{\gamma\}$.
- If $\gamma \cap \partial Y \neq \emptyset$, then after putting γ in a position so it has minimal intersection with ∂Y we identify parallel arcs of $\gamma \cap Y$ and define $\pi'_Y(\gamma)$ to be the union of these arcs and any closed curves in $\gamma \cap \partial Y \neq \emptyset$.

In the last case, it follows from the definition of $\mathcal{C}'(Y)$ that $\pi'_Y(\gamma)$ is a subset of $\mathcal{C}'_0(Y)$ with diameter at most one. Thus whenever $\pi'_Y(\gamma) \neq \emptyset$, it is a subset of $\mathcal{C}'(Y)$ of diameter at most one. Moreover, when $\xi(S) > 0$, there is a map $\phi_Y : \mathcal{C}'_0(Y) \to 2^{\mathcal{C}_0(Y)}$ which embeds the arc complex as a cobounded subset of the curve complex: the map ϕ_Y sends each arc to the boundary curves of a regular neighborhood of its union with ∂Y. For $Y \subset S$ with $\xi(Y) > 0$ we define $\pi_Y = \phi_Y \circ \pi'_Y : \mathcal{C}(S) \to \mathcal{C}(Y)$.

When $\xi(Y) = -1$ then any curve γ which crosses Y transversally has a lift $\widetilde{\gamma} \in \widetilde{Y}$ with at least one component which connects the two boundary components of $\widehat{Y}$. Together, the collection of lifts which connect the boundary components form a finite subset of $\mathcal{C}'(Y)$ with diameter at most 1, define $\pi_Y(\gamma)$ to be this set. If γ doesn't intersect Y or is the core curve of Y, then define $\pi_Y(\gamma) = \emptyset$. Also, for consistency of definitions, we define $\pi_Y : v \to \{v\}$ for $v \in \mathcal{C}_0(Y)$.

Since we often work with subsets of $\mathcal{C}(Y)$ rather than points, we use the following notation. First, for a set valued map $f : X \to 2^Y$, we adopt the notation $f(A) = \cup_{x \in A} f(x)$, thus allowing us to consider f as a map to Y. If $X \subseteq S$ and μ, ν are markings we define

$$d_X(\mu, \nu) = d_{\mathcal{C}(X)}(\pi_X(\mu), \pi_X(\gamma)).$$

Given sets $A, B \in \mathcal{M}(S)$ we set $d_X(A, B) = \min\{d_X(\alpha, \beta) : \alpha \in A \text{ and } \beta \in B\}$. Also, we write $\mathrm{diam_X(A)}$ to refer to the diameter of the set $\pi_{\mathcal{C}(X)}(A)$ and $\mathrm{diam_X(A, B)}$ for $\mathrm{diam_{\mathcal{C}(X)}(A \cup B)}$ in order to emphasize the symmetry between our use of minimal distance and diameter. (These same conventions apply to markings (as defined in the next section) as well as elements of the curve complex.)

An extremely useful result concerning these projections is the following (see [MM2] for the original proof, and [M] where the bound is corrected from 2 to 3):

Lemma 3.1.4. (Lipschitz projection, [MM2]). *Let Y be a subdomain of S. For any simplex ρ in $\mathcal{C}(S)$, if $\pi_Y(\rho) \neq \emptyset$ then* $\mathrm{diam_Y}(\rho) \leq 3$.

In light of this lemma, for any pair of subdomains $Y \subset Z$ of S we consider $\mathcal{C}(Z) \setminus B_1(\partial Y)$ with the path metric, i.e., the distance between $\gamma, \gamma' \in \mathcal{C}(Z) \setminus B_1(\partial Y)$ is the length of the shortest path connecting them in $\mathcal{C}(Z) \setminus B_1(\partial Y)$.

With the the path metric there is a *coarsely Lipschitz* map $\pi_{Z\to Y} : \mathcal{C}(Z) \setminus B_1(\partial Y) \to \mathcal{C}(Y)$, where a map is defined to be (K, C)-*coarsely Lipschitz* if and only if there exist a pair of constants K, C such that for each $a, b \in \mathcal{C}(Z) \setminus B_1(\partial Y)$ we have $d_Y(\pi_{Z\to Y}(a), \pi_{Z\to Y}(b)) \leq K d_{\mathcal{C}(Z)\setminus B_1(\partial Y)}(a, b) + C$. By the above lemma and the fact that $\pi_{Z\to Y}$ sends points of $\mathcal{C}(Z) \setminus B_1(\partial Y)$ to subsets of $\mathcal{C}(Y)$ of diameter at most 3, we have the following which we shall use in Chapter 4. (See Figure 3.2 for a cartoon of the behavior of this map near ∂Y.)

Corollary 3.1.5. *When $\mathcal{C}(Z) \setminus B_1(\partial Y)$ is endowed with the path metric, then for any domain $Y \subset Z$, we have*

$$\pi_{Z\to Y} : \mathcal{C}_1(Z) \setminus B_1(\partial Y) \to \mathcal{C}(Y)$$

is coarsely Lipschitz (with constants $K = 3$ and $C = 3$).

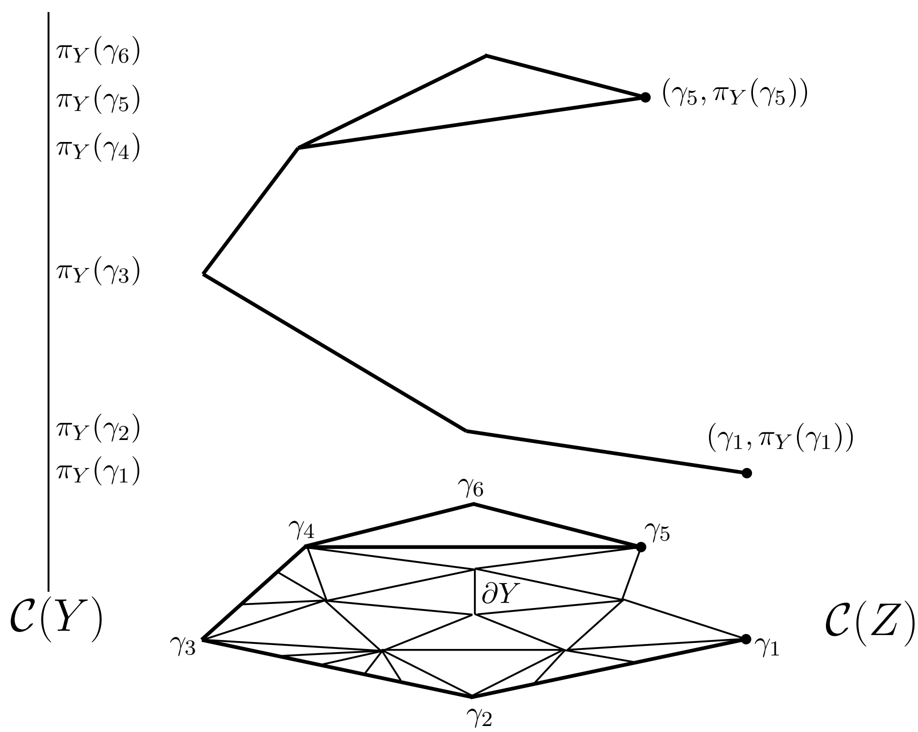

Figure 3.2: Letting $Y \subset Z$, the above is a caricature of the graph of $\pi_{Z\to Y} : \mathcal{C}(Z) \setminus B_1(\partial Y) \to \mathcal{C}(Y)$ near the point $\partial Y \in \mathcal{C}(Z)$.

In Figure 3.2, we observe that although $\pi_{Z\to Y}$ is Lipschitz, pairs of points like γ_1 and γ_5, although distance 2 in $\mathcal{C}(S)$, can be found which are arbitrarily far apart in the path metric on $\mathcal{C}(Z) \setminus B_1(\partial Y)$ and thus their distance when projected to Y may be very large.

Letting $\mathcal{C}_1(S)$ denote the one-skeleton of $\mathcal{C}(S)$, we state the following theorem which is one of the key technical results from [MM2].

Theorem 3.1.6. (Bounded geodesic image). *Let $Y \subsetneq S$ with $\xi(Y) \neq 0$ and let g be a geodesic segment, ray, or line in $\mathcal{C}_1(S)$, such that $\pi_Y(v) \neq \emptyset$ for every vertex v of g. There is a constant D depending only on $\xi(S)$ so that* $\mathrm{diam}_{\mathrm{Y}}(\mathrm{g}) \leq \mathrm{D}$.

3.2 Markings

In this section we describe the quasi-isometric model we will use for the mapping class group and explain the tools developed in [MM2] for computing with this model.

Definition 3.2.1. A *marking*, μ, on S is a collection of *base curves* to each of which we (may) associate a *transverse curve*. These collections are made subject to the constraints:

- The *base curves*, $\mathrm{base}(\mu) = \{\gamma_1, \ldots, \gamma_n\}$, consists of a simplex in $\mathcal{C}(S)$.
- The *transverse curve*, t, associated to a given base curve γ is either empty or an element of $\mathcal{C}_0(S)$ which intersects t once (or twice if $S = S_{0,4}$) and projects to a subset of diameter at most 1 in the annular complex $\mathcal{C}(\gamma)$.

When the transverse curve t is empty we say that γ doesn't have a transverse curve.

If the simplex formed by $\mathrm{base}(\mu)$ in $\mathcal{C}(S)$ is top dimensional and every curve has a transverse curve, then we say the marking is *complete*.

If a marking has the following two properties, then we say the marking is *clean*. First, for each γ, its transversal t is disjoint from the rest of the base curves. Second, for each γ and t as above their union $t \cup \gamma$ fills a surface denoted $F(t,\gamma)$, with $\xi(F(t,\gamma)) = 1$ and in which $d_{\mathcal{C}(F(t,\gamma))}(t,\gamma) = 1$. (See Figure 3.2.)

Let μ denote a complete clean marking with pairs $(\alpha_i, t(\alpha_i))$, we take as *elementary moves* the following two relations on the set of complete clean markings:

1. *Twist*: For some i, we replace $(\alpha_i, t(\alpha_i))$ by $(\alpha_i, t'(\alpha_i))$ where $t'(\alpha_i)$ is the result of one full (or half when possible) twist of $t(\alpha_i)$ around α_i. The rest of the pairs are left unchanged.

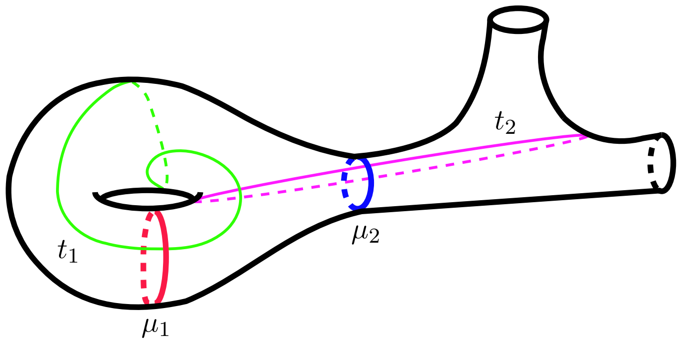

Figure 3.3: A complete clean marking $\mu \in \mathcal{M}(S_{1,2})$ with $\text{base}(\mu) = \mu_1 \cup \mu_2$.

2. *Flip*: For some i we swap the roles of the base and transverse curves between α_i and $t(\alpha_i)$. After doing this the complete marking may no longer be clean, one then needs to replace the new marking with a compatible clean one.

We say the clean marking μ is *compatible* with μ' if they have the same base curves, each base curve γ has a transverse curve t in one marking if and only if it has a transverse curve t' in the other marking, and when t exists then $d_\gamma(t, t')$ is minimal among all choices of t'.

In [MM2] it is shown that there exists a bound (depending only on the topological type of S) on the number of clean markings which are compatible with any other given marking. Thus even though the flip move is defined by choosing an arbitrary compatible complete clean marking, it is canonical up to some uniformly bounded amount of ambiguity.

One then defines the *marking complex*, denoted $\mathcal{M}(S)$, to be the graph formed by taking the complete clean markings on S as vertices and connecting two vertices by an edge if they differ by an elementary move. It is not hard to check that $\mathcal{M}(S)$ is a locally finite graph and that the mapping class group acts on it cocompactly and properly discontinuously. A consequence of this which we use throughout this work is:

Corollary 3.2.2. [MM2]. *$\mathcal{M}(S)$ is quasi-isometric to the mapping class group.*

For any subsurface $Y \subset S$ with $\mathcal{C}(Y) \neq \emptyset$ we considered in the previous section the *subsurface projection*, $\pi_Y : \mathcal{C}(S) \to 2^{\mathcal{C}(Y)}$. More generally one can consider subsurface projections from the marking complex; since this definition generalizes the above map we also denote it $\pi_Y : \mathcal{M}(S) \to 2^{\mathcal{C}(Y)}$. When $\mu \in \mathcal{M}(S)$ we define $\pi_Y(\mu)$ to be $\pi_Y(\text{base}(\mu))$ unless Y is an annulus around a curve $\gamma \in \text{base}(\mu)$, in which case we define $\pi_Y(\mu) = \pi_Y(t)$ where t is the transverse curve to γ. An important observation is that when μ is a complete marking on S then $\pi_Y(\mu) \neq \emptyset$ for each $Y \subset S$.

There is also a Lipschitz projection property for projections of markings.

Lemma 3.2.3. (Elementary move projection, [MM2]). *If $\mu, \mu' \in \mathcal{M}(S)$ differ by one elementary move, then for any domain $Y \subset S$ with $\xi(Y) \neq 0$,*

$$d_Y(\mu, \mu') \leq 4.$$

3.3 Hierarchies

Thurston's classification theorem for surface homeomorphisms [Th1] gives a layered structure for homeomorphisms based upon studying subsurfaces which are preserved (perhaps under an iterate of the homeomorphism). In [MM2], Masur and Minsky refine and further elucidate a layered structure for the mapping class group. Using the marking complex and the complex of curves, they provide a way to compare how relatively complicated two given mapping class group elements are on any subsurface (not just ones that are eventually periodic!). All the subsurface comparisons are tied together into one object called a hierarchy.

Definition 3.3.1. Let $Y \subset S$ with $\xi(Y) > 1$. A sequence of simplices $v_0, v_1, \ldots, v_n$ is called *tight* if:

1. For each $0 \leq i < j \leq n$, and vertices $v_i' \subset v_i$ and $v_j' \subset v_j$,

$$d_Y(v_i', v_j') = |i - j|.$$

2. For each $0 < i < n$ we have that v_i is the boundary of the subsurface filled by $v_{i-1} \cup v_{i+1}$.

When $\xi(Y) = 1$ we consider a sequence to be *tight* if and only if it is the vertex sequence of a geodesic.

When $\xi(Y) = -1$ we consider a sequence to be *tight* if and only if it is the vertex sequence of a geodesic where the set of endpoints on $\partial\hat{Y}$ of arcs representing the vertices equals the set of endpoints of the first and last arc.

We often use the following decorated version of a tight sequence.

Definition 3.3.2. A *tight geodesic* g in $\mathcal{C}(Y)$ consists of a tight sequence $v_0, v_1, \ldots, v_n$ and a pair of markings $\mathbf{I} = \mathbf{I}(g)$ and $\mathbf{T} = \mathbf{T}(g)$ (called the *initial* and *terminal markings* for g) such that v_0 is a vertex of base($\mathbf{I}$) and v_n is a vertex of base($\mathbf{T}$).

The integer n is called the length of g. The *domain* (sometimes called the *support*) of g refers to the surface Y, written $D(g) = Y$.

Below we explain a relationship between tight geodesics occurring in different subsurfaces.

When $\mu \in \mathcal{M}(S)$ we write $\mu|_Y$ to denote the restriction of this marking to Y, by which we mean:

- If $\xi(Y) = -1$, then $\mu|_Y = \pi_Y(\mu)$
- Otherwise, we let $\mu|_Y$ be the set of base curves which meet Y essentially, each taken with their associated transversal.

For a surface Y with $\xi(Y) \geq 1$ and a simplex $v \subset \mathcal{C}(Y)$, we say that X is a *component domain* of (Y, v) if either: X is a component of $Y \setminus v$ or X is an annulus whose core curve is a component of v. More generally, we say that a subsurface $X \subset Y$ is a *component domain of* g if for some simplex v_i in g, X is a component domain of $(D(g), v_i)$. Notice that v_i is uniquely determined by g and X.

Furthermore, when X is a component domain of g and $\xi(X) \neq 0$, define the *initial marking of* X *relative to* g:

$$\mathbf{I}(X, g) = \begin{cases} v_{i-1}|_X & v_i \text{ is not the first vertex (of } g) \\ \mathbf{I}(g)|_X & v_i \text{ is the first vertex} \end{cases}$$

Similarly define the *terminal marking of* X *relative to* g to be:

$$\mathbf{T}(X, g) = \begin{cases} v_{i+1}|_X & v_i \text{ is not the last vertex} \\ \mathbf{T}(g)|_X & v_i \text{ is the last vertex} \end{cases}$$

Observe that these are each markings, since ∂X is distance 1 in $\mathcal{C}(S)$ from each of $v_{i\pm1}$, or in the case where v_i is the first vertex, then ∂X is disjoint from base($\mathbf{I}(g)$) (similarly for the terminal markings).

When X is a component domain of g with $\mathbf{T}(X, g) \neq \emptyset$, then we say that X is *directly forward subordinate* to g, written $X \overset{d}{\searrow} g$. Similarly, when $\mathbf{I}(X, g) \neq \emptyset$ we say that X is *directly backward subordinate* to g, written $g \overset{d}{\swarrow} X$.

The definition generalizes to geodesics as follows.

Definition 3.3.3. Let g and h be tight geodesics. We say that g is *directly forward subordinate* to h, written $g \overset{d}{\searrow} h$, when $D(g) \overset{d}{\searrow} h$ and $\mathbf{T}(g) = \mathbf{T}(D(g), h)$. Similarly, h is *directly backward subordinate* to g, written $g \overset{d}{\swarrow} h$, when $g \overset{d}{\swarrow} D(h)$ and $\mathbf{I}(h) = \mathbf{I}(D(h), g)$.

We write *forward subordinate*, or $\searrow$, to denote the transitive closure of $\overset{d}{\searrow}$; similarly we define $\swarrow$.

We can now define the main tool which was introduced in [MM2].

Definition 3.3.4. A *hierarchy (of geodesics)* H, on S is a collection of tight geodesics subject to the following constraints:

1. There exists a tight geodesic whose support is S. This geodesic is called the *main geodesic* and is often denoted g_H. The initial and terminal markings of g_H are denoted $\mathbf{I}(H)$ and $\mathbf{T}(H)$.

2. Whenever there exists a pair of tight geodesics $g, k \in H$ and a subsurface $Y \subset S$ such that $g \overset{d}{\swarrow} Y \overset{d}{\searrow} k$ then H contains a unique tight geodesic h with domain Y such that $g \overset{d}{\swarrow} h \overset{d}{\searrow} k$.

3. For each geodesic $h \in H$ other than the main geodesic, there exists $g, k \in H$ so that $g \overset{d}{\swarrow} h \overset{d}{\searrow} k$.

Using an inductive argument, it is proved in [MM2] that given any two markings on a surface, there is a hierarchy which has initial marking one of them and terminal marking the other. The process begins with picking a base curve of the initial marking, one in the terminal marking, and a geodesic in $\mathcal{C}(S)$ between them (the main geodesic). Then, the second condition in the definition of a hierarchy forces certain proper subdomains to support geodesics: if there is a configuration $g \overset{d}{\swarrow} D \overset{d}{\searrow} k$ such that D does not support any geodesic h with $\mathbf{I}(h) = \mathbf{I}(D, g)$ and $\mathbf{T}(h) = \mathbf{T}(D, k)$, then we construct such a geodesic. When this geodesics is constructed, we can choose the initial vertex to be any element of base $\mathbf{I}(Y, g)$ (similarly for the terminal vertex). This process continues until enough geodesic are included in H so that the second and third conditions are satisfied. As the above suggests, there is in general not just one hierarchy, but many of them connecting any pair of markings. In our proof of Theorem 4.2.1 (Projection estimates) we will exploit this flexibility by building our hierarchy subject to certain additional constraints which we find useful.

When we consider several hierarchies at once, we use the notation $g_{H,Y}$ to denote the geodesic of H supported on Y, when this geodesic exists it is unique by Theorem 3.3.5.

For any domain $Y \subset S$ and hierarchy H, the *backward* and *forward sequences* are given respectively by

$$\Sigma_H^-(Y) = \{b \in H : Y \subseteq D(b) \; and \; I(b)|_Y \neq \emptyset\}$$

and

$$\Sigma_H^+(Y) = \{f \in H : Y \subseteq D(b) \; and \; T(f)|_Y \neq \emptyset\}.$$

The following theorem summarizes some results which are useful for making computations with hierarchies.

Theorem 3.3.5. (Structure of Sigma; [MM2]). *Let H be a hierarchy, and Y any domain in its support.*

1. *If $\Sigma_H^+(Y)$ is nonempty then it has the form of a sequence*

$$f_0 \stackrel{d}{\searrow} \cdots \stackrel{d}{\searrow} f_n = g_H,$$

where $n \geq 0$. Similarly, if $\Sigma_H^-(Y)$ is nonempty then it has the form of a sequence

$$g_H = b_m \stackrel{d}{\swarrow} \cdots \stackrel{d}{\swarrow} b_0,$$

where $m \geq 0$.

2. *If $\Sigma_H^{\pm}(Y)$ are both nonempty and $\xi(Y) \neq 0$, then $b_0 = f_0$, and Y intersects every vertex of f_0 nontrivially.*

3. *If Y is a component domain of a geodesic $k \in H$ and $\xi(Y) \neq 0$, then*

$$f \in \Sigma_H^+(Y) \iff Y \searrow f,$$

and similarly,

$$b \in \Sigma_H^-(Y) \iff b \swarrow Y.$$

If, furthermore, $\Sigma^{\pm}(Y)$ are both nonempty, then in fact Y is the support of the geodesic $b_0 = f_0$.

4. *Geodesics in H are determined by their supports. That is, if $D(h) = D(h')$ for $h, h' \in H$ then $h = h'$.*

Given a hierarchy, the following provides a useful criterion for determining when a domain is the support of some geodesic in that hierarchy.

Lemma 3.3.6. (Large link; [MM2]). *If Y is any domain in S and*

$$d_Y(I(H), T(H)) > M_2,$$

then Y is the support of a geodesic h in H, where M_2 only depends on the topological type of S.

The geodesics in a hierarchy admit a partial ordering, which generalizes both the linear ordering on vertices in a geodesic and the ordering coming from forward and backwards subordinacy. Below we recall the basic definitions and a few properties of this ordering.

Given a geodesic g in $\mathcal{C}_1(S)$ and a subsurface Y, define the *footprint* of Y in g, denoted $\phi_g(Y)$ to be the collection of vertices of g which are disjoint from Y.

It is easy to see that the diameter of this set is at most 2 and with a little work it can be shown that the footprint is always an interval of diameter at most 2. The proof that $\phi_g(Y)$ is interval uses the assumption that the geodesic is tight: this is the only place where tightness of geodesics gets used. That footprints form intervals is useful as it allows one to make the following definition.

Definition 3.3.7. For a pair of geodesics $g, k \in H$ we say *g precedes k in the time order*, or $g \prec_t k$, if there exists a geodesic $m \in H$ so that $D(g)$ and $D(k)$ are both subsets of $D(m)$ and

$$\max \phi_m(D(g)) < \min \phi_m(D(k)).$$

We call m the *comparison geodesic*.

Time ordering is a (strict) partial ordering on geodesics in a hierarchy, this is proven in [MM2]. It is worth remarking that when a pair of geodesics are time ordered, they are time ordered with respect to a unique comparison geodesic.

The following provides a way to use the time ordering to gain information about the hierarchy.

Lemma 3.3.8. (Order and projections; [MM2]). *Let H be a hierarchy in S and $h, k \in H$ with $D(h) = Y$ and $D(k) = Z$. Suppose that $Y \cap Z \neq \emptyset$ and neither domain is contained in the other. Under these conditions, if $k \prec_t h$ then $d_Y(\partial Z, I(H)) \leq M_1 + 2$ and $d_Z(T(H), \partial Y) \leq M_1 + 2$. The constant M_1 only depends on the topological type of S.*

Since we will often use the above hypothesis, we introduce the terminology that a pair of subsurfaces Y and Z of S *overlap* when $Y \cap Z \neq \emptyset$ and neither domain is contained in the other.

The next result provides a way to translate distance computations in the mapping class group to computations in curve complexes of subsurfaces.

Theorem 3.3.9. (Move distance and projections; [MM2]). *There exists a constant $t(S)$ such that for each $\mu, \nu \in \mathcal{M}$ and any threshold $t > t(S)$, there exists $K(= K(t))$ and $C(= C(t))$ such that:*

$$\frac{1}{K} d_{\mathcal{M}}(\mu, \nu) - C \leq \sum_{\substack{Y \subseteq S \\ d_Y(\mu,\nu) > t}} d_Y(\mu, \nu) \leq K d_{\mathcal{M}}(\mu, \nu) + C.$$

The importance of Theorem 3.3.9 can not be understated, as this theorem provides the crucial result that hierarchies give rise to quasi-geodesic paths in the mapping class group.

3.4 Teichmüller Space and the Pants Complex

For a detailed reference on Teichmüller space, consult [ImTa].

Definition 3.4.1. For a fixed open topological surface S, the *Teichmüller space of* S is the space of equivalence classes of pairs (X, f), where X is a finite area hyperbolic surface and $f : S \to X$ is a homeomorphism. A pair (X_1, f_1) and (X_2, f_2) are considered equivalent if there exists an isometry $h : X_1 \to X_2$ such that $h \circ f_2$ is homotopic to f_1.

A topology is obtained on Teichmüller space by infimizing over the distortion of maps $h : X_1 \to X_2$. Topologically Teichmüller space is fairly easily understood, as the following classical result indicates:

Theorem 3.4.2. *Teichmüller space is homeomorphic to* $\mathbb{R}^{6g-6+2p}$.

There are several natural metrics on Teichmüller space and its metric structure is far less transparent than its topological structure. Often of interest is the natural complex structure which Teichmüller space carries. What will be important in the sequel is the *Weil-Petersson metric* on Teichmüller space, which is a Kähler metric with negative sectional curvature. We will not need the integral form of the definition, so we just mention that the metric is obtained by considering a natural identification between the complex cotangent space of Teichmüller space and the space of holomorphic quadratic differentials, then defining the Weil-Petersson metric to be the one dual to the L^2-inner product on the space of holomorphic quadratic differentials. See [Wol] for a survey on the Weil-Petersson metric and its completion.

Definition 3.4.3. A *pair of pants* is a thrice punctured sphere. A *pants decomposition* of a surface $S_{g,p}$ is a maximal collection of pairwise non-parallel homotopy classes of simple closed curves; this decomposition obtains its name from the observation that such a curve system cuts $S_{g,p}$ into $2g + p - 2$ pairs of pants.

It is easy to verify that any pants decomposition consists of exactly $3g+p-3$ disjoint simple closed curves on S and is thus a maximal simplex in $\mathcal{C}(S)$. These decompositions have long been useful in the study of mapping class groups and Teichmüller space.

Originally defined by [Hat], the *pants complex of* S, denoted $\mathcal{P}(S)$, is a way of comparing all possible pants decompositions on a fixed surface. This complex consists of a vertex for every pants decomposition and an edge between each pair of decompositions which differ by an *elementary move*. Two pairs of pants $P = \{\gamma_1, \ldots, \gamma_{3g+p-3}\}$ and $P' = \{\gamma'_1, \ldots, \gamma'_{3g+p-3}\}$ differ by an *elementary move* if P and P' can be reindexed so that both:

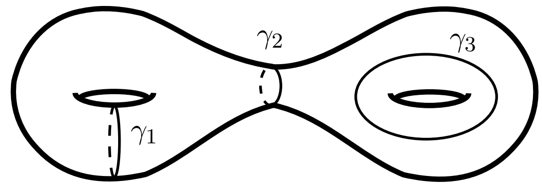

Figure 3.4: Example of a pants decomposition

1. $\gamma_i = \gamma_i'$ for all $2 \leq i \leq 3g+p-3$
2. In the component T of $S \backslash \cup_{2 \leq i \leq 3g+p-3} \gamma_i$ which is not a three-holed sphere (T is necessarily either a once punctured torus or a 4 punctured sphere) we have (see Figure 3.4):

$$d_{\mathcal{C}(T)}(\gamma_1, \gamma_1') = 1.$$

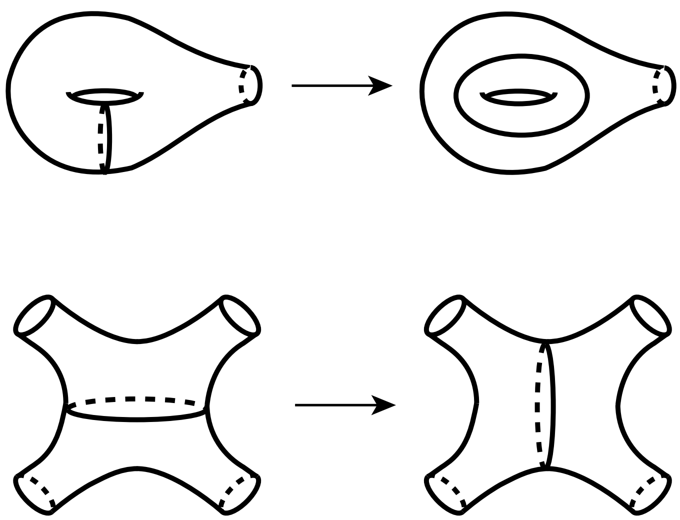

Figure 3.5: Elementary moves in the pants complex

$\mathcal{P}(S)$ is metrized by giving each edge the metric of the euclidean interval $[0,1]$. Our interest in this space comes out of the following remarkable theorem of Brock:

Theorem 3.4.4. ([Br]). *$\mathcal{P}(S)$ is quasi-isometric to the Teichmüller space of S with the Weil-Petersson metric.*

Noticing that a marking without transverse data is just a pants decomposition, in section 8 of [MM2] it is remarked that all of the constructions in their paper using hierarchies to obtain results about the marking complex can be replaced with analogous theorems for the pants complex. This is done by replacing markings (and hierarchies, etc) by markings without transverse data (and hierarchies-without-annuli, etc). The main result of this form we will use is:

Theorem 3.4.5. (Move distance and projections for $\mathcal{P}(S)$; [MM2]). *There exists a constant $t(S)$ such that for each $\mu, \nu \in \mathcal{P}$ (and any threshold $t > t(S)$) there exists $K(= K(t))$ and $C(= C(t))$ such that:*

$$\frac{1}{K} d_{\mathcal{P}}(\mu, \nu) - C \sum_{\substack{\text{non-annular } Y \subseteq S \\ d_Y(\mu,\nu) > t}} d_Y(\mu, \nu) \leq K d_{\mathcal{P}}(\mu, \nu) + C,$$

Chapter 4

Projection Estimates

In this section we address the following question: given a finite collection of subsurfaces $X_1, \ldots, X_N$ of S, what is the image of $\mathcal{M}(S)$ under the map which projects markings into the product of curve complexes of these subsurfaces, i.e., the image of $\mathcal{M}(S)$ as a subset of $\mathcal{C}(X_1) \times \cdots \times \mathcal{C}(X_N)$? For the case of two subsurfaces this question is answered in Theorems 4.2.1 and 4.2.2 using hierarchies in the curve complex as our main tool (see Section 3.1 for background). Theorem 4.3.6 generalizes the geometric description underlying Theorem 4.2.2 to arbitrary finite collections of subsurfaces by introducing the definition of a coarse tree-flat and using a delicate inductive step.

The first section of this chapter consists of a pair of toy examples which illustrates the dichotomy of Theorem 4.2.1 (Projection estimates), but which can be proven without use of the full hierarchy machinery. In Sections 4.2 we prove Theorem 4.2.1 and its geometric analogue Theorem 4.2.2. Finally, in Section 4.3 we prove Theorem 4.3.6 which is a generalization of the geometric picture in Theorem 4.2.2 to a description of the projection from the marking complex to the product of finitely many curve complexes.

4.1 A Motivating Example: the Dichotomy

Here we give two computations on the genus three surface. We exhibit a dichotomy between the behavior of projection maps into disjoint and intersecting pairs of subsurfaces. These examples illustrate the geometric meaning of Theorem 4.2.1 and provide a useful warm up for its proof, as the arguments we use here are a particularly easy form of what in general is a technical argument involving hierarchies.

Example 4.1.1. *(Disjoint subsurfaces).*

Let S be a closed genus three surface and let X and Y be subsurfaces of S which are each once-punctured tori and which can be realized disjointly on

S (figure 4.1.1).

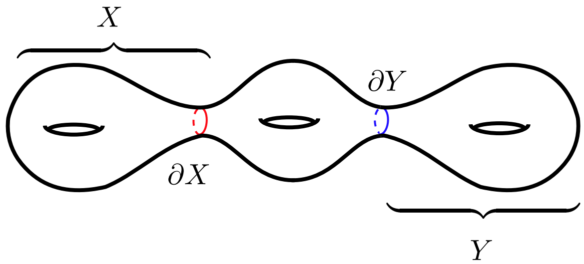

Figure 4.1: S and two disjoint subsurfaces

Let $\alpha \in \mathcal{C}(X)$ and $\beta \in \mathcal{C}(Y)$. Note that both these curves can also be considered as elements of $\mathcal{C}(S)$, then since $X \cap Y = \emptyset$ it follows that $d_{\mathcal{C}(S)}(\alpha, \beta) = 1$. Thus α and β are vertices of a common top dimensional simplex $\mu \subset \mathcal{C}(S)$. Taking a complete clean marking $\mu' \in \mathcal{M}(S)$ with base$(\mu') = \mu$, we have that the map

$$\mathcal{M}(S) \to \mathcal{C}(X) \times \mathcal{C}(Y)$$

is onto.

The next example provides the more interesting half of the dichotomy.

Example 4.1.2. *(Overlapping subsurfaces).* Let S be the closed genus three surface and consider two subsurfaces X and Y which each have genus two and one puncture. Moreover, suppose that X and Y overlap in a twice punctured torus, Z. Letting μ be an element of $\mathcal{C}(S)$ which intersects both X and Y. We shall show that both $d_X(\mu, \partial Y)$ and $d_Y(\mu, \partial X)$ can't simultaneously be large.

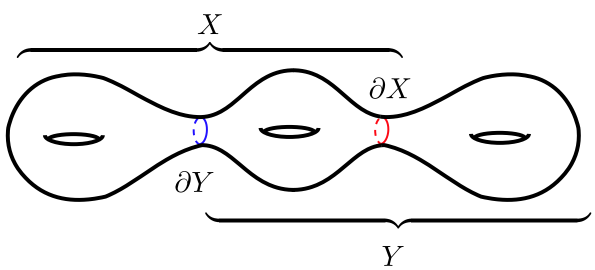

Figure 4.2: S and two overlapping subsurfaces

Let g be a geodesic in $\mathcal{C}_1(S)$ connecting μ to ∂X, with vertices $v_0, v_1, \ldots, v_n$ where $v_0 = \mu$ and $v_n = \partial X$. Since g is a geodesic in $\mathcal{C}(S)$, the curve v_{n-1} may be disjoint from X, but for all $k < n-1$ we have $v_k \cap \partial X \neq \emptyset$ and in particular,

$\pi_X(v_k) \neq \emptyset$ and $\pi_Y(v_k) \neq \emptyset$. Since $S = X \cup Y$ we have either $v_{n-1} \cap Y \neq \emptyset$ or $v_{n-1} \subset X$.

If $v_{n-1} \cap Y \neq \emptyset$, then every vertex of g intersects Y non-trivially. Theorem 3.1.6 (Bounded geodesic image; [MM2]) states that if g is a geodesic in $\mathcal{C}_1(S)$ with every vertex non-trivially intersecting a subsurface $Y \subset S$, then $\mathrm{diam_Y(\pi_Y(g))} < \mathrm{K}$ for a constant K depending only on the topological type of S. Accordingly, this theorem implies that we have $d_Y(\mu, \partial X) < K$ for a constant K depending only on the genus of S.

If $v_{n-1} \cap Y = \emptyset$, then it must be the case that $v_{n-1} \subset X$. In this case, we have $v_k \cap X \neq \emptyset$ for each $0 \leq k \leq n-1$. Applying the Bounded Geodesic Image Theorem tells us that $d_X(\mu, v_{n-1}) < K$. Since $v_{n-1} \cap Y = \emptyset$, we have in particular that $v_{n-1} \cap \partial Y = \emptyset$ which implies that $d_X(v_{n-1}, \partial Y) \leq 3$. Thus we may conclude that $d_X(\mu, \partial Y) < K + 3$.

Putting the two cases together, this example shows that there exists a constant $C = K+3$, for which any $\mu \in \mathcal{C}(S)$ must satisfy either $d_X(\mu, \partial Y) < C$ or $d_Y(\mu, \partial X) < C$.

If we take μ to be one of the base curves for some complete clean marking ν of S, we have shown that for any marking $\nu \in \mathcal{M}(S)$ either

$$d_X(\nu, \partial Y) < C \text{ or } d_Y(\nu, \partial X) < C.$$

We consider this a toy example, since the only tool we needed was the Bounded geodesic image Theorem. Difficulties in generalizing this example to arbitrary overlapping subsurfaces include problems such as: the boundaries of these subsurfaces may fill S or the union of these subsurfaces may not be all of S—both properties that were crucial to the above argument. To deal with these problems we use a powerful technical tool called a hierarchy, in which one considers not just a geodesic in the curve complex of S, but a family of geodesics each in the curve complex of some subsurface of S. The tools introduced by Masur and Minsky for calculating with these allow one to give the argument in the next section which is obtained by bootstrapping the key idea used in the above example.

4.2 Projection Estimates Theorem

Theorem 4.2.1. (Projection estimates). *Let Y and Z be two overlapping surfaces in S, with $\xi(Y) \neq 0 \neq \xi(Z)$, then for any $\mu \in \mathcal{M}(S)$:*

$$d_Y(\partial Z, \mu) > M \implies d_Z(\partial Y, \mu) \leq M,$$

for a constant M depending only on the topological type of S.

Proof:

Let M_1 and M_2 be given by Lemmas 3.3.6 and 3.3.8 (Large link Lemma and Order and projections Lemma; [MM2]). Define:

$$M = \max\{M_1 + 2, M_2 + 3\}.$$

We now fix a point $\mu \in \mathcal{M}$ for which $d_Y(\mu, \partial Z) > M$ and will show that this implies a bound on $d_Z(\mu, \partial Y)$.

Consider a hierarchy H with $\text{base}(I(H)) \supset \pi_Z(\mu) \cup \partial Z$ and $T(H) = \mu$ (when Z is an annulus take $\text{base}(I(H)) = \partial Z$ with transversal $\pi_Z(\mu)$). Furthermore, we assume H is built subject to the following two constraints. (See Section 3.3 for a discussion of the choices involved when building a hierarchy.)

- Choose the initial vertex of g_H to be an element of ∂Z which intersects Y. Call this vertex v_0.
- When adding a geodesic to H with a given initial marking $I(Y, g)$, if $v_0 \in \text{base}(I(Y, g))$ choose this as the initial vertex for the geodesic. If $v_0 \notin \text{base}(I(Y, g))$, then if any elements of ∂Z are in $\text{base}(I(h))$ choose one of these to be the initial vertex for the geodesic.

The proof is now broken into four steps. We show there exists a geodesic $k \in H$ with domain Z and a geodesic $h \in H$ with domain Y. Then we show $k \prec_t h$, from which the theorem follows as a consequence of Lemma 3.3.8 (Order and projections; [MM2]).

Step i (There exists a geodesic k with domain Z).

We start by considering the forward and backward sequences, $\Sigma^+_H(Z)$ and $\Sigma^-_H(Z)$, as defined in the discussion preceding Theorem 3.3.5 (Structure of Sigma; [MM2]).

Since μ is complete we have $\pi_Z(\mu) \neq \emptyset$; thus both the initial and terminal markings restrict to give nontrivial markings on Z (when Z is an annulus, we are using that $I(H)$ has a transverse curve). This implies for any subsurface $Q \supseteq Z$ that $\Sigma^-_H(Q)$ and $\Sigma^+_H(Q)$ both contain g_H and thus in particular are nonempty. Theorem 3.3.5 shows that when Q is a component domain of a geodesic in H and both $\Sigma^+_H(Q)$ and $\Sigma^-_H(Q)$ are non-empty, then Q must be the support of a geodesic in H.

Since the first vertex of g_H is $v_0 \in \partial Z$, we know it has a component domain $Q_1 \subsetneq S$ which contains Z. By the above observation, we know that $\Sigma^+_H(Q_1)$ and $\Sigma^-_H(Q_1)$ are both non-empty, and thus Q_1 supports a geodesic which we call k_1. If $Q_1 = Z$ we have produced a geodesic supported on Z and are done, otherwise since $Q_1 \supsetneq Z$ we can choose an element of ∂Z as the first vertex of the geodesic supported in Q_1.

Starting from the base case $Q_0 = S$ and $k_0 = g_H$, the above argument produces a sequence of properly nesting subsurfaces $S = Q_0 \supsetneq Q_1 \supsetneq \ldots \supsetneq Q_n = Z$ where for each $i > 0$ the subsurface Q_i is a component domain of a geodesic $k_{i-1} \in H$ supported in Q_{i-1} and each of the k_i has an element of ∂Z as an initial vertex.

Since S is a surface of finite type, the above sequence of nested surfaces must terminate with Z after finitely many steps.

Thus Z is a component domain of a geodesic in S and supports a geodesic which we call k.

Step ii (There exists a geodesic h with domain Y).

By hypothesis we have $d_Y(\partial Z, \mu) > M_2 + 3$ and thus:

$$d_Y(I(H), T(H)) = d_Y(\pi_Z(\mu) \cup \partial Z, \mu) > M_2.$$

Together with Lemma 3.3.6 (Large link Lemma; [MM2]), this implies Y is the support of a geodesic h in H.

Step iii ($k \prec_t h$).

In this step we will prove that the geodesic k precedes h in the time ordering on geodesics in H.

Since $D(k) = Z$ and $D(h) = Y$ and each are contained in $D(g_H) = S$: if $\max \phi_{g_H}(Z) < \min \phi_{g_H}(Y)$, then $k \prec_t h$ which is what we wanted to prove. In the general case, we we provide an inductive procedure to show that these geodesics have the desired time ordering.

We refer to the ordered vertices of g as $v_i(g)$. Recall that by the first part of the constraint we have $v_0(g_H) = v_0$, where v_0 was chosen to satisfy $v_0 \in \partial Z$ and $v_0 \cap Y \neq \emptyset$. Also, since g_H is a tight geodesic, $v_1(g_H) = \partial F(v_0(g_H), v_2(g_H))$ where $F(\alpha, \beta)$ denotes the surface filled by α and β.

Summarizing, we have:

- $Y \cap v_0 \neq \emptyset$, thus $v_0 \notin \phi_{g_H}(Y)$
- Z is contained in a component domain of v_0, thus $v_0 \in \phi_{g_H}(Z)$.
- Since g_H is a geodesic, we know $v_2(g_H)$ must intersect $v_0 (\in \partial Z)$. Since the diameter of a footprint is at most 2, it now follows that $\max \phi_{g_H}(Z) \leq v_1(g_H)$.

Together these imply:

$$\max \phi_{g_H}(Z) \leq v_1(g_H) \text{ and } \min \phi_{g_H}(Y) \geq v_1(g_H).$$

Now there are two mutually exclusive cases to consider:

1. $v_1(g_H) \cap Z \neq \emptyset$ or $v_1(g_H) \cap Y \neq \emptyset$

2. $v_1(g_H) \cap Z = \emptyset = v_1(g_H) \cap Y$.

In the first case, depending on which of the two sets is non-empty, we obtain $v_1(g_H) \notin \phi_{g_H}(Z)$ or $v_1(g_H) \notin \phi_{g_H}(Y)$, respectively. Either of these imply $\max \phi_{g_H}(Z) < \min \phi_{g_H}(Y)$, proving that in case (1) we have $k \prec_t h$.

Observing that in $S_{1,1}$ and $S_{0,4}$ the footprint of a domain must consist of at most one vertex, we have $v_1(g_H) \notin \phi_{g_H}(Z)$ and thus when S is either of these surfaces we are in case (1) and thus $k \prec_t h$. As the annulus does not contain any pair of surfaces which overlap we may now assume for the rest of the proof that $\xi(S) > 1$. In particular, for the remainder of the argument we need not consider the sporadic cases where $\mathcal{C}(S)$ has a special definition and thus we can assume two distinct homotopy classes of curves have distance one in the curve complex if and only if they can be realized disjointly.

For the remainder of this step we assume that $v_1(g_H) \in \phi_{g_H}(Z)$ and $v_1(g_H) \in \phi_{g_H}(Y)$ and will prove this implies $k \prec_t h$. Then, since $v_1(g_H) \cap Z = \emptyset$ there is a component of $S \setminus v_1(g_H)$ which contains Z and since $Z \cap Y \neq \emptyset$ the same component contains Y as well—we will call this component W_1. Since W_1 is a component domain of H which intersects both $I(H)$ and $T(H)$ (since $W_1 \supset Z$), Theorem 3.3.5 (Structure of Sigma) implies that it supports a geodesic l_1.

Since W_1 is a component domain of $v_1(g_H)$, we have $v_0 \in I(W_1, g_H)$. Thus, by our convention for choosing geodesics in H, we choose $v_0(l_1) = v_0$. As before we have $v_0(l_1) \in \phi_{l_1}(Z)$, $v_2(l_1) \notin \phi_{l_1}(Z)$, and $v_0(l_1) \notin \phi_{l_1}(Y)$. If $v_1(l_1)$ is in both $\phi_{l_1}(Z)$ and $\phi_{l_1}(Y)$ then we again restrict ourselves to the appropriate component W_2 of $W_1 \setminus v_1(l_1)$ and repeat the argument with W_2 and l_2. This gives a properly nested collection of subsurfaces each of which contains $Y \cup Z$, so the process terminates in a finite number of steps to produce a geodesic l_n with domain $W_n \supseteq Y \cup Z$ and which satisfies $v_0(l_n) = v_0 \in \phi_{l_n}(Z)$, $v_2(l_n) \notin \phi_{l_n}(Z)$, $v_0(l_n) \notin \phi_{l_n}(Y)$, and either $v_1(l_n) \notin \phi_{l_n}(Z)$ or $v_1(l_n) \notin \phi_{l_n}(Y)$. Hence $\max \phi_{l_n}(Z) < \min \phi_{l_n}(Y)$, and we have $k \prec_t h$ with comparison geodesic l_n.

Step iv (Conclusion).

We have now produced a hierarchy H with geodesics k and h with $D(h) = Y$, $D(k) = Z$, and $k \prec_t h$. Thus Lemma 3.3.8 (Order and projections; [MM2]) implies that $d_Z(\partial Y, T(H)) \leq M_1 + 2$.

Since $T(H) = \mu$, this implies $d_Z(\partial Y, \mu) \leq M_1 + 2$ which is exactly what we wanted to show.

□

The main case of the following result is an immediate corollary of Theorem 4.2.1 (Projection estimates); we give this restatement because it emphasizes the underlying geometry.

Theorem 4.2.2. (Projection estimates, geometric version). *For any distinct subsurfaces Y and Z of S, exactly one of the following holds for the map*

$$\pi_Y \times \pi_Z : \mathcal{M}(S) \to \mathcal{C}(Y) \times \mathcal{C}(Z)$$

1. *$Y \cap Z = \emptyset$ in which case the map is onto.*

2. *One of the surfaces, say Y, is contained inside the other. Here the image is contained in a radius 3 neighborhood of the set*

$$\mathcal{C}(Y) \times \pi_Z(\partial Y) \bigcup Graph(\pi_{Z \to Y}),$$

   where $\pi_{Z\to Y}$ is the projection map from $\mathcal{C}(Z) \setminus B_1(\partial Y)$ to $\mathcal{C}(Y)$.

3. *Y and Z overlap, in which case the image is contained in a radius $M = \max\{M_1 + 2, M_2\}$ neighborhood of the set*

$$\pi_Y(\partial Z) \times \mathcal{C}(Z) \bigcup \mathcal{C}(Y) \times \pi_Z(\partial Y).$$

   The constant M depends only on the topological type of S.

Proof:

Case 1 ($Y \cap Z = \emptyset$). Any pair consisting of a curve in $\mathcal{C}(Z)$ and a curve in $\mathcal{C}(Y)$ can be completed to a complete marking on S so the map is onto.

Case 2 ($Y \subset Z$).

For any $\mu \in \mathcal{M}$ we have that $d_Z(\pi'_Z(\mu), \pi_Z(\mu)) \leq 1$ and thus when we restrict these arcs (curves) to Y we get that $d_Y(\pi'_Y\pi'_Z(\mu), \pi'_Y\pi_Z(\mu)) \leq 1$. Applying Lemma 3.1.4 (Lipschitz projection; [MM2]) we then have that $d_Y(\psi\pi'_Y\pi'_Z(\mu), \psi\pi'_Y\pi_Z(\mu)) \leq 3$. Thus we have $d_Y(\mu, \pi_{Z\to Y} \circ \pi_Z(\mu)) \leq 3$.

This proves that if $Y \subset Z$ then $\pi_Y \times \pi_Z(\mathcal{M})$ is contained in a radius 3 neighborhood of $\mathcal{C}(Y) \times \pi_Z(\partial Y) \bigcup Graph(\pi_{Z\to Y})$.

Case 3 ($Y \cap Z$, $Y \not\subseteq Z$, and $Z \not\subseteq Y$).

If $\mu \in \mathcal{M}(S)$ projects to $\mathcal{C}(Y)$ with $d_Y(\partial Z, \mu) > M$ then Theorem 4.2.1 (Projection estimates) implies that $d_Z(\partial Y, \mu) \leq M$. Thus the image of the map $\mathcal{M}(S) \to \mathcal{C}(Y) \times \mathcal{C}(Z)$ is contained in the union of the radius M neighborhood of $\pi_Y(\partial Z) \times \mathcal{C}(Z)$ with the radius M neighborhood of $\mathcal{C}(Y) \times \pi_Z(\partial Y)$ as claimed.

□

Remark 4.2.3. In the above proposition we see that except in the case where $Y \cap Z = \emptyset$, there are uniform bounds for which the image of the map from $\mathcal{C}(S)$ to $\mathcal{C}(Y) \times \mathcal{C}(Z)$ lies in a neighborhood of the δ-hyperbolic space formed from the complexes of curves for Y and Z "joined together." The two complexes are glued together along a bounded diameter set in the non-nesting case. In the case of nesting $\mathcal{C}(Y)$ is glued to the link of a point in $\mathcal{C}(Z)$; essentially this is done by taking a "blow-up" of $\mathcal{C}(Z)$ at the point $\partial Y \in \mathcal{C}(Z)$. δ-hyperbolicity in the nesting case follows from Theorem 6.3.1.

Theorem 4.2.2 (Projection estimates, geometric version) is summarized in Figure 4.3. For a (slightly) more accurate picture of the behavior in the nesting case, see also Figure 3.2.

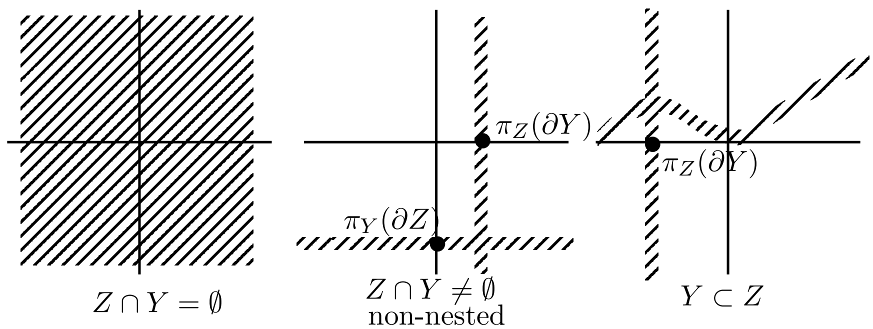

Figure 4.3: A cartoon of the possible images of projections from complete markings into $\mathcal{C}(Z) \times \mathcal{C}(Y)$

4.3 Coarse Tree-flats and the Marking Complex

In this section we generalize Theorem 4.2.2 to a form suitable for approaching questions on the structure of quasiflats in the mapping class group.

For a set Ω of subsurfaces of S we consider

$$\pi = \pi_\Omega : \mathcal{M}(S) \to \prod_{X \in \Omega} \mathcal{C}(X)$$

where for $\mu \in \mathcal{M}(S)$ we define $\pi(\mu) = \prod_X \pi_X(\mu)$.

Some notation we use to help organize the argument is:

If $U = \{Y_1, Y_2, \dots, Y_l\}$ is a set of homotopically distinct subsurfaces with disjoint representatives (henceforth simply referred to as disjoint subsurfaces) then we use the notation ∂U to denote the union of their boundary curves, $\cup_{1 \le i \le l} \partial Y_i$.

Definition 4.3.1. Let U denote a *maximal* subcollection of pairwise disjoint surfaces from a set of subsurfaces Ω; where maximal is defined to mean that no other surface from Ω is disjoint from all the surfaces in U. Define $\mathcal{U}_\Omega$ to be the set of all such maximal collections U which are built out of the surfaces in Ω. When no ambiguity arises, we will sometimes use the symbol U to refer to the union of all surfaces in U. More generally, if U is a collection of disjoint subsurfaces, we will say U is *maximal with respect to* Ω to mean that no element of Ω is disjoint from U; in this case U need not be made out of surfaces of Ω.

Note that each $U \in \mathcal{U}_\Omega$ is a set with $|U| \leq \min\{3g+p-3, |\Omega|\}$ (but $|U|$ need not be the same for each $U \in \mathcal{U}_\Omega$).

The useful property of a collection U being maximal with respect to a collection of surfaces Ω, is that this implies that each $X \in \Omega$ either satisfies $X \cap \partial U \neq \emptyset$ or there exist $Y \in U$ for which $X \subset Y$. This observation makes the following well defined.

Definition 4.3.2. Given a collection of subsurfaces U which is maximal with respect to Ω, we define the *C-Coarse tree-flat of* U, denoted $\mathcal{F}^C_\Omega(U)$, to be the set of points $(\sigma_x)_{x\in\Omega} \in \prod_{X\in\Omega} \mathcal{C}(X)$ satisfying the following two conditions for each $X \in \Omega$:

1. If $X \cap \partial U \neq \emptyset$, then $\mathrm{diam_X}(\sigma_\mathrm{X} \cup \pi_\mathrm{X}(\partial \mathrm{U})) \leq \mathrm{C}$.

2. For each $Y \in U$ if $X \subsetneq Y$, then $\mathrm{diam_X}(\sigma_\mathrm{X} \cup \pi_\mathrm{X}(\sigma_\mathrm{Y})) \leq \mathrm{C}$.

We refer to $|U|$ as the *rank* of this C-Coarse tree-flat.

We usually consider $\mathcal{F}_{\Omega^C(U)}$ when $U \in \mathcal{U}_\Omega$, although in the proof of Theorem 4.3.6 we will also consider the case that U is simply maximal with respect to Ω.

We will often consider the coarse tree-flat over U with $U \in \mathcal{U}_\Omega$.

Although this definition at first glance may appear somewhat opaque, it is fairly natural and arose out of a desire to generalize the phenomenon exhibited in Theorem 4.2.2.

Let us start by recasting the examples in Section 4.1 into this language. After reading these examples, we encourage the reader to re-examine Figure 4.3.

Example 4.1.1 Revisited. In this example, we show that when considering the product of the the complex of curves of two disjoint surfaces, any marking lies inside the coarse tree-flat over the pair.

In Example 4.1.1, the surfaces Q and W are disjoint; taking $\Omega = \{Q, W\}$ we have $U = \{Q, W\} \in \mathcal{U}_\Omega$, in this case U is the only element of $\mathcal{U}_\Omega$. Thus for each surface $Z \in \Omega$ we have $Z \cap \partial U = \emptyset$, hence both conditions in the definition of coarse tree flat are vacuous in this example. Thus for any constant C and each $\mu \in \mathcal{M}(S)$ we have $\pi_\Omega(\mu) \in \mathcal{F}^C_\Omega(U)$.

Example 4.1.2 Revisited. Here we consider again the setup of Example 4.1.2 where we considered two subsurfaces X and Y which overlap. We will show that in this case the image of the map $\pi : \mathcal{C}(X) \times \mathcal{C}(Y)$ lies in the union of the coarse tree-flat over X and the coarse tree-flat over Y.

Taking $\Omega = \{X, Y\}$, we have $\{X\} \in \mathcal{U}_\Omega$ and $\{Y\} \in \mathcal{U}_\Omega$. In this setting the second condition of coarse tree-flat is vacuous, but the first condition that $Z \cap \partial U \neq \emptyset$ for some $Z \in \Omega$ is satisfied for $Z = X$ when $U = \{Y\}$ and for $Z = Y$ when $U = \{X\}$.

Thus we would like to show that for any $\mu \in \mathcal{M}(S)$ either $\mathrm{diam_X}(\pi_\mathrm{X}(\mu) \cup \pi_\mathrm{X}(\partial \mathrm{U})) \leq \mathrm{C}$ or $\mathrm{diam_Y}(\pi_\mathrm{Y}(\mu) \cup \pi_\mathrm{Y}(\partial \mathrm{U})) \leq \mathrm{C}$. As long as C is larger than the constant M, as constructed in Theorem 4.2.2 (Projection Estimates) this follows from the overlapping case of that theorem. As a consequence this implies $\pi_\Omega(\mathcal{M}) \subset \mathcal{F}^C_\Omega(\{X\}) \cup \mathcal{F}^C_\Omega(\{Y\})$.

Example 4.3.3. **(Nested surfaces).**

Let S be the closed genus three surface and consider $\Omega = \{X, Y\}$, where X is a torus with one puncture, Y has genus two with one puncture, and $X \subset Y$ (see Figure 4.3.3).

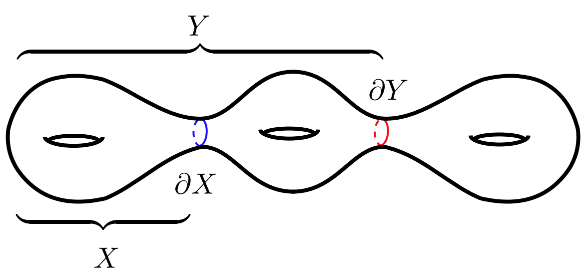

Figure 4.4: S and two nested subsurfaces

Let $U = \{Y\} \in \mathcal{U}_\Omega$. Since the first condition of coarse-tree flat is vacuously satisfied, it suffices to check the second condition. Thus showing $\pi(\mu) \in \mathcal{F}^M(U)$ reduces to proving $\mathrm{diam_X}(\pi_\mathrm{X}(\mu) \cup \pi_\mathrm{X}(\mu_\mathrm{Y})) \leq \mathrm{M}$. This inequality with $M = 3$ is the content of the nested case of Theorem 4.2.2 (compare Lemma 3.1.4).

The definition of coarse tree-flat was motivated by the following example.

Example 4.3.4. Let $U = \{A_1, \ldots, A_{3g+p-3}\}$ be a collection of disjoint essential annuli in S. We will now show that the notion of a course tree-flat provides a useful way to describe the image of the base marking μ_0 under the subgroup of the mapping class group generated by Dehn twists along the curves in U. Since U is a maximal collection of annuli, any subsurface X (which isn't a pair of pants) has $X \cap \partial U \neq \emptyset$. Furthermore, as a Dehn twist of base(μ_0) along the core curves $\{\gamma_1, \ldots, \gamma_{3g+p-3}\}$ of the annuli in U does not change the intersection numbers of base(μ_0) with γ_i for each i, we can easily obtain a bound on $\mathrm{diam_X(\partial U, f(\mu_0))}$ which is uniform over all such Dehn twists, f. Thus we have that for any finite set Θ of subsurfaces of S, there exist a constant C, for which all Dehn twists along curves in U lie in $\mathcal{F}^C_\Theta(U)$. C depends only on $|\Theta|$ and the choice of base marking.

We record the following elementary, but useful fact for later use.

Lemma 4.3.5. *Fix two collections, Ω and Ω', of subsurfaces of S. If $W \in \mathcal{U}_{\Omega\cup\Omega'}$ and we have both $\pi_\Omega(\mu) \in \mathcal{F}^c_\Omega(W)$ and $\pi_{\Omega'}(\mu) \in \mathcal{F}^{c'}_{\Omega'}(W)$ then $\pi_{\Omega\cup\Omega'}(\mu) \in \mathcal{F}^{\max\{c,c'\}}_{\Omega\cup\Omega'}(W)$.*

Proof:

To check $\pi_{\Omega\cup\Omega'}(\mu) \in \mathcal{F}^{\max\{c,c'\}}_{\Omega\cup\Omega'}(W)$, we need to show:

- if $X \cap \partial W \neq \emptyset$ then $\mathrm{diam_X(\pi_X(\mu) \cup \pi_X(\partial W)) \leq \max\{c, c'\}}$
- if $X \subset Y \in W$ then $\mathrm{diam_X(\pi_X(\mu) \cup \pi_X(\pi_Y(\mu))) \leq \max\{c, c'\}}$

Given that W is maximal, for each $X \in \Omega \cup \Omega'$ either $X \cap \partial W \neq \emptyset$ or $X \subset Y \in W$.

Since $\pi_\Omega(\mu) \in \mathcal{F}^c_\Omega(W)$ for each $X \in \Omega$ we either have

$$\mathrm{diam_X(\pi_X(\mu) \cup \pi_X(\partial W)) \leq c \leq \max\{c, c'\}}$$

or

$$\mathrm{diam_X(\pi_X(\mu) \cup \pi_X(\pi_Y(\mu))) \leq c \leq \max\{c, c'\}}$$

depending on whether $X \cap \partial W \neq \emptyset$ or $X \subset Y \in W$, respectively.

Similarly, for each $X \in \Omega'$.

Thus proving that $\pi_{\Omega\cup\Omega'}(\mu) \in \mathcal{F}^{\max\{c,c'\}}_{\Omega\cup\Omega'}(W)$.

□

We are now in a position to state and prove our main tool:

Theorem 4.3.6. (Coarse tree-flats carry markings). *For any integer $N \geq 1$ and collection of surfaces Θ with $|\Theta| = N$ the map $\pi = \prod_{X\in\Theta} \pi_X : \mathcal{M}(S) \to \prod_{X\in\Theta} \mathcal{C}(X)$ satisfies:*

$$Image(\pi) \subset \bigcup_{U \in \mathcal{U}_\Theta} \mathcal{F}^{C_N}_\Theta(U)$$

Where $M = max\{M_1 + 2, M_2 + 3\}$, $C_1 = M$, for $N > 1$ we set $C_N = (N-1)(M+9)$.

Proof:

The theorem reduces to showing that for each $\mu \in \mathcal{M}(S)$ there exists (at least) one $U \in \mathcal{U}_\Theta$ for which $\pi(\mu) \in \mathcal{F}_\Theta^{C_N}(U)$. We will prove this fact by induction on the number of surfaces being considered.

The base case of $N = 1$ is trivially satisfied since the conditions of coarse tree-flat are vacuous here. (The attentive reader will note that Examples 4.3, example:overlap:coarsetreeflat, and example:nested:coarsetreeflat combine to provide a complete proof of the $N = 2$ case.)

Induction step

Now that we have shown the base case we assume the following inductive hypothesis:

> For any Ω with $|\Omega| \leq N-1$ and any $\mu \in \mathcal{M}$ there exists (at least) one $U \in \mathcal{U}_\Omega$ for which $\pi_\Omega(\mu) \in \mathcal{F}_\Omega^{C_{|\Omega|}}(U)$.

For the remainder of the argument we shall fix an element $X \in \Theta$. We also fix a marking $\mu \in \mathcal{M}$ and denote by U an element of $\mathcal{U}_{\Theta \setminus X}$, as provided by the inductive hypothesis, for which $\pi_{\Theta\setminus X}(\mu) \in \mathcal{F}_{\Theta\setminus X}^{C_{N-1}}(U)$. (Note: we write $\Theta \setminus X$ to denote $\Theta \setminus \{X\}$ as we feel this should not be a source of confusion and makes the formulas easier to parse.)

We first deal with the two easiest cases, where $U \cap X = \emptyset$ or $X \subset S \in U$, before dealing with the main case where $X \cap \partial U \neq \emptyset$.

Case a: $(U \cap X = \emptyset)$.

Since the definition of the C-Coarse tree-flat of U doesn't induce any restrictions on the values of $\pi_Y(\mu)$ for $Y \in U$, we see that if $X \cap U = \emptyset$ then $\pi_X(\mu) \in \mathcal{F}_{\{X\}}^{C_{N-1}}(U \cup \{X\})$. The inductive assumption implies $\pi_{\Theta\setminus X}(\mu) \in \mathcal{F}_{\Theta\setminus X}^{C_{N-1}}(U)$, then Lemma 3.1.4 implies that $\pi_{\Theta\setminus X}(\mu) \in \mathcal{F}_{\Theta\setminus X}^{C_{N-1}+3}(U \cup \{X\})$. The two above computations combine via Lemma 4.3.5 to show that $\mu \in \mathcal{F}_\Theta^{C_{N-1}+3}(U \cup \{X\})$, and thus $\mu \in \mathcal{F}_\Theta^{C_N}(U \cup \{X\})$.

Case b: $(X \subset Y \in U)$.

If there exists $Y \in U$ for which $X \subset Y$ then the case of Theorem 4.2.1 (Projection estimates) for nested surfaces tells us that $\pi_X(\mu) \in \mathcal{F}_{\{X\}}^M(U)$ and thus invoking Lemma 4.3.5 we have in this case $\mu \in \mathcal{F}_\Theta^{C_N}(U)$.

Case c: $(X \cap \partial U \neq \emptyset)$**.**

After some preliminary definitions, we provide a sketch of the remaining argument. Define

$$\Delta = \{Z \in \Theta : \partial Z \cap X \neq \emptyset\}.$$

Another set which arises in the proof is:

$$\mathcal{V} = \{V \in \mathcal{U}_\Theta : Y \supseteq X \text{ for some } Y \in V\}.$$

We consider the set of surfaces which arise as elements in some collection of $\mathcal{V}$,

$$\Lambda = \{Y : Y \in V \text{ for some } V \in \mathcal{V}\}.$$

Notice that since Δ is the set of all elements of Θ which have boundary that intersect X and Λ is the set of all surfaces whose boundary is disjoint from X (thus $X \in \Lambda$) that we have $\Theta \setminus \Delta = \Lambda$.

The rest of the argument is organized with the following philosophy: we start by attempting to prove that $\pi_\Theta(\mu) \in \mathcal{F}_\Theta^{C_N}(U)$. If this is not the case, our argument will show that $d_Z(\mu, \partial V) \leq C_N$ for each $Z \in \Delta$ and $V \in \mathcal{V}$ with $\pi_Z(\partial V) \neq \emptyset$. Then invoking the inductive hypothesis, we know there exists a collection $V \in \mathcal{V}$ for which $\pi_\Lambda(\mu) \in \mathcal{F}_\Lambda^{C_{|\Lambda|}}(V)$. Since $\Theta \setminus \Lambda = \Delta$, the above two computations combine using Lemma 4.3.5 to show $\pi_\Theta(\mu) \in \mathcal{F}_\Theta^{C_N}(V)$, which will prove the theorem. The remainder of the proof is working out the details of this sketch including some necessary modifications to deal with the case of nested surfaces which were ignored in this overview.

We start with the following easy but useful lemma whose proof we provide at the end of this section.

Lemma 4.3.7. *If Q is a diameter 1 subset of $\mathcal{C}(S)$ with $\pi_X(Q) \neq \emptyset$ and ν is a marking, then either $d_X(Q, \nu) > M$ or $\mathrm{diam_X(Q, \nu) \leq M + 6}$.*

For our collection U, if $\mathrm{diam_X(\partial U, \mu) \leq M + 6}$ then coupling this with the inductive hypothesis we have that $\pi_\Theta(\mu) \in \mathcal{F}_\Theta^{C_{N-1}}(U)$. By Lemma 4.3.7 for the remainder of the argument it suffices to assume $d_X(\partial U, \mu) > M$.

The goal is to now show that even though $d_X(\partial U, \mu) > M$ this distance is either bounded above by C_{N-1} or there is a collection $V \in \mathcal{V}$, for which $\pi_\Theta(\mu) \in \mathcal{F}_\Theta^{C_N}(V)$.

We now analyze the three possibilities for surfaces $Z \in \Delta$.

- **Case i:** $\Delta_1 = \{Z \in \Delta : Z \subset X\}$**.**

  For any surface $Y \in V \in \mathcal{V}$ which contains X as a subsurface, we also have $Z \subset Y$. The nested surfaces case of Theorem 4.2.2 then implies $\pi_Z(\mu)$ is in the ball of radius M around $\pi_Z(\pi_Y(\mu))$.

  This shows that for each $V \in \mathcal{V}$, we have $\pi_{\Delta_1}(\mu) \in \mathcal{F}_{\Delta_1}^M(V)$.

- **Case ii:** $\Delta_2 = \{Z : \partial X \cap \partial Z \neq \emptyset$ and $d_X(\partial Z, \mu) > M\}$.

  By Theorem 4.2.1 (Projection estimates) if $Z \in \Delta_2$, then $\mathrm{diam_Z}(\partial \mathrm{X}, \mu) \leq \mathrm{M}$. Since the collections $V \in \mathcal{V}$ are those for which $\partial V \cap \partial X = \emptyset$, applying Lemma 3.1.4 tells us $\mathrm{diam_Z}(\partial \mathrm{V}, \mu) \leq \mathrm{M} + 3$ for every $V \in \mathcal{V}$ with $\pi_Z(\partial V) \neq \emptyset$. Thus we have for each $Z \in \Delta_2$ and each $V \in \mathcal{V}$ that $\pi_{\Delta_2}(\mu) \in \mathcal{F}^{M+9}_{\Delta_2}(V)$.

- **Case iii:** $\Delta_3 = \{Z : \partial X \cap \partial Z \neq \emptyset$ and $\mathrm{diam_X}(\partial \mathrm{Z}, \mu) \leq \mathrm{M} + 6\}$.

  In this case we can combine the inequality $\mathrm{diam_X}(\partial \mathrm{Z}, \mu) \leq \mathrm{M} + 6$, with the fact that $d_X(\partial U, \mu) > M$ and some computations involving $d_X(\partial U, Z)$ to get our desired conclusion.

  By Lemma 4.3.7 either $\mathrm{diam_X}(\partial \mathrm{U}, \partial \mathrm{Z}) \leq \mathrm{M} + 6$ or $d_X(\partial U, \partial Z) > M$, but this is useful information either way. If $\mathrm{diam_X}(\partial \mathrm{U}, \partial \mathrm{Z}) \leq \mathrm{M} + 6$ then together with $\mathrm{diam_X}(\mu, \partial \mathrm{Z}) \leq \mathrm{M} + 6$ we have that $\mathrm{diam_X}(\partial \mathrm{U}, \mu) \leq \mathrm{M} + 6 + \mathrm{M} + 6 = 2(\mathrm{M} + 6)$. Coupled with our induction hypothesis that $\pi_{\Theta \setminus X}(\mu) \in \mathcal{F}^{C_{N-1}}_{\Theta \setminus X}(U)$ this tells us that $\pi_\Theta(\mu) \in \mathcal{F}^{C_{N-1}}_{\Theta}(U)$. If this is the case, then we have finished proving the theorem. Thus, for the remainder of the proof we may assume $d_X(\partial U, \partial Z) > M$.

  Since $d_X(\partial U, \partial Z) > M$ and $\partial X \cap \partial Z \neq \emptyset$, Theorem 4.2.1 (Projection estimates) implies $d_Z(\partial U, \partial X) \leq M$, and thus $\mathrm{diam_Z}(\partial \mathrm{U}, \partial \mathrm{X}) \leq \mathrm{M} + 6$. When $Z \cap \partial U \neq \emptyset$ the induction hypothesis tells us $\mathrm{diam_Z}(\mu, \partial \mathrm{U}) \leq \mathrm{C_{N-1}}$. Together this gives us $\mathrm{diam_Z}(\mu, \partial \mathrm{X}) \leq \mathrm{C_{N-1}} + \mathrm{M} + 6$. So for any $V \in \mathcal{V}$ with $\pi_Z(\partial V) \neq \emptyset$ we have $\mathrm{diam_Z}(\mu, \partial \mathrm{V}) \leq (\mathrm{N} - 1)(\mathrm{M} + 6) + \mathrm{M} + 6 + 3 < \mathrm{N}(\mathrm{M}+9) = \mathrm{C_N}$ (the additional 3 here comes from applying Lemma 3.1.4, since ∂V need not be just a vertex, but rather may have diameter 1).

  We have thus shown, either there exists $Z \in \Delta_3$ with $\mathrm{diam_X}(\partial \mathrm{U}, \partial \mathrm{Z}) \leq \mathrm{M} + 9$ in which case $\pi_\Theta(\mu) \in \mathcal{F}^{C_{N-1}}_{\Theta}(U)$, or for every $Z \in \Delta_3$ we have $d_X(\partial U, \partial Z) > M$. In the latter case we showed that this implies $\pi_Z(\mu) \in \mathcal{F}^{C_N}_{\Delta_3}(V)$ for every $V \in \mathcal{V}$.

Combining the above three cases, since $\Delta = \Delta_1 \cup \Delta_2 \cup \Delta_3$, we have shown that either there exists a surface $Z \in \Delta_3$ with $\mathrm{diam_X}(\partial \mathrm{U}, \partial \mathrm{Z}) \leq \mathrm{M}+9$ (in which case $\mu \in \mathcal{F}^{C_N}_{\Theta}(U)$), or by Lemma 4.3.5 we have $\pi_\Delta(\mu) \in \mathcal{F}^{C_N}_{\Delta}(V)$ for every $V \in \mathcal{V}$. To finish, it remains to show in the latter case that $\pi_{\Theta \setminus \Delta}(\mu) \in \mathcal{F}^{C_N}_{\Theta \setminus \Delta}(V)$ for some $V \in \mathcal{V}$.

This leaves us to consider the set of surfaces with boundary disjoint from X, namely the set $\Theta \setminus \Delta = \Lambda$. Since $|\Lambda| < N$ (else X would be disjoint from all the surfaces in $\Theta \setminus X$, which was handled in case (a)) we can invoke the inductive hypothesis to conclude that there exists (at least one) $V' \in \mathcal{V}$ for which $\pi_\Lambda(\mu) \in \mathcal{F}_\Lambda^{C_{|\Lambda|}}(V')$. Since we above proved that for all $V \in \mathcal{V}$ we have $\pi_\Delta(\mu) \in \mathcal{F}_\Delta^{C_N}(V)$, Lemma 4.3.5 then implies that $\pi(\mu) \in \mathcal{F}_\Theta^{C_N}(V')$, finishing the argument.

$\square$

Proof of Lemma 4.3.7: Since Q and base ν are each diameter one subsets of $\mathcal{C}(S)$ Lemma 3.1.4 (Lipschitz projections; [MM2]) implies that $\mathrm{diam_X(Q)} \leq 3$ and $\mathrm{diam_X}(\nu) \leq 3$ (when X is an annulus around a base curve of ν we are using the fact that for each base curve, the transversals consist of a set of diameter at most one in the annular complex around that curve). Thus when $d_X(Q, \nu) \leq M$, we have $\mathrm{diam_X(\nu, Q) \leq M + 3 + 3}$ as desired.

$\square$

Chapter 5

Hyperbolicity of the Mapping Class Group and Teichmüller space in the Cases of Low Complexity

In this Chapter we give new proofs of the following theorems.

Theorem 5.0.8. $\mathcal{M}(S_{1,1})$ *and* $\mathcal{M}(S_{0,4})$ *are both* δ*-hyperbolic.*

Theorem 5.0.9. $\mathcal{P}(S_{1,2})$ *and* $\mathcal{P}(S_{0,5})$ *are both* δ*-hyperbolic.*

Theorem 5.0.8 is classical as these two mapping class groups are each isomorphic to $SL_2(\mathbb{Z})$, a group which is virtually free and thus easily seen to be δ-hyperbolic. Theorem 5.0.9 is a much deeper theorem, and was originally proved by J. Brock and B. Farb in [BrF]. Although the Weil-Petersson metric has been known for some time to have negative curvature, it was only recently shown by Z. Huang in [Hu] that the sectional curvature is not bounded away from zero, even for the case of $\mathcal{P}(S_{1,2})$ and $\mathcal{P}(S_{0,5})$, thereby prohibiting a proof of Theorem 5.0.9 by a comparison geometry argument. We find especially interesting the phenomenon that these proofs about two very different spaces are virtually identical.

We prove these two results simultaneously. This can be done since the property of low complexity which we will use is that any two proper subsurfaces (with nontrivial curve complex) must overlap: a property of $S_{1,1}$ and $S_{0,4}$ which is also true for $S_{1,2}$ and $S_{0,5}$ when one adds the assumption that the subsurfaces are not annuli, a natural assumption when one is considering pants decompositions instead of markings (compare Theorem 3.4.5). Thus the proof of Theorem 5.0.9 is obtained by rereading the proof for Theorem 5.0.8 reading "marking" as "pants decomposition," "hierarchy" as "hierarchy without annuli," and $\mathcal{M}(S)$ as $\mathcal{P}(S)$.

Our method of proof will be to show that for every hierarchy path there exist an "almost locally constant" map sending the marking (or pants) complex to the hierarchy path. That this implies hyperbolicity is a consequence of a general result from [MM1]; we also provide a new proof of this implication using asymptotic cones.

5.1 A Projection from Markings to Hierarchies

Fix a surface S. Also fix two complete clean markings I and T and a hierarchy H connecting them. By a hierarchy path, we mean a "resolution" of slices of the hierarchy into a sequence of markings separated by elementary moves; this is done by taking the slices of H and interpolating between them to get a path. The choice of a hierarchy path is not canonical, but they are each (K, C)-quasi-geodesics with the quasi-isometry constants depending only on the topological type of the surface. (See [MM2], especially Section 5 and the Efficiency of hierarchies Theorem of Section 6.)

We say that Y *appears as a large domain in* H if Y is in the set

$$\mathcal{G} = \{Y \subseteq S : d_Y(I(H), T(H)) > 6M + 4\delta'\}.$$

Where $M = \max\{M_1 + 2, M_2 + 3\}$ is the constant coming from Theorem 4.2.1 (Projection estimates), note in particular that this constant is sufficiently large to satisfy the hypothesis of Lemma 3.3.6 (Large link; [MM2]) and $\delta' = 4\delta + 5$ where δ denotes the maximum of the δ-hyperbolicity constants for $\mathcal{C}(Y)$ where $Y \subseteq S$. Alternatively, when dealing with the pants complex we consider:

$$\mathcal{G}' = \{Y \in \mathcal{G} : Y \text{ is not an annulus}\}.$$

Letting $g_{H,Y}$ denote the geodesic segment supported on Y in the hierarchy H, for each $Y \in \mathcal{G}$ we have a map $\mathcal{M}(S) \xrightarrow{\pi_Y} \mathcal{C}(Y) \xrightarrow{r_Y} g_{H,Y}$ where r_Y is the closest point(s) projection from $\mathcal{C}(Y)$ to $g_{H,Y}$. We denote the composition

$$p_Y = r_Y \circ \pi_Y : \mathcal{M}(S) \to g_{H,Y},$$

or just p when the surface Y is understood. Notice that δ-hyperbolicity of $\mathcal{C}(Y)$ (via Lemma 2.1.7) combined with the fact that $\mathrm{diam}_{\mathrm{Y}}(\pi_{\mathrm{Y}}(\mu)) \leq 3$ (Lemma 3.1.4) imply that for any $\mu \in \mathcal{M}(S)$, the set $p_Y(\mu)$ has diameter at most δ'. Accordingly, it follows that maps p_Y are coarsely distance decreasing in the sense that:

$$d_{\mathcal{C}(Y)}(p_Y(\mu), p_Y(\nu)) \leq d_{\mathcal{C}(Y)}(\pi_Y(\mu), \pi_Y(\nu)) + 2\delta'. \tag{5.1}$$

Define

$$\mathcal{L}(\mu) = \{Y \in \mathcal{G} : d_Y(p_Y(\mu), T(Y)) < 3M + 2\delta'\}$$

and

$$\mathcal{R}(\mu) = \{Y \in \mathcal{G} : d_Y(p_Y(\mu), I(Y)) < 3M + 2\delta'\}.$$

Usually the choice of marking μ will be fixed and we will drop the μ from the notation, just writing $\mathcal{L}$ and $\mathcal{R}$.

The next lemma summarizes the facts we will need about these sets.

Lemma 5.1.1. *1.* $\mathcal{L} \cap \mathcal{R} = \emptyset$

2. Let $Y, Z \in \mathcal{G}$*. If* $Y \prec_t Z$ *and* Y *and* Z *overlap, then either* $Y \in \mathcal{L}$ *or* $Z \in \mathcal{R}$*.*

3. If S *is* $S_{1,1}$ *or* $S_{0,4}$ *then the set* $\mathcal{G} \setminus (\mathcal{L} \cup \mathcal{R})$ *consists of at most one surface.*

4. If S *is* $S_{1,2}$ *or* $S_{0,5}$ *then* $\mathcal{G}' \setminus (\mathcal{L} \cup \mathcal{R})$ *consists of at most one surface.*

Proof:

1. Immediate. If $Y \in \mathcal{L}$ then $d_Y(p(\mu), T) < 3M$. If Y is also in $\mathcal{R}$ then $d_Y(p(\mu), I) < 3M$. But together these imply that the total distance H travels through Y is less than $6M$ contradicting our assumption that Y is a large domain.

2. Suppose not. Then we have $Y \prec_t Z$, $Y \notin \mathcal{L}$, and $Z \notin \mathcal{R}$.

   Since Y and Z are large domains with $Y \prec_t Z$, by Lemma 3.3.8 (Order and projection Lemma; [MM2]) we have that

   $$d_Y(\partial Z, T(H)) \leq M_1 + 2 \text{ and } d_Z(\partial Y, I(H)) \leq M_1 + 2.$$

   $Y \notin \mathcal{L}$ implies that

   $$d_Y(p(\mu), T(Y)) \geq 3M + 2\delta'.$$

   Similarly, $Z \notin \mathcal{R}$ implies that

   $$d_Z(p(\mu), I(Z)) \geq 3M + 2\delta'.$$

   Putting these facts together with the result from [MM2] that for any domain $d_D(I(D), I(H)) \leq M_1$, yields the two inequalities:

   $$d_Y(p(\mu), \partial Z) > 3M + 2\delta' - M - (M_1 + 2) > M + 2\delta'$$

and

$$d_Z(p(\mu), \partial Y) > 3M + 2\delta' - M - (M_1 + 2) > M + 2\delta'.$$

The fact that p_Y and p_Z are coarsely distance decreasing implies, via Equation 5.1, that both $d_Y(\mu, \partial Z) > d_Y(p(\mu), \partial Z) - 2\delta' > M$ and $d_Z(\mu, \partial Y) > d_Z(p(\mu), \partial Y) - 2\delta' > M$. But this is impossible, since together these inequalities contradict Theorem 4.2.1 (Projection estimates).

3. This follows from part 2, since in each of these surfaces any two subsurfaces overlap and thus are time ordered by Lemma 4.18 of [MM2].

4. Again this follows from part 2, since here any two non-annular surfaces must overlap.

□

Parts 3 and 4 of the above proposition suggests the following construction, which we give for the marking complexes for $S_{1,1}$ and $S_{0,4}$ and also for the pants complexes of $S_{1,2}$ and $S_{0,5}$ (parts of the construction generalize to general surfaces, but we won't use that generality here). Given any pair of markings $I, T \in \mathcal{M}(S)$ and a hierarchy H connecting them let $\widetilde{H}$ denote the set of markings associated to slices of H. Below we construct a map Φ which maps elements of $\mathcal{M}(S)$ to (uniformly) bounded diameter subsets of $\widetilde{H}$ (in the cases we are considering, $\widetilde{H}$ is metrized by a linear time ordering on slices). $\widetilde{H}$ be embedded into $\mathcal{M}(S)$ via a quasi-isometric embedding which can be extended by resolving the set of slices into a path (referred to as a *hierarchy path*), the map Φ induces a map to (uniformly) bounded diameter subsets of any hierarchy path from I to T. The map Φ is defined to be the identity on markings in $\widetilde{H}$; the following defines the map for $\mu \in \mathcal{M}(S) \setminus \widetilde{H}$:

1. If $\mathcal{G} \setminus (\mathcal{L} \cup \mathcal{R}) = \{A\} \neq \emptyset$, then define $\Phi(\mu) = (A, p_A(\mu))$.

2. If $\mathcal{G} \setminus (\mathcal{L} \cup \mathcal{R}) = \emptyset$, then consider

$$\Lambda = \{v \in g_H : d_{\mathcal{C}(S)}(v, \mu) \leq d_{\mathcal{C}(S)}(w, \mu) \text{ for any } w \in g_H\},$$

   i.e, the set of points on the main geodesic g_H which are closest to μ, namely $r_S(\mu)$. $\Phi(\mu)$ is defined to be the following sets of markings.

   (a) If any of the surfaces in $\mathcal{L}$ are component domains of g_H at a vertex in the set Λ, then denote by L the rightmost (the last to appear with respect to the time ordering) domain in $\mathcal{L}$, and define $\Phi(\mu)$ to contain each (L, v) where v is any vertex of $g_{H,L}$ within $3M + 3\delta'$ of the terminal marking T.

(b) If a component domain corresponding to a vertex of g_H in the set Λ is an element of $\mathcal{R}$ then choose the leftmost element of this set (call this R). Define $\Phi(\mu)$ to contain each of the markings (R, v) where v is any vertex of $g_{H,R}$ within $3M + 3\delta'$ of the initial marking I.

(c) For each domain D which is a component domain of a vertex in Λ and which is not in $\mathcal{G}$, if D supports a geodesic and it is time ordered after the L geodesic and before the R geodesic then define $\Phi(\mu)$ to contain each (D, v) where v is any vertex in $g_{H,D}$.

Before proceeding, a justification is needed as to why $(A, p_A(\mu))$, (L, v), and (R, v) give complete clean markings (in the text we abuse notation and often call them markings). In each case these consist of proper subsurface X and a point in $\mathcal{C}(X)$ (we unify the argument that these give markings by writing $(X, p_X(\mu))$). For $S_{1,1}$ and $S_{0,4}$, a point of $\mathcal{M}(S)$ consist of a curve in $\mathcal{C}(S)$ and a transversal to that curve, so it is clear that our prescription above indeed describes a marking since the only nontrivial proper subsurfaces are annuli so $(X, p_X(\mu))$ refers to the marking $(\partial X, p_X(\mu))$. The case of $\mathcal{P}(S)$ for the surfaces $S_{1,2}$ and $S_{0,5}$ requires (only slightly) more justification. In these cases an element of $\mathcal{P}(S)$ is a pair of disjoint curves. Recall that the only subsurfaces of S with nontrivial curve complex are once punctured tori and four punctured spheres (we ignore annuli when dealing with the pants complex). Thus $p_X(\mu) \in \mathcal{C}(X)$ is also an element of $\mathcal{C}(S)$. Furthermore, $p_X(\mu)$ and ∂X have distance 1 in $\mathcal{C}(S)$, so $(X, p_X(\mu))$ is taken to refer to the pants decomposition $(\partial X, p_X(\mu))$.

Lemma 5.1.2. *There is a uniform bound D (depending only on the topological type of S), so that for each $\mu \in \mathcal{M}$ the set $\Phi(\mu)$ has diameter less than D.*

Note that since the map from slices to paths is Lipschitz, this theorem yields a map which sends points in $\mathcal{M}$ to uniformly bounded diameter subsets of the hierarchy path, for any hierarchy path.
Proof:

First note that the set of points on the main geodesic g_H which are closest to μ has diameter at most δ' by Lemma 2.1.7. When $\mathcal{G} \setminus (\mathcal{L} \cup \mathcal{R}) \neq \emptyset$ there is only one such surface, as proven in Lemma 5.1.1. In this case it follows from the definition of Φ and δ-hyperbolicity of $\mathcal{C}(A)$ that $\Phi(\mu)$ is a subset of $\widetilde{H}$ of diameter at most δ'.

Note that under Φ the marking μ can not project to anything time ordered before the rightmost element $L \in \mathcal{L}$ as then Lemma 3.3.8 (Order and projections; [MM2]) would force $p_L(\mu)$ to lie near the initial marking of L and then we would have either $L \notin \mathcal{G}$ or $L \in \mathcal{R}$, either way contradicting Lemma 5.1.1

which proves that every surface in $\mathcal{L}$ is time ordered before every surface of $\mathcal{R}$; a similar argument gives the analogous result for $\mathcal{R}$.

So now $\Psi(\mu)$ consists of the rightmost element of $\mathcal{L}$ which we call L, the leftmost element of $\mathcal{R}$ which we call R, and all the rest of the small domains supporting geodesics which are time ordered between L and R (of which there are at most δ'). Φ was defined to be the union of the three set (any of which are possibly empty):

1. If $L \in \Lambda$ then (L, v) where v is any vertex of $g_{H,L}$ within $3M + 3\delta'$ of the terminal marking T
2. If $R \in \Lambda$ then (R, v) where v is any vertex of $g_{H,R}$ within $3M + 3\delta'$ of the initial marking I
3. If $D \in \Lambda$, $D \notin \mathcal{G}$, and $L \prec_t D \prec_t R$ then (D, v) where V is any vertex of $g_{H,D}$.

First, note that the diameter of the elements in item 1 is at most $B(3M + 3\delta')$, where B is the Lipschitz constant for the map from hierarchy slices to paths in the marking complex. (Similarly for the markings in item 2.) Since $\mathcal{C}(S)$ is δ-hyperbolic $\mathrm{diam}_{\mathcal{C}(\mathrm{S})}(\Lambda) < \delta'$, and thus if both L and R are in Λ then $d_{\mathcal{C}(S)}(L, R) < \delta'$. Also note that for any $D \in \Lambda$ as described in item 3 we have that the diameter in $\mathcal{M}(S)$ of the set of (D, v) is less than $B(6M + 4\delta')$ (since $6M + 4\delta'$ is the threshold for being in $\mathcal{G}$). Again, since the diameter in $\mathcal{C}(S)$ is bounded by δ' we then see that the set of all markings in item three has diameter in the marking complex bounded by $\delta' \cdot B(6M + 4\delta')$.

So now it follows that although $\Phi(\mu)$ consist of many slices:

$$\mathrm{diam}_{\mathcal{M}(\mathrm{S})}(\Phi(\mu) < \mathrm{2B(3M + 3\delta' + \delta'(6M + 4\delta'))}.$$

$\square$

Using the fact that the map from slices of a hierarchy to hierarchy paths is a Lipschitz map, the above shows that we can think of Φ as a map from $\mathcal{M}$ to subsets of hierarchy paths which has the property that points get sent to sets of uniformly bounded diameter.

In the next subsection we will show that these maps to hierarchy paths are "almost locally constant" off the hierarchy path, in the sense that for markings far from the hierarchy path the projections of extremely large diameter sets get mapped to uniformly small diameter subsets of the hierarchy path.

5.2 Coarsely Contracting Projections

The following will provide the framework for proving δ-hyperbolicity. It is a generalization to our context of Morse's Lemma on stability of quasi-geodesics in hyperbolic space.

Definition 5.2.1. In a space X, we say a family of paths $\mathcal{H}$ is *transitive* if every pair of points in X can be connected by a path in $\mathcal{H}$. We say $\mathcal{H}$ has the *coarsely contracting property* when $\mathcal{H}$ is a transitive family of paths in X with the property that for every $H \in \mathcal{H}$ there exists a map $\Phi_H : X \to H$ and constants b and c such that each for each $\mu, \mu' \in X$ satisfying $d_X(\mu, \mu') < b \cdot d_X(\mu, \Phi_H(\mu))$, then $\mathrm{diam}(\Phi_H(\mu), \Phi_H(\mu')) < \mathrm{c}$.

If considering $\mathcal{M}(S)$ then fix S to be either $S_{1,1}$ or $S_{0,4}$, if considering $\mathcal{P}(S)$ then fix S to be either $S_{1,2}$ or $S_{0,5}$ and use hierarchies-without-annuli.

Lemma 5.2.2.

- *When S is either $S_{1,1}$ or $S_{0,4}$, then the hierarchy paths form a coarsely contracting family of paths on $\mathcal{M}(S)$.*
- *When S is either $S_{1,2}$ or $S_{0,5}$, then the hierarchy-without-annuli paths form a coarsely contracting family of paths on $\mathcal{P}(S)$.*

Proof:

There exist a hierarchy connecting any pair of points in $\mathcal{M}(S)$, which shows that the set of hierarchy paths form a transitive path family (similarly in $\mathcal{P}(S)$ for hierarchy-without-annuli paths). Fix H a hierarchy path between two points in $\mathcal{M}(S)$ and use Φ to denote Φ_H.

We need to show: there exist constants b and c so that if $\mu, \mu' \in \mathcal{M}$ satisfy $d_{\mathcal{M}}(\mu, \mu') < b \cdot d_{\mathcal{M}}(\mu, \Phi(\mu))$, then $\mathrm{diam}(\Phi(\mu), \Phi(\mu')) < \mathrm{c}$.

First notice that a key property of Φ is that there exists a constant $Q = 3M + 3\delta'$ so that every domain A of a geodesic in H (including the main surface S) has $\mathrm{diam_A}(\Phi(\mu) \cup \mathrm{p_A}(\mu)) < \mathrm{Q}$; this fact is shown in the proof of Lemma 5.1.2.

Fix a hierarchy G from $\Phi(\mu)$ to μ. Let Z be a domain of G of length at least $4Q$, an assumption which implies that:

$$d_Z(\mu, \Phi(\mu)) > 3Q + M. \tag{5.2}$$

Let Y be a large domain of H. In order to compute $d_{\mathcal{M}}(\Phi(\mu), \Phi(\mu'))$ it is useful to calculate $d_Y(\mu, \mu')$. If Z is also a large domain of H then since Φ is defined so that in each domain Z, μ is sent via a closest point projection to a bounded diameter subset $\pi_Z(\Phi(\mu))$ of the $\mathcal{C}(Z)$ geodesic of H, we have that $d_Z(\mu, \partial Y) \geq d_Z(\mu, \Phi(\mu)) - 2\delta'$, by Equation 5.1. If Z is not a large domain of H, then $d_Z(\partial Y, \Phi(\mu)) < 6M + 4\delta'$; applying the triangle inequality then gives $d_Z(\mu, \partial Y) \geq d_Z(\mu, \Phi(\mu)) - \mathrm{diam_Z}(\Phi(\mu)) - \mathrm{d_Z}(\Phi(\mu), \partial \mathrm{Y}) - 2\delta'$.

Combining the two cases, independent of whether or not Z is a large domain in H, we have:

$$d_Z(\mu, \partial Y) \geq d_Z(\mu, \Phi(\mu)) - 3Q. \tag{5.3}$$

Combining the above with inequality 5.2 we have that $d_Z(\mu, \partial Y) > 3Q + M - 3Q = M$, which by the Projection estimates Theorem implies that $d_Y(\mu, \partial Z) < M$.

Now consider $d_Z(\mu', \partial Y)$: either $d_Z(\mu', \partial Y) > M$ for every large domain Y of H or for some domain Y of H we have $d_Z(\mu', \partial Y) < M$.

In the first case, from Theorem 4.2.1 (Projection estimates) we have that $d_Y(\mu', \partial Z) < M$, which combines with the paragraph above to give $d_Y(\mu', \mu) < 2M$. At which point it is easy to check that there is a uniform bound on $d_{\mathcal{M}}(\Phi(\mu), \Phi(\mu'))$—the bound comes since $d_Y(\mu', \mu) < 2M$ implies that the pairs $(\mathcal{L}(\mu), \mathcal{L}(\mu))$ and $(\mathcal{L}(\mu'), \mathcal{R}(\mu'))$ agree (except for a switch of at most one surface which then must have diameter not much larger than the threshold for $\mathcal{G}$) and that $\Lambda(\mu)$ is roughly the same as $\Lambda(\mu')$. Carrying out this computation, one obtains $d_{\mathcal{M}(S)}(\Phi(\mu), \Phi(\mu')) < 10M + 6\delta'$.

In the other case, when there is some domain Y of H for which $d_Z(\mu', \partial Y) < M$, we can combine this with equation 5.3 to give:

$$d_Z(\mu, \mu') \geq d_Z(\mu, \Phi(\mu)) - 3Q - M. \tag{5.4}$$

This tells us that whenever $d_Z(\Phi(\mu), \mu) > 2(3Q + M)$, we have

$$d_Z(\mu, \mu') > \frac{1}{2} d_Z(\Phi(\mu), \mu) \tag{5.5}$$

Thus, in domains Z of G for which $d_Z(\Phi(\mu), \mu)$ is sufficiently large, then unless $d_Z(\mu, \mu')$ is a definite fraction of $d_Z(\Phi(\mu), \mu)$ we have $d_Z(\mu', \partial Y) > M$ and as markings $\Phi(\mu)$ and $\Phi(\mu')$ are close.

If we can show that a definite proportion of the distance travelled in $\mathcal{M}$ between $\Phi(\mu)$ and μ takes place in domains larger than a threshold $6Q + 2M$ then we would be done by choosing the constant b less than this fraction, as then μ' has no chance to move far enough away from μ to have $d_Z(\mu', \partial Y) < M$ for a surface Z which is large in G and surface Y which is large in H.

That a definite fraction of distance occurs in large surfaces follows from the following counting argument. The fact that the the only domains of S which support geodesics in H are the surface S and component domains of g_H has as a consequence that if a hierarchy between I and T contains geodesics in $N+1$

different component domains, then $d_{\mathcal{C}(S)}(I,T) \geq N$. From this we see that either $\frac{1}{2}d_{\mathcal{M}(S)}(\Phi(\mu),\mu)$ occurs in large domains, or the sum of the lengths of geodesics shorter than $6Q+2M$ is larger than $\frac{1}{2}d_{\mathcal{M}(S)}(\Phi(\mu),\mu)$. In the latter case, we then have $|g_H| > \frac{d_{\mathcal{M}(S)}(\Phi(\mu),\mu)}{2(6Q+2M)}$. Thus we always have that at least $\frac{1}{2(6Q+2M)}d_{\mathcal{M}(S)}(\Phi(\mu),\mu)$ occurs in the large surfaces, as claimed.

Thus, choosing $b = \frac{1}{2(6Q+2M)}$ and $c > 10M + 6\delta'$, we have proved the theorem.

□

5.3 Hyperbolicity

In this section we provide proofs of Theorems 5.0.8 and 5.0.9 using the contracting properties of hierarchy paths provided by Lemma 5.2.2.

In [MM1] a contraction property is defined similar to Definition 5.2.1 and is used to prove δ-hyperbolicity of $\mathcal{C}(S)$. That property is what motivated the definition we gave in the previous section. In this section, using results about $\mathbb{R}$-trees we give a new proof that such contraction properties imply hyperbolicity, using different techniques than those of [MM1].

Theorem 5.3.1. *If X is a geodesic space with a family of (K,C)-quasi-geodesic paths $\mathcal{H}$ which have the coarsely contracting property, then X is δ-hyperbolic.*

Theorems 5.0.8 and 5.0.9 follow from this theorem since Lemma 5.2.2 proves that in the low complexity cases $\mathcal{M}(S)$ (and $\mathcal{P}(S)$) the hierarchy paths form a coarsely contracting family of paths (respectively hierarchy-without-annuli paths), and [MM2, Efficiency of hierarchies Theorem] implies that hierarchy paths are (K,C)-quasi-geodesics, with the constants depending only on the topological type of the surface.

Proof: The method of proof is to show that in the asymptotic cone, Φ induces a locally constant map to a certain family of bi-Lipschitz paths and then we show that this implies that $\mathrm{Cone}_\infty(\mathrm{X})$ is an $\mathbb{R}$-tree. In particular we will show for any pair of points $x,y \in \mathrm{Cone}_\infty(\mathrm{X})$ there exists a map $\Pi : \mathrm{Cone}_\infty(\mathrm{X}) \to [\mathrm{x},\mathrm{y}]$ such that:

1. $\Pi|_{[x,y]}$ is the identity

2. Π is locally constant outside $[x,y]$.

By our hypothesis, for each $H \in \mathcal{H}$ there exist a map $p_H : X \to X$ such that:

1. For each $\mu \in X$ we have $p_H(\mu) \in H$.

2. There exist constants $b, c > 0$ such that for any sequence of paths H_n and points μ_n with $r_n = d_X(\mu_n, \pi_H(\mu_n))$ growing linearly, the image $p_H(B_{b \cdot r_n}(\mu_n))$ has diameter less than or equal to c.

Fix three points $\bar{I} = \langle I_i \rangle_{i \in \mathbb{N}}, \bar{T} = \langle T_i \rangle_{i \in \mathbb{N}}, \bar{\mu} = \langle \mu_i \rangle_{i \in \mathbb{N}} \in \mathrm{Cone}_\infty(\mathrm{S})$. Let $\bar{H}$ denote the path connecting $\bar{I}$ to $\bar{T}$ given by a rescaled sequence of paths H_i connecting I_i to T_i. The definition of $\mathcal{H}$ having the coarsely contracting property (Definition 5.2.1) implies that when $d_X(\mu_i, H_i)$ grows linearly the ball with radius $b \cdot d_{\mathcal{M}(S)}(\mu_i, H_i)$ around μ_i maps to a ball of radius at most c in H_i. Then in the rescaled metric space $\frac{1}{i} \cdot X$ we have $d_{\frac{1}{i} \cdot X}(\mu_i, H_i)$ is a constant independent of i and in this metric the ball of radius b around μ_i maps to a set of diameter $\frac{c}{i}$ in H_i. Thus after taking ultralimits we have that the ball of radius b around any point $\bar{\mu} \in \omega\text{-}\lim_i \epsilon_i \cdot X$ not lying on $\bar{H}$ maps to a point on $\bar{H}$, which is to say the map is locally constant.

This proves that any embedded path connecting $\bar{I}$ and $\bar{T}$ must be a subset of $\bar{H}$. Since by hypothesis each $H \in \mathcal{H}$ is a (K, C)-quasi-geodesics, we have that $\bar{H}$ is the ultralimit as $i \to \infty$ of $(K, \frac{C}{i})$-quasi-geodesics and is thus a K-bi-Lipschitz embedded path. In particular it is homeomorphically embedded, so the only subpath connecting its endpoints is the path itself—thus proving that there is a unique arc between any pair of points in $\mathrm{Cone}_\infty(\mathrm{X})$. It is worth remarking that since $\mathrm{Cone}_\infty(\mathrm{X})$ is a geodesics space, the path $\bar{H}$ which *a fortiori* need only be bi-Lipschitz embedded is actually the geodesic between $\bar{I}$ and $\bar{T}$.

Mayer and Oversteegen's topological characterization of $\mathbb{R}$-trees [MaO] then tells us that in order to prove that $\mathrm{Cone}_\infty(\mathrm{X})$ is an $\mathbb{R}$-tree it suffices to show that it is uniquely arcwise connected and locally path connected. The first property was shown above, the second is true for any geodesic space.

Thus each $\mathrm{Cone}_\infty(\mathrm{X})$ is an $\mathbb{R}$-tree and thus Theorem 2.2.5 implies that X is δ-hyperbolic. □

Chapter 6

Evidence towards the Rank Conjecture

The goal of this chapter is to utilize the geometric picture developed in Theorem 4.3.6 (Coarse tree-flats carry markings) to develop further insight into the structure of quasiflats in the mapping class group. A connection between our characterization of the geometry of subsurfaces projection maps (Theorem 4.3.6) and Conjecture 6.3.6 (Rank Conjecture) comes from the Move distance and projection Theorem which states that distance in the mapping class group can be estimated by summing over the distance of certain subsurface projections. More precisely, for any sufficiently large constant t, there exist quasi-isometry constants $K(t)$ and $C(t)$, so that for any $\mu, \nu\mathcal{M}$ the distance in the marking graph agrees with $\sum_{\substack{Y\subseteq S \\ d_Y(\mu,\nu)>t}} d_Y(\mu,\nu)$ up to bounded additive and multiplicative errors (determined by $K(t)$ and $C(t)$). Note that this theorem does not give a quasi-isometric embedding into the product space of all curve complexes, since this distance estimates require only looking at sufficiently large distances (a finite set of subsurface distances) and thus differs from the l^1 metric on $\prod_{Y\subset S}\mathcal{C}(Y)$. In the next section we show that their construction is similar enough to a quasi-isometric embedding to yield useful information about the asymptotic cone of the mapping class group.

6.1 Ultralimits of Subsurfaces

In this section we remind the reader of some useful properties of ultralimits and provide a few preliminary results concerning ultralimits of subsurfaces. We incorporate these two goals with an introduction to the notation which we will use throughout this chapter.

Recall that for a fixed ultrafilter ω and surface S we defined Seq to be the set of sequences of homotopy classes of essential, non-peripheral subsurfaces $X \subset S$ with $\xi(X) \neq 0$, considered up to the relation $\sim$, where two sequences $\alpha = \langle\alpha_i\rangle_{i\in\mathbb{N}}$ and $\beta = \langle\beta_i\rangle_{i\in\mathbb{N}}$ satisfy $\alpha \sim \beta$ if and only if $\alpha_i = \beta_i$ for each i in some set $K \in \omega$.

We will use $\mu = \langle \mu_j \rangle_{j\in\mathbb{N}}$ to denote a point in $\mathcal{M}_\infty(S)$, which we recall is the space of equivalence classes of sequences of points $\mu_j \in \mathcal{M}$ with $\frac{1}{j} d_{\mathcal{M}}(0, \mu_j) < \infty$. As shorthand, we write $\mathcal{M}_\infty$ to denote $\mathrm{Cone}_\infty \mathcal{M}(\mathrm{S})$.

Note that once we have fixed a surface $S = S_{g,p}$, the subsurfaces of S are each one of a finite collection of topological types. Accordingly, for each $\alpha = \langle \alpha_j \rangle_{j\in\mathbb{N}} \in \mathrm{Seq}$ one obtains a partition of the indices into finitely many sets depending on the topological type of α_j. Since ω is an ultrafilter, in any decomposition of $\mathbb{N}$ into finitely many subsets, exactly one of these subsets is in ω. Thus the ultrafilter associates a (unique) topological type to α.

By the same principle as above, given any pair of sequences of subsurfaces (in this section these will be denoted by the first few lower case greek letters), $\alpha, \beta \in \mathrm{Seq}$, the ultrafilter tells us that a full ω-measure set of the pairs α_i, β_i are all either disjoint or intersecting. We remind the reader that we use the term intersection to refer to transverse intersection; for example a subsurface does not intersect its boundary. In the case where α and β intersect one can differentiate between the cases where they overlap and when they are nested. Accordingly we will simply say α and β are either disjoint, overlapping, or nested.

When $F \subset \mathrm{Seq}$ we use the notation F_i to refer to the set of i^{th} entries of F. A consequence of the above discussion is that there is a natural definition of *maximal collections of pairwise disjoint elements of* F; such collections will be denoted either by Δ or Λ. For clarity, we define this explicitly: we say $\Delta = \langle U_i \rangle_{i\in\mathbb{N}}$ is a maximal collection of pairwise disjoint elements of F if and only if for a full ω-measure set of indices, U_i is a maximal disjoint collection of elements of F_i as defined in Definition 4.3.1. It follows from Lemma 2.2.6 that $|\Delta| = \omega\text{-}\lim_i |U_i|$. As in Definition 4.3.1, we define $\mathcal{U}_F$ as the set of all maximal collection of disjoint elements of F.

Let us fix a marking $0 \in \mathcal{M}(S)$ to use as a base point and introduce the notation $0_Y = \pi_Y(0)$. Theorem 3.1.2 tells us that each curve complex is δ-hyperbolic. Since there are only finitely many topological types of subsurfaces of S we need to consider only finitely many isometry types of curve complexes, and thus we can choose a uniform δ which works for all of them. Since $\frac{1}{i}(\mathcal{C}(\alpha_i), 0_{\alpha_i})$ is $\frac{\delta}{i}$-hyperbolic, $\omega\text{-}\lim_i \frac{1}{i}(\mathcal{C}(\alpha_i), 0_{\alpha_i})$ is an ultralimit of spaces with hyperbolicity constant going to zero; Theorem 2.2.4 then implies that this space is an $\mathbb{R}$-tree which we henceforth denote T_α. As an example to keep in mind, notice that if α is a constant sequence, then T_α is just the asymptotic cone of the complex of curves of that particular subsurface. Furthermore, since ultralimits commute with finite products: for any finite set $F \subset \mathrm{Seq}$ we have

$$\omega\text{-}\lim_i \frac{1}{i}(\prod_{\alpha_i \in F_i} \mathcal{C}(\alpha_i), \prod_{\alpha_i \in F_i} 0_i) = \prod_{\alpha \in F} T_\alpha.$$

When $\partial\beta$ intersects α we consider two possibilities for the point

$$\langle \pi_{\alpha_i}(\partial\beta_i) \rangle \in \omega\text{-}\lim_i \frac{1}{i}\mathcal{C}(\alpha_i).$$

First, if $\langle \pi_{\alpha_i}(\partial\beta_i) \rangle$ is in T_α (the connected component of $\omega\text{-}\lim_i \frac{1}{i}\mathcal{C}(\alpha_i)$ containing 0_α), then we denote this point by $\pi_\alpha(\beta)$. The second case is when $\langle \pi_{\alpha_i}(\partial\beta_i) \rangle$ is in a connected component not containing 0_α, here we set $\pi_\alpha(\beta) = \emptyset$. Usually we will provide a discussion of $\pi_\alpha(\beta)$ assuming it is a point in T_α, but the arguments always work equally well in the case that this is the empty set.

When $\alpha \subset \beta$ we consider the map $\pi_{\beta \to \alpha} : T_\beta \setminus \pi_\beta(\alpha) \to T_\alpha$, which is defined as the ultralimit of $\pi_{\beta_i \to \alpha_i} : T_{\beta_i} \setminus B_1(\pi_{\beta_i}(\partial\alpha_1)) \to T_{\alpha_i}$. The following result is a strong version of Corollary 3.1.5 (Projections coarsely Lipschitz), providing more than just a non-coarse version of that theorem.

Lemma 6.1.1. (Projections locally constant) *For any S and any $\alpha, \beta \in$* Seq*, the map*

$$\pi_{\beta \to \alpha} : T_\beta \setminus \pi_\beta(\alpha) \to T_\alpha$$

is locally constant.

Proof: Fix a pair of points μ, ν in the same connected component of $T_\beta \setminus \pi_\beta(\alpha)$.

Since T_β is an $\mathbb{R}$-tree, it is uniquely arc connected. This implies in particular that the geodesic between μ_i and ν_i is uniformly bounded away from the radius 1 neighborhood of $\pi_{\beta_i}(\alpha_i)$. Thus, by Theorem 3.1.6 (Bounded geodesic image; [MM2]) we have $\mathrm{diam}_{\mathcal{C}(\beta_\mathrm{i})}(\pi_{\beta_\mathrm{i} \to \alpha_\mathrm{i}}(\mu_\mathrm{i} \cup \nu_\mathrm{i}) < \mathrm{K}$ for a uniform constant K depending only on the topological type of S.

Thus, $\mathrm{diam}_{\mathrm{T}_\beta}(\pi_{\beta \to \alpha}(\mu \cup \nu)) = 0$ and we see that $\pi_{\beta \to \alpha}(\mu) = \pi_{\beta \to \alpha}(\nu)$, i.e., the map locally constant. □

Recall, the complexity function $\xi(S) = 3g + p - 3$ which when $\xi(S_{g,p}) > 0$ gives the maximal torsion free rank of an abelian subgroup of $\mathcal{MCG}(S_{g,p})$.

ξ generalizes to a complexity function on disjoint unions of non-annular subsurfaces via $\xi(U) = 3g + p - 3c$ where $c = |U|$ is the number of connected components of U. One sees that this has similar properties to the complexity function as defined on connected surfaces, so we maintain the use of the name ξ. Recall, that subsurfaces are always assumed to be essential and non-peripheral;

also they are considered up to isotopy, so parallel annuli are not considered disjoint.

We end this section by remarking that this can be extended to a complexity function defined on all disjoint unions of subsurfaces, moreover this extension is additive.

Let U denote a collection of disjoint subsurfaces of S. Define $\bar{\xi}(U)$ to be the maximal number of distinct homotopy classes of essential and non-peripheral curves which can simultaneously be realized on U plus the number of annular components in U.

Proposition 6.1.2. *The function $\bar{\xi}(U)$ satisfies:*

1. *For a connected subsurface $X \subset S$, $\bar{\xi}(X) = |\xi(X)|$.*
2. *If U is a collection of disjoint surfaces inside X then $\bar{\xi}(U) \leq \bar{\xi}(X)$.*
3. *Given a collection U of disjoint subsurfaces, then if $U = U_1 \sqcup U_2$ then $\bar{\xi}(U) = \bar{\xi}(U_1) + \bar{\xi}(U_2)$.*

Proof:

(1) When X is not an annulus: a maximal collection of disjoint, essential, non-peripheral curves in X yields a pants decomposition; one can check using the Euler characteristic that such a decomposition consists of $3g+p-3$ curves. Thus $\bar{\xi}(X) = \xi(X)$. When X is an annulus, then $\xi(X) = -1$, so $\bar{\xi}(X) = |\xi(X)|$.

(2) If $U \subset X$, then any homotopy class of an essential, non-peripheral curve in U remains so in X. Furthermore, since the annuli are required to be non-peripheral any annuli in U yields an essential curve in X which is distinct from the non-peripheral curves in U, thus yielding the result.

(3) From Part 2, we have that each curve (and annulus) in each of the U_i is an essential and non-peripheral curve in X. Since we require curves to be non-peripheral, no curve in U_1 can be homotoped into a non-peripheral curve of U_2. Then, since the same homotopy class of annuli can not appear in both the U_i, we have that $\bar{\xi}(U) = \bar{\xi}(U_1) + \bar{\xi}(U_2)$.

□

By the discussion earlier in this section, it follows from the above proposition that the quantity $\bar{\xi}(\Delta)$ is well defined, where Δ is any maximal collection of disjoint elements of Seq,

6.2 A Lipschitz Map to Trees

In this section and the next we will define the relevant terms and then prove the following result which is a combination of Proposition 6.2.2 and Theorem 6.3.2.

By taking ultralimits of the projection maps studied in the previous chapters, we will define a map $\widehat{\psi} : \mathcal{M}_\infty \to \prod_\alpha T_\alpha$ and prove that the following holds.

Theorem 6.2.1. *The map $\widehat{\psi} : \mathcal{M}_\infty \to \prod_\alpha T_\alpha$ is Lipschitz. Moreover any projection of $\widehat{\psi}(\mathcal{M}_\infty)$ to finitely many products of trees is a finite union of tree-flats each of rank less than or equal to $3g + p - 3$.*

We now define the main character in this chapter. For each $\alpha \in \mathrm{Seq}$ we let:

$$\widehat{\psi_\alpha} : \mathrm{Cone}_\infty(\mathcal{M}, 0) \to \mathrm{T}_\alpha$$

be the map defined by $\widehat{\psi_\alpha} : \langle \mu_i \rangle \mapsto \omega\text{-}\lim_i \frac{1}{i}(\pi_{\alpha_i}(\mu_i), 0_{\alpha_i})$. Note that for any $\mu, \nu \in \mathcal{M}_\infty$ we have $d_{T_\alpha}(\widehat{\psi}_\alpha(\mu), \widehat{\psi}_\alpha(\nu)) = \omega\text{-}\lim_i \frac{1}{i} d_{Y_{\alpha_i}}(\mu, \nu)$.

Although it follows from the next theorem, we will give a quick argument here to show that the map $\widehat{\psi}$ is well defined. Since $\langle \mu_i \rangle \in \mathcal{M}_\infty$ implies that $\omega\text{-}\lim_i \frac{1}{i} d_{\mathcal{M}}(0, \mu_i) < \infty$, by Theorem 3.3.9 (Move distance and projections; [MM2]) we have $\omega\text{-}\lim_i \frac{1}{i} \sum_{\substack{Y \subseteq S \\ d_Y(\mu,\nu) > t}} d_{\mathcal{C}(Y)}(0_Y, \mu) < \infty$. In particular this shows that $\omega\text{-}\lim_i \frac{1}{i} d_{\mathcal{C}(\alpha_i)}(0_{\alpha_i}, \mu) < \infty$ for each $\alpha \in \mathrm{Seq}$. From this we have that the map $\widehat{\psi}_\alpha$ is well defined.

Giving $\prod_{\alpha \in \mathrm{Seq}} T_\alpha$ the l^1 metric, we consider

$$\widehat{\psi} = \prod_{\alpha \in \mathrm{Seq}} \widehat{\psi}_\alpha : \mathcal{M}_\infty \to \prod_{\alpha \in \mathrm{Seq}} T_\alpha.$$

The next theorem proves the first half of Theorem 6.2.1.

Proposition 6.2.2. *The map*

$$\widehat{\psi} : \mathcal{M}_\infty \to \prod_{\alpha \in \mathrm{Seq}} T_\alpha$$

is Lipschitz.

Proof: Fix an ultrafilter ω, and two points $\langle \mu_i \rangle, \langle \nu_i \rangle \in \mathrm{Cone}_\omega \mathcal{M}$. By definition of the asymptotic cone $d_{\mathrm{Cone}_\omega(\mathcal{M})}(\bar{\mu}, \bar{\nu}) = \omega\text{-}\lim_i \frac{1}{i} d_{\mathcal{M}}(\mu_i, \nu_i)$, and thus Theorem 3.3.9 implies that there exist a threshold $t(S)$, so that for each

$t > t(S)$, there exists constants $K = K(t) > 1$ and $B = B(t) > 0$ so that the distance in the asymptotic cone is bounded below by

$$\omega\text{-}\lim \frac{1}{i}\left(-B + \frac{1}{K}\sum_{\substack{Y\subseteq S\\ d_Y(\mu_i,\nu_i)>t}} d_Y(\mu_i,\nu_i)\right).$$

Since ultralimits commute with addition and multiplication, the above expression is equal to $\frac{1}{K}\omega\text{-}\lim \frac{1}{i}\sum_{\substack{Y\subseteq S\\ d_Y(\mu_i,\nu_i)>t}} d_Y(\mu_i,\nu_i)$.

To show that $\widehat{\psi} : \mathrm{Cone}_\omega\mathcal{M} \to \prod_{\alpha\in\mathrm{Seq}} \mathrm{T}_\alpha$ is a Lipschitz map, it suffices to show that

$$\sum_{\alpha\in\mathrm{Seq}} d_{T_\alpha}(\widehat{\psi}_\alpha(\mu), \widehat{\psi}_\alpha(\nu)) \leq \omega\text{-}\lim \frac{1}{i}\sum_{\substack{Y\subseteq S\\ d_Y(\mu_i,\nu_i)>t}} d_Y(\mu_i,\nu_i). \tag{6.1}$$

To prove this inequality we first show that it holds when Seq is replaced on the left-hand side by a finite set of distinct sequences $\Gamma \subset \mathrm{Seq}$.

We have the following inequalities

$$\begin{aligned}\sum_{\gamma\in\Gamma} d_{T_\gamma}(\widehat{\psi}_\gamma(\mu), \widehat{\psi}_\gamma(\nu)) \quad &= \omega\text{-}\lim \frac{1}{i}\sum_{\gamma\in\Gamma} d_{\gamma_i}(\mu_i,\nu_i) &(6.2)\\ &\leq \omega\text{-}\lim \frac{1}{i}\sum_{\substack{Y\subseteq S\\ d_Y(\mu_i,\nu_i)>t}} d_Y(\mu_i,\nu_i). &(6.3)\end{aligned}$$

The first line is the definition of distance in T_γ. The second line comes from the following two facts. First, if $d_{T_\gamma}(\widehat{\psi}_\gamma(\mu), \widehat{\psi}_\gamma(\nu)) = C > 0$, then for any $\epsilon > 0$ there exists a set $I \in \omega$, where for each $i \in I$ we have $d_{\gamma_i}(\mu_i,\nu_i) > i(C-\epsilon)$. This implies that if the distance on the left hand side is not zero, then when i is sufficiently large we have $d_\gamma(\mu_i,\nu_i) > t$ and thus this distance is included in the summation on line 6.3. Second, by Lemma 2.2.6, we have that pairwise distinct sequences in Γ yield distinct terms in the last summation of the above inequality.

Since an infinite sum is the limit of the finite approximating sums, we have:

$$\begin{aligned}\sum_{\alpha\in\mathrm{Seq}} d_{T_\alpha}(\widehat{\psi}_\alpha(\mu), \widehat{\psi}_\alpha(\nu)) \quad &= \sup_{\substack{\Gamma\subset\mathrm{Seq}\\ |\Gamma|<\infty}} \sum_{\gamma\in\Gamma} d_{Y_\gamma}(\widehat{\psi}_\gamma(\mu), \widehat{\psi}_\gamma(\nu))\\ &\leq \omega\text{-}\lim \frac{1}{i}\sum_{\substack{Y\subseteq S\\ d_Y(\mu_i,\nu_i)>t}} d_Y(\mu_i,\nu_i)\end{aligned}$$

which is what we wanted to show.

□

6.3 Dimension and Rank

There are certain subsets of $\mathcal{M}_\infty$ on which $\widehat{\psi}$ is bi-Lipschitz. The plan for this section is to study the image of such subsets. Then in the remaining sections of this chapter we work to understand in an inductive manner the subsets of $\mathcal{M}_\infty$ where this map collapses large diameter sets and explain the beginnings of a process to tie this information together in order to control the geometry of the entire asymptotic cone. (We hope to elaborate on this process in a future paper.)

We now combine the map $\widehat{\psi}$ constructed in the previous section with Theorem 4.3.6 (Coarse tree-flats carry markings) to gain information about the asymptotic geometry of the mapping class group and of Teichmüller space.

By Theorem 4.3.6, we know that for any finite collection of surfaces Θ we have $\pi_\Theta(\mathcal{M})$ is contained in a finite union of coarse tree-flats, we now give a version of that theorem in the asymptotic cone of $\mathcal{M}$.

For a set $F \subset \mathrm{Seq}$ we write:

$$\widehat{\psi}_F : \mathcal{M}_\infty \to \prod_{\alpha \in F} T_\alpha.$$

For $\Delta = \langle U_i \rangle_{i\in\mathbb{N}}$ a sequence of maximal disjoint collections of F, we define a *tree flat in* $\prod_{\alpha\in F} T_\alpha$ *of rank* $\bar{\xi}(\Delta)$ to be:

$$\mathcal{T}(\Delta) = \omega\text{-}\lim_i \frac{1}{i}(\mathcal{F}_{F_i}(\Delta_i), 0_{F_i}).$$

In the following theorem, we show that for any finite set $F \in \mathrm{Seq}$, we have $\widehat{\psi}_F(\mathcal{M})$ is a finite union of tree-flats, and these tree-flats each have topological dimension at most $\bar{\xi}(\Delta)$.

In the Appendix we provide a brief summary of the relevant background in dimension theory, which the reader may want to consult at this time.

Theorem 6.3.1. *Given a surface S and $\alpha, \beta \in \mathrm{Seq}$, then exactly one of the following occurs:*

1. *$\alpha \cap \beta = \emptyset$, in which case*

$$\widehat{\psi}_{\{\alpha,\beta\}}(\mathcal{M}_\infty) = T_\alpha \times T_\beta.$$

   Moreover, $\dim(\widehat{\psi}_{\{\alpha,\beta\}}(\mathcal{M}_\infty)) = 2$

2. *α and β overlap, in which case*

$$\widehat{\psi}_{\{\alpha,\beta\}}(\mathcal{M}_\infty) \subset T_\alpha \times \pi_\beta(\alpha) \bigcup \pi_\alpha(\beta) \times T_\beta,$$

   where at most one of $\pi_\beta(\alpha)$ or $\pi_\alpha(\beta)$ is allowed to be empty. Moreover, $\widehat{\psi}_{\{\alpha,\beta\}}(\mathcal{M}_\infty)$ is an $\mathbb{R}$-tree and thus in particular has dimension 1.

3. *One of the elements, say β, is nested in the other. Here*

$$\widehat{\psi}_{\{\alpha,\beta\}}(\mathcal{M}_\infty) \subset \pi_\alpha(\beta) \times T_\beta \bigcup Graph(\pi_{\alpha\to\beta}),$$

where $\pi_{\alpha\to\beta}$ is a locally constant map with domain $T_\alpha \setminus \pi_\alpha(\beta))$ and range T_β. $\pi_\alpha(\beta)$ is allowed to be the empty set, in which case $\pi_\alpha(\beta) \times T_\beta = \emptyset$ and $\pi_{\alpha\to\beta}$ is a constant map from T_α to T_β. Moreover, $\widehat{\psi}_{\{\alpha,\beta\}}(\mathcal{M}_\infty)$ is an $\mathbb{R}$-tree and thus in particular has dimension 1.

Proof:

We use the three cases of Theorem 4.2.2 (Projection estimates, geometric version) to establish the theorem.

In each case, we will use Δ to denote a maximal disjoint collection of elements in $F = \{\alpha, \beta\}$.

(Disjoint surfaces).

When α and β are disjoint we have $\Delta = \{\alpha, \beta\}$. Hence

$$\pi_{F_i}(\mathcal{M}) \subset \mathcal{F}^{C_2}_{F_i}(\Delta_i) = \mathcal{C}(\alpha_i) \times \mathcal{C}(\beta_i),$$

for all i in a full ω-measure set, and thus

$$\begin{aligned} \omega\text{-}\lim_i \tfrac{1}{i}\pi_{F_i}(\mathcal{M}) \subset \ & \omega\text{-}\lim_i \frac{1}{i}(\mathcal{F}^{C_2}_{F_i}(\Delta_i), 0) \\ & = \omega\text{-}\lim_i \frac{1}{i}(\mathcal{C}(\alpha_i) \times \mathcal{C}(\beta_i), 0_{\alpha_i} \times 0_{\beta_i}) \\ & = T_\alpha \times T_\beta. \end{aligned}$$

By Corollary A.7 (a consequence of work of Mayer and Oversteegen, [MaO]) we have $\dim(\mathcal{T}(\mathrm{T}_\alpha \times \mathrm{T}_\beta)) = 2$ and thus Theorem 4.2.2 (Projection estimates; geometric version) implies $\dim(\mathcal{T}(\Delta)) = \dim(\widehat{\psi}_{\mathrm{F}}(\mathcal{M}_\infty)) = 2$.

(Overlapping surfaces).

We now consider the case where α and β overlap. Here by Theorem 4.2.2 in the case of overlapping surfaces, we know that

$$\widehat{\psi}_{F_i}(M) \subset \mathcal{F}^{C_2}_{F_i}(\{\alpha_i\}) \cup \mathcal{F}^{C_2}_{F_i}(\{\beta_i\}).$$

We start with the first of these sets:

$\mathcal{F}_{F_i}^{C_2}(\{\alpha_i\}) = B_M(\mathcal{C}(\alpha_i) \times \pi_{\beta_i}(\partial\alpha_i))$. Since $\mathcal{C}(\alpha_i)$ is δ-hyperbolic, we have

$$\begin{aligned}
\mathcal{T}_F(\alpha) &= \omega\text{-}\lim_i \frac{1}{i}\widehat{\psi}(\mathcal{F}_{F_i}^{C_2}(\{\alpha_i\})) \\
&= \omega\text{-}\lim_i \frac{1}{i} B_M((\mathcal{C}(\alpha_i) \times \pi_{\beta_i}(\partial\alpha_i))) \\
&= T_\alpha \times \pi_\beta(\alpha),
\end{aligned}$$

which is an $\mathbb{R}$-tree sitting as a coordinate factor of $T_\alpha \times T_\beta$ and thus is closed. If $\pi_\beta(\alpha) = \emptyset$, then the above set $\mathcal{T}_F(\{\alpha\})$ is the empty set.

Analysis of $\mathcal{T}_F(\{\beta\})$ is the same as above, so we have $\mathcal{T}_F(\beta) = \pi_\alpha(\beta) \times T_\beta$. Again, it is possible that $\pi_\alpha(\beta) = \emptyset$, which would render $\mathcal{T}_F(\beta) = \emptyset$.

A key observation is that Theorem 4.2.1 (Projection estimates) implies that if $\omega\text{-}\lim_i \frac{1}{i} d_{\alpha_i}(\partial\beta_i, 0) = \infty$, then $d_{\beta_i}(\partial\alpha_i, 0) < M$ for a full ω-measure set of indices. Thus if $d_\alpha(\beta, 0) = \infty$, then $d_\beta(\alpha, 0) = 0$ and thus in particular, $\pi_\beta(\alpha) \neq \emptyset$. Since the same logic holds reversing the roles of α and β we see that at most one of $\pi_\beta(\alpha)$ or $\pi_\alpha(\beta)$ is empty.

Since $\widehat{\psi}_F(\mathcal{M}_\infty) \subset \mathcal{T}_F(\{\alpha\}) \cup \mathcal{T}_F(\{\beta\})$, we have finished the first half of the assertion. That this is an $\mathbb{R}$-tree follows immediately, since $\widehat{\psi}_F(\mathcal{M}_\infty)$ is the union of two $\mathbb{R}$-trees along a point (or when one of the base points is empty, it is just an $\mathbb{R}$-tree union the empty set).

(Nested surfaces).

We now assume $\beta \subset \alpha$. For nested surfaces Theorem 4.2.2 tells us that

$$\begin{aligned}
\pi_{F_i}(\mathcal{M}) \subset \mathcal{F}_{F_i}^{C_2}(\{\alpha_i\}) = \;& \pi_{\alpha_i}(\partial\beta_i)) \times B_{C_2}(\mathcal{C}(\beta_i)) \\
& \bigcup B_{C_2}(Graph(\pi_{\alpha_i \to \beta_i})),
\end{aligned}$$

where B_{C_2} denotes the ball of radius C_2 (as computed in Theorem 4.2.2) and $\pi_{\alpha_i \to \beta_i}$ has domain $\mathcal{C}(\alpha_i) \setminus B_1(\pi_{\alpha_i}(\beta_i))$ and range $\mathcal{C}(\beta_i)$.

We now consider these two subsets separately. First, we note

$$\omega\text{-}\lim_i \frac{1}{i} B_{C^2}(\mathcal{C}(\beta_i) \times \pi_{\alpha_i}(\partial\beta_i)) = T_\beta \times \pi_\alpha(\beta).$$

This is an $\mathbb{R}$-tree and forms a closed subset of $T_\alpha \times T_\beta$. The second subset is the neighborhood of a graph with domain $\mathcal{C}(\alpha_i) \setminus B_1(\pi_{\alpha_i}(\partial\beta_i))$ and range $\mathcal{C}(\beta_i)$. Thus we have:

$$\omega\text{-}\lim_i \frac{1}{i}\mathcal{F}_F^{C_2}(\{\alpha_i\})) = \pi_\alpha(\beta) \times T_\beta \cup Graph(\pi_{\alpha \to \beta}), \tag{6.4}$$

where $\pi_{\alpha\to\beta}$ has domain $T_\alpha \setminus \pi_\alpha(\beta)$ and range T_β.

Lemma 6.1.1 shows that $\pi_{\alpha\to\beta} : T_\alpha \setminus \pi_\alpha(\beta) \to T_\beta$ is locally constant, and thus being the graph over an $\mathbb{R}$-tree (minus the point $\pi_\alpha(\beta)$ if $\pi_\alpha(\beta) \neq \emptyset$) is a union of disjoint $\mathbb{R}$-trees each incomplete at a point of $T_\beta \times \pi_\alpha(\beta)$ (when $\pi_\alpha(\beta) = \emptyset$ then we just have one $\mathbb{R}$-tree). Since the $\pi_{\alpha\to\beta}$ is locally constant, the closure of each of these $\mathbb{R}$-trees contains exactly one new point, and that point is in $T_\beta \times \pi_\alpha(\beta)$ In particular, this shows that the set described in Equation 6.4 is closed and is an $\mathbb{R}$-tree.

Theorem A.6 then implies that since the above is an $\mathbb{R}$-tree we have $\dim(\widehat{\psi}_{\mathrm{F}}(\mathcal{M}_\infty)) = 1$.

$\square$

We will now prove that second part of Theorem 6.2.1.

Theorem 6.3.2. *Given a surface S and a finite subset $F \subset \mathrm{Seq}$, then $\widehat{\psi}_F(\mathcal{M}_\infty)$ is closed and is contained in a finite union of tree flats. In particular,*

$$\dim(\widehat{\psi}_{\mathrm{F}}(\mathcal{M}_\infty)) \leq 3\mathrm{g} + \mathrm{p} - 3.$$

Proof:

That $\widehat{\psi}_F(\mathcal{M}_\infty)$ is a finite union of tree-flats is an immediate consequence of the finiteness in Theorem 4.3.6 (Coarse tree-flats carry markings).

The key step to our proof of this theorem is to show that for any finite set $F \subset \mathrm{Seq}$ and for each $\Delta = \langle U_i \rangle_{i\in\mathbb{N}}$, a maximal disjoint collection of elements of F, the tree-flat $\mathcal{T}(\Delta)$ is closed and has dimension at most $\bar{\xi}(\Delta) \leq 3g + p - 3$.

Once this bound is established for the dimension of the tree-flats, the dimension bound for $\widehat{\psi}_F(\mathcal{M}_\infty)$ follows from the fact that dimension does not increase by taking finite unions of closed sets (see Theorem A.3).

Taking Theorem 6.3.1 as the base case, we proceed by induction on the number of elements in F.

We assume the inductive hypothesis: for any subset $G \subsetneq F$ and any $\Delta \in \mathcal{U}_G$ the set $\mathcal{T}_G(\Delta)$ is closed and has dimension at most $\bar{\xi}(\Delta)$.

For each $\delta \in \Delta$, define $F_\delta = \{\sigma \in F : \sigma \subset \delta\}$. Note that for each $\sigma \in F_\delta$ one has $\pi_{\delta_i}(\partial\sigma_i) \neq \emptyset$, although as per the usual discussion, $\pi_\delta(\sigma)$ may still be empty if $\omega\text{-}\lim_i \frac{1}{i} d_{\mathcal{M}}(\pi_{\delta_i}(\partial\sigma_i), 0_{\delta_i}) = \infty$.

For each $\delta \in \Delta$ we let $P_\delta = \{\pi_\delta(\sigma)$ for some $\sigma \in F_\delta\}$, noting that $|P_\delta| \leq |F| < \infty$. The $p \in P_\delta$ can be used to partition F_δ into finitely many disjoint sets $F_\delta(p) = \{\sigma \in F_\delta : \pi_\delta(\sigma) = p\}$.

Notice that if σ and σ' are disjoint then $\pi_\delta(\sigma) = \pi_\delta(\sigma')$. Thus each maximal collection of disjoint elements of F_δ lies in some $F_\delta(p)$.

We now consider $\widehat{\psi}_{F_\delta(p)\cup\{\delta\}}(\mathcal{M}_\infty)$ for some fixed δ and p. By definition, if $\widehat{\psi}_\delta(\mu) = p$ then $\widehat{\psi}_{F_\delta(p)\cup\{\delta\}} \in \{p\} \times \widehat{\psi}_{F_\delta(p)}(\mathcal{M}_\infty)$. By the inductive hypothesis, $\widehat{\psi}_{F_\delta(p)}(\mathcal{M}_\infty) \subset \cup_{\Lambda \in \mathcal{U}_{F_\delta(p)}} \mathcal{T}_{F_\delta(p)}(\Lambda)$.

Theorem 4.3.6 and Lemma 6.1.1 then implies that $\widehat{\psi}_{F_\delta(p)\cup\{\delta\}}(\mathcal{M}_\infty)$ is the union of $\{p\} \times \cup_{\Lambda\in\mathcal{U}_{F_\delta(p)}} \mathcal{T}_{F_{\delta(p)}}(\Lambda)$ and the graph of a locally constant function with domain $T_\delta \setminus \{p\}$ and range $\cup_{\Lambda\in\mathcal{U}_{F_\delta(p)}} \mathcal{T}_{F_\delta(p)}(\Lambda)$.

Define $X_p = \cup_{\Lambda\in\mathcal{U}_{F_\delta(p)}} \mathcal{T}_{F_{\delta(p)}}(\Lambda) \subset \prod_{\sigma\in F_\delta(p)} T_\sigma$. By the inductive hypothesis, it follows that this set is closed and by Lemma 6.1.2 has $\mathrm{dim(X_p)} = \max_{\Lambda\in\mathcal{U}_{\mathrm{F}_\delta(\mathrm{p})}} \bar{\xi}(\Lambda) \leq \bar{\xi}(\delta)$.

Lemma 6.1.1 implies that for any $p' \neq p$, any point $\mu \in \mathcal{M}_\infty$ with $\pi_{F_\delta(p)\cup\delta}(\mu) \in X_p$, and any $\sigma \in F_\delta(p')$ we have $\pi_\sigma(\mu) = \pi_{\delta\to\sigma}(p)$. Accordingly we now define the following subset of $\prod_{\sigma\in F_\delta\cup\delta} T_\sigma$,

$$\widehat{X_p} = \{p\} \times \textstyle\prod_{\sigma\in F_\delta\setminus F_\delta(p)} \pi_{\delta\to\sigma}(p) \times X_p$$

We thus obtain that $\widehat{\psi}_{F_\delta\cup\{\delta\}}(\mathcal{M}_\infty)$ is contained in the set:

$$Graph\left(f : T_\delta \setminus P \to \prod_{p\in P_\delta} X_p\right) \cup \bigcup_{p\in P} \widehat{X_p}, \tag{6.5}$$

where as above, the function f is a locally constant function. Since Λ is a collection of disjoint elements of Seq nested inside δ, by the properties of $\bar{\xi}$ enumerated in Lemma 6.1.2, we have $\bar{\xi}(\Lambda) \leq \bar{\xi}(\delta)$. Thus the above set is the union of two sets: one with dimension 1 (the graph over a tree) and the other closed with dimension at most $\bar{\xi}(\delta)$, the second dimension bound comes from Theorem A.2 (Dimension is subadditive) and the fact that dimension can not increase by taking finite unions of closed sets with uniformly bounded dimension (Theorem A.3). Thus the set described on line 6.5 has dimension at most $\bar{\xi}(\delta)$. This set is also closed, since the limit points of the graph are precisely the other set that we union it with.

Since Δ is maximal, any $\sigma \notin \Delta \cup_{\delta\in\Delta} F_\delta$ must overlap with Δ, thus we now see that $\mathcal{T}_F(\Delta)$ is given by

$$\prod_{\sigma\in F\setminus(\Delta\cup_{\delta\in\Delta}F_\delta)} (\pi_\sigma(\Delta)) \times \prod_{\delta\in\Delta} \widehat{\psi}_{F_\delta\cup\{\delta\}}(\mathcal{M}_\infty).$$

This tree-flat is homeomorphic to $\prod_{\delta\in\Delta} \widehat{\psi}_{F_\delta\cup\{\delta\}}(\mathcal{M}_\infty)$. By the above we know that for each $\delta \in \Delta$, we have $\dim(\widehat{\psi}_{\mathrm{F}_\delta\cup\{\delta\}}(\mathcal{M}_\infty)) \leq \bar{\xi}(\delta)$. Since Proposition 6.1.2 proved that $\sum_{\delta\in\Delta} \bar{\xi}(\delta) = \bar{\xi}(\Delta)$, using Theorem A.2, we have

$$\dim(\mathcal{T}_{\mathrm{F}}(\Delta)) \leq \sum_{\delta\in\Delta} (\dim(\widehat{\psi}_{\mathrm{F}_\delta\cup\{\delta\}}(\mathcal{M}_\infty))) \leq \sum_{\delta\in\Delta} \bar{\xi}(\delta) = \bar{\xi}(\Delta).$$

This proves the desired dimension bound for each tree-flat. Since $\widehat{\psi}_F(\mathcal{M}_\infty)$ is a finite union of tree-flats, each of which is closed and has dimension at most $\bar{\xi}(S) = 3g + p - 3$, the theorem now follows from Theorem A.3.

□

In the remainder of this section, we show that the dimension bound on $\widehat{\psi}_F(\mathcal{M}_\infty)$ from Theorem 6.3.2 is useful in the study of $\mathcal{M}_\infty$. We first remind the reader of a very general phenomenon.

Lemma 6.3.3. *Fix a continuous embedding $\phi : X \to (\prod_{\gamma\in\Gamma} T_\gamma, 0, l^1)$, where X is a compact and separable metric space, T_γ are metric spaces, and Γ is an index set of arbitrary infinite cardinality. For any $\epsilon > 0$ there exists $M \in \mathbb{N}$ so that $\phi_M : X \to \prod_{i=1}^{M(\epsilon)} T_i$ is an ϵ-mapping.*

Proof:

For two points $a, b \in \prod_{\gamma\in\Gamma} T_\gamma$ we use the notation $d_\gamma(a, b)$ to denote the distance between a and b in the T_γ coordinate. Since we are using the l^1 metric, for any pair of points $a, b \in \prod_{\gamma\in\Gamma} T_\gamma$ we have

$$d_{\prod_{\gamma\in\Gamma} T_\gamma}(a, b) = \sum_{\gamma\in\Gamma} d_\gamma(a, b).$$

Since $\phi^{-1} : \phi(X) \to X$ is a continuous function on a compact set, it is uniformly continuous. Thus, for each $\epsilon > 0$ there exists $\epsilon' > 0$, such that $d_{\prod_{\gamma\in\Gamma}}(a, b) < \epsilon'$ implies $d_X(\phi^{-1}(a), \phi^{-1}(b)) < \epsilon$

For each $q \in \phi(X)$ and any $\epsilon' > 0$, $\sum_{\gamma\in\Gamma} d_\gamma(0, q) < \infty$ implies that there exists a finite set $M_q \subset \Gamma$ for which $\sum_{\gamma\notin M_q} d_\gamma(0, q) < \frac{\epsilon'}{4}$. Letting $N(q)$ be the $\frac{\epsilon'}{4}$ neighborhood of q, we see that for each $q' \in N(q)$ one has $\sum_{\gamma\notin M_q} d_\gamma(0, q') < \frac{\epsilon'}{2}$. Compactness of X implies that there exists a finite collection of points $Q \subset \phi(X)$ so that $\cup_{q\in Q} N(q)$ cover $\phi(X)$. Taking $M = \cup_{q\in Q} M_q$ we have that each $x \in X$ has $\sum_{\gamma\notin M} d_\gamma(0, \phi(x)) < \frac{\epsilon'}{2}$ and moreover for each $x, y \in X$ we have $\sum_{\gamma\notin M} d_\gamma(\phi(y), \phi(x)) < \epsilon'$. Thus for each point $z \in \prod_{\gamma\in M} T_\gamma$ we have that the set of points in $\phi(X)$ which agree with z in their projection to $\prod_{\gamma\in M} T_\gamma$ has diameter less than ϵ'. Thus the uniform continuity mentioned above implies that any pair of points in this set have $d_X(x, y) < \epsilon$, and thus for each $z \in \prod_{\gamma\in M} T_\gamma$ we have $\mathrm{diam_X}(\phi_\mathrm{M}^{-1}(\mathrm{z})) < \epsilon$, i.e. ϕ_M is an ϵ-mapping.

□

An appeal to some standard dimension theory leads us to the following partial result towards Conjecture 6.3.6.

Theorem 6.3.4. *Let X be a compact separable subset of $\mathcal{M}_\infty(S)$. If $\widehat{\psi}$ restricted to X is bi-Lipschitz, then* $\mathrm{dim(X)} \leq \mathrm{3g + p - 3}$.

Proof:

Since X is compact and separable, Lemma 6.3.3 shows that for each $\epsilon > 0$ there exists a constant $N(\epsilon) \in \mathbb{N}$ for which $\widehat{\psi}$ induces an ϵ-mapping from X into $\prod_{i=1}^{N(\epsilon)} T_i$, note that this is the map $\widehat{\psi}_{N(\epsilon)}$. In Theorem 6.3.2 we proved that $\dim(\widehat{\psi}_{\mathrm{N}(\epsilon)}(\mathcal{M}_\infty)) \leq 3\mathrm{g} + \mathrm{p} - 3$.

Taking a sequence of positive real numbers $\epsilon \to 0$, we obtain a sequence of ϵ-mappings of X into spaces with dimension $\leq 3g + p - 3$. Proposition A.5 then implies that $\dim(\mathrm{X}) \leq 3\mathrm{g} + \mathrm{p} - 3$. $\square$

At this point it is worth mentioning a speculation on the importance of what has already been shown. We feel the following is a reasonable conjecture concerning quasiflats in the mapping class group.

Conjecture 6.3.5. *On each quasiflat in $\mathcal{M}(S_{g,p})$ of rank $\geq 3g+p-3$ the map $\widehat{\psi}$ restricts to a bi-Lipschitz embedding from the asymptotic cone over this set into $\prod_{\alpha \in \mathrm{Seq}} T_\alpha$.*

Recall that if we fix a quasi-isometric embedding of $\phi : \mathbb{R}^n \hookrightarrow \mathcal{M}$ this induces a bi-Lipschitz embedding $\widehat{\phi} : \mathbb{R}^n \hookrightarrow \mathcal{M}_\infty$. Also note that a bi-Lipschitz embedding is in particular a topological embedding. These observations and Theorem 6.3.4 prove that an affirmative resolution to Conjecture 6.3.5 would imply the following Rank Conjecture for the mapping class group.

Conjecture 6.3.6. **[BrF]** *$\mathcal{M}(S)$ admits quasi-isometric embeddings of $\mathbb{R}^n$ if and only if $n \leq 3g + p - 3$.*

The "if" direction of the conjecture has already been established, since Dehn twists along a maximal collection of simple closed curves yield a quasi-isometric embedding of $\mathbb{R}^{3g+p-3}$ into $\mathcal{M}$. This fact was first proven for punctured surfaces in [Mo], then in full generality in [FLM]. We remark that it also follows directly by applying Theorem 3.3.9 (the analogous argument to produce lower bounds on the geometric rank of the pants complex has been written out explicitly in [BrF]).

Also relevant is the Rank Conjecture for the Weil-Petersson metric on Teichmüller space:

Conjecture 6.3.7. *Teichmüller space with the Weil-Petersson metric admits quasi-isometric embeddings of $\mathbb{R}^n$ if and only if $n \leq g - 1 + \lfloor \frac{g+p}{2} \rfloor$.*

A proof of this could follow much the same argument as for the mapping class group. The main difference is that instead of using Theorem 3.3.9 to estimate distances in the marking complex by sums of projections in the curve complex one should use Theorem 3.4.5, its variant for the pants complex.

Theorem 3.4.5 differs from Theorem 3.3.9 in that instead of summing over all sufficiently large subsurface projection distances, we only sum over those large distances which occur in curve complexes of non-annular subsurface. This then produces an 'almost' quasi-isometric embedding into the product of all curve complexes of non-annular subsurfaces of S. Again this 'almost' quasi-isometry induces a Lipschitz map from the $\mathrm{Cone}_\infty(\mathcal{P})$ into a product of trees. The topological count of how many such subsurfaces of S can be simultaneously disjoint shows the rank of a maximal tree-flat is $g - 1 + \lfloor \frac{g+p}{2} \rfloor$ instead of the $3g - 3 + p$ one obtains for the mapping class group. Thereby providing an upper bound for the dimension of subsets of the asymptotic cone on which this map is bi-Lipschitz.

6.4 A Stratification of $\mathcal{M}_\infty$

Before explaining why $\widehat{\psi}$ fails to induce a bi-Lipschitz embedding from all of $\mathcal{M}_\infty$ into $\prod_{\alpha \in \mathrm{Seq}} T_\alpha$, we will first examine some subsets of $\mathcal{M}_\infty$ from which $\widehat{\psi}$ is a bi-Lipschitz embedding onto its image.

Example 6.4.1. Since $\mathcal{M}$ is equipped with a basepoint 0, for any element $\mu = \langle \mu_i \rangle \in \mathcal{M}_\infty$ and each $\alpha \in \mathrm{Seq}$ we can consider the growth of $d_{\alpha_i}(\mu_i, 0)$ as $i \to \infty$. Consider the subset $L_{C,C'} \subset \mathcal{M}_\infty$ consisting of those μ for which $|\{Y \subset S : d_Y(\mu_i, 0) > C\}| < C'$ for all $i \in \mathbb{N}$. The set $L_{C,C'}$ provides an example of a subset of $\mathcal{M}_\infty$ on which $\widehat{\psi}$ is bi-Lipschitz. This observation is a consequence of the fact that the asymptotic cone functor commutes with finite products (Proposition 2.2.7); alternatively a direct argument can readily be given by simply invoking the pigeonhole principle.

The next example generalizes the one above.

Example 6.4.2. Recall, $d_{\mathcal{M}_\infty}(\mu, \nu)$ is defined as $\omega\text{-}\lim_i \frac{1}{i}(\mu_i, \nu_i)$, which by Theorem 3.3.9 is estimated (up to a bounded multiplicative error, K, which depends on T) by $\sum_{\substack{Y \subseteq S \\ d_Y(\mu_i,\nu_i) > T}} d_Y(\mu_i, \nu_i))$, for a constant T chosen sufficiently large. In a similar vein to the above example, one can consider L_{linear} to be the maximal subset of $\mathcal{M}_\infty$ containing 0 for which any pair of points $\mu, \nu \in L_{\text{linear}}$ satisfy $d_{\mathcal{M}_\infty}(\mu, \nu) = K \sum_{\alpha \in \mathrm{Seq}} \omega\text{-}\lim_i \frac{1}{i} d_{\alpha_i}(\mu_i, \nu_i)$.

This clearly generalizes the previous example which is the subset of L_{linear} with only finitely many non-zero terms in the sum. It is tautological that this subset of the asymptotic cone bi-Lipschitz embeds into $\prod_{\alpha \in \mathrm{Seq}} T_\alpha$. More importantly though, as we discuss in more detail below, this example indicates a way in which the asymptotic cone may fail to embed into $\prod_{\alpha \in \mathrm{Seq}} T_\alpha$.

The second example suggests that $\widehat{\psi}$ is not a bi-Lipschitz embedding on all of $\mathcal{M}_\infty$ because there are sequences of mapping class elements whose distance grows linearly in the marking complex, but in the curve complex of each subsurface they differ by an amount that grows sublinearly. Indeed, such examples can be produced even in the case of the torus, moreover they can be made with the further restriction that the growth is not only sublinear everywhere, but actually asymptotically bounded. We now sketch a construction of one such example which for concreteness we give on the once punctured torus; generalizations of this to an arbitrary surface can readily be seen.

Example 6.4.3. Let $S = S_{1,1}$ and fix a marking $I \in \mathcal{M}$ with base curve γ_0 and transverse curve $t_0(0)$. As discussed in section 3.2 there are two elementary moves in $\mathcal{M}(S)$: twist and flip, denoted W and F, respectively.

The following makes sense for $n \in \{n : n = k^2 \text{ for some } k \in \mathbb{N}\}$. If we let ω be an ultrafilter for which $\{n : n = k^2 \text{ for some } k \in \mathbb{N}\} \in \omega$, then an argument which works for such integers is all we need to draw conclusions concerning $\mathrm{Cone}_\omega\mathcal{M}$. For ease of notation, we suppress this technicality from further mention.

For each $n \in \mathbb{N}$ consider the mapping class given by $(F \circ W^{\sqrt{n}})^{\sqrt{n}}$. We think of this map as being factored into $\sqrt{n}$ iterations of $F \circ W^{\sqrt{n}}$: where $W^{\sqrt{n}}$ fixes the previous base annuli γ_i and changes $t_i(j)$ to $t_i(j+1)$, whereas F changes γ_i to γ_{i+1} and renames the transversal to $t_{i+1}(0)$. For simplicity of notation we refer to the marking given by $(F \circ W^{\sqrt{n}})^{\sqrt{n}}(I)$ as T_n.

Since $\mathcal{M}(S_{1,1})$ is the Farey graph, once may verify that when $n > 2$ there is a hierarchy from I to T_n whose only large domains are the γ_i and the surface S. Moreover, $d_{\gamma_i}(t_i(0), t_i(\sqrt{n})) = \sqrt{n} - 1$ and $d_{\mathcal{C}(S)}(I, T_n) = \sqrt{n}$. Thus by Theorem 3.3.9 we have that up to some quasi-isometry constants which are independent of n and which we omit from this discussion: $d_{\mathcal{M}(S)}(I, T_n) = \sqrt{n} + \sqrt{n} \cdot (\sqrt{n} - 1)$. Then, taking the asymptotic cone of the mapping class group based at I, we have

$$d_{\mathcal{M}_\infty}(\langle T_n \rangle, \langle I \rangle) = \omega\text{-}\lim_n \frac{\sqrt{n} \cdot \sqrt{n}}{n} = 1.$$

But, $d_{\mathcal{C}(S)}(I, T_n) = \sqrt{n} = d_{\mathcal{C}(\gamma_i)}(I, T_n) + 1$ and since we these are the only large domains in the hierarchy from I to T_n, the distance between I and T_n is uniformly bounded in the curve complex of any other subsurface. From this it follows that for any $\alpha \in \mathrm{Seq}$ we have

$$\omega\text{-}\lim_i \frac{1}{i} d_{\mathcal{C}(\alpha_i)}(I, T_n) = 0.$$

Thus we have constructed an example of two points $\langle T_n \rangle, \langle I \rangle \in \mathcal{M}_\infty$ which are distance one in $\mathcal{M}_\infty$, but which are identified in $\prod_{\alpha \in \mathrm{Seq}} T_\alpha$.

The previous two examples indicate a dichotomy between points in the asymptotic cone corresponding to sequences of markings where the distance is completely given by linear growth in curve complex distance for certain sequences of subsurfaces versus sequences of markings with sublinear growth in every sequence of curve complexes. The following definition isolates the key aspect of this phenomenon.

Given any subsurface $Z \subsetneq S$, we can define a projection $\pi_{\mathcal{M}(Z)} : \mathcal{M}(S) \to 2^{\mathcal{M}(Z)}$, which sends elements of $\mathcal{M}(S)$ to uniformly bounded diameter subsets of $\mathcal{M}(Z)$. Given any $\mu \in \mathcal{M}(S)$ we can build a marking in the following way: pick an element $\gamma_1 \in \pi_Z(\mu)$ and let this be a base curve of the marking. Then letting $Z_2 = Z \setminus \gamma_1$, choose an element $\gamma_2 \in \pi_{Z_2}(\mu)$ as another base curve. Repeat this process until Z_{n+1} is a disjoint union of thrice punctured spheres (this happens when $n = 3g(Z) + p(Z) - 3$). Now for each γ_i define its transversal to be $\pi_{\gamma_i}(\mu)$. This process creates an element of $\mathcal{M}(Z)$, but arbitrary choices were made along the way. Define $\pi_{\mathcal{M}(Z)}(\mu)$ to be the union of all possible markings built following this process, we now show that this is a bounded diameter subset of $\mathcal{M}(Z)$. Fix $\nu, \nu' \in \pi_{\mathcal{M}(Z)}(\mu)$. For each $Y \subseteq S$ and each $1 \leq i, j \leq 3g + p - 3$, since $\gamma_i \in \text{base}(\nu)$ and $\gamma'_j \in \text{base}(\nu')$ are obtained by projecting μ to a subsurface, Lemma 3.1.4 (Lipschitz projection) implies that $d_S(\gamma_i, \gamma'_j) \leq 3$. Thus by Theorem 3.3.9 the $d_{\mathcal{M}(Z)}(\nu, \nu')$ is bounded and we have shown that $\pi_{\mathcal{M}(Z)}(\mu)$ is a bounded diameter subset of $\mathcal{M}(Z)$.

As in the case of the maps π_Z, we abbreviate $d_{\mathcal{M}(Z)}(\pi_{\mathcal{M}(Z)}(\mu_i), \pi_{\mathcal{M}(Z)}(\nu_i))$ by writing $d_{\mathcal{M}(Z)}(\mu_i, \nu_i)$.

Definition 6.4.4. Define a pair of points (μ, ν) in $\mathcal{M}_\infty(S) \times \mathcal{M}_\infty(S)$ to be *sublinear of the k^{th} order* if:

$$\omega\text{-}\lim_i \frac{\sup\{d_{\mathcal{M}(Z)}(\mu_i, \nu_i) : Z \subsetneq S \text{ with } \xi(Z) < \xi(S) - k\}}{i} = 0.$$

For our fixed basepoint $0 \in \mathcal{M}(S)$, let F_k denote the set of ν such that $(0, \nu)$ is sublinear of the k^{th} order. We use the term *strongly sublinear* to refer to the elements of F_0.

In this notation Example 6.4.3 shows that $|F_0| > 1$. Then since $\widehat{\psi}$ maps each point of F_0 to the basepoint of $\prod_{\alpha \in \text{Seq}} T_\alpha$ we have the following.

Proposition 6.4.5. *$\widehat{\psi} : \mathcal{M}_\infty \to \prod_{\alpha \in \text{Seq}} T_\alpha$ is not an embedding.*

Our goal is to use to use the notion of sublinearity to provide a stratification of $\mathcal{M}_\infty$ into understandable pieces.

In Chapter 5 we used projections from $\mathcal{M}(S)$ to hierarchy paths, Φ_H, and showed that these maps had a strong contraction property. In fact, these maps were sufficiently contraction so that for each hierarchy path, in the asymptotic cone we obtained a locally constant map from $\mathcal{M}_\infty$ to a bi-Lipschitz path. Then from this we concluded that in these cases $\mathcal{M}_\infty$ was an $\mathbb{R}$-tree.

We expect in a way analogous to the construction of Φ_H, we could show that for any surface S there exists a locally constant map $\Phi : \mathcal{M}_\infty \to F_0$. Following this, in a forthcoming paper we hope to show:

Conjecture 6.4.6. *F_0 satisfies the following properties:*

1. *F_0 is geodesically convex.*
2. *Given two points $\mu, \nu \in F_0$ they are connected by a unique path in $\mathcal{M}_\infty$. In particular, F_0 is an $\mathbb{R}$-tree.*

This conjecture would imply that any quasiflat in $\mathcal{M}_\infty$ intersects F_0 in at most a point, which would be a step towards producing the desired stratification of $\mathcal{M}_\infty$.

Appendix

A Dimension Theory

We summarize some standard definitions and results in dimension theory. (For further details consult [HW], [N], or [En].)

The following is widely considered to be the definition which most deserves to be called topological dimension; henceforth "dimension" will refer to *covering dimension.*

Definition A.1. Fix a topological space X. If for any open covering $\mathcal{U}$ of X there exists an open cover $\mathcal{V}$ refining $\mathcal{U}$ such that every collection $V_0, V_1, \dots V_k \in \mathcal{V}$ for which $\cap_{i=0}^{k} V_i \neq \emptyset$ must have $k \leq N$, then we say that X has *covering dimension* $\leq N$. We denote this $\mathrm{dim(X)} \leq \mathrm{N}$. When X has $\mathrm{dim(X)} \leq \mathrm{N}$ but not $\mathrm{dim(X)} \leq \mathrm{N} - 1$, then we say $\mathrm{dim(X)} = \mathrm{N}$.

Note in particular that $\mathrm{dim(X)} = \mathrm{N}$ does not imply that X has dimension N at every point.

Theorem A.2. *For topological spaces R and S,*

$$dim(R \times S) \leq dim(R) + dim(S)$$

Theorem A.3. *Fix a topological space X and two subsets K and L, one of which is closed. If $dim(K) \leq n$ and $dim(L) \leq n$, then $dim(K \cup L) \leq n$*

Definition A.4. A continuous mapping f of a space R into a space S is called an *ϵ-mapping* if the inverse image $f^{-1}(q)$ of each point q of S has diameter $< \epsilon$.

We recall the following characterization of dimension which has origins in Brouwer's 1911 proof of Invariance of Domain.[1]

Proposition A.5. *Let R be an n-dimensional compact metric space. Then there exists a positive number ϵ such that R cannot be mapped by an ϵ-mapping onto any metric space of dimension $\leq n - 1$.*

[1] I would like to thank M. Bestvina for suggesting to me the relevance of this proposition.

The main specialized result in dimension theory that we use is:

Theorem A.6. (Mayer and Oversteegen; [MaO]). *An $\mathbb{R}$-tree has dimension 1.*

Corollary A.7. *Let T_1 and T_2 be $\mathbb{R}$-trees. Then $dim(T_1 \times T_2) = 2$*

Proof: Theorems A.6 and A.2 imply that $dim(T_1 \times T_2) \leq 2$. Since $\mathbb{R}^2$ topologically embeds into $T_1 \times T_2$, we have $dim(T_1 \times T_2) \geq 2$. □

Bibliography

[BaS] D. Bachman and S. Schleimer. Surface bundles versus Heegaard splittings. 2002. arXiv:math.GT/0212104.

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

[BeH] M. Bestvina and M. Handel. Train tracks and automorphisms of free groups. *Ann. of Math.* **135** (1992), 1–51.

[BiLMc] J. S. Birman, A. Lubotzky, and J. McCarthy. Abelian and solvable subgroups of the mapping class groups. *Duke Math. J.* **50** (1983) no. 4, 1107–1120.

[Bo1] B. Bowditch. Intersection numbers and the hyperbolicity of the curve complex. *Preprint*.
Available at, `http://www.maths.soton.ac.uk/staff/Bowditch`.

[Bo2] B. Bowditch. Tight geodesics in the curve complex. Preprint.
Available at, `http://www.maths.soton.ac.uk/staff/Bowditch`.

[BriH] M. Bridson and A. Haefliger. *Metric Spaces of Non-positive Curvature*. Springer, 1999.

[Br] J. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. 2001. arXiv:math.GT/0109048.

[BrCM] J. Brock, R. Canary, and Y. Minsky. The classification of Kleinian surface groups II: the ending lamination conjecture. *In preparation.*

[BrF] J. Brock and B. Farb. Curvature and rank of Teichmüller space. 2002. arXiv:math.GT/0109045.

[Bu] J. Burillo. Dimension and fundamental groups of asymptotic cones. *Ph.D. Thesis*, University of Utah. June 1996.

[D] C. Druţu. Quasi-isometry invariants and asymptotic cones. *Int. J. of Algebra and Computation.* **12** (2002), no. 1 and 2, 99–135.

[En] R. Engelking. *Theory of Dimensions Finite and Infinite. Sigma Series in Pure Mathematics* vol. 10. Heldermann Verlag. (1995).

[EF] A. Eskin and B. Farb. Quasi-flats and rigidity in higher rank symmetric spaces. *Journal Amer. Math. Soc.* **10** (1997), 653–692.

[F] B. Farb. Revised preprint. 1998. *Geom. Funct. Analysis.* **8** (1998), 810–840.

[FLM] B. Farb, A. Lubotzky, and Y. Minsky. Rank one phenomena for mapping class groups. *Duke Math. J.* **106** (2001), 581–597.

[FMo] B. Farb and L. Mosher. A rigidity theorem for the solvable Baumslag-Solitar groups. *Invent. Math.* **131** (1998) no. 2, 419–451, Appendix by D. Cooper.

[FaLaP] A. Fathi, F. Laudenbach, V. Ponerau, et al. *Travaux de Thurston sur les surfaces.* Séminaire Orsay, Société Mathématique de France, Paris. Astérisque. **66–67** (1979).

[GhHar] E. Ghys and P. de la Harpe. *Sur les groupes hyperboliques d'aprés Mikhael Gromov.* Birkhäuser, 1990.

[G1] M. Gromov. Groups of polynomial growth and expanding maps. Inst. Hautes Études Sci. Publ. Math. **53** (1981), 53-73.

[G2] M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in Group Theory*, number 8 of *MSRI Publications*. Springer-Verlag, 1987.

[G3] M. Gromov. Asymptotic invariants of infinite groups. In G. Niblo and M. Roller, editors, *Geometric Group Theory*, volume 182 of *LMS Lecture Notes*. Cambridge Univ. Press, 1993.

[Ham] U. Hamenstädt. *Personal Communication.*

[Hare] J. L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.* **84** (1986), 157–176.

[Harv] W. J. Harvey. Boundary structure of the modular group. In I. Kra and B. Maskit, editors, *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference*, volume 97 of *Ann. of Math. Stud.* Princeton, 1981.

[Hat] A. Hatcher. Pants decompositions of surfaces. *Preprint. Available online at* `http://www.math.cornell.edu/~hatcher`.

[Hem] J. Hempel. 3-manifolds as viewed from the curve complex. *Topology.* **40** (2001) no. 3, 631–657.

[Hu] Z. Huang. Asymptotic Flatness of the Weil-Petersson Metric on Teichmüller space. *Preprint.*

[HW] W. Hurewicz and H. Wallman. *Dimension Theory. Princeton University Press*, 1948.

[ImTa] Y. Imayoshi and M. Taniguchi. *An Introduction to Teichmüller Spaces. Springer-Verlag*, 1992.

[Ir] E. Irmak. Superinjective Simplicial Maps of Complexes of Curves and Injective Homomorphisms of Subgroups of Mapping Class Groups. *Topology.* **43** (2004), 513–541.

[Iv] N. V. Ivanov. Mapping class groups. In R. Daverman and R. Sher, editor, *Handbook of Geometric Topology.* Elsevier, 2001, 523–633.

[Ka] M. Kapovich. Lectures on the Geometric Group Theory. University of Utah, Spring 2003.

[KaLe] M. Kapovich and B. Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. *Invent. Math.* **128** (1997) no. 2, 393–416.

[KlLe] B. Kleiner and B. Leeb. Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings. *Inst. Hautes Études Sci. Publ. Math.* **No. 86** (1997), 115–197.

[Mar] D. Margalit. The automorphism group of the pants complex. 2002. arXiv:math.GT/0201319. *Duke Math. J.* **121** (2004) no. 3, 457–479.

[MM1] H. Masur and Y. Minsky. Geometry of the complex of curves I: hyperbolicity. *Invent. Math.* **138** (1999), 103–149.

[MM2] H. Masur and Y. Minsky. Geometry of the complex of curves II: hierarchical structure. *Geom. & Funct. Anal.* **10** (2000), 902–974.

[MaO] J. Mayer and L. Oversteegen. A topological Characterization of ℝ-Trees. *Trans. Amer. Math. Soc.* **320** (1990) no. 1, 395–415.

[M] Y. Minsky. The classification of Kleinian surface groups I: models and bounds. 2002. arXiv:math.GT/0302208.

[Mo] L. Mosher. Mapping class groups are automatic. *Ann. of Math.* **142** (1995), 303–384.

[Mo2] L. Mosher. Homology and dynamics in quasi-isometric rigidity of once-punctured mapping class groups. 2003. arXiv:math.GR/0308065.

[N] J. Nagata. *Modern Dimension Theory. Interscience Publishers*, 1965.

[R] T. Riley. Higher connectedness of asymptotic cones. *Topology* **42** (2003), 1289–1352.

[Sc] R. Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.* **82** (1995) 133–168.

[Th1] W. P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)* **19** (1988) no. 2, 417–431.

[Th2] W. P. Thurston. The Geometry and Topology of 3-manifolds. Princeton University Press. 1997.

[Tho] S. Thomas. Cayley Graphs of Finitely Generated Groups. *Preprint. Available online at* `http://www.math.rutgers.edu/~sthomas`.

[VW] L. Van den Dries and A. Wilkie. On Gromov's theorem concerning groups of polynomial growth and elementary logic. *Journ. of Algebra* **89** (1984), 349–374.

[Wo] K. Wortman. Quasiflats with holes in reductive groups. 2004. arXiv:math.GT/0401360.

[Wol] K. Wortman. Geometry of the Weil-Petersson completion of Teichmueller space. In S. T. Yau, editor, *Surveys in Differential Geometry, VIII: Papers in Honor of Calabi, Lawson, Siu and Uhlenbeck*. International Press. 2003.

Index