Thurston's skinning map and curves on surfaces

BY

JONAH GASTER
B.S., University of Wisconsin at Madison, 2002
M.S., University of Illinois at Chicago, 2010

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Chicago, 2014

Chicago, Illinois

Defense committee:

David Dumas, Chair and Advisor
Marc Culler
Daniel Groves
Peter Shalen
Richard Kent, University of Wisconsin at Madison

To my mom, Laurie Sucher, for her boundless love and support. She wouldn't have been particularly interested in the material contents of this thesis, but she loved beauty, exploration, and truth. With some helpful interpretation, perhaps she would have found some remnant of those qualities in this work.

ACKNOWLEDGMENTS

There are many people who deserve thanks for making my time in graduate school alternately possible, enjoyable, and fruitful. To all of these people, I offer a sincere thank you.

Thanks to my lovely extended family: to my dad, Michael Gaster, for writing those arithmetic worksheets when I was five years old, and for generally being excited about mathematics; to Gabe, Toby, Barak, Noam, and Jeremiah, for their unending enthusiasm for argument and sharpening of ideas; to Cheryl, Zvika, Linda, Juan, Sophie, Nancy, Tova, Miriam, and Thea, for their love and encouragement; to Susan, Janet, Cat, Stacy, Diana, Roberta, and David, for good conversation. Special thanks to Elly, for her love, support, and understanding, for putting up with me in the past year, and for making me laugh on a regular basis.

Thanks to my colleagues and friends (in no particular order): Michael S., Matt, Ellie, Yen, Mike C., Filya, Phil, Brad, Hao, Natalie, Gabe, Deniz, Ali, Csar, Richard, Holly, Paul, Luigi, Chih-Chi, Xudong, and Lei, and the rest of the past and present MSCS graduate students at UIC who made possible countless learning experiences throughout Chicago. Your company has made graduate school less lonely, and sometimes even fun.

Thanks to Professors Alex Furman, Laura DeMarco, Peter Shalen, Marc Culler, Kevin Whyte, Izzet Coskun, and Daniel Groves, for teaching righteous mathematics at UIC that affected me tremendously; and to the entire MSCS staff at UIC, in particular to Maureen Madden, and the Director of Graduate Studies, Professor Ramin Takloo-Bighash, for being

constantly helpful and understanding. Thanks as well to my high school math teacher Doug O'Roark, whose wit and clarity were an early source of inspiration in my career.

For innumerable stimulating and challenging conversations, thanks to Brice Loustau, Andres Sambarino, Jeff Danciger, Michael Hull, Andy Sanders, Will Cavendish, Subhojoy Gupta, Sebastian Hensel, Simion Filip, Professors Gabriel Calsamiglia, Chris Leininger, Jason DeBlois, Richard Canary, Jeff Brock, Javier Aramayona, Martin Bridgeman, Richard Kent, Moira Chas, and Curt McMullen. Thanks as well to Tarik Aougab, for the fruitful discussions, and for the collaboration that forms the second part of this thesis.

Thanks are of course due to my committee, for agreeing to donate their time and energy to the evaluation of this dissertation, and for a lot of helpful advice over the years. Finally, many thanks are due to my advisor, David Dumas. I am grateful for his patience, his financial and moral support, his numerous careful suggestions and comments, and his constant time and energy for mathematics. Above all, I will continue to take inspiration from his uncanny ability to simplify and enrich an idea. This thesis would not have been possible without his help.

The financial support of the National Science Foundation, research grants DMS-0952869, 1107452, 1107263, 1107367 "RNMS: GEometric structures And Representation varieties" (the GEAR Network), is gratefully acknowledged.

TABLE OF CONTENTS

CHAPTER			PAGE
1	**INTRODUCTION**		1
	1.1	Thurston's geometrization and the skinning map	1
	1.2	Mapping class group orbits among maximal complete 1-systems on surfaces	4
2	**A FAMILY OF NON-INJECTIVE SKINNING MAPS WITH CRITICAL POINTS**		9
	2.1	Background and notation	9
	2.2	Skinning maps of finite covers of M	15
	2.3	The Example 3-manifold	17
	2.4	Deforming the example	23
	2.5	Rhombic 4-punctured spheres	27
	2.6	The convex core boundaries of Γ_t	34
	2.7	Non-injectivity of $t \mapsto Z_t$	42
	2.8	Computational lemmas	51
	2.9	A family of finite covers of M	57
	2.10	Appendix: Numerical lemmas	59
3	**MAPPING CLASS GROUP ORBITS OF MAXIMAL COMPLETE 1-SYSTEMS AND THEIR DUAL CUBE COMPLEXES**		60
	3.1	Background and motivation	60
	3.2	Sageev's construction	68
	3.3	Mapping class group orbits of collections of curves	77
	3.4	Recognizing n-cubes dual to a 1-system	78
	3.5	A family of maximal complete 1-systems	90
	3.6	Restricting mapping class group orbits via the dual cube complex	98
	3.7	Appendix: Finite-dimensionality of the cube complex dual to a system of curves	102
	APPENDICES		105
	Appendix A		106
	CITED LITERATURE		107
	VITA		111

LIST OF FIGURES

FIGURE		PAGE
1	Core curves of annuli in P	18
2	Generators for $\pi_1(H_2, x)$	18
3	Generators for $\pi_1(\Sigma, x)$	18
4	Another view of $\pi_1(\Sigma, x)$	19
5	The octahedron $\mathcal{O}$ with side identifications labelled by α and β.	22
6	Global coordinates on a $\{\xi, \eta\}$-rhombic $X \in \mathcal{T}(\Sigma)$.	29
7	The transversal δ	29
8	The quadrilateral Q_X, with vertical sides shaded.	33
9	Lifts for $\xi \in \pi_1\Sigma$	35
10	Maximal support planes for Γ_t, and the bending angle $\theta(t)$	37
11	A support plane for $\Gamma_{\frac{1}{2}}$ fails to be a support plane for $\hat{\Gamma}_{\frac{1}{2}}$	39
12	The quadrilateral $Q_t \subset Z_t$ and the curve ξ	44
13	A picture of the quasi-Fuchsian projective structure on Q_t. The horizontal sides are labelled h_1 and h_2.	46
14	The normalizing transformation $Q_t \to R_t$	47
15	The non-monotonicity of $Mod\,(Q_t)$. The quadrilateral R_t is pictured rotated so that its vertical sides appear horizontal.	50
16	The fundamental domain F for the covering $\Sigma' \to \Sigma$.	58
17	The cover M_4, the pared 3-manifold (H_4, P_4). Note that P_4 has two connected components.	58

LIST OF FIGURES (Continued)

FIGURE		PAGE
18	A vertex v of $\widetilde{G(\lambda)}$ determines a map from the 1-skeleton of the square to $\widetilde{G(\lambda)}^*$	64
19	Changing the realization λ by a Reidemeister type III move changes the isomorphism type of $\mathcal{S}(\lambda)$.	68
20	A point in $\alpha_j \cap \beta_0'$ corresponds to a point in $\alpha_j \cap \beta_0$.	83
21	'New' intersections of β_0' with α_j, violating the assumption that α_1 and α_j are in minimal position.	83
22	A 'localized realization' of β: $\beta \cap C$ is connected, and $p_i \in C$.	84
23	The triangle T_1 realizes (p_0, p_2, p_1), forming a localized realization of β.	86
24	The triangle T_1 realizes (p_0, p_1, p_2), where no localized realization is yet assured.	86
25	Given that the triangle T_1 realizes (p_0, p_1, p_2), the arcs $\beta \cap C$ may appear as pictured.	87
26	Given triangle T_1 realized as (p_0, p_1, p_2), the darkly shaded triangle T_2 cannot have a corner at p_0 as pictured.	87
27	The remaining case where triangle T_2 realizes (p_0, p_2, p_3).	88
28	Applying a homotopy to β using T_2, getting 'closer' to a localized realization.	88
29	The complex $\mathcal{C}(\{\alpha, \gamma\})$ is not a subcomplex of $\mathcal{C}(\{\alpha, \beta, \gamma\})$.	90
30	The generators of $\pi_1(S, x)$ in one handle.	91
31	The curves $\alpha_1, \alpha_2 \in A$	92
32	The curves $\beta_1, \beta_2 \in B$	92
33	The curves $\alpha_3, \alpha_4 \in A$ and $\beta_5, \beta_6 \in B$	93
34	The curve δ_0	93

LIST OF FIGURES (Continued)

FIGURE **PAGE**

35 The system of curves $\Gamma(1,1,1,1,1)$ 94

36 The system of curves $\Gamma(1,-1,1,1,-1)$ 95

37 The system of curves $\Gamma(1,1,1,1,1,1,1)$ 96

38 The system of curves $\Gamma(1,-1,1,-1,-1,1,-1)$ 97

39 The inductive step for Lemma 3.7.1. 103

SUMMARY

The 'deformation space' of a given geometric structure on a fixed smooth manifold is a major theme in low-dimensional geometry. In this thesis we present work on two problems that fit into such a framework for 2 and 3-dimensional hyperbolic manifolds. The first concerns the behavior of the *skinning map* associated to the family of infinite-volume hyperbolic structures on the interior of a 3-manifold with boundary, and the second concerns counting *maximal complete 1-systems* on a closed surface.

Thurston's revolutionary Geometrization Program left many paths for exploration, including both quantitative and qualitative questions about the tools he created. In Thurston's approach to finding hyperbolic structures, when a 3-manifold is irreducible, atoroidal, and Haken—when it contains an embedded π_1-injective surface but no such tori—in many cases the 'cutting' along the surface allows an inductive procedure that locates the hyperbolic structure. In Chapter 2, we show that an important tool in this approach, the skinning map, may fail to be an immersion.

The skinning map we examine is associated to a 3-manifold M homeomorphic to a genus-2 handlebody, and to a family of hyperbolic structures with two rank-1 cusps. We exploit an orientation-reversing isometry of M to conclude that the skinning map of M preserves a one-dimensional submanifold of the deformation space, and use estimates on extremal length functions to show the existence of a critical point in this submanifold. A family of finite covers of M produces examples of non-immersion skinning maps on the Teichmüller spaces of surfaces in each even genus, and with either four or six punctures.

SUMMARY (Continued)

In analogy with the examination of an immersed surface in a 3-manifold, the study of immersions of 1-manifolds on surfaces (i.e. closed curves) has borne considerable fruit. One avenue of study in this vein is concerned with 'maximal topological designs' of a collection of curves, where the notion of maximality depends on context. In Chapter 3 we advance this story by presenting the first bounds for the number of mapping class group orbits of maximal complete 1-systems. In the process, we develop tools making Sageev's general approach of dual cube complexes more accessible in our setting. While our linear lower bounds only hold in odd genus, we hope to improve them to exponential bounds in any genus in the near future.

The results of both Chapters 2 and 3 leave several directions for further exploration. From Chapter 2, the main question left unanswered concerns the nature and exact location of the critical point of the skinning map we studied. An answer to this question might illuminate some criterion for existence of critical points. From Chapter 3, many questions remain about the set of maximally complete 1-systems.

Chapter 2 has been accepted for publication by the *Transactions of the AMS*, published by the American Mathematical Society. The work presented in Chapter 3 was developed jointly with Tarik Aougab, whose permission for the inclusion of the work is gratefully acknowledged (see p. 106).

CHAPTER 1

INTRODUCTION

1.1 Thurston's geometrization and the skinning map

Thurston introduced the 'skinning map' as a tool for locating hyperbolic structures on certain closed 3-manifolds, an integral part of his celebrated proof of the Hyperbolization Theorem for Haken manifolds (see (Thurston, 1986), (Otal, 1998), (Kapovich, 2001)). Apart from a class of simple cases, explicit examples of skinning maps remain unexplored. One reason for this is that the definition of the skinning map uses the deformation theory of Ahlfors and Bers to pass back and forth from a conformal structure on the boundary of a 3-manifold to a hyperbolic structure on its interior. Like the uniformization theorem of Riemann surfaces, an explicit formula for the resulting map is typically out of reach.

Let M be an orientable compact 3-manifold with nonempty incompressible boundary $\Sigma := \partial M$, such that the interior of M admits a geometrically finite hyperbolic structure. Given a point $X \in \mathcal{T}(\Sigma)$ in the Teichmüller space of Σ, the 3-manifold M has a quasi-Fuchsian cover, corresponding to $\pi_1\Sigma \hookrightarrow \pi_1 M$, with conformal structures X and $\overline{Y} \in \mathcal{T}(\overline{\Sigma})$ at infinity, where $\bar{\cdot}$ indicates a reversal of orientation. The skinning map of M, σ_M, is given by $\sigma_M(X) = \overline{Y}$. Thurston's key result about skinning maps is the Bounded Image Theorem: when M is acylindrical (when there are no non-trivial conjugacies in $\pi_1 M$ between boundary parallel curves), then $\sigma_M(\mathcal{T}(\Sigma))$ is contained in a compact set.

Thurston described how to locate hyperbolic structures on a class of closed Haken 3-manifolds by iteration of $\tau \circ \sigma_M$, where τ is a certain isometry from $\mathcal{T}(\overline{\Sigma})$ to $\mathcal{T}(\Sigma)$. There are further questions, and results, about how 'effectively' iteration of $\tau \circ \sigma_M$ locates fixed points. Results in this direction include: The diameter of the image of σ_M is controlled by constants depending only on the volume of M (Kent, 2010). If M is without rank-1 cusps then σ_M is non-constant (Dumas and Kent, 2009). This result has an improvement, which shows that σ_M is open and finite-to-one (Dumas, 2012).

When M is an interval bundle over a surface, the map σ_M is a global diffeomorphism, easily described as a switch of coordinates using the Ahlfors-Bers parametrization. However, no similarly explicit descriptions of σ_M are known in all other cases. Thus it is consistent with the current literature to ask:

Question 1. *Are skinning maps always diffeomorphisms onto their images? Are they always immersions?*

In this paper, we present a negative answer to both questions. The reader should note that we work in the category of pared 3-manifolds and hyperbolic structures with rank-1 cusps. This allows low-dimensional calculations, in which σ_M is 'simply' a holomorphic map $\sigma_M : \mathbb{H} \to \overline{\mathbb{H}}$.

Theorem A. *There exists a non-injective skinning map with a critical point.*

Briefly, the proof of Theorem A (see Theorem 2.7.6) examines a certain pared 3-manifold M, using an orientation-reversing isometry of M to conclude that the skinning map sends a certain real 1-dimensional submanifold of $\mathcal{T}(\Sigma)$ to itself. Analysis of this submanifold together

with extremal length estimates show that this restricted map is not one-to-one, and hence the skinning map has a critical point. That is, this proof produces a (real) 1-parameter path of quasi-Fuchsian groups $\{Q(X_t, \sigma_M(X_t))\}$, where X_t and $\sigma_M(X_t)$ are both contained in a line in $\mathcal{T}(\Sigma)$, such that $t \mapsto X_t$ is injective and $t \mapsto \sigma_M(X_t)$ is not.

By passing to finite covers of the manifold from Theorem A (see Corollary 2.9.1) this example implies a family of examples of skinning maps with critical points, on Teichmüller spaces of arbitrarily high dimension.

This work owes a debt of inspiration to numerical experiments developed by Dumas-Kent which suggested the presence of a critical point of the skinning map associated to the manifold in Theorem 2.7.6. Their work also examines some other closely related manifolds and skinning maps, at least one of which appears to be a diffeomorphism onto its image. It would be interesting to have a unified understanding of these different behaviors.

The reader may note that the proof of Theorem A is only one of existence. It does not identify a critical point, determine the number of critical points, or determine any local degrees of σ_M near critical points. However, for this skinning map there is some evidence of a unique simple critical point, both in the experiments of Dumas-Kent and in numerical tests conducted by the author. The latter tests further allow analysis of the critical point itself. Curiously, it appears that the critical point occurs where M has hexagonal convex core boundary (that is, as a point in moduli space the convex core boundary is isometric to the Poincaré metric of $\mathbb{C} \setminus \{0, 1, e^{\pi i/3}\}$). This would seem to imply some geometrical significance to the critical point,

which for the moment remains elusive. Based on this numerical evidence, we frame this as a conjecture:

Conjecture. *For M as in Theorem 2.7.6, there is a unique simple critical point of σ_M that occurs when the convex core boundary of M is isometric to $\mathbb{C} \setminus \{0, 1, e^{\pi i/3}\}$.*

1.2 Mapping class group orbits among maximal complete 1-systems on surfaces

The results described in this section, detailed in Chapter 3, represent a collaboration with Tarik Aougab (see p. 106).

A significant amount of research in Teichmüller space has been motivated by, and has grown out of, a detailed understanding of the structure of the space of simple closed curves on a surface of finite type. One such direction of exploration has focused on combinatorial design questions. A notable example is a question of Farb-Leininger:

Question 2. *What is the maximum number of distinct homotopy classes of curves on a closed surface that can be made to mutually intersect at most once?*

The question above has resisted a complete solution, though (Malestein et al., 2014) have presented lower and upper bounds (quadratic and exponential, respectively, in genus) [1] for the answer. In addition, they answered a simpler version of Question 2 precisely. Let a *1-system* refer to a collection of homotopy classes of curves that pairwise intersect at most once, and call a 1-system *complete* if every pair of curves intersects exactly once.

---

[1] (Przytycki, 2014) has recently improved these upper bounds to cubic in genus.

Theorem 1.2.1. *(Malestein et al., 2014, Thm. 1.4) On a compact surface Σ of genus g, a maximal complete 1-system has size $2g+1$.*

This size is achieved for a simple example: There are $2g+1$ diagonals of a $(4g+2)$-gon which, when opposite sides of the polygon are glued, form a maximal complete 1-system on a surface of genus g. Choosing any orientations for the curves, the algebraic intersection number of the oriented curves must be ± 1, so it is clear that a maximal complete 1-system injects into $H_1(\Sigma, \mathbb{Z}/2\mathbb{Z})$. Maximality here is shown by examining the image in $H_1(\Sigma, \mathbb{Z}/2\mathbb{Z})$ using the algebraic intersection pairing that makes $H_1(\Sigma, \mathbb{Z}/2\mathbb{Z})$ into a symplectic vector space.

Let $N(g)$ indicate the number of extended mapping class group orbits of maximal complete 1-systems on a compact surface of genus g. In addition to indicating that $N(2) = 1$, (Malestein et al., 2014, p. 11) posed the question of whether $N(g) = 1$ for all g. This has been resolved negatively (Aougab, 2012): $N(g) \geq 2$ for each $g \geq 3$. In joint work with Aougab (see Theorem 3.5.4, Chapter 3) we show:

Theorem B. *We have the upper bound $(4g^2+2g)! \geq N(g)$. For g odd, we have the lower bound $N(g) \geq g/2$.*

Towards the lower bound, we exhibit a construction of exponentially many maximal complete 1-systems. Though we expect that these lie in exponentially many distinct extended mapping class group orbits (see Conjecture 3.5), we use the dual cube complex construction of (Sageev, 1995) to distinguish linearly many orbits in this collection. Through our analysis, we are lead to an examination of properties of the dual cube complex, which we now highlight:

Theorem C. *The extended mapping class group orbit of a filling system of curves is characterized by the isomorphism type of its dual cube complex.*

This could be compared to a familiar statement for '0-systems': When Γ_1 and Γ_2 are collections of homotopy classes of pairwise disjoint curves, then they differ by a mapping class of the surface if and only if the dual graph to the systems are isomorphic. When the curves of Γ_i are not disjoint, the dual graph is not a well-defined invariant of the collection of homotopy classes. However, Theorem 3.3.1 can be seen as an appropriate generalization of this idea.

Sageev's construction has the benefit of being quite general, but the same generality may make it difficult to work with in specific examples. (See §3.1 and §3.2 for a presentation of Sageev's construction in the setting of curves on surfaces). In our setting of 1-systems of curves on a surface, we develop a simple, accessible method for detecting the presence of n-cubes in the dual cube complex:

Theorem D. *A set of n curves in a 1-system forms an n-cube in the dual cube complex if and only if every trio of curves in minimal position has a triangle in its complement.*

Note that 'having a triangle in its complement when in minimal position' is an invariant of a trio of homotopy classes of curves (see Lemma 3.4.1). The assumption that the curves form a 1-system is crucial, and the conclusion is false without it.

Among various natural questions in this domain, we highlight one that we find particularly curious. Recall that the *dimension* of a cube complex is the supremum of the dimensions of the cubes in the complex. By Theorem D, the dual cube complex of the maximal complete 1-system formed from a gluing of a $(4g+2)$-gon is $(2g+1)$-dimensional. This is the maximal

possible dimension of the dual cube complex to a 1-system, so it is tempting to speculate about the corresponding minimal dimension:

Question 3. *Does there exist a maximal complete 1-system with a 2-dimensional dual cube complex? Does the minimum dimension of the dual cube complexes of maximal complete 1-systems grow with genus?*

This question fits into a broader framework:

Question 4. *What are restrictions on the dual cube complex of a maximal complete* 1*-system?*

A slightly different perspective on maximal complete 1-systems is provided by a *stabilization* process: Given an arc on a surface that intersects each of the curves of a maximal complete 1-system once, one can sew in an annulus at the endpoints of this arc, and thread the arc through the annulus to form a new closed curve. Together with the curve twisted around the core of the annulus, one obtains a new maximal complete 1-system.

Question 5. *Is every maximal complete* 1*-system obtainable by a sequence of stabilizations of the maximal complete* 1*-system on a surface of genus 2?*

A positive answer would imply a negative answer to the first part of Question 3.

An entirely different method of inquiry, which is natural in the context of Theorem B, concerns pursuing analogies of maximal complete 1-systems with more foundational Teichmüller-theoretic tools. The *curve graph* $\mathcal{C}(S)$ of S is a locally infinite, infinite diameter graph whose vertices are in correspondence with homotopy classes of simple closed curves, and whose edges

correspond to pairs of curves that can be realized disjointly. The curve graph has been an essential object of study in Teichmüller theory, mapping class groups, and hyperbolic 3-manifolds. Relatedly, the *pants graph* $\mathcal{P}(S)$ can be formed by forming vertices for maximal complete subgraphs of $\mathcal{C}(S)$, or *pants decompositions*, with edges that correspond to maximal complete subgraphs that share all but one of their vertices, and so that those differing vertices have either one or two intersections.

The *systole graph*, originally considered by Schaller, has vertices corresponding to non-separating curves, with edges that correspond to pairs of curves that have exactly one intersection. It is not hard to see that the identity map on vertices, from the systole graph to $\mathcal{C}(S)$, is bi-Lipschitz, from which it follows that much of the rich complexity of $\mathcal{C}(S)$ carries over to the systole graph. In analogy with the formation of $\mathcal{P}(S)$ from $\mathcal{C}(S)$, maximal complete 1-systems determine maximal complete subgraphs of the systole graph.

It is tempting to seek a completion of this analogy, but there are significant obstacles. For instance, it is not true that maximal complete 1-systems are precisely the maximal complete subgraphs of the systole graph: There exists a triangle in the systole graph of the genus 2 surface that is not contained in a complete graph on 5 vertices.

Question 6. *What is the 'right' analogy of* $\mathcal{P}(S)$ *in the setting of Schaller's systole graph?*

CHAPTER 2

A FAMILY OF NON-INJECTIVE SKINNING MAPS WITH CRITICAL POINTS

In §2.1, we introduce background and notation. The proof that critical points of skinning maps persist under certain finite covers is found in §2.2. In §2.3 and §2.4 we introduce the 3-manifold relevant to Theorem A, and a path of geometrically finite structures on it, noting that this path maintains some important symmetry. In §2.5 we study the implications of this symmetry on the geometry of a four-punctured sphere. We apply some of these implications in §2.6 to the convex core boundaries of our path. This information is used in §2.7 which collects the main ideas for the proof of Theorem 2.7.6, deferring some computations to §2.8 and Appendix 2.10. Finally, §2.9 presents a family of finite covers of the example from Theorem 2.7.6, which proves Corollary 2.9.1.

2.1 Background and notation

Let $\Sigma = \Sigma_{g,n}$ be a smooth surface of genus g with n punctures.

Let $\mathcal{S}$ denote the set of non-peripheral unoriented simple closed curves on Σ and $\mathcal{ML}(\Sigma)$ the space of measured geodesic laminations. Recall the natural inclusion $\mathbb{R}_+ \times \mathcal{S} \hookrightarrow \mathcal{ML}(\Sigma)$, which has dense image. The quotient of $\mathcal{ML}(\Sigma)$ under the action of multiplication by positive weights will be written $\mathbb{PML}(\Sigma)$. See (Casson and Bleiler, 1988) and (Fathi et al., 1979) for details.

Recall that the *Teichmüller space* of Σ, denoted $\mathcal{T}(\Sigma)$, is the set of marked complex structures (equivalently, marked hyperbolic structures, via uniformization) on Σ up to marking equivalence. Recall that $\mathcal{T}(\Sigma)$ is naturally a complex manifold homeomorphic to $\mathbb{C}^{6g-6+2n}$.

For a smooth manifold N, possibly with boundary, define $MCG^*(N) := {}^{\mathrm{Diff}(N)}/_{\mathrm{Diff}_0(N)}$, the *extended mapping class group* of N. Pre-composition of a marking with a diffeomorphism of the surface Σ descends to an action of $MCG^*(\Sigma)$ on $\mathcal{T}(\Sigma)$. Note that we allow orientation-reversing diffeomorphisms. When Σ is the boundary of a 3-manifold M, there is a restriction map $r : MCG^*(M) \to MCG^*(\Sigma)$. See (Lehto, 1987), (Farb and Margalit, 2012), and (Fathi et al., 1979) for details about $\mathcal{T}(\Sigma)$ and $MCG^*(\Sigma)$.

A *Kleinian group* Γ is a discrete subgroup of $\mathrm{PSL}(2,\mathbb{C})$. The action of Γ on $\mathbb{CP}^1$ has a domain of discontinuity Ω_Γ, and its complement is the *limit set*, Λ_Γ. One component of the complement of a totally geodesic plane in $\mathbb{H}^3$ is a *supporting half-space* for a connected component of Ω_Γ—or, less specifically, a supporting half-space for Γ—if it meets the conformal boundary $\mathbb{CP}^1$ in that component, and the closure of the intersection with $\mathbb{CP}^1$ intersects Λ_Γ in at least two points. The totally geodesic boundary of a supporting half-plane for Γ is a *support plane* for Γ. The *convex hull* of Γ is the complement of the union of all of its supporting half-spaces and the *convex core* of Γ is the quotient of the convex hull by Γ. When an ϵ-neighborhood of the convex core is finite volume, for some $\epsilon > 0$, we say that Γ is *geometrically finite*. See (Maskit, 1988) and (Canary et al., 2006) for details.

When N is a topological space with finitely-generated fundamental group, $\mathrm{Hom}(\pi_1 N, \mathrm{PSL}_2\,\mathbb{C})$ naturally has the structure of an algebraic variety. There is a conjugation action of $\mathrm{PSL}_2\,\mathbb{C}$ on

$\mathrm{Hom}(\pi_1 N, \mathrm{PSL}_2\,\mathbb{C})$, and the quotient, interpreted in the sense of geometric invariant theory, is the $\mathrm{PSL}_2\,\mathbb{C}$-*character variety of* N. We will denote the $\mathrm{PSL}_2\,\mathbb{C}$-character variety of N by $\mathcal{X}(N)$. Recall that there is a natural holomorphic structure that $\mathcal{X}(N)$ inherits from $\mathrm{PSL}_2\,\mathbb{C}$. See (Goldman, 1984) and (Goldman, 2009) for details.

A Kleinian group Γ is *Fuchsian* if it is conjugate, in $\mathrm{PSL}_2\,\mathbb{C}$, into $\mathrm{PSL}_2\,\mathbb{R}$. In this case, Γ stabilizes a totally geodesic plane in $\mathbb{H}^3$, and Λ_Γ is contained in a round circle on $\mathbb{CP}^1$. If Λ_Γ is a quasi-circle and Γ stabilizes each of the components of Ω_Γ, then Γ is *quasi-Fuchsian*. In this case, the hyperbolic manifold $\mathbb{H}^3/\Gamma$ has a conformal compactification $^{\mathbb{H}^3 \sqcup \Omega_\Gamma}/_\Gamma$, with two conformal structures on finite-type Riemann surfaces 'at infinity'. We denote by $\mathcal{QF}(\Sigma)$ the locus of $\mathcal{X}(\Sigma)$ consisting of faithful representations with quasi-Fuchsian image. Recall Bers' Theorem (Bers, 1960),

Theorem 2.1.1 (Simultaneous Uniformization)**.** *There exists an identification* $\mathrm{AB}_\Sigma : \mathcal{QF}(\Sigma) \cong \mathcal{T}(\Sigma) \times \mathcal{T}(\overline{\Sigma})$*, a biholomorphism of complex manifolds.*

When ρ is faithful, we identify $[\rho] \in \mathcal{X}(N)$ with the conjugacy class of its image $\rho(\pi_1 N)$ in $\mathrm{PSL}_2\,\mathbb{C}$. We write Q for the map AB_Σ^{-1}, so that, for $X, Y \in \mathcal{T}(\Sigma)$, we have $Q(X, \overline{Y}) \in \mathcal{QF}(\Sigma)$.

A *pared 3-manifold* is a pair, (H_0, P), where H_0 is a compact, atoroidal, irreducible 3-manifold with boundary and P is a disjoint union of incompressible annuli and tori in ∂H_0, such that:

- P contains all of the tori in ∂H_0

· Any embedded cylinder with boundary components in P is homotopic, relative to the boundary, into P

By the 'boundary' of a pared 3-manifold, $\partial(H_0, P)$, we mean $\partial H_0 \setminus P$. See (McMullen, 1990, p.434) or (Canary and McCullough, 2004, §5) for details. The pared 3-manifolds we consider will satisfy a slightly stronger condition, so that the skinning map will take values in $\mathcal{T}(\Sigma)$ (McMullen, 1990, p. 435):

($\star$) Any embedded cylinder with one boundary component in P is homotopic, relative to the boundary, into P

This latter condition is motivated by consideration of accidental parabolics, which we discuss briefly below. An *accidental parabolic* of $\mathbb{H}^3/\Gamma$ is a curve homotopic into a cusp of $\mathbb{H}^3/\Gamma$ but which descends to a non-boundary parallel curve on a component of the boundary Ω_Γ/Γ.

When $M = (H_0, P)$ is a pared 3-manifold, we denote by $\mathcal{X}(M) \subset \mathcal{X}(H_0)$ the locus of points for which the image of curves homotopic into P are parabolic. The subset $\mathcal{GF}(M) \subset \mathcal{X}(M)$ consists of $[\rho] \in \mathcal{X}(M)$ such that $\rho(\pi_1 H_0)$ is a geometrically finite Kleinian group, the image $\rho(\gamma)$ is parabolic if and only if γ is homotopic into P, and such that ρ is induced by a homeomorphism $\mathbb{H}^3/\rho(\pi_1 H_0) \cong H_0^\circ$. It is a theorem of Thurston that a pared 3-manifold M is *hyperbolizable*, i.e. $\mathcal{GF}(M) \neq \emptyset$, so that M carries a complete geometrically finite hyperbolic structure (see (Canary and McCullough, 2004, p.106)). Let $\mathrm{Diff}(H_0, P)$ be the diffeomorphisms of H_0 that preserve the set P and $MCG^*(M) := \mathrm{Diff}(H_0,P)/\mathrm{Diff}_0(H_0,P)$. Note that $MCG^*(M)$ acts on $\mathcal{X}(M)$ preserving $\mathcal{GF}(M)$.

Since $\mathcal{QF}(\Sigma)$ identifies with $\mathcal{GF}(\Sigma \times \mathbb{R})$, it is natural to seek a generalization of Bers' Uniformization Theorem to 3-manifolds other than $\Sigma \times \mathbb{R}$. Such a theorem is provided by the deformation theory developed by Ahlfors, Bers, Marden, Sullivan, and others. We refer to this identification as the 'Ahlfors-Bers' parameterization:

Theorem 2.1.2 (Ahlfors-Bers Parameterization)**.** *For a pared 3-manifold $M = (H_0, P)$ with incompressible boundary, there exists an identification $\mathrm{AB}_M : \mathcal{GF}(M) \cong \mathcal{T}(\partial M)$, a biholomorphism of complex manifolds.*

Note that ∂M may be disconnected, in which case $\mathcal{T}(\partial M)$ is the product of the Teichmüller spaces of the components.

For the remainder of the section, let $M = (H_0, P)$ be a hyperbolizable pared 3-manifold, satisfying $(\star)$, such that $\Sigma = \partial M$ is nonempty, incompressible, and connected. Let $[\hat{\Gamma}] \in \mathcal{GF}(M)$. Fix a choice of basepoints for M and Σ, and a path connecting them, so that inclusion induces the well-defined homomorphism $\iota_* : \pi_1\Sigma \hookrightarrow \pi_1 M$ and the restriction morphism ι^* on representations. Let $\Gamma = \iota^*\hat{\Gamma}$.

Lemma 2.1.3. *Under the assumptions above, we have $[\Gamma] \in \mathcal{QF}(\Sigma)$.*

Proof. The choices above fix an identification of universal covers $\widetilde{M} \cong \mathbb{H}^3$ and $\widetilde{\Sigma} \cong U_0$, for some component $U_0 \subset \Omega_{\hat{\Gamma}}$. In this case, $\Gamma = \mathrm{Stab}_{\hat{\Gamma}}(U_0)$ and (Marden, 1974, Cor. 6.5.) shows that Γ must be geometrically finite. Because Σ is incompressible in M, the domain U_0 must be simply-connected.

The condition ($\star$) implies that $\rho(\pi_1\Sigma)$ has no accidental parabolics: Suppose a curve γ on Σ has parabolic image, but is *not* homotopic into a cusp in the boundary. Because γ is homotopic into P, the Annulus Theorem (**?**) would yield a cylinder between a non-pared curve and a pared curve in ∂H_0, violating ($\star$). □

Note that the assumption that ∂M is nonempty rules out a fibered hyperbolic 3-manifold, where the Kleinian group corresponding to the fiber is a geometrically infinite surface subgroup of a geometrically finite Kleinian group, and the conclusion of Lemma 2.1.3 fails.

Definition. For M as above and $[\hat{\Gamma}] \in \mathcal{GF}(M)$ with $\Gamma = \iota^*\hat{\Gamma}$, the domains Ω_Γ and $\Omega_{\hat{\Gamma}}$ share a connected component, which we refer to as the *top* of Γ. By Lemma 2.1.3, the domain Ω_Γ has one other component, which we refer to as the *bottom*. When clear from context, we may refer to the quotient of the top component of Ω_Γ as the top of Γ.

Lemma 2.1.3 ensures the skinning map will take values in $\mathcal{T}(\overline{\Sigma})$. Below, we assume the connectedness of Σ (for Σ disconnected, see (Otal, 1998)).

Definition. The skinning map $\sigma_M : \mathcal{T}(\Sigma) \to \mathcal{T}(\overline{\Sigma})$ fits into the following commutative diagram:

$$\begin{array}{ccccccc} \mathcal{T}(\Sigma) & \xrightarrow{\mathrm{AB}_M} & \mathcal{GF}(M) & \xrightarrow{\iota^*} & \mathcal{QF}(\Sigma) & \xrightarrow{\mathrm{AB}_\Sigma} & \mathcal{T}(\Sigma) \times \mathcal{T}(\overline{\Sigma}) \\ & \searrow^{\sigma_M} & & & & & \downarrow proj_2 \\ & & & & & & \mathcal{T}(\overline{\Sigma}) \end{array}$$

In words, $\iota^* \circ \mathrm{AB}_M : \mathcal{T}(\Sigma) \to \mathcal{QF}(\Sigma)$ associates to $X \in \mathcal{T}(\Sigma)$ a quasi-Fuchsian structure $Q(X, \overline{Y}) \in \mathcal{QF}(\Sigma)$, where the top of $Q(X, \overline{Y})$ is X, and the bottom is $\overline{Y}$. The skinning map is given by $\sigma_M(X) = \overline{Y}$.

Because each of the maps above is holomorphic, σ_M is holomorphic. As a result of the naturality of the Ahlfors-Bers identifications σ_M is equivariant for the actions of $MCG^*(M)$, $MCG^*(\Sigma)$, and the restriction map $r : MCG^*(M) \to MCG^*(\Sigma)$. As a consequence,

Proposition 2.1.4. *For each $\phi \in MCG^*(M)$, the fixed point set* $\operatorname{Fix} r(\phi)$ *is preserved by the skinning map, i.e. $\sigma_M(\operatorname{Fix} r(\phi)) \subset \operatorname{Fix} r(\phi)$.*

The two copies of Fix $r(\phi)$ in Proposition 2.1.4 technically lie in Teichmüller spaces of surfaces with opposite orientations. Which one is intended will be clear from context. Because there is a canonical anti-holomorphic isometry $\bar{\cdot} : \mathcal{T}(\overline{\Sigma}) \to \mathcal{T}(\Sigma)$, this distinction is not crucial.

The reader should note that the existence of any nontrivial mapping class of the 3-manifold M is not assured. In the 3-manifold from Theorem 2.7.6 this existence allows a reduction of dimension, an essential tool in our analysis.

2.2 Skinning maps of finite covers of M

Let $p : (M', \Sigma') \to (M, \Sigma)$ be a finite covering of manifolds with boundary. We will denote the restriction $p|_{\Sigma'}$ by p as well.

Proposition 2.2.1. *Suppose that σ_M has a critical point at $X \in \mathcal{T}(\Sigma)$. Then $\sigma_{M'}$ has a critical point at $p^*X \in \mathcal{T}(\Sigma')$.*

Proof. Denote the image $\mathcal{E}_M := \iota^* \mathcal{GF}(M) \subset \mathcal{QF}(\Sigma)$, and the *Bers slice* through $Y \in \mathcal{T}(\Sigma)$ by $\mathcal{B}_Y := \{Q(X, \overline{Y}) \mid X \in \mathcal{T}(\Sigma)\} \subset \mathcal{QF}(\Sigma)$ (that is, quasi-Fuchsian manifolds with $\overline{Y} \in \mathcal{T}(\overline{\Sigma})$ on the bottom). Note that $\mathcal{E}_M$ is the graph of σ_M in the Ahlfors-Bers coordinates for $\mathcal{QF}(\Sigma)$. From this point of view, $\mathcal{B}_Y$ is a 'vertical line', and $\overline{Y} = \sigma_M(X)$ if and only if $\mathcal{E}_M$ and $\mathcal{B}_Y$ intersect

at $Q(X, \overline{Y})$. If $\overline{Y} = \sigma_M(X)$, then σ_M has a critical point at X if and only if $\mathcal{E}_M$ and $\mathcal{B}_Y$ have a tangency at $Q(X, \overline{Y})$.

As Σ' is incompressible in M', we may consider the following diagram:

$$\begin{array}{ccccccc} \mathcal{T}(\Sigma') & \xrightarrow{\mathrm{AB}_{M'}} & \mathcal{GF}(M') & \xrightarrow{\iota^*} & \mathcal{QF}(\Sigma') & \xrightarrow{\mathrm{AB}_{\Sigma'}} & \mathcal{T}(\Sigma') \times \mathcal{T}(\overline{\Sigma}') \\ \uparrow{\scriptstyle p^*} & & \uparrow{\scriptstyle p^*} & & \uparrow{\scriptstyle p^*} & & \uparrow{\scriptstyle p^*} \\ \mathcal{T}(\Sigma) & \xrightarrow[\mathrm{AB}_M]{} & \mathcal{GF}(M) & \xrightarrow[\iota^*]{} & \mathcal{QF}(\Sigma) & \xrightarrow[\mathrm{AB}_\Sigma]{} & \mathcal{T}(\Sigma) \times \mathcal{T}(\overline{\Sigma}) \end{array} \tag{2.1}$$

The maps p^* above each correspond to lifting structures via p. For example, the natural map induced by restriction $\mathcal{GF}(M) \to \mathcal{GF}(M')$ can be interpreted as lifting hyperbolic structures via p.

In the center of the diagram, each map is given by restricting a representation to a subgroup. This is induced by the commutative diagram of subgroups,

$$\begin{array}{ccc} \pi_1\Sigma' & \subset & \pi_1 M' \\ \cap & & \cap \\ \pi_1\Sigma & \subset & \pi_1 M \end{array}$$

and hence the center of diagram (Equation 2.1) commutes.

The outside pieces of diagram (Equation 2.1) commute because the map AB_M is natural with respect to passing to finite covers: Given a point $[\Gamma] \in \mathcal{GF}(M)$, let $\Gamma' := p^*(\Gamma)$, let $\mathrm{AB}_M^{-1}[\Gamma] = X$, and let $X' := \mathrm{AB}_{M'}^{-1}[\Gamma']$. Because $\Gamma' < \Gamma$ is finite-index, the tops of the domains of discontinuity of Γ and Γ' are the same set. Thus X' holomorphically covers X, with compatible markings so that $p^*(X) = X'$.

Consider a critical point $X \in \mathcal{T}(\Sigma)$, with $\overline{Y} = \sigma_M(X)$, i.e. a tangency between $\mathcal{B}_Y$ and $\mathcal{E}_M$ in $\mathcal{QF}(\Sigma)$. By commutativity of the diagram, $p^*\mathcal{B}_Y \subset \mathcal{B}_{Y'}$, where $Y' = p^*(Y)$, and $p^*\mathcal{E}_M \subset \mathcal{E}_{M'}$. Thus, in order to see a critical point of $\sigma_{M'}$ at p^*X, it is enough to observe that $p^* : \mathcal{QF}(\Sigma) \to \mathcal{QF}(\Sigma')$ is an immersion: A tangency between $\mathcal{B}_Y$ and $\mathcal{E}_M$ at $Q(X, \overline{Y})$ lifts to a tangency between $\mathcal{B}_{Y'}$ and $\mathcal{E}_{M'}$ at $p^*Q(X, \overline{Y}) = Q(p^*X, p^*\overline{Y})$.

In order to see that $p^* : \mathcal{QF}(\Sigma) \to \mathcal{QF}(\Sigma')$ is an immersion, we must consider the smooth structure on $\mathcal{X}(\Sigma)$. One interpretation of this structure identifies $T_{[\rho]}\mathcal{X}(\Sigma)$ with $H^1\left(\Sigma, (\mathrm{sl}_2\mathbb{C})_\rho\right)$, the vector space of 1-forms with values in the flat $\mathrm{sl}_2\mathbb{C}$-bundle on Σ associated to ρ. (See (Goldman, 1984) for details about the infinitesimal deformation theory of $\mathcal{X}(\Sigma)$). The restriction map $p^* : \mathcal{X}(\Sigma) \to \mathcal{X}(\Sigma')$ naturally induces the pullback map $p^* : H^1\left(\Sigma, (\mathrm{sl}_2\mathbb{C})_\rho\right) \to H^1\left(\Sigma', (\mathrm{sl}_2\mathbb{C})_{p^*\rho}\right)$ on tangent spaces.

Finite covering maps induce injective pullback maps on cohomology groups: If the pullback $p^*\phi$ is a coboundary df, one may average f over the finite sheets of the cover to obtain a form that descends, showing that ϕ was also a coboundary. Hence $d_{[\rho]}p^* : T_{[\rho]}\mathcal{X}(\Sigma) \to T_{p^*[\rho]}\mathcal{X}(\Sigma')$ is injective, $p^* : \mathcal{X}(\Sigma') \to \mathcal{X}(\Sigma)$ is an immersion, and $\sigma_{M'}$ has a critical point at p^*X. □

2.3 The Example 3-manifold

Consider the following pared 3-manifold $M = (H_2, P)$: Let H_2 denote the closed genus-2 handlebody and P the union of the annuli obtained as regular neighborhoods of the curves pictured in Figure 1. Fix a choice of basepoint $x \in \partial H_2 \backslash \overline{P}$, and the presentation $\pi_1(H_2, x) = \langle A, B \rangle$

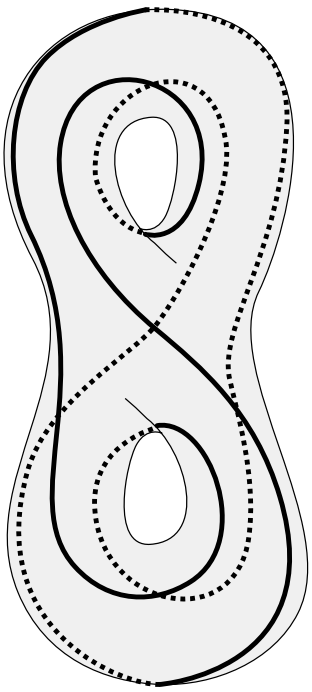

Figure 1. Core curves of annuli in P

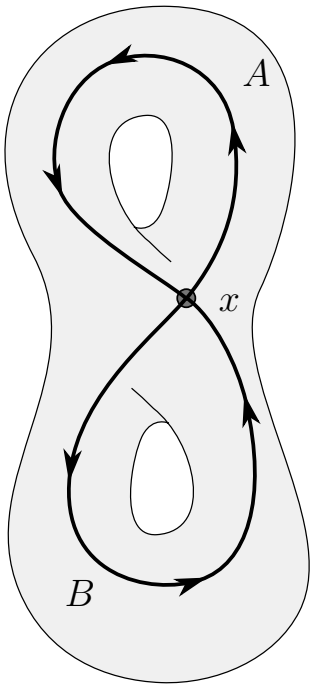

Figure 2. Generators for $\pi_1(H_2, x)$

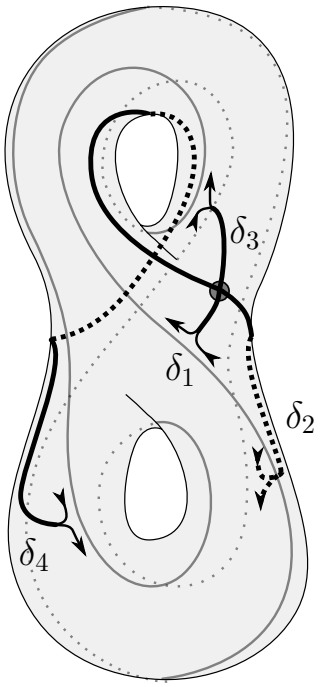

Figure 3. Generators for $\pi_1(\Sigma, x)$

(see Figure 2). For a natural choice of presentation for $\pi_1(\Sigma_{2,0}, x)$, given by $\langle a_1, b_1, a_2, b_2 | [a_1, b_1][a_2, b_2] = 1 \rangle$ the conjugacy classes of the core curves of P are $\{[b_1 a_1 b_1 b_2 [a_1, b_1]], [b_2 a_2 b_2 b_1]\}$[1].

Since the annuli in P are *disk-busting* in the genus-2 surface, (Otal, 1988, Lem. 1.15.) guarantees that the boundary is incompressible and (H_2, P) is acylindrical. Since the annuli in P are non-separating, $\Sigma = \partial M$ is a 4-holed sphere. Fix the presentation $\pi_1(\Sigma, x) = \langle \delta_1, \delta_2, \delta_3, \delta_4 \mid \prod \delta_i = 1 \rangle$, with the δ_i as pictured in Figure 3: δ_i takes the path drawn to a boundary component and winds around it counter-clockwise. Deleting P from ∂H_2 we arrive at a topological picture of Σ, shown in Figure 4.

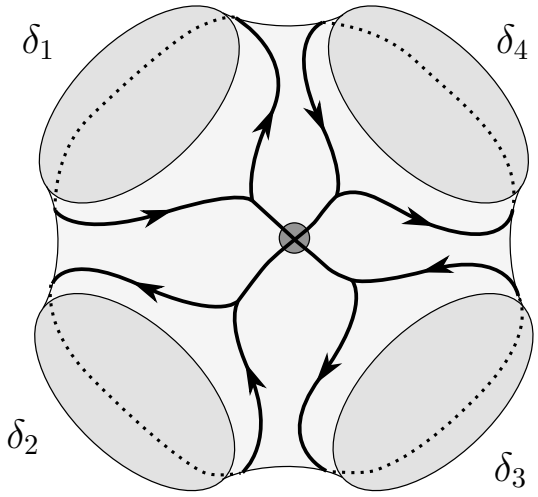

Figure 4. Another view of $\pi_1(\Sigma, x)$

---

[1]The reader is warned of the notational offense that the choices above force $\pi_1(\Sigma_{2,0}, x) \to \pi_1(M, x)$ to be given by $b_1 \mapsto A$ and $b_2 \mapsto B$. Though inconvenient, a_i and b_i will play no further role in our analysis, and the reader may ignore $\pi_1(\Sigma_{2,0}, x)$.

Choosing the constant path from x to x, we record the homomorphism induced by inclusion $\iota_* : \pi_1(\Sigma, x) \hookrightarrow \pi_1(M, x)$ in the chosen generators:

$$\iota_*(\delta_1) = A^{-2}B \tag{2.2}$$

$$\iota_*(\delta_2) = B^{-2}A^{-1} \tag{2.3}$$

$$\iota_*(\delta_3) = AB^{-1}A \quad = A\ \iota_*(\delta_1)^{-1}\ A^{-1} \tag{2.4}$$

$$\iota_*(\delta_4) = A^{-1}B^2A^2 \ = A^{-2}\ \iota_*(\delta_2)^{-1}\ A^2 \tag{2.5}$$

We suppress the basepoint x and the notation ι_* in what follows, and view $\pi_1(\Sigma)$ as a subgroup of $\pi_1(M)$. In everything that follows, we fix the notation $M = (H_2, P)$ and Σ for the 4-holed sphere boundary. We now describe a geometrically finite hyperbolic structure on M. Details about similar structures (gluings of regular right-angled ideal polyhedra) can be found in (Chesebro and DeBlois, 2012).

For $v, w \in \left(\mathbb{CP}^1\right)^3$, each with pairwise distinct entries, let $m(v, w)$ be a lift to $\mathrm{SL}_2\,\mathbb{C}$ (which one is immaterial) of the unique Möbius transformation taking the triple v to the triple w. Fix an identification $\partial_\infty \mathbb{H}^3 \cong \mathbb{CP}^1$, and let $\mathcal{O}$ be the ideal octahedron in $\mathbb{H}^3$ with totally geodesic triangular faces whose vertices are $\{1, 0, -1, i, -i, \infty\}$.

Consider the following representation of $\pi_1 H_2$:

$$\tilde{\rho}_1 : \langle A, B \rangle \ \rightarrow \ \mathrm{SL}_2\,\mathbb{C} \tag{2.6}$$

$$A \quad \mapsto \quad m((-1,i,0),(0,-i,1)) \quad = \ \frac{1+i}{2}\begin{pmatrix} 1 & 1 \\ 1+2i & 1 \end{pmatrix} \tag{2.7}$$

$$B \quad \mapsto \quad m((-1,-i,\infty),(1,\infty,i)) \ = \ -\frac{1+i}{2}\begin{pmatrix} i & -1+2i \\ 1 & i \end{pmatrix} \tag{2.8}$$

Let $\rho_1 : \langle A, B \rangle \rightarrow \mathrm{PSL}_2\,\mathbb{C}$ be the representation induced by $\tilde{\rho}_1$.

Geometrically, $\rho_1(A)$ and $\rho_1(B)$ perform face identifications on $\mathcal{O}$ with hyperbolic isometries (see Figure 5), and so ρ_1 is faithful by a standard Ping-Pong Lemma argument. Let $\hat{\Gamma}_1 = \rho_1(\langle A, B \rangle)$. By the Poincaré Polyhedron Theorem (Maskit, 1988, p. 75), $\hat{\Gamma}$ is discrete, and $\mathbb{H}^3/\hat{\Gamma}$ is a hyperbolic structure on the genus-2 handlebody. Since all of the dihedral angles of $\mathcal{O}$ are $\frac{\pi}{2}$, the non-paired faces meet flush. Thus the quotient $\mathcal{O}/\hat{\Gamma}_1$ is a hyperbolic structure on the genus-2 handlebody, with totally geodesic boundary, homotopy equivalent to $\mathbb{H}^3/\hat{\Gamma}_1$. Horoball neighborhoods of the six ideal vertices glue to two rank-1 cusp neighborhoods in the complete hyperbolic manifold $\mathbb{H}^3/\hat{\Gamma}_1$.

Denote the convex core of $\hat{\Gamma}_1$ by $\mathcal{C}$. Since $\mathcal{O}$ is the convex hull of points that are in the limit set, $\mathcal{O}/\hat{\Gamma}_1 \subset \mathcal{C}$. Since $\mathcal{C}$ is minimal among convex subsets of $\mathbb{H}/\hat{\Gamma}_1$ for which inclusion is a homotopy equivalence, we see that $\mathcal{C} = \mathcal{O}/\hat{\Gamma}_1$. Since $\mathcal{O}$ has finite volume and nonempty interior, $\hat{\Gamma}_1$ is geometrically finite.

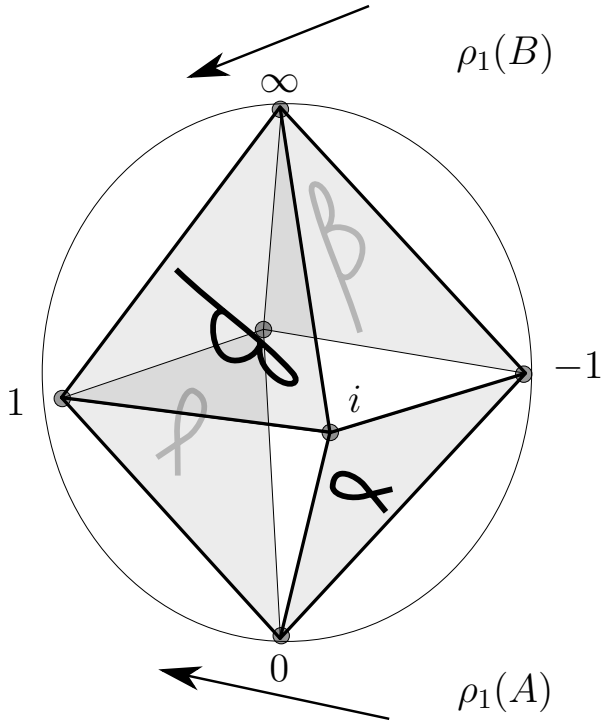

Figure 5. The octahedron $\mathcal{O}$ with side identifications labelled by α and β.

The ρ_1-images of the core curves of P are distinct maximal parabolic conjugacy classes in $\hat{\Gamma}_1$. Any parabolics in $\hat{\Gamma}_1$ must be conjugate to parabolics stabilizing the equivalence class of a vertex of $\mathcal{O}$ (Maskit, 1988, p. 131). Since P consists of two annuli, and since the vertices split into two equivalence classes, ρ_1 determines a one-to-one correspondence between components of P and conjugacy classes of maximal parabolic subgroups of $\hat{\Gamma}_1$. This means we are in the setting of (Chesebro and DeBlois, 2012, Lem. 2.6.) and ρ_1 is induced by a homeomorphism from the convex core of $\hat{\Gamma}_1$, ${}^{\mathcal{O}}/_{\hat{\Gamma}_1}$, to $H_2 \setminus P$. We denote this point $[\rho_1] \in \mathcal{GF}(M) \subset \mathcal{X}(M)$.

The octahedron has an evident self-map which respects the gluings: Reflect through an equatorial plane and rotate by 90° around the axis perpendicular to the plane. In the coordinates chosen in Figure 5, this is the anti-Möbius map $z \mapsto i/\bar{z}$. This self-map of $\mathcal{O}$ thus descends to an orientation-reversing isometry of the hyperbolic 3-manifold $\mathbb{H}^3/\hat{\Gamma}_1$ (cf. §2.4).

Before analyzing this symmetry of (M, Σ) (cf. §2.5), we examine deformations of the above hyperbolic structure. In particular, the octahedron $\mathcal{O}$ provides a 'hands-on' method to deform the representation ρ_1.

2.4 Deforming the example

Consider now the path of representations:

$$\begin{aligned}
\tilde{\rho}_t : \langle A, B\rangle &\to \mathrm{SL}_2\,\mathbb{C} \\
A &\mapsto m((-1, it, 0), (0, -it, 1)) = \frac{1}{t(t-i)} \begin{pmatrix} t^2 & t^2 \\ 1+2it & t^2 \end{pmatrix} \\
B &\mapsto m((-1, -it, \infty), (1, \infty, it)) = -\frac{i}{t+i} \begin{pmatrix} it & -1+2it \\ 1 & it \end{pmatrix}
\end{aligned} \tag{2.9}$$

Once again, let $\rho_t : \langle A, B\rangle \to \mathrm{PSL}_2\,\mathbb{C}$ be the representations induced by $\tilde{\rho}_t$. Taking $t = 1$, we recover the representation ρ_1 from §2.3. For $t \neq 1$, ρ_t performs identifications of the triangular totally geodesic faces of an ideal octahedron with vertices $\{1, 0, -1, it, -it, \infty\}$. Note that the octahedron is not regular, and the dihedral angles are no longer 90°.

For each $t \in \mathbb{R}$, let $\hat{\Gamma}_t := \rho_t(\langle A, B\rangle)$ and $\Gamma_t := \rho_t(\pi_1\Sigma)$.

Lemma 2.4.1. *The maps* $[\rho_t] : (0,\infty) \to \mathcal{X}(M)$ *and* $\left[\rho_t|_{\pi_1(\Sigma)}\right] : (0,1] \to \mathcal{X}(\Sigma)$ *are injective. For* $t > 0$, $[\rho_t|_{\pi_1(\Sigma)}]$ *is in the real locus of* $\mathcal{X}(\Sigma)$ *if and only if* $t = 1$.

Proof. Let $A_t := \rho_t(A)$ and $B_t := \rho_t(B)$. Noting that the square of the trace is a well-defined function on $\mathrm{PSL}_2\,\mathbb{C}$, for $i \in \{1,2,3,4\}$ we have

$$tr^2\rho_t(\delta_i) = 4.$$

Thus $[\rho_t] \in \mathcal{X}(M) \subset \mathcal{X}(H_0)$.

In order to show injectivity of the paths $[\rho_\cdot]$ and $[\rho_\cdot|_{\pi_1\Sigma}]$ on $(0,\infty)$ and $(0,1)$, respectively, we show that tr^2A_t and $tr^2\rho_t(\delta_1\delta_2)$ are one-to-one functions of t on the respective intervals.

We compute:

$$tr^2A_t = \left(\frac{2t}{t-i}\right)^2$$

Note that $\frac{2t}{t-i} - 1 = \frac{t+i}{t-i}$. For $t > 0$, the quantity $\frac{t+i}{t-i}$ is evidently a parameterization of the upper hemisphere of the unit circle centered at 0. Thus $\{tr^2A_t : t > 0\}$ is a parameterization of the square of the upper hemisphere of the unit circle centered at 1. In particular, it is one-to-one on $(0,\infty)$.

For brevity, let $f(t) := tr\tilde{\rho}_t(\delta_1\delta_2)$. A computation shows

$$f(t) = \frac{2t^2(t^4 - 22t^2 - 7)}{(1+t^2)^3} + i\frac{(t^2-1)(5t^2+1)^2}{t(1+t^2)^3}.$$

For $t \in (0,1)$, the quantity $\operatorname{Im} f(t) < 0$. Thus $f(t)^2$ is one-to-one on $(0,1)$ if and only if $f(t)$ is one-to-one on $(0,1)$. We compute

$$\frac{d}{dt} \operatorname{Im} f(t) = \frac{-(1+5t^2)(5t^6 - 35t^4 + 7t^2 - 1)}{t^2(1+t^2)^4}.$$

We now estimate, for $t \in (0,1)$,

$$\begin{aligned} 5t^6 - 35t^4 + 7t^2 - 1 = \quad & 5(t^2-1)^3 \; - \quad 20(t^2-1)^2 \; - \quad 48(t^2-1) \; - \quad 64 \\ < \quad & 0 \quad + \quad 0 \quad + \quad 48 \quad - \quad 64 \\ < \quad & 0 \end{aligned}$$

Thus $\frac{d}{dt} \operatorname{Im} f(t) > 0$, and $\operatorname{Im} f(t)$ is monotone increasing for $t \in (0,1)$. In particular, the function $tr^2\rho_t(\delta_1\delta_2) : (0,1) \to \mathbb{C}$ is one-to-one.

Finally, since the Kleinian group $\hat{\Gamma}_1$ has totally geodesic boundary, the boundary subgroup $[\Gamma_1]$ is in the real locus of $\mathcal{X}(\Sigma)$. As well, it is clear that $Im\ f(t) = 0$, for $t > 0$, if and only $t = 1$. Thus $[\hat{\Gamma}_t]$ is in the real locus of $\mathcal{X}(\Sigma)$ if and only if $t = 1$. □

The 'symmetry' of $\hat{\Gamma}_1$ (cf. §2.3) persists along the path $t \mapsto [\rho_t]$. (In fact, this path parametrizes the full fixed subset of $\mathcal{X}(M)$ preserved by this symmetry, but this is not relevant to our analysis). In Lemma 2.4.2 we check that conjugation by the anti-Möbius map $\Psi_t(z) = it/\bar{z}$ descends to an isometry of $\mathbb{H}^3/\hat{\Gamma}_t$. This isometry will be instrumental in our analysis of the path $[\rho_t]$ (cf. §6 and §7).

Lemma 2.4.2. *For all t, we have $\Psi_t\hat{\Gamma}_t\Psi_t^{-1} = \hat{\Gamma}_t$, and $[\Psi_t\Gamma_t\Psi_t^{-1}] = [\Gamma_t]$. Thus Ψ_t induces a mapping class $\Psi \in MCG^*(M)$, with $[\hat{\Gamma}_t] \in \text{Fix } \Psi$ and $[\Gamma_t] \in \text{Fix } \Psi|_\Sigma$. The action of $\Psi|_\Sigma$ on the punctures is an order four permutation, and $\Psi|_\Sigma$ preserves the two simple closed curves $\delta_1\delta_3$ and $\delta_2\delta_4$.*

Proof. A calculation using the definition of ρ_t (Equation 2.9) shows:

- $\Psi_t A_t \Psi_t^{-1} = B_t$
- $\Psi_t B_t \Psi_t^{-1} = A_t^{-1}$

Thus $\Psi_t\hat{\Gamma}_t\Psi_t^{-1} = \hat{\Gamma}_t$, and Ψ_t induces a mapping class $\Psi \in MCG^*(M)$ with $[\hat{\Gamma}_t] \in \text{Fix } \Psi$. We will drop the restriction map notation and consider Ψ as simultaneously an element of $MCG^*(M)$ and $MCG^*(\Sigma)$. We compute:

- $\Psi_t\rho_t(\delta_1)\Psi_t^{-1} = A_t^2 \cdot \rho_t\left(\delta_4^{-1}\right) \cdot A_t^{-2}$
- $\Psi_t\rho_t(\delta_2)\Psi_t^{-1} = A_t^2 \cdot \rho_t\left(\delta_1^{-1}\right) \cdot A_t^{-2}$
- $\Psi_t\rho_t(\delta_3)\Psi_t^{-1} = A_t^2 \cdot \rho_t\left(\delta_1\delta_3\delta_4\right) \cdot A_t^{-2}$
- $\Psi_t\rho_t(\delta_4)\Psi_t^{-1} = A_t^2 \cdot \rho_t\left(\delta_1\delta_2\delta_4\right) \cdot A_t^{-2}$

Thus $\Psi_t\Gamma_t\Psi_t^{-1}$ is conjugate to Γ_t, and Ψ is an orientation-reversing mapping class of Σ with $[\Gamma_t] \in \text{Fix } \Psi$.

Because $\delta_1\delta_3\delta_4 = \delta_1\delta_2^{-1}\delta_1^{-1} \sim \delta_2^{-1}$, (and similarly $\delta_1\delta_2\delta_4 \sim \delta_3^{-1}$), we see that the action of Ψ on the conjugacy classes of punctures is a cyclic permutation. The fact that Ψ preserves the geodesic representatives of the simple closed curves $\delta_1\delta_3$ and $\delta_2\delta_4$ is immediate. □

Note that, by Proposition 2.1.4, the fixed set of the mapping class Ψ is preserved by the skinning map. Namely:

$$\sigma_M\left(\text{Fix } \Psi|_\Sigma\right) \subset \text{Fix } \Psi|_\Sigma \tag{2.10}$$

Since $\mathcal{GF}(M) \subset \mathcal{X}(M)$ is open (Marden, 1974, Thm. 10.1.), there is some open interval about 1 so that $[\rho_t] \in \mathcal{GF}(M)$ and, by Lemma 2.1.3, $\left[\rho_t|_{\pi_1(\Sigma)}\right] \in \mathcal{QF}(\Sigma)$. Denote the maximal such open interval around 1 by U.

The path of quasi-Fuchsian groups $t \mapsto \Gamma_t$, for $t \in U$, induces two paths in $\mathcal{T}(\Sigma)$ corresponding to the top and bottom of Γ_t. For $t \in U$, fix notation $[\Gamma_t] = Q(X_t, \overline{Z_t})$, so that $\sigma_M(X_t) = \overline{Z_t}$ and $X_t, Z_t \in \text{Fix } \Psi|_\Sigma$. Lemma 2.4.1 implies, in particular, that $\{X_t \mid t \in U\}$ is an injective path in $\mathcal{T}(\Sigma)$. In fact, the remainder of Chapter 2 is devoted to checking that $\{Z_t \mid t \in U\}$ is a non-injective path whose image is confined to the real one-dimensional submanifold Fix $\Psi|_\Sigma \subset \mathcal{T}(\Sigma)$. Our strategy will be to examine the convex core boundary surfaces and bending laminations of Γ_t. First, we need to better understand the set Fix $\Psi|_\Sigma$. This is the subject of §2.5.

2.5 Rhombic 4-punctured spheres

In this section, we examine the symmetries of 4-punctured spheres, and collect some useful facts about a special symmetrical set in $\mathcal{T}(\Sigma)$.

It is well-known that for all $X \in \mathcal{T}(\Sigma)$ there exists a Klein 4-group of conformal automorphisms that acts trivially on $\mathcal{S}$, whose non-trivial elements are involutions that exchange punctures in pairs. Lemma 2.5.1 characterizes a more restrictive orientation-reversing symmetry of a subset of $\mathcal{T}(\Sigma)$.

Definition. If $X \in \mathcal{T}(\Sigma)$ may be written as the complement of the vertices of a Euclidean rhombus in $\mathbb{CP}^1$, with ξ and η as pictured (see Figure 6), then X is $\{\xi, \eta\}$-*rhombic*.

Let $\mathcal{R}_{\{\xi,\eta\}} := \{X \in \mathcal{T}(\Sigma) \mid X \text{ is } \{\xi,\eta\}\text{-rhombic}\}$. Recall *Fenchel-Nielsen* coordinates for $\mathcal{T}(\Sigma)$: Let ξ be the pants curve. Using the coordinate pictured in Figure 6, for any $s < 0$, let the transversal δ be the homotopy class of the curve $\operatorname{Im} z = s \cdot \operatorname{Re} z$ (see Figure 7). By the classical work of Fenchel and Nielsen, we obtain a diffeomorphism $\mathcal{T}(\Sigma) \cong \{(\ell, \theta) \mid \ell \in \mathbb{R}^+,\ \theta \in \mathbb{R}\}$.

Lemma 2.5.1. *Let $X \in \mathcal{T}(\Sigma)$, $\xi, \eta \in \mathcal{S}$. The following are equivalent:*

1. *X is $\{\xi, \eta\}$-rhombic.*
2. *In the Fenchel-Nielsen coordinates above, $X \in \{\theta = \pi/2\}$.*
3. *There exists an orientation-reversing $\psi \in MCG^*(\Sigma)$ that acts as an order four cyclic permutation of the punctures, preserves ξ and η, and $X \in \operatorname{Fix} \psi$.*

Proof. (1) $\Rightarrow$ (2): In order to measure the twisting coordinate relative to the choice of transversal δ, we observe that the four points of intersection of ξ with the real and imaginary axes divide ξ into four arcs. These arcs are cyclically permuted by the isometry $z \mapsto \frac{is}{\bar{z}}$, and hence

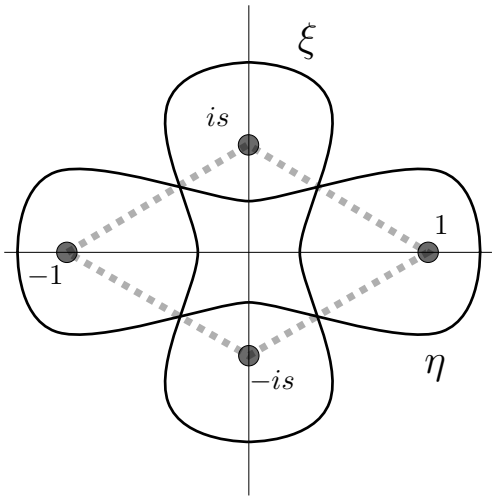

Figure 6. Global coordinates on a $\{\xi, \eta\}$-rhombic $X \in \mathcal{T}(\Sigma)$.

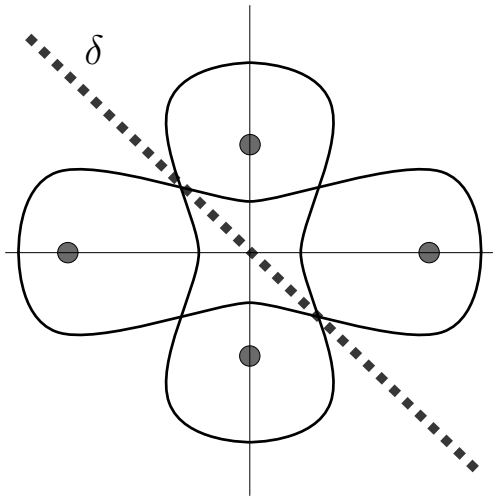

Figure 7. The transversal δ

are of equal length. This shows that the twisting coordinate is in $\mathbb{Z} \cdot \frac{\pi}{2}$, and the choice of δ guarantees the twisting coordinate is $\frac{\pi}{2}$.

(2) $\Rightarrow$ (3): There is a reflection/twist symmetry of X that preserves ξ and η and acts on punctures as desired.

(3) $\Rightarrow$ (1): By the existence of holomorphic involutions exchanging the punctures in pairs, there exists an automorphism α such that $\phi := \alpha \circ \psi$ is a simple transposition of the punctures that preserves ξ and η. These two fixed punctures must lie in the same component of the complement of η, since ϕ preserves η. With an appropriate Möbius transformation, take these to 1 and -1. Anti-conformal involutions of $\mathbb{CP}^1$ are inversions through a circle, and their fixed point sets are the circle they involve through. Since ϕ fixes the other two punctures, they lie

on the fixed circle for ϕ. We may thus apply another Möbius transformation, which fixes 1 and -1 and takes this fixed circle to the imaginary axis. Apply an elliptic Möbius transformation fixing 1 and -1, centering the imaginary punctures about 0, and multiplying by i and rescaling if necessary, and X has the desired form. □

Following Lemma 2.5.1, we refer to the defining symmetry ψ as $\{\xi, \eta\}$-*rhombic* as well. Though it is not involved in our analysis, the reader may recall that there is a natural identification $\mathcal{T}(\Sigma) \cong \mathbb{H}$. Each element of $\mathcal{T}(\Sigma)$ is covered by $\mathbb{C} \setminus \mathbb{Z} \oplus \tau\mathbb{Z}$ for some $\tau \in \mathbb{H}$, with covering group $\langle z \mapsto z + 2, z \mapsto z + 2\tau, z \mapsto -z \rangle$, such that the line $\{\operatorname{Im} z = \frac{1}{2} \operatorname{Im} \tau\}$ projects to ξ and the line $\{\operatorname{Im} z(\operatorname{Re} \tau - \frac{1}{2}) = \operatorname{Im} \tau(\operatorname{Re} z - \frac{1}{4})\}$ projects to η. Relative to this parametrization the reader may identify $\mathcal{R}_{\{\xi,\eta\}}$ as the line $\{\operatorname{Re} z = \frac{1}{2}\}$: The orientation-reversing symmetry $z \mapsto 1 - \bar{z}$ preserves the lattices $\mathbb{Z} \oplus \mathbb{Z}(\frac{1}{2} + yi)$ and preserves ξ and η, fulfilling (3).

We look to the action of $MCG^*(\Sigma)$ on $\mathbb{PML}(\Sigma)$. As it turns out, the fixed points of the action of $\{\xi, \eta\}$-rhombic isometries are easy to characterize.

Lemma 2.5.2. *A $\{\xi, \eta\}$-rhombic mapping class $\psi \in MCG^*(\Sigma)$, acting on $\mathbb{PML}(\Sigma)$, has* $\operatorname{Fix} \psi = \{\xi, \eta\}$.

Proof. There is an identification $(\mathbb{PML}(\Sigma), \mathcal{S}) \cong (\mathbb{RP}^1, \mathbb{QP}^1)$ which is equivariant with respect to a homomorphism $MCG^*(\Sigma) \to \mathrm{PGL}_2 \mathbb{Z}$ (cf. (Farb and Margalit, 2012, p. 60)). For the four-punctured sphere there is a $MCG^*(\Sigma)$-equivariant map from rays in $\mathcal{ML}(\Sigma)$ to lines in $H_1(\Sigma, \mathbb{R})$, which sends rays of laminations supported on simple closed curves to lines in $H_1(\Sigma, \mathbb{Q}) \subset H_1(\Sigma, \mathbb{R})$. Because the action of an element of $MCG^*(\Sigma)$ on $\mathbb{PML}(\Sigma)$ is given by the action of

an element of $\mathrm{PGL}_2\,\mathbb{Z}$ on $\mathbb{RP}^1$, each element of $MCG^*(\Sigma)$ which acts non-trivially on $\mathbb{PML}(\Sigma)$ has at most two fixed points. □

For a curve $\gamma \in \mathcal{S}$, recall the real-valued functions ℓ_γ and Ext_γ on $\mathcal{T}(\Sigma)$: For $X \in \mathcal{T}(\Sigma)$, the quantity $\ell_\gamma(X)$ is the hyperbolic length of the geodesic representative of γ and the quantity $Ext_\gamma(X)$ is the extremal length of the family of curves homotopic to γ. See (Farb and Margalit, 2012) and (Ahlfors, 2006) for details.

Lemma 2.5.3. *The maps ℓ_ξ, ℓ_η, Ext_ξ, and Ext_η each provide a diffeomorphism from $\mathcal{R}_{\{\xi,\eta\}}$ to $\mathbb{R}^+$.*

Proof. Fix pants decomposition with pants curve ξ and transversal δ, as in Figure 6 and Figure 7. By Lemma 2.5.1, $\ell_\xi|_{\mathcal{R}_{\{\xi,\eta\}}}$ is injective. Since one may construct $X \in \mathcal{R}_{\{\xi,\eta\}}$ with pants curve of specified hyperbolic length, and since length functions are smooth, the lemma is clear for ℓ_ξ. The proof for ℓ_η is the same.

Along a Teichmüller geodesic, the extremal lengths of the vertical and horizontal foliations each provide diffeomorphisms to $\mathbb{R}^+$ (this follows from (Gardiner and Masur, 1991, Lem. 5.1), Gardiner's formula (Gardiner, 1984, Thm. 8), and the inverse function theorem). Thus it is enough to check:

1. The set $\mathcal{R}_{\{\xi,\eta\}}$ is a Teichmüller geodesic.

2. The projective classes of its foliations are given by $[\xi], [\eta] \in \mathbb{PML}(\Sigma)$.

The first follows from the fact that fixed point sets of isometries, in uniquely geodesic spaces such as $(\mathcal{T}(\Sigma), d_{\mathcal{T}})$, are convex, while the second follows because the two foliations must be preserved by the rhombic symmetry, and Lemma 2.5.2 applies. □

In fact, it will be more direct to deal with a conformal invariant closely related to extremal length, the *modulus* of a quadrilateral. Recall that the the modulus of a quadrilateral is the extremal length of the family of curves that connect the vertical sides (see (Lehto, 1987), (Ahlfors, 2006) for details).

For $X \in \mathcal{R}_{\{\xi,\eta\}}$, consider the natural coordinate on X provided by Lemma 2.5.1. Let Q_X be the quadrilateral given by the first quadrant in $\mathbb{CP}^1$, with vertices $\{is, 0, 1, \infty\}$ and vertical sides given by the arcs $(0, 1)$ and (is, ∞) (see Figure 8). Denote the modulus of a quadrilateral or annulus by $Mod(\cdot)$.

Lemma 2.5.4. *We have* $Ext_\xi X = 4\, Mod(Q_X)$.

Proof. The quantity $Ext_\xi X$ can also be computed as the modulus of the unique maximal modulus annulus containing ξ as its core curve. Consider this annulus A_ξ. Since ξ is preserved by the anti-holomorphic maps $z \mapsto \overline{z}$ and $z \mapsto -\overline{z}$, the annulus A_ξ is preserved by these maps. This is enough to ensure that $A_\xi = \mathbb{C} \setminus ((-\infty, -1) \cup (1, \infty) \cup (-is, is))$. (Alternatively, in (Ahlfors, 2006, p. 23, II.), there is a classical extremal length problem whose solution—due to Teichmüller—implies that A_ξ has the given form).

Since the curves

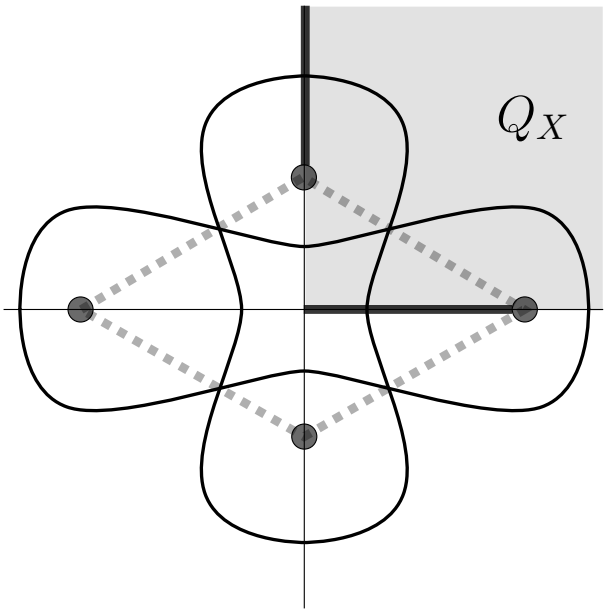

Figure 8. The quadrilateral Q_X, with vertical sides shaded.

- $\{i(s+\tau) \mid \tau \in (0,\infty)\}$
- $\{-i(s+\tau) \mid \tau \in (0,\infty)\}$
- $\{\tau \mid \tau \in (0,1)\}$
- $\{-\tau \mid \tau \in (0,1)\}$

are all fixed by $z \mapsto \overline{z}$ and $z \mapsto -\overline{z}$, they must be vertical geodesics in the Euclidean metric for A_ξ. By (Ahlfors, 2006, p. 16), the modulus $Mod(A_\xi)$ is equal to the sum of the moduli of the four quadrilateral pieces obtained after cutting along these vertical geodesics. Since these four quadrilateral pieces are each congruent to Q_X, we have $Mod(A_\xi) = 4Mod(Q_X)$. □

2.6 The convex core boundaries of Γ_t

We return to the goal of understanding Γ_t via its convex core boundary. Recall from Lemma 2.4.2 that Ψ is an order-4 orientation-reversing isometry of Γ_t, cyclically permuting the punctures and preserving the simple closed curves $\delta_1\delta_3$ and $\delta_2\delta_4$. From here on we fix the the notation $\xi := \delta_1\delta_3$ and $\eta := \delta_2\delta_4$. The reader may notice that the isometry Ψ is $\{\xi, \eta\}$-rhombic. (Equation 2.10) now becomes

$$\sigma_M\left(\mathcal{R}_{\{\xi,\eta\}}\right) \subset \mathcal{R}_{\{\xi,\eta\}}. \tag{2.11}$$

Below, we present the computation of the bending angles for the convex core of Γ_t. We choose the branch $[0, \pi]$ for $\cos^{-1}$ and $[0, 2\pi)$ for arg. Note that the existence of a 'top' and 'bottom' of Γ_t, relative to $\hat{\Gamma}_t$, implies that there is a 'top' and 'bottom' bending lamination. The computation of the bending measure for the top lamination is computable by the same methods presented as for the bottom. Recall that U is the maximal open interval, containing 1, so that Γ_t is quasi-Fuchsian for $t \in U$ (see §2.4).

Lemma 2.6.1. *For $t \in U$, the top and bottom surfaces of the convex core boundary of Γ_t are both in $\mathcal{R}_{\{\xi,\eta\}}$. The bending lamination on the bottom of the convex core boundary is $\theta(t) \cdot \xi$, for*

$$\theta(t) = \cos^{-1}\left(\frac{2t^3(3-t^2)}{(1+t^2)^2}\right). \tag{2.12}$$

Proof. By Lemma 2.4.2, Ψ_t is a $\{\xi, \eta\}$-rhombic isometry of Γ_t. It thus sends the convex core to itself, so preserves the convex core boundary pleated surfaces, which are hence in $\mathcal{R}_{\{\xi,\eta\}}$.

Moreover, Ψ_t preserves the bending laminations $\{\lambda^+, \lambda^-\}$, where λ^+ is on the top and λ^- is on the bottom. By Lemma 2.5.2, $\{[\lambda^+], [\lambda^-]\} \subset \{\xi, \eta\}$, where $[\cdot]$ denotes the projective class of the measured lamination. Before computing $\theta(t)$, and determining that $\xi = [\lambda^-]$, we describe the method in words.

To compute the bending angle $\theta(t)$, it is necessary to find a pair of distinct maximal support planes that intersect along the axis of the hyperbolic Möbius map $\rho_t(\xi)$, i.e. two half-planes that meet along the axis of $\rho_t(\xi)$. This can be achieved by considering the two other lifts of ξ which are the axes of $\rho_t\left(\delta_1^{-1}\xi\delta_1\right)$ and $\rho_t\left(\delta_4\xi\delta_4^{-1}\right)$: One can check easily that $\left[\delta_1^{-1}\cdot\xi\right]$ and $[\delta_4\cdot\xi]$—the homotopy classes of the concatenations—are simple closed curves. Thus, for the pair $\left\{\xi, \delta_1^{-1}\cdot\xi\right\}$ for example, there is a connected fundamental domain for the action of $\pi_1\Sigma$ that has the lifts of these curves on its boundary (motivating the term 'neighbors' for this pair, see Figure 9). The same is true, of course, for $\{\xi, \delta_4\cdot\xi\}$.

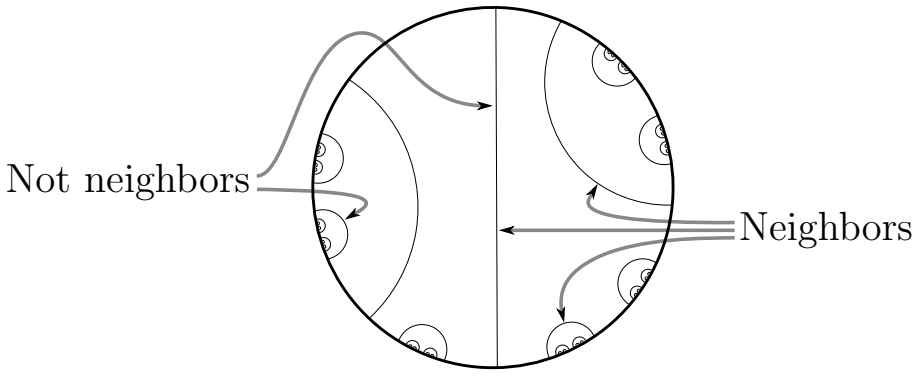

Figure 9. Lifts for $\xi \in \pi_1\Sigma$

As we know that ξ is the support of a bending lamination for Γ_t, the four fixed points of each of the pairs $\left\{\xi, \delta_1^{-1} \cdot \xi\right\}$ and $\{\xi, \delta_4 \cdot \xi\}$ lie on a round circle which intersects Λ_{Γ_t} at these four points. Each such circle is the frontier of a maximal support plane for Γ_t. The two pairs $\left\{\xi, \delta_1^{-1} \cdot \xi\right\}$ and $\{\xi, \delta_4 \cdot \xi\}$ thus determine two maximal support planes intersecting in $\mathbb{H}^3$ along the axis of $\rho_t(\xi)$, as desired.

We thus consider three pairs of fixed points, corresponding to the three chosen lifts of ξ: Fix $\rho_t\left(\delta_1^{-1}\xi\delta_1\right)$, Fix $\rho_t(\xi)$, and Fix $\rho_t\left(\delta_4\xi\delta_4^{-1}\right)$. The first four points determine one half-plane in $\mathbb{H}^3$, and the last four determine a second. We may now compute the bending angle between these half-planes using the argument of a cross-ratio.

There is an ambiguity in our calculation of support planes: we do not know which side of the totally geodesic plane through a pair of neighboring axes is actually a supporting half-space for either of the groups Γ_t or $\hat{\Gamma}_t$. In order, then, to distinguish between the top and the bottom laminations we will use the following consequence of the definitions: if there exists a support plane for a component $\Omega_0 \subset \Omega_{\Gamma_t}$ which is not simultaneously a support plane for $\hat{\Gamma}_t$, then Ω_0 is the bottom. Such a support plane is produced explicitly below at $t = \frac{1}{2}$, for the domain facing the convex core boundary component with bending supported on ξ.

In fact, for our purposes, it is enough to distinguish the top lamination from the bottom lamination at $t = \frac{1}{2}$: The bending lamination map from $\mathcal{GF}(M)$ to $\mathcal{ML}(\Sigma)$ is continuous (Keen and Series, 1995, Thm. 4.6.), and its image along the path $\{[\rho_t] \mid t \in U\}$, by the argument above, is confined to the subset of $\mathcal{ML}(\Sigma)$ supported on either ξ or η. This subset is a pair of rays, intersecting only at the 'zero' lamination. Thus, for instance, if ξ switches at some point

from the support for the bottom lamination of Γ_t to the support for the top lamination of Γ_t, the Kleinian group $\hat{\Gamma}_t$ must have the 'zero' bending lamination, i.e. totally geodesic boundary. In this case, Γ_t would be Fuchsian at this value of t. By Lemma 2.4.1, $[\rho_t]$ is not Fuchsian for $t < 1$, hence the lamination on top at $t = \frac{1}{2}$ must remain on top for all $t < 1$.

We proceed with the calculation of $\theta(t)$. Denote the fixed points of $\rho_t(\xi)$ by $p_t^{\pm}$, where the choice is fixed by asking that the root in the expression of the fixed points be positive for p^+ and negative for p^-. The fixed points of $\rho_t\left(\delta_1^{-1}\xi\delta_1\right)$ and $\rho_t\left(\delta_4\xi\delta_4^{-1}\right)$ are thus, respectively, $\rho_t\left(\delta_1^{-1}\right)\cdot p_t^{\pm}$, and $\rho_t(\delta_4)\cdot p_t^{\pm}$. The root below refers to the positive one.

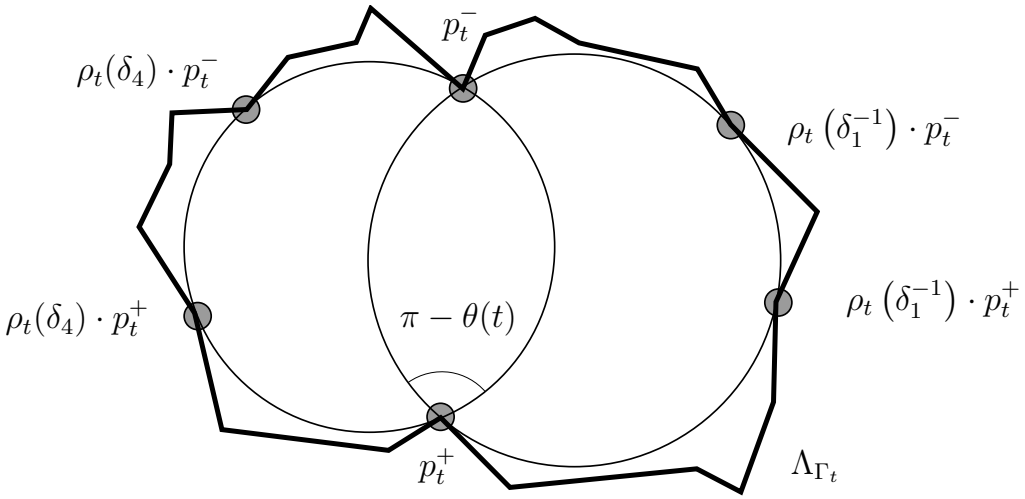

Figure 10. Maximal support planes for Γ_t, and the bending angle $\theta(t)$

We use the notation for the cross-ratio $[a:b:c:d] = \frac{d-a}{d-c}\frac{b-c}{b-a}$.

$$
\begin{aligned}
&[\ p_t^+ : \rho_t(\delta_4)\cdot p_t^+ : p_t^- : \rho_t\left(\delta_1^{-1}\right)\cdot p_t^+\ ] = \frac{1}{2t(-i+t)(i+t)^3}\cdot \\
&\left(-i-2t-2it^2-13it^4+2t^5+4it^6+\right. \\
&\left.\left(-i-2t-it^2+2t^3\right)\sqrt{1+6t^2+17t^4-4t^6}\right) \\
&= \frac{1+5t^2+\sqrt{1+6t^2+17t^4-4t^6}}{2t(t^2+1)^3}\cdot \\
&\left(2t^3(t^2-3)+\ i\ (1-t^2)\sqrt{1+6t^2+17t^4-4t^6}\right).
\end{aligned}
$$

It is easy to see that the imaginary part of this cross-ratio, for $t \in (0,1)$, is positive (cf. (Equation 2.14)). Thus $\theta(t) = \pi - \arg\ \left[p_t^+ : \rho_t(\delta_4)\cdot p_t^+ : p_t^- : \rho_t\left(\delta_1^{-1}\right)\cdot p_t^+\right]$. Taking the argument of the expression above gives

$$\theta(t) = \cos^{-1}\left(\frac{2t^3(3-t^2)}{(1+t^2)^2}\right).$$

It remains to distinguish the top of the convex core boundary from the bottom. Following the outline described above, we fix some notation at the point $t = \frac{1}{2}$. Let S be the circle through the points $p_{1/2}^+$, $p_{1/2}^-$, and $\rho_{1/2}(\delta_1^{-1})\cdot p_{1/2}^+$, let its center be c and its radius r. Noting that $B_{1/2}$ is loxodromic, let its fixed points be $p_1, p_2 \in \Lambda_{\hat{\Gamma}_{1/2}}$. The following computations are straightforward. The root is chosen with positive imaginary part.

$$c = -\frac{1}{2} - i\ ; \quad r = \frac{\sqrt{5}}{2}\ ; \quad p_1 = \left(\frac{1+2i}{5}\right)\sqrt{7+i}\ ; \quad p_2 = -\left(\frac{1+2i}{5}\right)\sqrt{7+i}$$

$$|c - p_1|^2 = \frac{5}{4} + \sqrt{2} + \operatorname{Re}\sqrt{7+i} > \frac{5}{4} = r^2$$
$$|c - p_2|^2 = \frac{5}{4} + \sqrt{2} - \operatorname{Re}\sqrt{7+i} < \frac{5}{4} = r^2$$

(Note that $\operatorname{Re}\sqrt{7+i} = \sqrt{\frac{7+5\sqrt{2}}{2}} > \sqrt{2}$.)

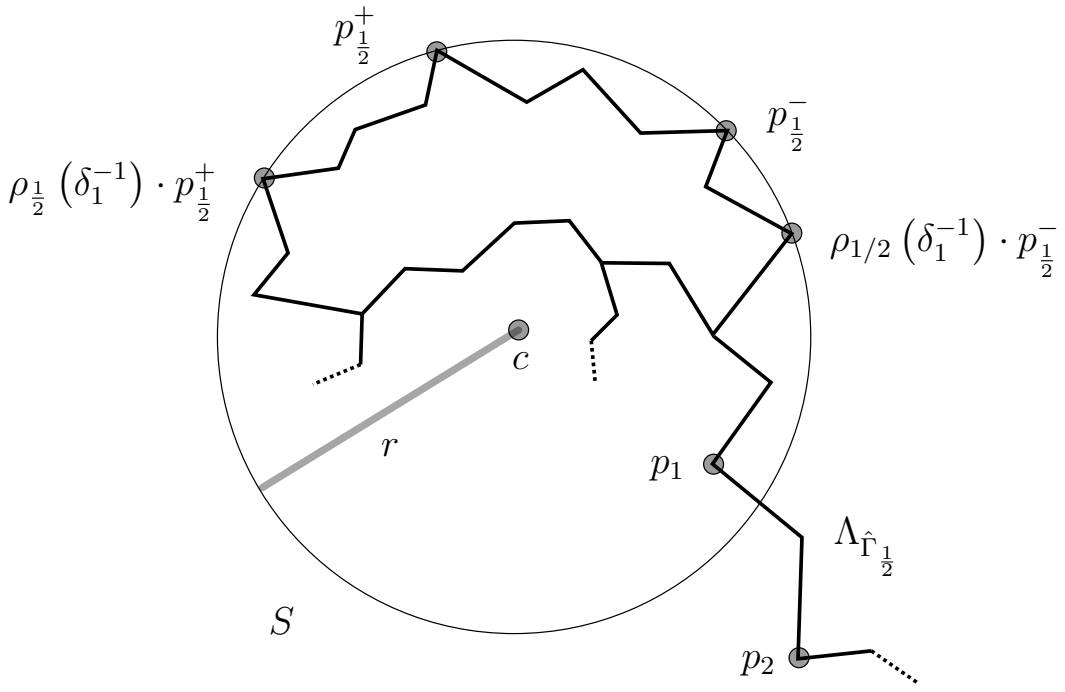

Figure 11. A support plane for $\Gamma_{\frac{1}{2}}$ fails to be a support plane for $\hat{\Gamma}_{\frac{1}{2}}$

A support plane for $\hat{\Gamma}_{1/2}$ must be the boundary of a supporting half-space, which must intersect $\mathbb{CP}^1$ in the domain of discontinuity. Because there are points $p_1, p_2 \in \Lambda_{\hat{\Gamma}_{1/2}}$ in the

interior and exterior of the round disk S, the geodesic plane which meets $\mathbb{CP}^1$ in S cannot be the boundary of a supporting half-space for $\hat{\Gamma}_{1/2}$. Because S is a support plane for $\Gamma_{1/2}$ but not for $\hat{\Gamma}_{1/2}$, S is a support plane for the bottom of $\Gamma_{1/2}$. (See Figure 11 for a schematic where the exterior of S is a supporting half-space for the bottom of $\Gamma_{1/2}$). Since S is on the side of $\Gamma_{1/2}$ with bending lamination ξ, this implies that ξ is the support of the bending lamination on the bottom. □

For $\gamma \in \pi_1(M)$, let $\ell(\gamma, \Gamma_t)$ denote the hyperbolic translation length of $\rho_t(\gamma)$ in $\mathbb{H}^3$. Using Lemma 2.6.1 we may now find $\inf U$ explicitly.

Lemma 2.6.2. *The bottom (resp. top) convex core boundary surface of Γ_t is the element of $\mathcal{R}_{\{\xi,\eta\}}$ determined by $\ell(\xi, \Gamma_t)$ (resp. $\ell(\eta, \Gamma_t)$), where*

$$\ell(\xi, \Gamma_t) = 2\ \cosh^{-1}\left(\frac{1 + 8t^2 + 21t^4 - 2t^6}{2t^2(1+t^2)^2}\right) \tag{2.13}$$

and

$$\ell(\eta, \Gamma_t) = 2\ \cosh^{-1}\left(\frac{-1 + 4t^2 + 74t^4 + 196t^6 - t^8}{(1+t^2)^4}\right).$$

Moreover, $\frac{1}{2}\left(5 + 3\sqrt{3} - \sqrt{44 + 26\sqrt{3}}\right) = \inf U$.

Proof. For a hyperbolic isometry $A \in \mathrm{PSL}_2\,\mathbb{C}$, its translation length is given by $\ell(A) = 2\cosh^{-1}\left(\frac{|trA|}{2}\right)$. Thus (Equation 2.13) can be computed directly. For $t \in U$, the hyperbolic length of the curve which supports the bending lamination of Γ_t, on the hyperbolic surface on the top or the bottom of the convex core boundary, is precisely the length of the curve in

the quasi-Fuchsian manifold $\mathbb{H}^3/\Gamma_t$. By Lemma 2.5.3, this quantity determines the convex core boundary surface.

Let $t_0 = \inf U$. Because Σ is incompressible in M, by (Bonahon, 1986, Thm. A) the hyperbolic manifold $\mathbb{H}^3/\hat{\Gamma}_{t_0}$ is *geometrically tame*. In particular, it has an end invariant λ, a geodesic lamination which is the union of limits of simple closed geodesics exiting the end and/or simple closed curves that are parabolic in $\hat{\Gamma}_{t_0}$. Since $\hat{\Gamma}_{t_0} \in \text{Fix } \Psi$, the invariant λ must also be preserved by the rhombic symmetry. Because the natural map from $\mathbb{PML}(\Sigma)$ to geodesic laminations is equivariant for the action of $MCG^*(\Sigma)$, Lemma 2.5.2 implies that $\lambda \in \{\xi, \eta\}$.

Thus $\hat{\Gamma}_{t_0}$ is geometrically finite with end invariant either ξ or η.

We check below the following computational facts:

1. $|tr\tilde{\rho}_t(\xi)| > 2$ for all $t \in (0, 1]$.
2. For $1 > t > \frac{1}{2}\left(5 + 3\sqrt{3} - \sqrt{44 + 26\sqrt{3}}\right)$, we have $|tr\tilde{\rho}_t(\eta)| > 2$, and $tr^2\rho_t(\eta) = 4$ at $t = \frac{1}{2}\left(5 + 3\sqrt{3} - \sqrt{44 + 26\sqrt{3}}\right)$.

As a result of (1), the end invariant of $\hat{\Gamma}_{t_0}$ must be η. As a result of (2), we may compute explicitly the lower bound $t_0 = \frac{1}{2}\left(5 + 3\sqrt{3} - \sqrt{44 + 26\sqrt{3}}\right)$.

(1): For all $t \in (0, 1]$,

$$\begin{aligned} 2 + tr\tilde{\rho}_t(\xi) &= \frac{-1 - 8t^2 - 21t^4 + 2t^6}{t^2(1+t^2)^2} + 2 \\ &= \frac{-1 - 6t^2 - 17t^4 + 4t^6}{t^2(1+t^2)^2} \\ &= \frac{-1 - 6t^2 - 13t^4 - 4t^4(1-t^2)}{t^2(1+t^2)^2} < 0. \end{aligned} \tag{2.14}$$

(2): For all $t \in (t_0, 1]$,

$$\begin{aligned} 2 + tr\tilde{\rho}_t(\eta) &= \frac{2(1 - 4t^2 - 74t^4 - 196t^6 + t^8)}{(1+t^2)^4} + 2 \\ &= \frac{4(1 - 34t^4 - 96t^6 + t^8)}{(1+t^2)^4} \\ &= \frac{4\left((1+t^2)^2 - 2t(1+5t^2)\right)\left((1+t^2)^2 + 2t(1+5t^2)\right)}{(1+t^2)^4}. \end{aligned}$$

One may check directly (using the explicit equation for the zeros of a quartic) that $(1+t^2)^2 - 2t(1+5t^2)$ has two real zeros, the smaller of which is $t_0 = \frac{1}{2}\left(5 + 3\sqrt{3} - \sqrt{44 + 26\sqrt{3}}\right)$. The larger zero, $\frac{1}{2}\left(5 + 3\sqrt{3} + \sqrt{44 + 26\sqrt{3}}\right)$, is greater than 1, and for $t = 1$, $(1+t^2)^2 - 2t(1+5t^2)$ is negative. Thus $2 + tr\tilde{\rho}_t(\eta) < 0$ for all $t \in (t_0, 1]$, and $tr^2\rho_{t_0}(\eta) = 4$. □

For $t \in (t_0, 1]$, recall the notation $[\Gamma_t] = Q(X_t, Z_t)$, so that in particular $\sigma_M(X_t) = Z_t$. Denote the bottom surface of the convex core boundary of Γ_t by Y_t. That is, recalling Lemma 2.5.3,

$$Y_t = \left(\ell_\xi|_{\mathcal{R}_{\{\xi,\eta\}}}\right)^{-1}\left(\ell(\xi, \Gamma_t)\right).$$

2.7 Non-injectivity of $t \mapsto Z_t$

In this section we use a description of Z_t as a grafted surface to show the non-monotonicity of the function $Ext_\xi Z_t : (t_0, 1] \to \mathbb{R}^+$. Since $Z_t \in \mathcal{R}_{\{\xi,\eta\}}$ by Lemma 2.4.2, and since $Ext_\xi : \mathcal{R}_{\{\xi,\eta\}} \to \mathbb{R}^+$ is a diffeomorphism by Lemma 2.5.3, the non-monotonicity of $Ext_\xi Z_t$ implies the non-injectivity of $t \mapsto Z_t$, for $t \in (t_0, 1]$.

Recall that *grafting* is a map (see (Kamishima and Tan, 1992) and (McMullen, 1998) for details)

$$gr : \mathcal{ML}(\Sigma) \times \mathcal{T}(\Sigma) \to \mathcal{T}(\Sigma).$$

Briefly, for $\gamma \in \mathcal{S}$, the Riemann surface $gr(\tau \cdot \gamma, X)$ is obtained by cutting open X along the geodesic representative for γ, and inserting a Euclidean annulus of height τ.

For a geometrically finite hyperbolic 3-manifold N, with conformal boundary $X_\infty \in \mathcal{T}(\partial N)$, convex core boundary $X_{cc} \in \mathcal{T}(\partial N)$, and bending lamination $\lambda \in \mathcal{ML}(\partial N)$, grafting provides the description:

$$X_\infty = gr(\lambda, X_{cc}).$$

Employing Lemmas 2.6.1 and 2.6.2, we thus have

$$Z_t = gr(\theta(t) \cdot \xi, Y_t). \tag{2.15}$$

We note that there is also a projective version of grafting, Gr, that 'covers' conformal grafting (Kamishima and Tan, 1992). This provides the natural (quasi-Fuchsian) projective structure $\mathcal{Z}_t = Gr(\theta(t) \cdot \xi, Y_t)$ which the surface Z_t inherits as the boundary of $\mathbb{H}^3/\Gamma_t$.

To each $X \in \mathcal{R}_{\{\xi,\eta\}}$, recall the quadrilateral Q_X (see §2.5). For brevity, let $Q_t := Q_{Z_t}$ (see Figure 12).

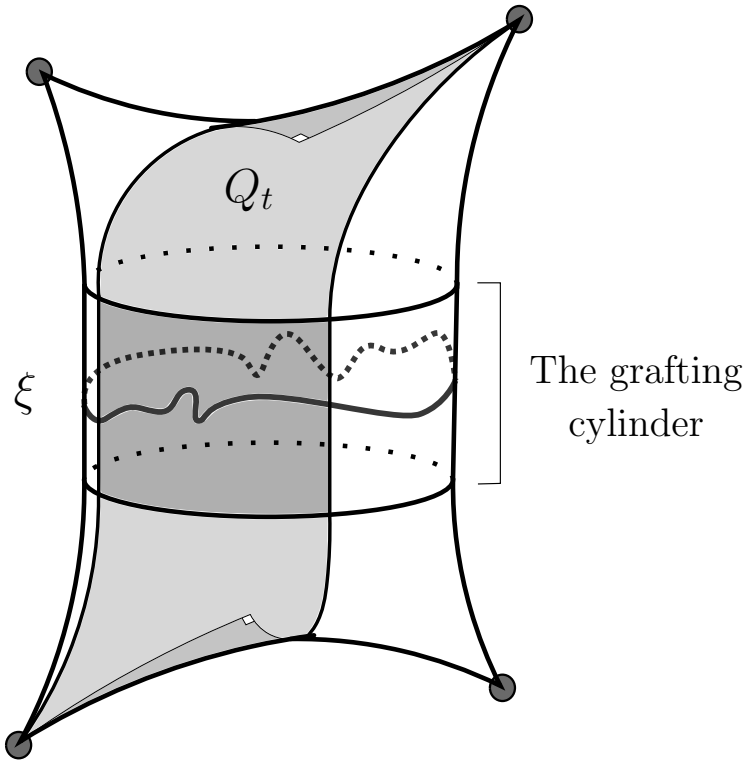

Figure 12. The quadrilateral $Q_t \subset Z_t$ and the curve ξ

In Lemma 2.7.1, we present the natural projective structure on Q_t, given by the projective grafting description of the projective structure $\mathcal{Z}_t$ on Z_t. Let $D(z, r)$ be the open disk in $\mathbb{C}$ centered at z of radius r, let E^c indicate the complement of a set $E \subset \mathbb{C}$, and let $L(t) := \frac{1}{4}\ell(\xi, \Gamma_t)$.

Lemma 2.7.1. *The projective structure on the quadrilateral Q_t is given by the region (see Figure 13)*

$$Q_t \cong \overline{D(c_1, r_1)} \cap \left(\bigcap_{i=2}^{4} D(c_i, r_i)^c \right)$$

with 'vertical' sides given by the two concentric arcs, such that:

1. $(c_1, r_1) = \left(0, e^{L(t)}\right)$

2. $(c_2, r_2) = (0, 1)$

3. $(c_3, r_3) = \left(e^{L(t)} \cosh(L(t)), e^{L(t)} \sinh(L(t))\right)$

4. $(c_4, r_4) = \left(e^{i(\pi+\theta(t))} \cosh(L(t)), \sinh(L(t))\right)$

Proof. By (Equation 2.15), the projective structure induced by Γ_t on Z_t is given by the projective grafting of the uniformized hyperbolic surface Y_t along the measured lamination $\theta(t) \cdot \xi$. We form the picture presented in Figure 13 by developing Y_t onto $\mathbb{H}^2$ as follows: Choose an orientation for ξ. Consider the intersection of the geodesic representative of ξ with Q_{Y_t}. Develop this arc onto the imaginary axis of $\mathbb{H}^2$, so that its endpoints lift to i and $e^{L(t)}i$. In this case, the sides of Q_{Y_t} are contained in the boundaries of the disks $D(0,1)$, $D\left(e^{L(t)} \cosh L(t), e^{L(t)} \sinh L(t)\right)$, $D\left(-\cosh L(t), \sinh L(t)\right)$, and $D\left(0, e^{L(t)}\right)$. With the chosen coordinate, the process of projective grafting along $\theta(t) \cdot \xi$ inserts a wedge of angle $\theta(t)$. □

We now have a function, $t \mapsto Mod(Q_t)$, for $t \in (t_0, 1]$, whose non-monotonicity would suffice to show the non-injectivity of $t \mapsto Z_t$. As a result of Lemma 2.7.1, it is possible to produce numerical estimates of the modulus that suggest non-monotonicity: One may use conformal maps to 'open' the two punctures to right angles, and a carefully scaled logarithm to make the quadrilateral look 'nearly rectangular'. These produce the rough estimates:

- $Mod\left(Q_1\right) \approx 2.3$
- $Mod\left(Q_{.52}\right) \approx 1.8$

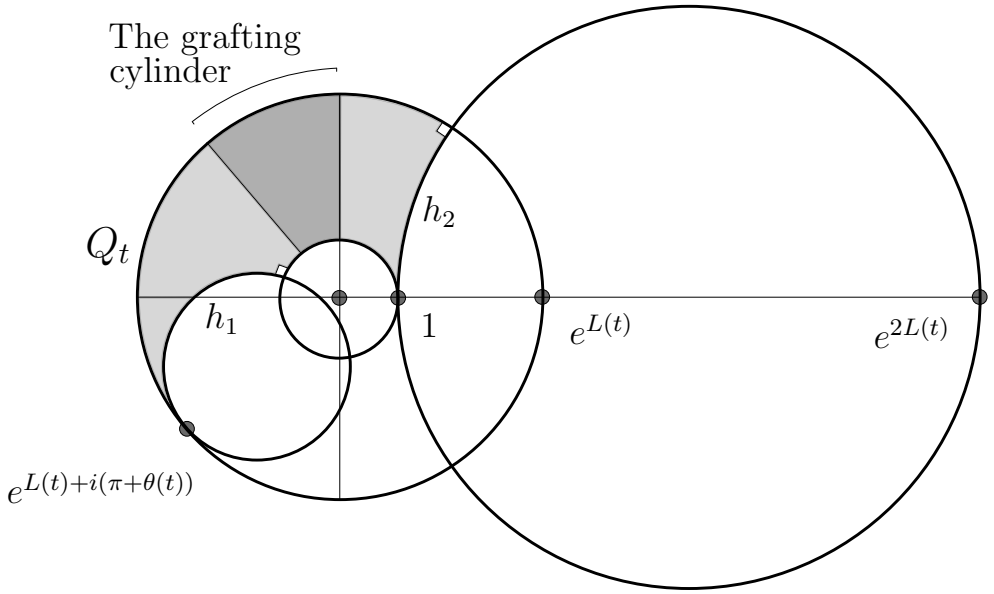

Figure 13. A picture of the quasi-Fuchsian projective structure on Q_t. The horizontal sides are labelled h_1 and h_2.

· $Mod\,(Q_{.39}) \approx 2.0$

(Recall that $t_0 = \frac{1}{2}\left(5 + 3\sqrt{3} - \sqrt{44 + 26\sqrt{3}}\right)$, and note that $.39 > t_0$). Unfortunately, controlling the error in these approximations quickly becomes very delicate, so we pursue a different approach. In Lemma 2.7.2, we apply a normalizing transformation to the conformal structure given in Lemma 2.7.1 which allows a direct comparison of moduli of quadrilaterals.

While there remain some involved calculations, this method reduces the computational difficulties considerably. After Lemma 2.7.2, we have a reasonable list of verifications to check, involving integers and simple functions. We defer this list of verifications to §2.8 and to the Appendix, as they contribute no new ideas to the proof.

Choose a branch of the logarithm $\arg(z) \in (0, 2\pi)$. Applying $z \mapsto \log z$ takes Q_t into a vertical strip, with its vertical sides sent into a pair of vertical lines. Its horizontal sides can be described as the graphs of two explicit functions over the interval between the vertical sides. Following this map by $z \mapsto \frac{1}{L(t)}(z - \frac{\pi+\theta(t)}{2}i)$ we arrive at R_t, a certain conformal presentation of Q_t (see Figure 14). The proof of the following lemma is a straightforward calculation.

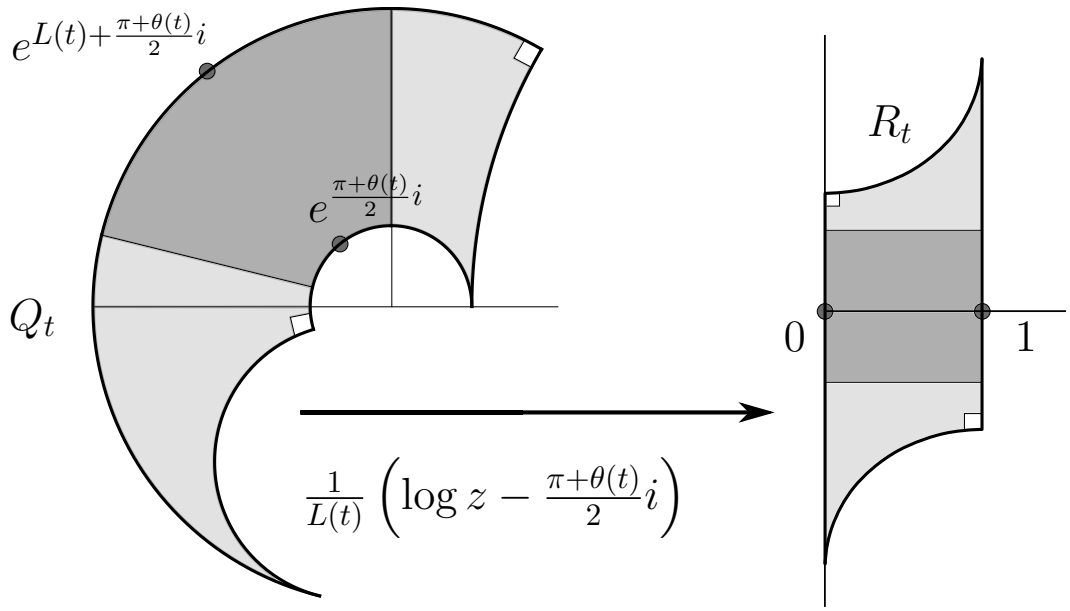

Figure 14. The normalizing transformation $Q_t \to R_t$

Lemma 2.7.2. *The image under the transformation* $z \mapsto \frac{1}{L(t)}\left(\log z - \frac{\pi+\theta(t)}{2}i\right)$ *of* Q_t *is given by* R_t *(see Figure 14). The quadrilateral* R_t *has vertical sides in the vertical lines* $\{\operatorname{Re} z = 0\}$,

$\{\mathrm{Re}\, z = 1\}$, *and horizontal sides given as the graphs of the functions* $x \mapsto F(x,t)$ *and* $x \mapsto -F(1-x,t)$ *over* $x \in [0,1]$. *The function* $F(x,t)$ *is given by:*

$$F(x,t) = \alpha(t) - \beta(x,t) \tag{2.16}$$

where α *and* β *are given by:*

$$\alpha(t) = \frac{\pi + \theta(t)}{2L(t)} \tag{2.17}$$

$$\beta(x,t) = \frac{1}{L(t)} \cos^{-1}\left(\frac{\cosh\left(xL(t)\right)}{\cosh\left(L(t)\right)}\right) \tag{2.18}$$

The useful aspect of Lemma 2.7.2 is that $Mod(Q_t)$ now depends only on the graph of the function $x \mapsto F(x,t)$ over $x \in [0,1]$, which forms a kind of 'profile' for the quadrilateral Q_t. There is some geometric intuition to the two terms $\alpha(t)$ and $\beta(x,t)$ of this profile function. Note that $F(x,t)$ depends on x only through β, as α is independent of x. Let $\widetilde{Z_t}$ denote the cover of Z_t corresponding to the subgroup $\langle\xi\rangle < \pi_1(\Sigma)$. Since the maximal modulus annulus with core curve homotopic to ξ can be lifted to $\widetilde{Z_t}$, one obtains an estimate relating $Ext_\xi \widetilde{Z_t}$ and $Ext_\xi Z_t$ (cf. (McMullen, 1998, p. 21)). The reader may interpret the $\alpha(t)$ term in $F(x,t)$ as determining $Ext_\xi \widetilde{Z_t}$, while the $\beta(x,t)$ term determines the error in this estimate.

Precisely, one has

$$Ext_\xi \widetilde{Z_t} = \frac{4L(t)}{\pi + \theta(t)} = \frac{2}{\alpha(t)}$$

and

$$Ext_\xi Z_t \leq Ext_\xi \widetilde{Z_t}.$$

In other words,

$$Mod(Q_t) \ \leq \ \frac{1}{2\alpha(t)}. \tag{2.19}$$

Since $\frac{1}{2\alpha(t)}$ is the modulus of the rectangle sharing its vertical sides with Q_t and its horizontal sides contained in the lines $\{\operatorname{Im} z = \alpha(t)\}$ and $\{\operatorname{Im} z = -\alpha(t)\}$, $\beta(x,t)$ accounts for the discrepancy in inequality (Equation 2.19).

Some computational details, deferred to §2.8, provide the following:

Lemma 2.7.3. *We have the containment of quadrilaterals $R_1 \subset R_{\frac{1}{2}}$, with the vertical sides of R_1 contained in the vertical sides of $R_{\frac{1}{2}}$.*

Lemma 2.7.4. *We have the containment of quadrilaterals $R_{\frac{2}{5}} \subset R_{\frac{1}{2}}$, with the vertical sides of $R_{\frac{2}{5}}$ contained in the vertical sides of $R_{\frac{1}{2}}$.*

These Lemmas directly imply non-monotonicity of $t \mapsto Mod(Q_t)$.

Proposition 2.7.5. *The function $t \mapsto Mod\,(Q_t)$ is not monotone on the interval $(t_0, 1]$.*

Proof. Note that $t_0 < \frac{2}{5}$, by Lemma 2.10.1. By Lemmas 2.7.3 and 2.7.4 (see Figure 15), we have:

$$Mod\,(Q_1) \leq Mod\left(Q_{\frac{1}{2}}\right)$$
$$Mod\left(Q_{\frac{2}{5}}\right) \leq Mod\left(Q_{\frac{1}{2}}\right)$$

□

Figure 15. The non-monotonicity of $Mod\,(Q_t)$. The quadrilateral R_t is pictured rotated so that its vertical sides appear horizontal.

The non-monotonicity of Proposition 2.7.5 implies the existence of a critical point for the skinning map of the 3-manifold from §2.3, so we are now ready to give the proof of Theorem A.

Theorem 2.7.6. *The pared 3-manifold $M = (H_2, P)$ has a skinning map σ_M that is not one-to-one and has a critical point.*

Proof. The map σ_M, by (Equation 2.10) and Lemma 2.5.1, sends the real 1-dimensional submanifold $\mathcal{R} := \mathcal{R}_{\{\xi,\eta\}}$ to itself. By Lemma 2.5.3 and Proposition 2.7.5, $(Ext_\xi \circ \sigma_M)\,|_{\mathcal{R}}$ is a non-monotonic continuously differentiable function, and thus $(Ext_\xi \circ \sigma_M)\,|_{\mathcal{R}}$ has a critical point. By Lemma 2.5.3, $Ext_\xi|_{\mathcal{R}}$ is a diffeomorphism, and thus σ_M has a critical point. □

2.8 Computational lemmas

The goal of this section is to prove Lemmas 2.7.3 and 2.7.4, which together imply non-monotonicity of $Mod(Q_t)$ (see Proposition 2.7.5). Since the vertical sides of R_t are contained in the lines $\{\mathrm{Re}\, z = 0\}$ and $\{\mathrm{Re}\, z = 1\}$, and the horizontal sides are graphs of functions over $x \in [0,1]$, the containments in these Lemmas are a consequence of inequalities involving the functions $x \mapsto F(x,t)$. In particular, we seek uniform estimates such as $F(x,t_1) < F(x,t_2)$, for $t_1, t_2 \in (t_0, 1]$, and for all $x \in [0,1]$.

Essentially, it is the non-monotonicity of $\alpha(t)$ that accounts for such uniform estimates. However, while α is independent of x, F is not, and in fact we may have $\alpha(t_1) > \alpha(t_2)$ while $F(x_0,t_1) < F(x_0,t_2)$, for some $x_0 \in [0,1]$. We must therefore take more care in estimates on $F(x,t_1) - F(x,t_2)$ by accounting for the effect of β.

While β does depend on x, we can obtain control over

$$\max_{x\in[0,1]} |\beta(x,t_1) - \beta(x,t_2)|.$$

This allows comparisons between $F(x,t_1)$ and $F(x,t_2)$ over all x. Some computational pieces that are easily checked with a computer are collected in Appendix 2.10, where a Mathematica notebook containing a demonstration of these verifications is also indicated.

For ease in presentation, we introduce some notation for functions that will be used throughout this section:

$$p_1(t) = 1 + 8t^2 + 21t^4 - 2t^6 \tag{2.20}$$

$$p_2(t) = 2t^2(1 + t^2)^2 \tag{2.21}$$

$$p_3(t) = 4t^5(3 - t^2) \tag{2.22}$$

In terms of these polynomials, (Equation 2.12) and (Equation 2.13) become:

$$\theta(t) = \cos^{-1}\left(\frac{p_3(t)}{p_2(t)}\right) \tag{2.23}$$

$$L(t) = \frac{1}{2}\cosh^{-1}\left(\frac{p_1(t)}{p_2(t)}\right) \tag{2.24}$$

Recall (Equation 2.17) and (Equation 2.18):

$$\alpha(t) = \frac{\pi + \theta(t)}{2L(t)} \tag{2.25}$$

$$\beta(x,t) = \frac{1}{L(t)}\cos^{-1}\left(\frac{\cosh\left(x\ L(t)\right)}{\cosh\left(L(t)\right)}\right) \tag{2.26}$$

The interested reader can refer to (Equation 2.20–Equation 2.26) to check the computations below. The computational strategy is to first turn inequalities involving the $\cos^{-1}$ and $\cosh^{-1}$ functions into inequalities involving logarithms and constants consisting of algebraic numbers and π. These statements are then reduced to explicit inequalities involving large powers of algebraic numbers, which can be checked with a computer. Indeed, for the patient reader, they could all be checked by hand.

Lemma 2.8.1. *We have* $\alpha(\frac{1}{2}) > \alpha(1)$.

Proof. Calculating $\alpha(1)$ and $\alpha\left(\frac{1}{2}\right)$, and estimating $\cos^{-1}\left(\frac{11}{25}\right)$ using Lemma 2.10.2, we have:

$$\alpha(1) = \frac{\pi}{\cosh^{-1}\left(\frac{7}{2}\right)}$$
$$\alpha\left(\frac{1}{2}\right) = \frac{\pi + \cos^{-1}\left(\frac{11}{25}\right)}{\cosh^{-1}\left(\frac{137}{25}\right)} > \frac{\frac{27\pi}{20}}{\cosh^{-1}\left(\frac{137}{25}\right)}$$

Now referring to Lemma 2.10.6,

$$\left(\frac{7+3\sqrt{5}}{2}\right)^{27} > \left(\frac{137+36\sqrt{14}}{25}\right)^{20}.$$

Taking logarithms and re-arranging, we find

$$\frac{\frac{27}{20}}{\log\left(\frac{137+36\sqrt{14}}{25}\right)} > \frac{1}{\log\left(\frac{7+3\sqrt{5}}{2}\right)}.$$

Since, for $x \geq 1$, $\cosh^{-1} x = \log\left(x + \sqrt{x^2-1}\right)$, we have

$$\frac{\frac{27}{20}}{\cosh^{-1}\left(\frac{137}{25}\right)} > \frac{1}{\cosh^{-1}\left(\frac{7}{2}\right)}.$$

This implies $\alpha\left(\frac{1}{2}\right) > \alpha(1)$, as desired. □

The proof of Lemma 2.8.2 proceeds along similar lines, so we suppress commentary.

Lemma 2.8.2. *We have* $\alpha(\frac{1}{2}) - \alpha(\frac{2}{5}) > \beta(0, \frac{1}{2}) - \beta(0, \frac{2}{5})$.

Proof.

$$\alpha\left(\frac{1}{2}\right)-\alpha\left(\frac{2}{5}\right)>\frac{\frac{27\pi}{20}}{\cosh^{-1}\left(\frac{137}{25}\right)}-\frac{\frac{17\pi}{12}}{\cosh^{-1}\left(\frac{43897}{6728}\right)}\quad\text{by Lemmas 2.10.2 and 2.10.3}$$

$$\begin{aligned}\beta\left(0,\frac{1}{2}\right)-\beta\left(0,\frac{2}{5}\right)&=\frac{2\cos^{-1}\left(\frac{5}{9}\right)}{\cosh^{-1}\left(\frac{137}{25}\right)}-\frac{2\cos^{-1}\left(\frac{2}{3}\right)}{\cosh^{-1}\left(\frac{7}{2}\right)}\\&<\frac{\frac{19\pi}{30}}{\cosh^{-1}\left(\frac{137}{25}\right)}-\frac{\frac{13\pi}{20}}{\cosh^{-1}\left(\frac{43897}{6728}\right)}\quad\text{by Lemmas 2.10.4 and 2.10.5}\end{aligned}$$

$$\left(\frac{137+36\sqrt{14}}{25}\right)^{46}<\left(\frac{43897+225\sqrt{37169}}{6728}\right)^{43}\quad\text{by Lemma 2.10.7}$$

$$\begin{aligned}&\implies\frac{46}{43}=\frac{\frac{13}{20}-\frac{17}{12}}{\frac{19}{30}-\frac{27}{20}}<\frac{\log\left(\frac{43897+225\sqrt{37169}}{6728}\right)}{\log\left(\frac{137+36\sqrt{14}}{25}\right)}=\frac{\cosh^{-1}\left(\frac{43897}{6728}\right)}{\cosh^{-1}\left(\frac{137}{25}\right)}\\&\implies\frac{\frac{27\pi}{20}}{\cosh^{-1}\left(\frac{137}{25}\right)}-\frac{\frac{17\pi}{12}}{\cosh^{-1}\left(\frac{43897}{6728}\right)}>\frac{\frac{19\pi}{30}}{\cosh^{-1}\left(\frac{137}{25}\right)}-\frac{\frac{13\pi}{20}}{\cosh^{-1}\left(\frac{43897}{6728}\right)}\end{aligned}$$

The inequality follows. □

Lemma 2.8.3. *L is monotone decreasing on $(0,1]$.*

Proof. Recall that $L(t) = \frac{1}{2}\cosh^{-1}\left(\frac{p_1(t)}{p_2(t)}\right)$ (see (Equation 2.20), (Equation 2.21), and (Equation 2.24)). Since $t^2 > t^4$ for $t \in (0,1)$, it is clear that $p_1(t) > p_2(t) > 0$. As $\cosh^{-1}$ is monotone increasing on $(1,\infty)$, the lemma will follow from checking that $\frac{p_1(t)}{p_2(t)}$ is monotone decreasing on $(0,1)$. We have:

$$\frac{d}{dt}\left(\frac{p_1(t)}{p_2(t)}\right) = -\frac{(1+5t^2)(1-2t^2+5t^4)}{t^3(1+t^2)^3}$$

As $t^4 - t^2 \geq -\frac{1}{4}$, it is straightforward to check that $\frac{d}{dt}\left(\frac{p_1(t)}{p_2(t)}\right) < 0$ for $t > 0$. □

Lemma 2.8.4. *For $t_1, t_2 \in (0,1)$ such that $t_2 > t_1$, and for all $x \in (0,1)$, we have*

$$0 < \beta(x,t_2) - \beta(x,t_1) < \beta(0,t_2) - \beta(0,t_1).$$

Proof. We compute:

$$\frac{\partial^2\beta}{\partial x \partial t}(x,t) = \left[\frac{-L'(t)\sinh(2L(t))\cosh(xL(t))}{2\left(\cosh^2 L(t) - \cosh^2(xL(t))\right)^{\frac{3}{2}}}\right] \cdot [x\tanh L(t) - \tanh(xL(t))]$$

By Lemma 2.8.3, $L'(t) < 0$ for all $t \in (0,1)$ and the first bracketed term is positive.

The function $x^{-1}\tanh x$ is decreasing for $x > 0$, so for $x \in (0,1)$ we have

$$\frac{\tanh(xL(t))}{xL(t)} > \frac{\tanh L(t)}{L(t)}$$

and the second bracketed term is negative. We conclude that $\frac{\partial^2\beta}{\partial x\partial t}(x,t) < 0$ for all $t, x \in (0,1)$.

Thus on the domain $(0,1) \times (0,1)$, we have:

(i) For fixed x, the function $\frac{\partial \beta}{\partial x}(x,t)$ is decreasing in t.

(ii) For fixed t, the function $\frac{\partial \beta}{\partial t}(x,t)$ is decreasing in x.

A straightforward computation shows that

$$\lim_{x \to 1} \frac{\partial \beta}{\partial t}(x,t) = 0.$$

By (ii), this limit is the infimum of $\frac{\partial \beta}{\partial t}(x,t)$ over $x \in (0,1)$. Thus $\frac{\partial \beta}{\partial t}(x,t) > 0$ and $\beta(x,t)$ is increasing in t.

Suppose that $t_2 > t_1$. By (i),

$$\frac{\partial}{\partial x}\left(\beta(x,t_2) - \beta(x,t_1)\right) < 0$$

for all $x \in (0,1)$. We now have that $\beta(x,t_2) - \beta(x,t_1)$ is positive and monotone decreasing in x. The result follows. □

Proof of Lemma 2.7.3. For any $x \in (0,1)$, combining Lemmas 2.8.1 and 2.8.4,

$$\begin{aligned} F\left(x, \frac{1}{2}\right) - F(x,1) &= \left(\alpha\left(\frac{1}{2}\right) - \alpha(1)\right) + \left(\beta(x,1) - \beta\left(x, \frac{1}{2}\right)\right) \\ &> \left(\alpha\left(\frac{1}{2}\right) - \alpha(1)\right) \;>\; 0. \end{aligned}$$

Referring to the description of R_t in Lemma 2.7.2, this inequality implies the containment $R_1 \subset R_{\frac{1}{2}}$ with containment of vertical sides as desired. □

Proof of Lemma 2.7.4. For any $x \in (0,1)$, combining Lemmas 2.8.2 and 2.8.4,

$$\begin{aligned} F\left(x,\frac{1}{2}\right) - F\left(x,\frac{2}{5}\right) &= \left(\alpha\left(\frac{1}{2}\right) - \alpha\left(\frac{2}{5}\right)\right) - \left(\beta\left(x,\frac{1}{2}\right) - \beta\left(x,\frac{2}{5}\right)\right) \\ &> \left(\alpha\left(\frac{1}{2}\right) - \alpha\left(\frac{2}{5}\right)\right) - \left(\beta\left(0,\frac{1}{2}\right) - \beta\left(0,\frac{2}{5}\right)\right) \\ &> \ 0. \end{aligned}$$

The containment $R_{\frac{2}{5}} \subset R_{\frac{1}{2}}$ follows as above. □

2.9 A family of finite covers of M

Using Theorem 2.7.6 and Proposition 2.2, we construct a family of pared 3-manifolds with non-immersion skinning maps, on Teichmüller spaces of arbitrarily high dimension.

Corollary 2.9.1. *There exists a family of pared 3-manifolds $\{M_n = (H_n, P_n)\}_{n=2}^{\infty}$, each admitting a geometrically finite hyperbolic structure, satisfying:*

- *For n even, the paring locus P_n consists of two rank-1 cusps and the boundary has topological type $\partial M_n \cong \Sigma_{n-2,4}$.*
- *For n odd, the paring locus P_n consists of three rank-1 cusps and the boundary has topological type $\partial M_n \cong \Sigma_{n-3,6}$.*

For each integer $n \geq 2$, the skinning map σ_{M_n} has a critical point.

Proof. Recall that H_g indicates the closed genus g handlebody, and recall the identification $\pi_1 H_2 \cong \langle A, B\rangle$ as in Figure 2. For each genus $g \geq 2$, consider the cyclic cover $H_g \to H_2$ corresponding to the subgroup $\langle A, B^{g-1}\rangle$. The family $\{M_n\}_{n=0}^{\infty}$ is generated from this family of covers, as we now describe.

The annuli in P lift to annuli $P_g \subset \Sigma_g$. Denote $\Sigma' := \Sigma_g \setminus P_g$ (and recall $\Sigma = \Sigma_2 \setminus P$). We choose a fundamental domain for the covering $\Sigma' \to \Sigma$, as pictured in Figure 16, and denote it F. Precisely, choose a two-holed torus fundamental domain for $\Sigma_g \to \Sigma_2$, include only one of the two boundary components, and take the complement of P_g. Because F is connected and Σ' is the union of finitely many copies of F, it is clear that the cover Σ' is connected.

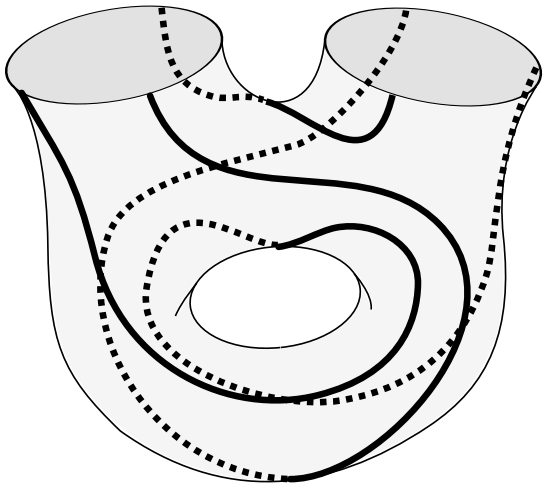

Figure 16. The fundamental domain F for the covering $\Sigma' \to \Sigma$.

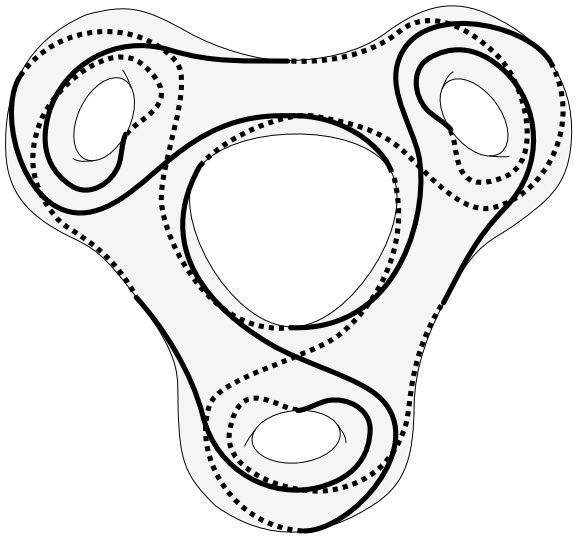

Figure 17. The cover M_4, the pared 3-manifold (H_4, P_4). Note that P_4 has two connected components.

The intersection of P_g with F consists of three connected components. The gluing pattern of $g-1$ copies of F, in order to build Σ', matches these components in a straightforward fashion. Upon inspection it is clear that P_{2g} has two components and P_{2g+1} has three. (See Figure 17 for a picture of P_4).

We consider the pared 3-manifolds $M_n := (H_n, P_n)$. The boundary ∂M_{2g} is a genus $2g$ surface, with two non-separating disk-busting annuli deleted, and is thus homeomorphic to $\Sigma_{2g-2,4}$. Similarly, $\partial M_{2g+1} \cong \Sigma_{2g-2,6}$. Applying Proposition 2.2.1, the corollary follows. □

2.10 Appendix: Numerical lemmas

Lemmas 2.10.1 through 2.10.7, as they can clearly be verified by a finite list of computations involving integers and standard functions, are available in a Mathematica notebook[1].

Lemma 2.10.1. $\frac{1}{2}\left(5+3\sqrt{3}-\sqrt{44+26\sqrt{3}}\right) < \frac{2}{5}$

Lemma 2.10.2. $\frac{11}{25} < \cos(\frac{7\pi}{20})$

Lemma 2.10.3. $\frac{1136}{4205} > \cos(\frac{5\pi}{12})$

Lemma 2.10.4. $\frac{5}{9} > \cos(\frac{19\pi}{60})$

Lemma 2.10.5. $\frac{116}{225} < \cos(\frac{13\pi}{40})$

Lemma 2.10.6. $25^{20}(7+3\sqrt{5})^{27} > 2^{27}(137+36\sqrt{14})^{20}$

Lemma 2.10.7. $6728^{43}(137+36\sqrt{14})^{46} < 25^{46}(43897+225\sqrt{37169})^{43}$

---

[1] as an ancillary file at arxiv.org/abs/1212.6210

CHAPTER 3

MAPPING CLASS GROUP ORBITS OF MAXIMAL COMPLETE 1-SYSTEMS AND THEIR DUAL CUBE COMPLEXES

This chapter presents work that was developed jointly with Tarik Aougab (see p. 106).

After motivating the dual cube complex with a discussion of a dual square complex of a collection of curves in §3.1, we describe the construction in detail in §3.2. Two general theorems about dual cube complexes, Theorems C and D, are proved in §3.3 and §3.4, respectively. In §3.5, we describe a general construction of exponentially many maximal complete 1-systems. Towards Theorem B, we examine mapping class group orbits within our construction, postponing the proof of an important technical step (Proposition 3.5.2) to §3.6.

3.1 Background and motivation

Let $S = S_g$ be a compact oriented surface of genus g, and let $\Lambda = \{\gamma_1, ..., \gamma_n\}$ be a collection of free homotopy classes of closed curves on S. Recall that a collection of curves form a *bigon* if there is an embedded disk in S whose boundary is the union of two arcs of the curves. A *minimal position realization* of Λ is a set $\lambda = \{\eta_1, ..., \eta_n\}$ such that:

(i) Each $\eta_i : S^1 \to S$ is a smooth immersion in the free homotopy class γ_i.

(ii) The union $\bigcup_i \eta_i(S^1)$ forms no bigons.

(iii) The immersed submanifolds $\{\eta_1(S^1), ..., \eta_n(S^1)\}$ intersect only at transverse double points.

We will refer to minimal position realizations simply as *realizations*. In everything that follows, we suppose that $\lambda = \{\eta_1, ..., \eta_n\}$ is a realization of Λ. Condition (ii) above implies that λ minimizes the sum of the pairwise geometric intersection numbers of the curves in Λ (Farb and Margalit, 2012, Prop. 1.7, p. 31). See (Farb and Margalit, 2012, Ch. 1) for background on curves on surfaces.

Let $\widetilde{\lambda}$ indicate the set of the lifts of the elements of λ to the universal cover $\widetilde{S}$, so that the elements of $\widetilde{\lambda}$ are curves in $\widetilde{S}$. When clear from context, we may also use $\widetilde{\lambda}$ to refer to the union of the elements of $\widetilde{\lambda}$. The union of the curves in $\widetilde{\lambda}$ may be considered as an embedded graph $\widetilde{G(\lambda)} \subset \widetilde{S}$. Condition (iii) above guarantees that each vertex of this graph has valence four, and condition (ii) implies that every pair of curves in $\widetilde{\lambda}$ intersect at most once (see (Farb and Margalit, 2012, Lemma 1.8, p. 30)). As a consequence, a pair of lifts in $\widetilde{\lambda}$ intersect in $\widetilde{S}$ if and only if the pairs of limit points they determine in the boundary $\partial\widetilde{S}$ are *linked*, though we will not utilize this point of view. (A pair of pairs of points of S^1 are linked when one pair of points is in distinct components of the complement of the other pair).

The complementary regions of $\widetilde{S} \setminus \widetilde{G(\lambda)}$ are diffeomorphic to polygons (possibly infinite-sided and unbounded), where the sides of the boundary polygons are arcs of curves in $\widetilde{\lambda}$. Let $\widetilde{G(\lambda)^*}$ denote the dual graph of $\widetilde{G(\lambda)}$, i.e. the graph whose vertices are in correspondence with the complementary regions of $\widetilde{G(\lambda)}$, and whose edges correspond to arcs of the boundary polygons shared between complementary regions.

As each vertex $v \in \widetilde{G(\lambda)}^{(0)}$ is valence four, there is a map of the 1-skeleton of the square $\overline{v} : ([-1,1] \times [-1,1])^{(1)} \to \widetilde{G(\lambda)^*}$. The map $\overline{v}$ takes the four sides of the square to the four

edges of $\widetilde{G(\lambda)^*}$ that correspond to the arcs of curves in $\widetilde{\lambda}$ meeting at v, as in Figure 18. We construct a cell complex by using each $\overline{v}$ to glue the square $[-1,1]\times[-1,1]$ to $\widetilde{G(\lambda)^*}$ along its 1-skeleton, for each vertex $v \in \widetilde{G(\lambda)}^{(0)}$:

$$\widetilde{\mathcal{S}(\lambda)} := \left(\widetilde{G(\lambda)^*} \cup \bigsqcup_{v\in\widetilde{G(\lambda)}^{(0)}} [-1,1]\times[-1,1] \right) / \overline{v}$$

When Λ consists of disjoint curves, $\widetilde{G(\lambda)}^{(0)}$ is empty, and $\widetilde{\mathcal{S}(\lambda)}$ is the dual tree $\widetilde{G(\lambda)^*}$.

By choosing a point in each complementary component of $\widetilde{S}\setminus\widetilde{\lambda}$, we can construct an embedding of $\widetilde{G(\lambda)^*}\to\widetilde{S}$, transverse to $\widetilde{\lambda}$. Extend this map over the 2-skeleton of $\widetilde{\mathcal{S}(\lambda)}$ so that it is a local homeomorphism on the interior of the squares $[-1,1]\times[-1,1]$ of $\widetilde{\mathcal{S}(\lambda)}$. We will refer to this as 'the natural embedding' $\widetilde{\mathcal{S}(\lambda)}\to\widetilde{S}$.

From the point of view of this embedding, each edge of $\widetilde{G(\lambda)^*}$ has an element of $\widetilde{\lambda}$ passing through its midpoint, and each square in $\widetilde{\mathcal{S}(\lambda)}$ has a pair of lifts joining opposite sides and intersecting in the center. In fact, $\widetilde{\mathcal{S}(\lambda)}$ is an example of a topological space with slightly more structure than a cell complex, as its 2-cells are all copies of the square $[-1,1]\times[-1,1]$. This is an example of a cube complex.

Recall that a *cube* (of dimension n) is the cell complex $\square^n := [-1,1]^n$, and a *cube complex* is the quotient of disjoint cubes by a gluing map which is a Euclidean isometry on each cell. In such a complex, a *maximal cube* is a cube which is not contained in a higher dimensional cube. The *dimension* of a cube complex is the supremum of the dimensions of its cubes. A

square complex is a two-dimensional cube complex. See (Sageev, 2012) for an introduction to cube complexes.

The *local hyperplanes* of an n-cube are the intersections of $\square^n$ with coordinate planes. In the quotient cube complex, the local hyperplanes are well-defined: As the gluing maps are Euclidean isometries, they preserve the locus of points halfway between two faces. We introduce an equivalence relation on local hyperplanes in a cube complex: Two local hyperplanes are *hyperplane equivalent* if they intersect along the 1-skeleton of the cube complex, and a *hyperplane* of a cube complex is the hyperplane equivalence class of a local hyperplane. An isomorphism of cube complexes is a homeomorphism that is a cellular Euclidean isometry, which thus preserves hyperplanes. We will denote the hyperplanes of a cube complex $\mathcal{C}$ by $\mathcal{H}_{\mathcal{C}}$.

Notice that 1-dimensional cube complexes are graphs and trees, whose hyperplanes (and local hyperplanes) are midpoints of edges. The complex $\widetilde{\mathcal{S}(\lambda)}$ constructed above is an example of a square complex, in which the hyperplanes $\mathcal{H}_{\widetilde{\mathcal{S}(\lambda)}}$ of $\widetilde{\mathcal{S}(\lambda)}$ are naturally in correspondence with $\widetilde{\lambda}$ (see Lemma 3.1.2).

We note a simple useful fact, which we present as a lemma to make it easily recalled:

Lemma 3.1.1. *If a group acts by automorphisms on a cube complex so that any element preserving a cube fixes it pointwise, then the quotient is a cube complex as well.*

The action of $\pi_1 S$ on $\widetilde{S}$ by deck transformations induces an action of $\pi_1 S$ on the graph $\widetilde{G(\lambda)}$ and on its complementary components $\widetilde{S} \setminus \widetilde{G(\lambda)}$. This induces an action of $\pi_1 S$ on $\widetilde{G(\lambda)^*}$, the 1-skeleton of $\widetilde{\mathcal{S}(\lambda)}$, which extends over the 2-skeleton by construction. We arrive at an action of $\pi_1 S$ on $\widetilde{\mathcal{S}(\lambda)}$ by cube complex automorphisms, making the natural embedding $\pi_1 S$-equivariant.

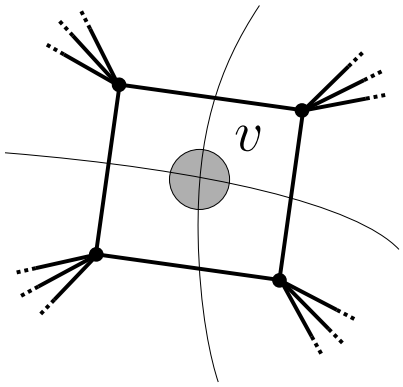

Figure 18. A vertex v of $\widetilde{G(\lambda)}$ determines a map from the 1-skeleton of the square to $\widetilde{G(\lambda)^*}$

Lemma 3.1.2. *The natural embedding $\widetilde{\mathcal{S}(\lambda)} \to \widetilde{S}$ induces a $\pi_1 S$-equivariant correspondence of the hyperplanes $\mathcal{H}_{\widetilde{\mathcal{S}(\lambda)}}$ with $\widetilde{\lambda}$.*

Proof. Under the natural embedding of $\widetilde{\mathcal{S}(\lambda)}$ into $\widetilde{S}$, each 1 and 2-cube of $\widetilde{\mathcal{S}(\lambda)}$ has local hyperplanes that naturally correspond to curves in $\widetilde{\lambda}$: Each 1-cube intersects a unique curve, and each 2-cube has two local hyperplanes corresponding to the curves of $\widetilde{\lambda}$ running through the 2-cube. This provides a $\pi_1 S$-equivariant map from the local hyperplanes of $\widetilde{\mathcal{S}(\lambda)}$ to the curves in $\widetilde{\lambda}$. Since hyperplane equivalent local hyperplanes correspond to the same curve in $\widetilde{\lambda}$, this map induces a $\pi_1 S$-equivariant map $\Psi : \mathcal{H}_{\widetilde{\mathcal{S}(\lambda)}} \to \widetilde{\lambda}$. Since the 1-skeleton of $\widetilde{\mathcal{S}(\lambda)}$ is the dual graph $\widetilde{G(\lambda)^*}$, the map Ψ is surjective.

If a 1-cube is maximal, its unique local hyperplane corresponds to a curve disjoint from the other curves of $\widetilde{\lambda}$. As the curve intersects no other edges of $\widetilde{G(\lambda)^*}$, there are no other local hyperplanes that determine this curve. If there are two 2-cubes that have local hyperplanes that

both correspond to γ, then one may follow the curve γ from one of these 2-cubes to the other, passing through finitely many 2-cubes on the way. At each step, the local hyperplanes that correspond to γ meet along the 1-skeleton of $\widetilde{\mathcal{S}(\lambda)}$. This implies that these local hyperplanes are hyperplane equivalent in $\widetilde{\mathcal{S}(\lambda)}$, and Ψ is injective. □

Lemma 3.1.3. *The action of $\pi_1 S$ on $\widetilde{\mathcal{S}(\lambda)}$ satisfies:*

1. *The stabilizer of any 2-cube is trivial.*
2. *If $g \in \pi_1 S$ stabilizes a 1-cube in $\widetilde{\mathcal{S}(\lambda)}$, then g fixes the 1-cube pointwise.*

Proof. The first fact follows since an element stabilizing a 2-cube must fix the corresponding vertex of $\widetilde{G(\lambda)}$, violating the freeness of the action of $\pi_1 S$ on $\widetilde{S}$. As for the second, a Euclidean isometry of a 1-cube either fixes it pointwise or exchanges its endpoints. Each 1-cube in $\widetilde{\mathcal{S}(\lambda)}$ is dual to a curve in $\widetilde{\lambda}$. The subgroup of $\pi_1 S$ that fixes this curve $\gamma \in \widetilde{\lambda}$ is cyclic, and none of its elements flip the half-spaces $\widetilde{S} \setminus \gamma$, as S is orientable. □

Note that Lemma 3.1.3 implies that the only source of non-trivial stabilizers of $\pi_1 S$ on $\widetilde{\mathcal{S}(\lambda)}$ are from the action of $\pi_1 S$ on the 1-skeleton $\widetilde{G(\lambda)^*}$, as when Λ consists of disjoint curves. By Lemma 3.1.1, the quotient $\mathcal{S}(\lambda) := \widetilde{\mathcal{S}(\lambda)}/\pi_1 S$ is a square complex. As the hyperplanes $\mathcal{H}_{\widetilde{\mathcal{S}(\lambda)}}$ are in $\pi_1 S$-equivariant correspondence with the curves in $\widetilde{\lambda}$, the hyperplanes of the quotient $\mathcal{S}(\lambda)$ are in correspondence with the elements of λ.

Recall that the collection of curves Λ is *filling* if the complement of the curves in S, in any realization, is a disjoint union of disks.

Lemma 3.1.4. *If Λ is filling, the action of $\pi_1 S$ on $\widetilde{\mathcal{S}(\lambda)}$ is free.*

Proof. When the collection Λ is filling, the complementary components of λ in S are all finite-sided polygons, and the complementary components of $\widetilde{S} \setminus \widetilde{\lambda}$ are all finite-sided polygons as well. Suppose that $g \in \pi_1 S$ fixed a vertex $v \in \widetilde{\mathcal{S}(\lambda)}$. By definition this implies that g preserves the corresponding compact polygon in $\widetilde{S} \setminus \widetilde{\lambda}$, and by Brouwer's fixed point theorem this contradicts freeness of the action of $\pi_1 S$ on $\widetilde{S}$. □

Proposition 3.1.5. *If the collection of curves Λ is filling, then, for any realization λ for Λ, there is a homeomorphism of the square complex $\mathcal{S}(\lambda)$ with S taking hyperplanes of $\mathcal{S}(\lambda)$ to λ.*

Proof. When Λ is filling, every maximal cube of the square complex $\widetilde{\mathcal{S}(\lambda)}$ is 2-dimensional: A 1-dimensional maximal cube of $\widetilde{\mathcal{S}(\lambda)}$ is dual to a unique curve in $\widetilde{\lambda}$ that is necessarily disjoint from the other curves in $\widetilde{\lambda}$. In this case, the image of this curve on S fills some subsurface of S, whose boundary curves would produce essential annuli in the complement of λ.

This implies that every 1-cube in $\widetilde{\mathcal{S}(\lambda)}$ bounds two 2-cubes: Suppose e is a 1-cube, i.e. an edge in $\widetilde{G(\lambda)^*}$. Because e is not maximal, it bounds at least one 2-cube. This 2-cube corresponds to an intersection v of two lifts in $\widetilde{\lambda}$, where e is dual to one of the arcs incident to v. We denote the ray rooted at v, in the direction of e, by β. Since β covers a curve in λ infinitely many times, there is a 'first' vertex of $\widetilde{G(\lambda)^*}$ along the ray β after e. This implies that e is incident to two 2-cubes.

As these 2-cubes necessarily map to opposite sides of e, the natural embedding $\widetilde{\mathcal{S}(\lambda)} \to \widetilde{S}$ is now a $\pi_1 S$-equivariant embedding that is a local homeomorphism, which is thus a $\pi_1 S$-equivariant homeomorphism.

By Lemma 3.1.4, the action of $\pi_1 S$ on $\widetilde{\mathcal{S}(\lambda)}$ is free. In fact, when the action of a group by automorphisms on a locally finite cube complex is free, it is necessarily properly discontinuous: A compact set K intersects finitely many maximal cubes, by local finiteness. As cube complex automorphisms take maximal cubes to maximal cubes, infinitely many group elements g satisfying $g \cdot K \cap K \neq \emptyset$ imply the existence of a non-trivial group element preserving a maximal cube, and Brouwer's fixed point theorem again applies.

Thus $\pi_1 S$ acts on $\widetilde{\mathcal{S}(\lambda)}$ freely and properly discontinuously. Since $\widetilde{\mathcal{S}(\lambda)}$ is homeomorphic to $\widetilde{S}$ and $\pi_1 S$ acts freely and properly discontinuously, the quotient complex is homeomorphic to the surface S. By Lemma 3.1.2, this homeomorphism takes hyperplanes of $\mathcal{S}(\lambda)$ to λ. □

The square complex $\mathcal{S}(\lambda)$ is not an invariant of Λ, as different realizations may yield non-isomorphic complexes. For example, the presence of a triangle in the complement of λ allows a Reidemeister type III move, creating a new realization but changing the isomorphism types of $\widetilde{G(\lambda)^*}$ and $\widetilde{\mathcal{S}(\lambda)}$, as shown in Figure 19.

While the square complex $\mathcal{S}(\lambda)$ depends essentially on the realization λ, we now describe the construction of a related cube complex 'dual' to Λ which is independent of realization. Originally due to Sageev, this produces an isometric action of $\pi_1 S$ on a finite-dimensional cube complex, which we denote by $\widetilde{\mathcal{C}(\Lambda)}$, with quotient $\mathcal{C}(\Lambda) := \widetilde{\mathcal{C}(\Lambda)}/\pi_1 S$, a cube complex with finitely many maximal cubes. Though it will be unnecessary in this work, in fact $\mathcal{C}(\Lambda)$ is non-positively curved, and this construction can be placed in a considerably more general context (see (Sageev, 1995), (Sageev, 2012), and (Chatterji and Niblo, 2005)).

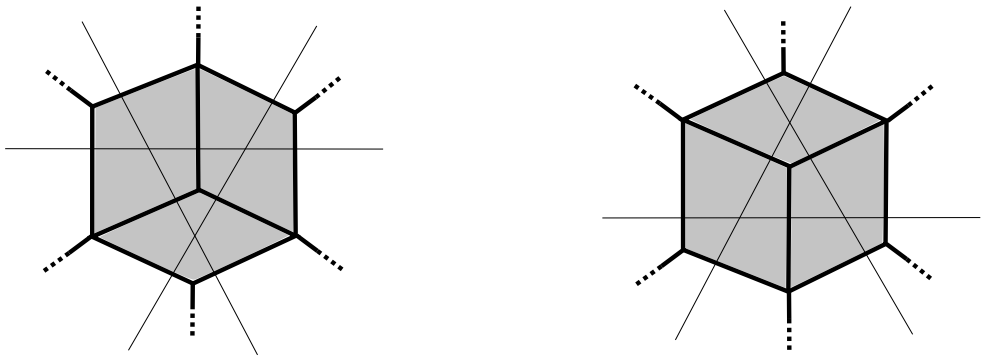

Figure 19. Changing the realization λ by a Reidemeister type III move changes the isomorphism type of $\mathcal{S}(\lambda)$.

3.2 Sageev's construction

Fix a choice of realization λ of Λ, with lifts $\widetilde{\lambda}$ as before. We now describe the construction of a cube complex $\widetilde{\mathcal{C}(\lambda)}$, which we show below is independent of the choice of realization. Each element $\gamma \in \widetilde{\lambda}$ separates $\widetilde{S} \setminus \gamma$ into two connected components. Fix a choice of identification of these two half-spaces with $\{1, -1\}$ for each $\gamma \in \widetilde{\lambda}$. We will refer to a choice of one of these two as a *labeling* of γ, and a labeling of $\widetilde{\lambda}$ is a labeling for each of the curves of $\widetilde{\lambda}$. Identifying 2 with $\{1, -1\}$, we say a labeling $v \in 2^{\widetilde{\lambda}}$ is *admissible* if the half-spaces $v(\alpha)$ and $v(\beta)$ intersect for every pair of curves $\alpha, \beta \in \widetilde{\lambda}$. Note that the coordinates of a labeling are naturally identified with the curves in $\widetilde{\lambda}$.

Let $\mathcal{V}_0 \subset 2^{\widetilde{\lambda}}$ denote the collection of admissible labelings. For each connected region $U \subset \widetilde{S} \setminus \widetilde{\lambda}$, there is an admissible labeling $v_\lambda(U) \in \mathcal{V}_0$: For each curve in $\widetilde{\lambda}$, $v_\lambda(U)$ chooses the half-space containing U. Let Y be the graph whose vertex set is $\mathcal{V}_0$, where two labelings are

joined by an edge when they differ on exactly one element of $\widetilde{\lambda}$. Hence the edges of Y are in correspondence with the elements of $\widetilde{\lambda}$.

Choose a connected region $U \subset \widetilde{S} \setminus \widetilde{\lambda}$, and let $\widetilde{\mathcal{C}(\lambda)}^{(1)}$ denote the connected component of $v_\lambda(U)$ in Y. This will be the 1-skeleton of the cube complex $\widetilde{\mathcal{C}(\lambda)}$. The choice of connected region $U \subset \widetilde{S} \setminus \widetilde{\lambda}$ is not essential:

Lemma 3.2.1. *For any connected region $V \subset \widetilde{S} \setminus \widetilde{\lambda}$, the labeling $v_\lambda(V)$ is a vertex of $\widetilde{\mathcal{C}(\lambda)}^{(1)}$.*

Proof. Consider a compact arc α from U to V, transverse to $\widetilde{\lambda}$. Each intersection of α with $\widetilde{\lambda}$ corresponds to an edge in $\widetilde{\mathcal{C}(\lambda)}$. These finitely many edges join $v_\lambda(U)$ to $v_\lambda(V)$. □

Suppose there is a set of distinct labelings $\nu = \{v_1, ..., v_{2^k}\} \subset \widetilde{\mathcal{C}(\lambda)}^{(0)}$, for $k \geq 2$, that are equal on $\widetilde{\lambda} \setminus \{\gamma_1, ..., \gamma_k\}$. The coordinates of the v_i on the set $\{\gamma_1, ..., \gamma_k\}$ provide an identification of ν with the vertices of the k-cube $\left(\square^k\right)^{(0)}$. By design $\left(\square^k\right)^{(1)}$ has an edge between two vertices precisely when the corresponding labelings v_i and v_j are joined by an edge in $\widetilde{\mathcal{C}(\lambda)}^{(1)}$ – namely, when they differ in exactly one coordinate. This provides an embedding of the 1-skeleton of a k-cube $\overline{\nu} : \left(\square^k\right)^{(1)} \to \widetilde{\mathcal{C}(\lambda)}^{(1)}$.

In analogy with the construction of $\widetilde{\mathcal{S}(\lambda)}$, we may now construct a cube complex by using the maps $\overline{\nu}$ to glue in k-cubes to $\widetilde{\mathcal{C}(\lambda)}$. For brevity below we let $E_k \subset \widetilde{\mathcal{C}(\lambda)}^{(0)}$ denote the set of collections of 2^k distinct labelings that are equal in the complement of k coordinates. For $\nu \in E_k$, we refer to these k distinguished coordinates as the *support* of ν, $\mathrm{supp}(\nu)$. Let

$$\widetilde{\mathcal{C}(\lambda)} := \left(\widetilde{\mathcal{C}(\lambda)}^{(1)} \cup \bigsqcup_{\substack{\nu \in E_k, \\ k \geq 2}} \square^k \right) / \sim .$$

When useful for clarity we may write $\square^k_\nu$ for the k-cube of $\mathcal{C}(\lambda)$ whose support is ν. The equivalence relation $\sim$ in the definition of $\widetilde{\mathcal{C}(\lambda)}$ identifies $\left(\square^k_\nu\right)^{(1)}$ to $\widetilde{\mathcal{C}(\lambda)}^{(1)}$ via the map $\overline{\nu}$ for each $\nu \in E_k$, and the following inclusion relation: Suppose we are given $\nu \in E_k$ and $\mu \in E_l$, satisfying $\mu \subsetneq \nu$. In this case, there are $k-l$ coordinates in $\mathrm{supp}(\nu) \setminus \mathrm{supp}(\mu)$. The coordinates of μ on $\mathrm{supp}(\nu) \setminus \mathrm{supp}(\mu)$ define a unique l-cube face of $\square^k_\nu$, which we identify with $\square^l_\mu$ via the equivalence relation $\sim$.

The action of $\pi_1 S$ on $\widetilde{\lambda}$ induces a permutation action of $\pi_1 S$ on $2^{\widetilde{\lambda}}$: For $g \in \pi_1 S$ and admissible labeling $v \in \widetilde{\mathcal{C}(\lambda)}^{(0)}$, the labeling $g \cdot v \in \widetilde{\mathcal{C}(\lambda)}^{(0)}$ is given by $(g \cdot v)(\gamma) = g^{-1} \cdot v(g \cdot \gamma)$ for each $\gamma \in \widetilde{\lambda}$. Since the action on $\widetilde{S}$ is by deck transformations, it is straightforward to check that the action of $\pi_1 S$ preserves the admissible set $\mathcal{V}_0$ and induces an action by graph automorphisms on the 1-skeleton $\widetilde{\mathcal{C}(\lambda)}^{(1)}$. Finally, this action extends to $\widetilde{\mathcal{C}(\lambda)}$ via a permutation action of $\pi_1 S$ on $\bigsqcup_{k=2}^{\infty} \bigsqcup_{\nu \in E_k} \square^k_\nu$: For $g \in \pi_1 S$, we have $\nu \in E_k$ if and only if $g \cdot \nu \in E_k$, and we let $g \cdot \square^k_\nu = \square^k_{g \cdot \nu}$. This is an isomorphism of cube complexes that descends to $\widetilde{\mathcal{C}(\lambda)}$ by construction.

Lemma 3.2.2. *The square complex $\widetilde{\mathcal{S}(\lambda)}$ is a $\pi_1 S$-invariant square subcomplex of $\widetilde{\mathcal{C}(\lambda)}$.*

Proof. Recall the dual graph $\widetilde{G(\lambda)^*}$ as in §3.1, whose vertices are in correspondence with the connected components of $\widetilde{S} \setminus \widetilde{G(\lambda)}$. By Lemma 3.2.1, these components determine vertices of $\widetilde{\mathcal{C}(\lambda)}^{(0)}$, producing a map from $\left(\widetilde{G(\lambda)^*}\right)^{(0)}$ to $\widetilde{\mathcal{C}(\lambda)}^{(0)}$, to which we refer below.

Recall also that $\widetilde{\mathcal{S}(\lambda)}$ is formed by gluing 2-cubes to the graph $\widetilde{G(\lambda)^*}$ along squares in the graph that correspond to vertices of $\widetilde{G(\lambda)}$. The map above extends to a $\pi_1 S$-equivariant embedding of $\widetilde{G(\lambda)^*}$ to the 1-skeleton of $\widetilde{\mathcal{C}(\lambda)}$: Edges of $\widetilde{G(\lambda)^*}$ correspond to regions $U, V \subset \widetilde{S} \setminus \widetilde{\lambda}$ such that $v_\lambda(U)$ and $v_\lambda(V)$ differ in exactly one coordinate, in which case these vertices are

connected by an edge in $\widetilde{\mathcal{C}(\lambda)}^{(1)}$. This map extends further, $\pi_1 S$-equivariantly, over the 2-cubes in $\widetilde{\mathcal{S}(\lambda)}$: A square of $\widetilde{\mathcal{S}(\lambda)}$ has vertices in $\widetilde{G(\lambda)^*}$ (see Figure 18) that are sent to four labelings that pairwise differ on at most 2 elements of $\widetilde{\lambda}$, which determines a square in $\widetilde{\mathcal{C}(\lambda)}$. □

Recall that $\mathcal{H}_{\widetilde{\mathcal{C}(\lambda)}}$ indicates the set of hyperplanes of $\widetilde{\mathcal{C}(\lambda)}$. As with Lemma 3.1.2, we have:

Proposition 3.2.3. *There is a $\pi_1 S$-equivariant correspondence of $\mathcal{H}_{\widetilde{\mathcal{C}(\lambda)}}$ with $\widetilde{\lambda}$. A set of k hyperplanes mutually intersect in a k-cube of $\widetilde{\mathcal{C}(\lambda)}$ if and only if the corresponding k curves in $\widetilde{\lambda}$ mutually intersect.*

Proof. In the above construction a k-cube $\square^k_\nu \subset \widetilde{\mathcal{C}(\lambda)}$ arises from k curves of $\widetilde{\lambda}$ forming the support of ν, a set of 2^k admissible labelings. Each local hyperplane of $\square^k_\nu$ corresponds to one coordinate of the cube being held constant, i.e. one curve in the support of ν. Consider the map Ψ_0 that sends a local hyperplane of a k-cube $\square^k_\nu$ to the unique curve in $\mathrm{supp}(\nu)$ whose coordinate is constant on the local hyperplane.

Suppose two local hyperplanes are hyperplane equivalent in $\widetilde{\mathcal{C}(\lambda)}$, i.e. they meet along the 1-skeleton of $\widetilde{\mathcal{C}(\lambda)}$. The edge e they meet along crosses each local hyperplane, and in doing so there is a unique curve in $\widetilde{\lambda}$ whose labeling changes, which is thus the coordinate on which the local hyperplanes were constant. Thus Ψ_0 induces a well-defined $\pi_1 S$-equivariant map $\Psi : \mathcal{H}_{\widetilde{\mathcal{C}(\lambda)}} \to \widetilde{\lambda}$. It is immediate that this map is surjective, as each curve in $\widetilde{\lambda}$ is a hyperplane of a 1-cube in $\widetilde{G(\lambda)^*}$.

Injectivity is slightly more subtle. In analogy with the proof of Lemma 3.1.2, if two local hyperplanes both have Ψ_0-image γ, then they can be 'traced' to each other by 'following' the

curve γ in $\widetilde{S}$. We will conclude that these two local hyperplanes are hyperplane equivalent, so that there is in fact a unique hyperplane of $\widetilde{\mathcal{C}(\lambda)}$ corresponding to γ, and Ψ is injective. Note that in the proof of Lemma 3.1.2 we had the natural embedding $\widetilde{\mathcal{S}(\lambda)} \to \widetilde{S}$ to work with, which is absent in this higher-dimensional setting.

Suppose two local hyperplanes H_ν and H_μ, contained in cubes $\square^k_\nu$ and $\square^l_\mu$ respectively, both have $\Psi_0(H_\nu) = \Psi_0(H_\mu) = \gamma$. In this case, $\nu \in E_k$ and $\mu \in E_l$ have $\gamma \in \text{supp}(\nu) \cap \text{supp}(\mu) \subset \widetilde{\lambda}$. The set $\text{supp}(\nu) \cup \text{supp}(\mu) \setminus \{\gamma\} \subset \widetilde{\lambda}$ is a finite set of lifts intersecting γ. Choose an identification of γ with $\mathbb{R}$. Since every pair of curves in $\widetilde{\lambda}$ intersect at most once, we can find points p^+ and p^- in γ, in the complement of the rest of the curves in $\widetilde{\lambda} \setminus \{\gamma\}$, so that all intersections of γ with the curves in $\text{supp}(\nu) \cup \text{supp}(\mu) \setminus \{\gamma\}$ are contained in the interval $[p^-, p^+]$. Enumerate the intersections in increasing order as $p_1, p_2, ..., p_n$ and the curves of $\widetilde{\lambda} \setminus \{\gamma\}$ that they lie on as $\gamma_1, ..., \gamma_n$.

Among $\gamma_1, ..., \gamma_n$ there is some i and j so that $\gamma_i \in \text{supp}(\nu)$ and $\gamma_j \in \text{supp}(\mu)$. Let $\square^2_i$ and $\square^2_j$ indicate the 2-cubes of $\widetilde{\mathcal{S}(\lambda)}$ dual to p_i and p_j, respectively. Since both $\gamma, \gamma_i \in \text{supp}(\nu)$, by Lemma 3.2.2 the square $\square^2_i$ is a sub-cube of $\square^k_\nu$, and likewise for $\square^2_j$ and $\square^l_\mu$. The local hyperplanes of a k-cube are hyperplane equivalent to the intersections of these local hyperplanes with its m-cube faces, for $m \geq 1$. Thus H_ν is hyperplane equivalent in $\widetilde{\mathcal{C}(\lambda)}$ to the local hyperplane corresponding to γ of $\square^2_i$, and similarly for H_μ and the local hyperplane corresponding to γ of $\square^2_j$. These local hyperplanes corresponding to γ are hyperplane equivalent in $\widetilde{\mathcal{S}(\lambda)}$ by Lemma 3.1.2, so that by Lemma 3.2.2 they are hyperplane equivalent in $\widetilde{\mathcal{C}(\lambda)}$. Thus H_μ and H_ν are hyperplane equivalent, and Ψ is a $\pi_1 S$-equivariant identification of $\mathcal{H}_{\widetilde{\mathcal{C}(\lambda)}}$ with $\widetilde{\lambda}$.

Finally, we consider the second assertion of the proposition, relating a cube of $\widetilde{\mathcal{C}(\lambda)}$ to mutually intersecting curves in $\widetilde{\lambda}$. Given a k-cube $\square^k_\nu \subset \widetilde{\mathcal{C}(\lambda)}$, because $|\nu| = 2^k$, every choice of half-space for the k curves in $\mathrm{supp}(\nu)$ must extend to an admissible labeling of $\widetilde{\lambda}$. In particular, the k curves corresponding to the local hyperplanes of $\square^k_\nu$ all pairwise intersect: If two curves were disjoint, there would be labeling in ν that chooses disjoint half-spaces for these two curves, which is thus not admissible.

Conversely, given k mutually intersecting curves $\{\gamma_1, ..., \gamma_k\} \subset \widetilde{\lambda}$, there is a k-cube in $\widetilde{\mathcal{C}(\lambda)}$ containing the corresponding k hyperplanes: Choose a component $U \subset \widetilde{S} \setminus \widetilde{\lambda}$ contained in a bounded component of $\widetilde{S} \setminus \{\gamma_1, ..., \gamma_k\}$, and note that by Lemma 3.2.1, $v_\lambda(U) \in \widetilde{\mathcal{C}(\lambda)}^{(0)}$. Consider the 2^k labelings $\nu_1, ..., \nu_{2^k}$ formed by allowing any value in the $\gamma_1, ..., \gamma_k$ coordinates, while fixing all other coordinates of $v_\lambda(U)$. These labelings are all admissible: For any other lift $\gamma \in \widetilde{\lambda}$, in either case where γ and γ_i do or do not intersect, changing the γ_i coordinate of $v_\lambda(U)$ doesn't change admissibility. We thus have a set of 2^k distinct admissible labelings agreeing on $\widetilde{\lambda} \setminus \{\gamma_1, ..., \gamma_k\}$, so that there is a corresponding k-cube in $\widetilde{\mathcal{C}(\lambda)}$. □

Note that the cube constructed above from k mutually intersecting lifts may not be uniquely determined by the k lifts in $\widetilde{\lambda}$.

Lemma 3.2.4. *The complex $\widetilde{\mathcal{C}(\lambda)}$ is simply-connected.*

This is (Chatterji and Niblo, 2005, Lemma 2.6, p. 8). We include a sketch of a proof:

Proof sketch. A loop in $\widetilde{\mathcal{C}(\lambda)}$ can be homotoped so that it lies in the 1-skeleton, so we assume we have a path $\alpha : [0, 1] \to \widetilde{\mathcal{C}(\lambda)}^{(1)}$, with $\alpha(0) = \alpha(1)$. Towards contradiction, we assume further

that α is an embedding, and that α doesn't 'backtrack on maximal cubes': That is, there exist distinct maximal cubes $C_1, ..., C_n$ and $0 = s_0 < s_1 < ... < s_n < 1$ so that $\alpha(s_i) \subset C_i \setminus C_{i-1}$, for each $i = 0, ..., n-1$.

Suppose that the vertices visited by α are $v_1, ..., v_m \in \widetilde{\mathcal{C}(\lambda)}^{(0)}$, in order, and the curve dual to the edge e_i from v_i to v_{i+1} is γ_i. Each edge e_i represents a switch of choice of half-space at γ_i. Note that, as α is a loop, each γ_i must be dual to two edges, which make opposite switches. Moreover, there exists a subsequence $(e_{i_1}, ..., e_{i_n})$ of $(e_1, ..., e_m)$, chosen so that $e_{i_j} \in C_j \setminus C_{j-1}$, in which case γ_i and γ_{i+1} do not intersect. By the Jordan Curve Theorem, this implies the following nesting condition: The edge that crosses γ_{i+1} backwards must occur before the edge that crosses γ_i backwards. This will guarantee the existence of backtracking in the 1-skeleton. □

In fact, $\widetilde{\mathcal{C}(\lambda)}$ is $CAT(0)$ (Chatterji and Niblo, 2005, Lemmas 2.5 and 2.6), so it is contractible. We will not require this stronger statement, however.

Lemma 3.2.5. *The action of $\pi_1 S$ on $\widetilde{\mathcal{C}(\lambda)}$ satisfies:*

1. *The stabilizer of an n-cube is trivial for $n \geq 2$.*
2. *If $g \in \pi_1 S$ stabilizes a 1-cube in $\widetilde{\mathcal{C}(\lambda)}$, then g fixes the 1-cube pointwise.*
3. *If λ is filling, the action of $\pi_1 S$ is free and properly discontinuous.*

Proof. Since $\pi_1 S$ is torsion-free, if a non-trivial element stabilizes an n-cube for $n \geq 2$, then it has some power that is a non-trivial element that fixes two of the hyperplanes, and the proof is as in Lemma 3.1.3. The second statement also follows by the same argument in Lemma 3.1.3, with $\widetilde{\mathcal{S}(\lambda)}$ replaced by $\widetilde{\mathcal{C}(\lambda)}$.

Towards the third statement, suppose that λ is filling. In order to see that the action of $\pi_1 S$ is free and properly discontinuous, it suffices to show that the action is free, and that $\widetilde{\mathcal{C}(\lambda)}$ is locally finite (as in the proof of Lemma 3.1.5). In fact, it is enough to check that each vertex of $\widetilde{\mathcal{C}(\lambda)}$ is finite valence: If $g \in \pi_1 S$ stabilizes a vertex, then finite valence of the vertex ensures that g stabilizes a maximal cube containing the vertex. Stabilizers of higher-dimensional cubes are trivial, so this maximal cube is an edge, and the curve dual to this edge would produce an annulus disjoint from λ, contradicting the assumption that λ is filling.

In order to check finite valence of the vertices of $\widetilde{\mathcal{C}(\lambda)}$, we use the $\pi_1 S$-invariant subcomplex $\widetilde{\mathcal{S}(\lambda)} \subset \widetilde{\mathcal{C}(\lambda)}$ guaranteed by Lemma 3.2.2. Restricted to the 1-skeletons, this inclusion preserves combinatorial distances (cf. (Chatterji and Niblo, 2005, p. 6)): Given two vertices $v_\lambda(U_1)$ and $v_\lambda(U_2)$ of $\widetilde{\mathcal{S}(\lambda)}$, any path in $\widetilde{\mathcal{C}(\lambda)}$ from $v_\lambda(U_1)$ to $v_\lambda(U_2)$ must flip the labels of each of the curves in $\widetilde{\lambda}$ that separate U_1 from U_2. An arc in $\widetilde{S}$ from U_1 to U_2, intersecting $\widetilde{\lambda}$ minimally and transversally, provides a path in $\widetilde{\mathcal{S}(\lambda)}$ whose length is precisely the number of curves in $\widetilde{\lambda}$ that separate U_1 from U_2.

Now an infinite valence vertex of $\widetilde{\mathcal{C}(\lambda)}$ would provide infinitely many maximal cubes of $\widetilde{\mathcal{C}(\lambda)}$ whose union is bounded diameter. On the other hand, $\widetilde{\mathcal{S}(\lambda)}$ intersects each maximal cube of $\widetilde{\mathcal{C}(\lambda)}$, and a bounded diameter set of $\widetilde{\mathcal{S}(\lambda)}$ may only intersect finitely many squares. $\square$

The quotient $\mathcal{C}(\lambda) := \widetilde{\mathcal{C}(\lambda)}/\pi_1 S$ is again a cube complex, by Lemma 3.1.1. By Lemma 3.2.3, the hyperplanes $\mathcal{H}_{\mathcal{C}(\lambda)}$ are in correspondence with the curves in λ. By Lemma 3.2.2, the square complex $\mathcal{S}(\lambda)$ is a subcomplex of $\mathcal{C}(\lambda)$. See Appendix 3.7 for a demonstration that $\widetilde{\mathcal{C}(\lambda)}$ consists of finitely many $\pi_1 S$-orbits of maximal cubes.

Finally, we check that $\widetilde{\mathcal{C}(\lambda)}$ is independent of realization λ:

Lemma 3.2.6. *If λ and ξ are realizations of Λ, then there is a $\pi_1 S$-equivariant isomorphism of cube complexes $\widetilde{\mathcal{C}(\lambda)} \cong \widetilde{\mathcal{C}(\xi)}$.*

Proof. Note that the construction of $\widetilde{\mathcal{C}(\lambda)}$ is independent of an isotopy of the union $\widetilde{\lambda}$. By (de Graaf and Schrijver, 1997, Thm. 1) the realizations λ and ξ are homotopic through isotopies and finitely many Reidemeister type III moves. Hence it suffices to show the above equality when λ and ξ differ by exactly one Reidemeister type III move. From now on, we assume that λ and ξ agree outside of a disk on S containing the Reidemeister triangles.

Note that, because λ and ξ consist of homotopic curves, there is a natural bijection of the lifts $\psi : \widetilde{\lambda} \to \widetilde{\xi}$, and an associated bijection from admissible labelings of $\widetilde{\lambda}$ to admissible labelings of $\widetilde{\xi}$. In order to check that this defines a map on the 0-skeletons $\widetilde{\mathcal{C}(\lambda)}^{(0)} \to \widetilde{\mathcal{C}(\xi)}^{(0)}$, we need only check that this map sends the component of $v_\lambda(U)$ to $v_\lambda(V)$ in the graphs of admissible labelings, for some connected regions $U \subset \widetilde{S} \setminus \widetilde{\lambda}$ and $V \subset \widetilde{S} \setminus \widetilde{\xi}$. Choose a connected region $U \subset \widetilde{S} \setminus (\widetilde{\lambda} \cup \widetilde{\xi})$ in one of the lifts of the region on S on which λ and ξ agree. By design, we have $v_\lambda(U)(\gamma) = v_\xi(U)(\psi(\gamma))$ for $\gamma \in \widetilde{\lambda}$.

This yields a map of 0-skeletons $\widetilde{\mathcal{C}(\lambda)}^{(0)} \to \widetilde{\mathcal{C}(\xi)}^{(0)}$ which extends straightforwardly: If 2^k labelings of $\widetilde{\lambda}$ differ on at most k curves of $\widetilde{\xi}$, then the corresponding 2^k labelings of $\widetilde{\xi}$ differ on at most k curves of ξ. As the bijection $\psi : \widetilde{\lambda} \to \widetilde{\xi}$ is $\pi_1 S$-equivariant, the isomorphism described above is as well. □

As a consequence of Lemma 3.2.6, the isomorphism type of the quotient $\mathcal{C}(\lambda)$ is independent of realization λ for Λ. We refer to this isomorphism type of cube complex as $\mathcal{C}(\Lambda)$. In what follows, we do not return to the combinatorial definition of $\mathcal{C}(\Lambda)$. The reader aware of the correspondence in Proposition 3.2.3 will understand the essence of our analysis.

3.3 Mapping class group orbits of collections of curves

We now characterize a filling curve system from the isomorphism type of its dual cube complex. Recall that the elements of Λ are in correspondence with the hyperplanes of $\mathcal{C}(\Lambda)$. Recall the *extended mapping class group* $MCG^*(S)$ (see §2.1). The following is a restatement of Theorem C:

Theorem 3.3.1. *Two filling curve systems Λ_1 and Λ_2 are equivalent under the action of $MCG^*(S)$ if and only if there is an isomorphism of cube complexes $\mathcal{C}(\Lambda_1) \cong \mathcal{C}(\Lambda_2)$. The induced set map from Λ_1 to Λ_2 corresponds to the induced map between hyperplanes of $\mathcal{C}(\Lambda_1)$ and hyperplanes of $\mathcal{C}(\Lambda_2)$.*

Proof. One direction is straightforward: Suppose $\phi \cdot \Lambda_1 = \Lambda_2$ for $\phi \in MCG^*(S)$. Choose a realization λ_1 for Λ_1, and a homeomorphism ϕ' realizing ϕ. In this case ϕ' induces a $\pi_1 S$-equivariant isomorphism $\widetilde{\phi'} : \widetilde{\mathcal{C}(\lambda_1)} \cong \widetilde{\mathcal{C}(\phi \cdot \lambda_1)}$, where the induced map of hyperplanes corresponds to the set map induced by ϕ' from λ_1 to $\phi' \cdot \lambda_1$. Since $\phi' \cdot \lambda_1$ is a realization of Λ_2, we have $\mathcal{C}(\Lambda_1) \cong \mathcal{C}(\lambda_1)$ and $\mathcal{C}(\Lambda_2) \cong \mathcal{C}(\phi \cdot \lambda_1)$. The result follows.

On the other hand, suppose $\Phi : \mathcal{C}(\Lambda_1) \cong \mathcal{C}(\Lambda_2)$ is an isomorphism of cube complexes. Choose realizations λ_1 and λ_2 for Λ_1 and Λ_2. By Lemmas 3.2.4 and 3.2.5, we have that $\widetilde{\mathcal{C}(\Lambda_i)}$ is simply-connected and the action of $\pi_1 S$ is free and properly discontinuous, for each $i = 1, 2$. Thus $\widetilde{\mathcal{C}(\Lambda_i)}$

is the universal cover of $\mathcal{C}(\Lambda_i)$ and there are isomorphisms $\pi_1 S \cong \pi_1 \mathcal{C}(\Lambda_i)$ for each i. Composing these isomorphisms with Φ_*, we find an automorphism $\phi_* : \pi_1 S \to \pi_1 S$. By construction, we may lift Φ to a ϕ_*-equivariant isomorphism of cube complexes $\widetilde{\Phi} : \widetilde{\mathcal{C}(\Lambda_1)} \to \widetilde{\mathcal{C}(\Lambda_2)}$, in the sense that

$$\widetilde{\Phi}(g \cdot x) = \phi_*(g) \cdot \widetilde{\Phi}(x)$$

for each $x \in \widetilde{\mathcal{C}(\Lambda_1)}$ and $g \in \pi_1 S$. This induces a corresponding equivariant map on hyperplanes, which in turn induces a correspondence of the collections of conjugacy classes of stabilizers of hyperplanes of $\widetilde{\mathcal{C}(\Lambda_1)}$ with those of $\widetilde{\mathcal{C}(\Lambda_2)}$.

Since Proposition 3.2.3 guarantees that there is an identification of the collection of conjugacy classes of hyperplane stabilizers of $\pi_1 S$ acting on $\widetilde{\mathcal{C}(\Lambda_i)}$ with the collection of conjugacy classes of stabilizers of curves in $\widetilde{\lambda}_i$, we have a correspondence of the collection of conjugacy classes of stabilizers of curves in $\widetilde{\lambda_1}$ with those of $\widetilde{\lambda_2}$. The conjugacy classes of the stabilizers of curves in $\widetilde{\lambda}_i$ are naturally identified with the collection of conjugacy classes of Λ_i in $\pi_1 S$, from which it follows that the automorphism ϕ_* of $\pi_1 S$ takes the conjugacy classes determined by Λ_1 to those of Λ_2. By the Dehn-Nielsen-Baer Theorem (Farb and Margalit, 2012, Ch. 8, Thm. 8.1), there is $\phi \in MCG^*(S)$ inducing ϕ_*, so that $\phi \cdot \Lambda_1 = \Lambda_2$. □

3.4 Recognizing n-cubes dual to a 1-system

There remains the problem of recognizing quantitative information about $\mathcal{C}(\Lambda)$ from a given set of curves Λ. In this section, we prove Theorem D, giving a criterion for recognizing the dimensions of cubes in the complex dual to a 1-system.

A realization of a collection of curves forms a *triangle* if there is an embedded disk on S whose boundary components are three arcs of the curves, and so that these arcs intersect pairwise exactly once on the boundary of the disk. A collection of homotopy classes of curves forms a triangle if there is a realization of the curves that forms a triangle.

Lemma 3.4.1. *If a collection of homotopy classes of closed curves forms a triangle then every realization of the curves forms a triangle.*

Proof. By (de Graaf and Schrijver, 1997, Thm. 1) any two realizations are homotopic through isotopies and finitely many Reidemeister type III moves. Neither of these changes the existence of a triangle in the complement. □

Lemma 3.4.2. *If* $\Gamma = \{\gamma_1, \gamma_2, \gamma_3\}$ *is a realization of simple closed curves on* S*, then* $\dim \mathcal{C}(\Gamma) = 3$ *if and only if* $S \setminus \Gamma$ *has a connected component that is a triangle.*

Proof. If $\dim \mathcal{C}(\Gamma) = 3$ and Γ is in minimal position, then by Lemma 3.2.3 there are three mutually intersecting lifts of Γ in $\widetilde{S}$. As Γ is a realization, the intersections of these lifts are not concurrent, and there is a triangle T in their complement. Since T is compact, the intersection $T \cap \widetilde{\Gamma}$ consists of finitely many arcs. If an arc doesn't form a bigon with the boundary of a triangle, then the complement of the arc has a triangular component. Thus there is a triangular component $T' \subset T \setminus \widetilde{\Gamma}$ (cf. with *innermost* bigon, (Farb and Margalit, 2012, p. 31)).

Let $\pi : \widetilde{S} \to S$ be the covering map. If $int(T')$ does not embed under π, then there is an element $g \in \pi_1 S$ so that $int(T') \cap g \cdot int(T') \neq \emptyset$. In this case, by the Jordan Curve Theorem

there is an intersection $p \in \partial T' \cap g \cdot \partial T'$. Since $\partial T'$ consists of arcs from lifts of Γ, this violates $T' \subset T \setminus \widetilde{\Gamma}$. Thus $int(T')$ embeds in S, and we have found a triangle in the complement of Γ.

Conversely, suppose Γ forms a triangle T. Lift this topological disk to $\widetilde{S}$, and observe that the arcs of the boundary are contained in curves that pairwise intersect. Thus by Lemma 3.2.3 there is a 3-cube corresponding to this collection (see §3.2). By Lemma 3.2.3 again, a 4-cube would yield four lifts of curves in Γ that pairwise intersect. Since $|\Gamma| = 3$, two of these lifts are in the same $\pi_1 S$-orbit. As each of the γ_i are simple, this is impossible. □

Recall that a *1-system* is a collection of homotopy classes of simple closed curves whose pairwise geometric intersection number is at most one.

Theorem 3.4.3. *Suppose* $\Gamma = \{\gamma_1, ..., \gamma_n\}$ *is a* 1*-system on* S*. Then the dimension of* $\mathcal{C}(\Gamma)$ *is* n *if and only if the dimension of* $\mathcal{C}(\Gamma')$ *is* 3*, for every triple* $\Gamma' \subset \Gamma$*.*

One direction is straightforward: If $\dim \mathcal{C}(\Gamma) = n$, then by Lemma 3.2.3 there is a set of n lifts of the curves of Γ which mutually intersect. In this case, every trio of these lifts mutually intersect, and, by Lemma 3.2.3 again, form a 3-cube in the complex dual to that trio. Towards the other direction, using Lemma 3.4.2 we assume that each trio forms a triangle.

We define an *almost-realization* of a curve system to be a minimal position realization (see §3.1) with one slight change: the words 'only at transverse double points' in condition (iii) should be replaced by 'transversally'. (The terms *pairwise minimal position realization* would be more descriptive, but less economical). Note that (Farb and Margalit, 2012, Prop. 1.7, p. 31)

still guarantees that almost-realizations minimize pairwise geometric intersection numbers, as with realizations.

Proposition 3.4.4. *If $\{\gamma_1, ..., \gamma_n\}$ is a 1-system so that every trio $\{\gamma_i, \gamma_j, \gamma_k\}$ forms a triangle, then there is an almost-realization $\{\alpha_1, ..., \alpha_n\}$ of $\{\gamma_1, ..., \gamma_n\}$ so that the curves $\alpha_1, ..., \alpha_n$ have a single common intersection point.*

Note that if the curves form a 1-system and every trio forms triangles, then each pair intersect exactly once, which we assume below. It is interesting to note that our proof of Proposition 3.4.4 makes essential use of the orientation of S.

Proof. We proceed by induction on n. For $n = 3$, choose a realization of the three curves that forms a triangle. Note that since the curves form a 1-system, there is one boundary arc of the triangle contained in each of the curves. Fixing the curves outside of a disk containing the triangle, homotope one of the curves across the triangle so that it transversally crosses the other two curves at their intersection point. This produces an almost-realization of the curves, since no bigons have been created.

Assume now that we have $n + 1$ curves $\{\gamma_1, ..., \gamma_{n+1}\}$ so that every trio forms a triangle. By inductive hypothesis, there is an almost-realization $\{\alpha_1, ..., \alpha_n\}$ of $\{\gamma_1, ..., \gamma_n\}$ so that the α_i all have a unique common intersection point. Our strategy of proof will be to first choose a curve in the homotopy class of γ_{n+1} that is in minimal position with the α_i, and then to use the intersection properties of the curves – namely the fact that each pair intersects exactly once, and that every trio forms triangles – to 'weave' this curve around the other curves, achieving the

desired arrangement. Note that in this process the curves α_i, for $i = 1, ..., n$, remain unchanged. A straightforward adaptation of (Farb and Margalit, 2012, Prop. 1.7, p. 31) gives:

Lemma 3.4.5. *If $\{\alpha_1, ..., \alpha_n\}$ is an almost-realization of a collection of simple closed curves, then for any other simple closed curve β, there is a curve β', homotopic to β, so that $\{\alpha_1, ..., \alpha_n, \beta'\}$ is an almost-realization.*

Proof. By (Farb and Margalit, 2012, Prop. 1.7, p. 31), two curves are in minimal position if and only if they form no bigons. One half of this proof shows that if a bigon is formed (Farb and Margalit, 2012, p. 32), one may homotope one of these curves across the bigon while the other remains fixed. Choose any realization β_0 of β. If β_0 is in minimal position with α_1, then let $\beta_1 = \beta_0$. If β_0 is not in minimal position with α_1, find a bigon that β_0 forms with α_1. Use this bigon to find a curve β_0', homotopic to β_0, so that $|\beta_0' \cap \alpha_1| < |\beta_0 \cap \alpha_1|$.

It is crucial to observe that this step may be done so that $|\beta_0' \cap \alpha_j| \leq |\beta_0 \cap \alpha_j|$, for each j: The new curve $\beta_0' = \beta_+ \cup \beta_-$ is composed of two arcs, where β_+ follows β_0 and β_- follows α_1. If an arc of α_j intersects β_0' in the arc β_+, then there is a corresponding point of intersection of α_j with β_0. If an arc of α_j intersects β_0' in the arc β_-, then this arc of α_j enters the bigon formed by β_0 and α_1 along the side contained in α_1. In this case, α_j must exit the bigon through the side contained in β_0, since α_1 and α_j form no bigons by hypothesis (see Figure 20 and Figure 21). Thus there is again a point of intersection of β_0 with this arc of α_j that corresponds to the intersection of α_j with β_0'. We conclude that $|\beta_0' \cap \alpha_j| \leq |\beta_0 \cap \alpha_j|$, for $j \neq 1$.

We apply these finitely many homotopies across bigons to β_0, one for each of the bigons formed by β_0 and α_1. The result is a curve β_1, homotopic to β, so that β_1 and α_1 are in minimal

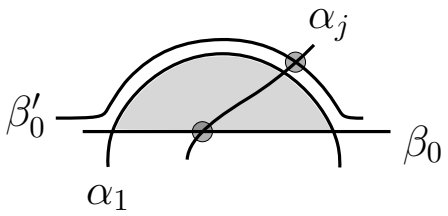

Figure 20. A point in $\alpha_j \cap \beta_0'$ corresponds to a point in $\alpha_j \cap \beta_0$.

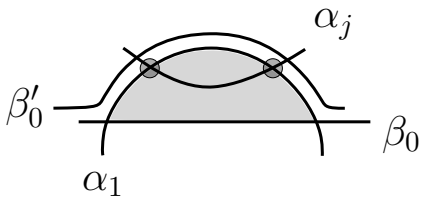

Figure 21. 'New' intersections of β_0' with α_j, violating the assumption that α_1 and α_j are in minimal position.

position, and so that $|\beta_1 \cap \alpha_j| \leq |\beta_0 \cap \alpha_j|$, for $j \neq 1$. Do this one-by-one for each α_k, producing a curve β_k, homotopic to β, so that β_k is in minimal position with $\alpha_1, ..., \alpha_k$. We terminate with a curve β', homotopic to β, so that β is in minimal position with α_j for each $j = 1, ..., n$. Thus $\{\alpha_1, ..., \alpha_n, \beta'\}$ is an almost-realization. □

Lemma 3.4.5 implies that we may find a curve β, homotopic to γ_{n+1}, so that $\{\alpha_1, ..., \alpha_n, \beta\}$ is an almost-realization. After possibly renaming the curves $\alpha_1, ..., \alpha_n$, we may assume that these curves are arranged around their unique common intersection in the counter-clockwise order $\alpha_1, ..., \alpha_n$. Let p_0 indicate the common intersection point of the α_i, and choose a disk neighborhood C of p_0, embedded on S, so that $\alpha_i \cap C$ is connected for each $i = 1, ..., n$.

Let $p_i := \beta \cap \alpha_i$. Replacing β with a homotopic curve if necessary, we may assume that all of the p_i are contained in the arcs $\alpha_i \cap C$, and that $\beta \cap C$ consists only of arcs that intersect an arc $\alpha_i \cap C$. We now have a picture where β weaves in and out of C, intersecting each of the arcs $\alpha_i \cap C$ precisely once. Outside of C the curves $\{\alpha_1, ..., \alpha_n, \beta\}$ are disjoint. The argument

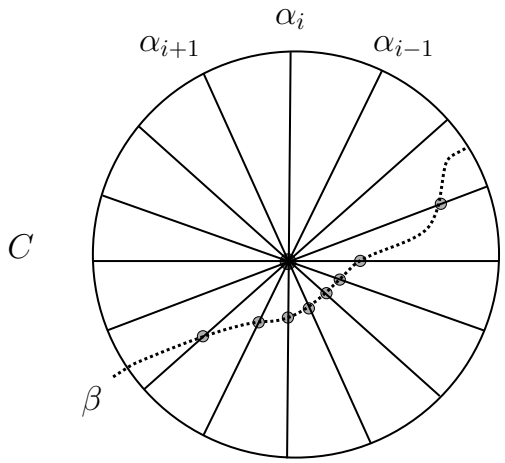

Figure 22. A 'localized realization' of β: $\beta \cap C$ is connected, and $p_i \in C$.

that follows applies various homotopies to β, maintaining the fact that β and α_i are in minimal position, for each $i = 1, ..., n$. Note that as we apply such homotopies, the p_i move as well.

In fact, for each p_i there is a homotopy of the curve β that slides the intersection p_i along α_i, out of C, until p_i "reappears" on the other side of p_0 on $\alpha_i \cap C$. Since the curves α_i do not intersect outside of C, this homotopy leaves the collection $\{\alpha_1, ..., \alpha_n, \beta\}$, pairwise, in minimal position. This homotopy of β will be exploited in the following argument, where we will refer to it as the 'slide move applied to p_i'.

Our first goal is to apply a homotopy to β to achieve a special formation, in which intersections of β with α_i have been 'localized' to C: We will say a curve is a *localized realization* of β if it is homotopic to β, $\beta \cap C$ is connected, and $p_i \in C$, for each i. See Figure 22 for an illustrated example. The bulk of the proof of this inductive step concerns existence of a localized realization of β. We will proceed by analyzing the triangles formed by β with α_i and

α_{i+1}, as i goes from 1 to $n-1$. At each step, we either find a localized realization of β, or we apply a homotopy to β so that the number of connected components of $\beta \cap C$ is at most $n-i$.

Given vertices a_1, a_2, a_3 of the triangle T, we will say that T 'realizes (a_1, a_2, a_3)' if the counter-clockwise orientation of ∂T induces the cyclic order (a_1, a_2, a_3). We will refer to a sufficiently small neighborhood of a_i as a 'corner' of T, or 'the corner at a_i' when the points are distinct, where the sufficiency is fulfilled when the intersection of T with this neighborhood is connected and disjoint from the side opposite the vertex.

Suppose that two curves η and δ in minimal position, together with a third curve, form a triangle T, so that T has a corner at the intersection point $p \in \eta \cap \delta$. In this case, the complement of $\eta \cup \delta$ in a small enough neighborhood of p consists of four components. Note that these four components correspond to the four possibilities for the placement of a corner of T at p. (When η and δ are simple curves it is straightforward to check that these possibilities are mutually exclusive). This is exploited repeatedly below, for the placement of corners of triangles at p_0.

By hypothesis, and by Lemma 3.4.1, there is a triangular component T_1 of $S \setminus \{\alpha_1, \alpha_2, \beta\}$. Because the curves form a 1-system, the vertices of T_1 are necessarily given by p_0, p_1, and p_2. Of the four ways for there to be a corner formed by α_1 and α_2 at p_0, two of them – namely, the choice of corner so that T realizes (p_0, p_2, p_1) (see Figure 23) – would achieve a localized realization: The composition, if necessary, of slide moves applied to p_1 and p_2 with a homotopy of β across the triangle T_2 would form a localized realization. We thus assume that T realizes (p_0, p_1, p_2), in which case, applying slide moves to p_1 and p_2 if necessary, we have the situation

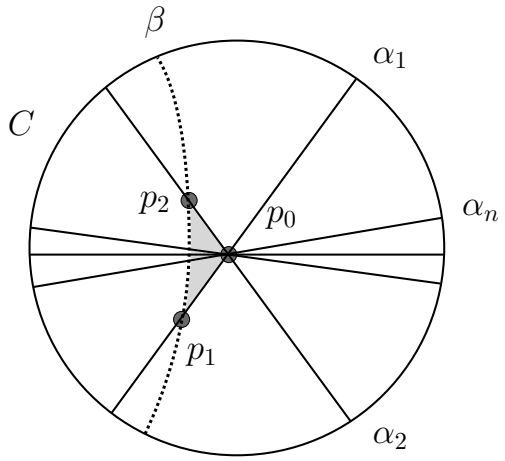

Figure 23. The triangle T_1 realizes (p_0, p_2, p_1), forming a localized realization of β.

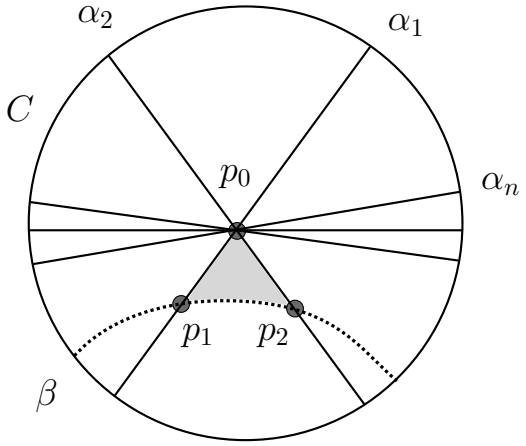

Figure 24. The triangle T_1 realizes (p_0, p_1, p_2), where no localized realization is yet assured.

pictured in Figure 24. Applying further slide moves when necessary, we may now assume that the p_i, for $i > 2$, all lie on the same side of $\alpha_1 \cap C$ inside C.

We now consider a triangular component T_2 of $S \setminus \{\alpha_2, \alpha_3, \beta\}$, with vertices p_0, p_2, and p_3. As before, there are four possible placements of the corner of T_2 at p_0. Two of them correspond to a situation in which T_2 realizes (p_0, p_3, p_2), again allowing a homotopy of β that forms a localized realization. Assuming then that T_2 realizes (p_0, p_2, p_3), we now describe why there is only one possible placement of the corner of T_2 at p_0.

First suppose the corner of T_2 at p_0 is placed as pictured in Figure 26. That is, suppose T_2 does not share the side between p_0 and p_2, contained in the arc $\alpha_2 \cap C$, with T_1. Because T_2 realizes (p_0, p_2, p_3), the counter-clockwise orientation of ∂T_2 turns left from α_2 to β at p_2. There are two ways to make a left turn from α_2 to β at p_2. Of them, the one so that T_2 does

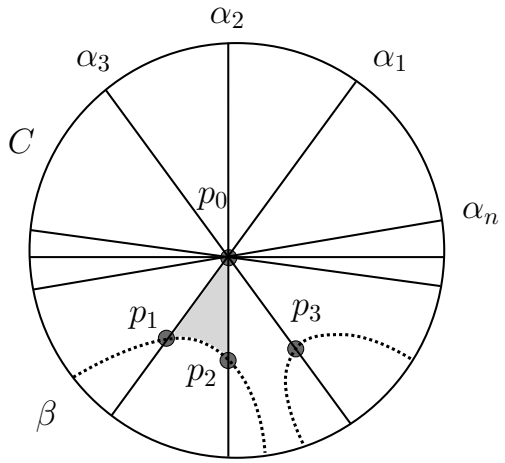

Figure 25. Given that the triangle T_1 realizes (p_0, p_1, p_2), the arcs $\beta \cap C$ may appear as pictured.

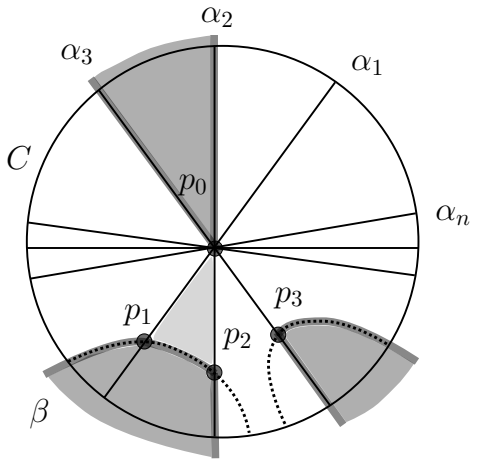

Figure 26. Given triangle T_1 realized as (p_0, p_1, p_2), the darkly shaded triangle T_2 cannot have a corner at p_0 as pictured.

not share the side between p_0 and p_2 with T_1 contains the arc of the curve β that lies between the vertices p_2 and p_3 and containing p_1. Consequently, one side of the arc $\alpha_1 \cap C$ leaves C inside T_2, while the other leaves C outside of T_2. Since α_1 may only intersect the sides of T_2 inside C, this is impossible.

This leaves only one possibility for the placement of the corner of T_2 at p_0 in which T_2 realizes (p_0, p_2, p_3) (see Figure 27). This case allows us to apply a homotopy to β supported in the triangle T_2, leaving the intersections p_i inside C, and ensuring that the number of connected components of $\beta \cap C$ is at most $n - 2$ (see Figure 28).

Similarly, at the kth step, the triangular component T_k of $S \setminus \{\alpha_k, \alpha_{k+1}, \beta\}$ either provides a homotopy of β that achieves a localized realization, or provides a homotopy of β, supported

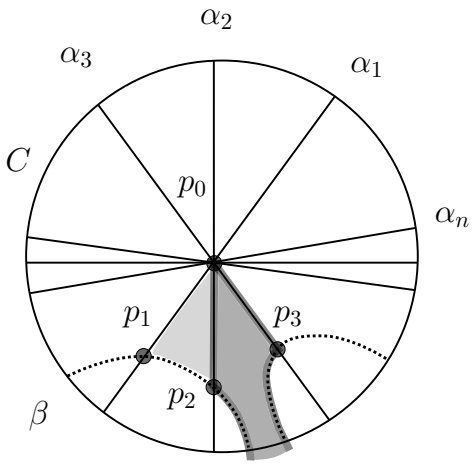

Figure 27. The remaining case where triangle T_2 realizes (p_0, p_2, p_3).

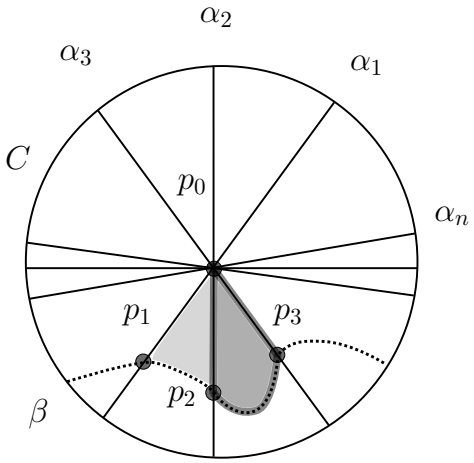

Figure 28. Applying a homotopy to β using T_2, getting 'closer' to a localized realization.

inside C, that ensures that the number of connected components of $\beta \cap C$ is at most $n - k$. When $k = n - 1$, we have ensured the existence of a localized realization of β.

Finally, after replacing β with its localized realization, we assume that $\beta \cap C$ is connected, and that $p_i \in C$ for each i. Note that this implies that $\beta \cap \partial C$ consists of exactly two points, which we denote p_- and p_+. By construction, there are exactly $2n$ components of $\partial C \setminus \{\alpha_1, ..., \alpha_n\}$, which we refer to as 'sectors', arranged cyclically around ∂C. We claim that p_- and p_+ are in diametrically opposed sectors of ∂C.

After renumbering, we may assume that p_- is in one of the sectors between α_n and α_1, and that, reading from p_- to p_+ along β, the intersections with α_i occur in the order $p_1, ..., p_n$. In this case, reading from p_1 to p_+, β stays in one component of $C \setminus \alpha_1$. Suppose that p_+ is in a sector between α_k and α_{k+1}, for $k < n$. In this case, if β intersects α_{k+1}, then it must intersect

at least twice, by the Jordan Curve Theorem. As β is in minimal position with α_{k+1}, this does not happen. Thus p_+ is in the sector between α_n and α_1, opposite the sector containing p_-. We may now apply a homotopy to β making it into a diameter passing through p_0, completing the inductive step. □

Proof of Theorem 3.4.3. By Proposition 3.4.4, the curves $\gamma_1, ..., \gamma_n$ have an almost-realization as $\alpha_1, ..., \alpha_n$, so that the α_i have a single common intersection point. Choose a lift of this point to $\widetilde{S}$, and consider the lifts of the α_i passing through this point. Since the α_i form an almost-realization, there are no bigons formed by the curves, and the chosen lifts pairwise intersect in exactly one point. This implies that in any realization of the curves $\gamma_1, ..., \gamma_n$ these lifts will pairwise intersect. By Lemma 3.2.3, there is an n-cube dual to these curves. □

A word of caution: It is not generally true that if $\Gamma' \subset \Gamma$ then the complex $\mathcal{C}(\Gamma')$ is a subcomplex of $\mathcal{C}(\Gamma)$ (see Figure 29). The following corollary is a weaker version of such a statement that will suffice for our application. Given a collection of curves Γ, we will say that a subset of n curves $\Gamma' \subset \Gamma$ *form an n-cube* in $\mathcal{C}(\Gamma)$ if there are n hyperplanes corresponding to the curves of Γ' intersecting in an n-cube of $\mathcal{C}(\Gamma)$.

Corollary 3.4.6. *If Γ is a 1-system of curves, and $\{\gamma_1, ..., \gamma_n\} = \Gamma' \subset \Gamma$, then the curves of Γ' form an n-cube in $\mathcal{C}(\Gamma)$ if and only if every triple of curves from Γ' form a 3-cube.*

Proof. Using Proposition 3.2.3, the curves of Γ' form an n-cube in $\mathcal{C}(\Gamma)$ if and only if there is a choice of lifts $\{\widetilde{\gamma_1}, ..., \widetilde{\gamma_n}\}$ so that these lifts pairwise intersect, which in turn occurs if and only if $\dim \mathcal{C}(\Gamma') = n$, at which point we apply Theorem 3.4.3. □

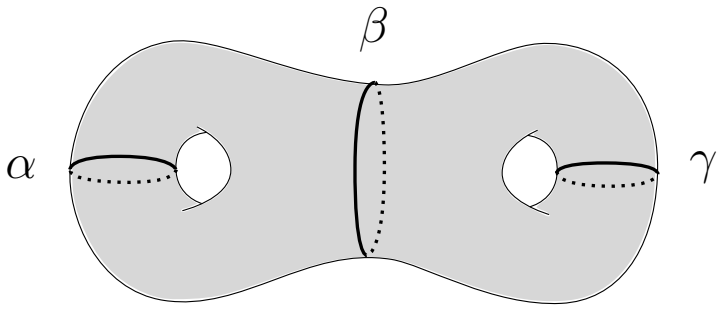

Figure 29. The complex $\mathcal{C}(\{\alpha, \gamma\})$ is not a subcomplex of $\mathcal{C}(\{\alpha, \beta, \gamma\})$.

3.5 A family of maximal complete 1-systems

We construct now a curve system on a surface S of any odd genus $g = 2n + 1$. Consider the genus g surface as the unit sphere in $\mathbb{R}^3$ with g handles attached at evenly spaced disks centered on an equator, making the order g homeomorphism σ that cyclically permutes the handles apparent. Let x be a point fixed by this homeomorphism. Consider the presentation $\pi_1(S, x) = \langle r_i, s_i | \prod_{i=1}^{g} [r_i, s_i] \rangle$, where the generators r_1 and s_1 are as pictured in Figure 30, and $r_i = \sigma(r_{i-1})$ and $s_i = \sigma(s_{i-1})$ for $i = 2, ..., g$.

Let α_1, β_1, and δ_0 be given by

- $\alpha_1 = [r_{n+2} r_{n+3} ... r_{2n+1} s_1]$,
- $\beta_1 = \left[r_{n+2} r_{n+3} ... r_{2n+1} s_1^{-1} \prod_{i=1}^{n+1} [s_i, r_i] \right]$, and
- $\delta_0 = [r_1 r_2 ... r_{2n+1}]$.

The orbit of α_1 under σ gives g curves which we denote by $\{\alpha_1, \alpha_3, ..., \alpha_{2g-1}\}$. Similarly, we denote the σ-orbit of β_1 by $\{\beta_1, \beta_3, ..., \beta_{2g-1}\}$. We complete these collections to sequences

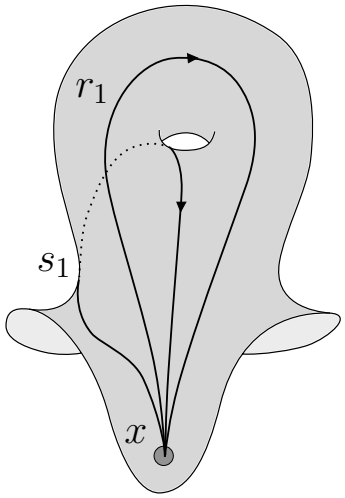

Figure 30. The generators of $\pi_1(S, x)$ in one handle.

$\{\alpha_i\}_{i=1}^{2g}$ and $\{\beta_i\}_{i=1}^{2g}$ by defining $\alpha_{2i} = \tau_i(\alpha_{2i-1})$ and $\beta_{2i} = \tau_i^{-1}(\beta_{2i-1})$ for $i = 1, ..., g$, and where τ_i is the right Dehn twist around r_i. See Figure 31, Figure 32, Figure 33, and Figure 34 for illustrative examples.

Grouping these curves together, we will refer to $A = \{\alpha_1, ..., \alpha_{2g}\}$ as the set of 'up' curves and $B = \{\beta_1, ..., \beta_{2g}\}$ as the set of 'down' curves. We will refer to the pair of up curves (resp. down curves) α_{2i-1} and α_{2i} (resp. β_{2i-1} and β_{2i}) as 'partners', for $i = 1, ..., g$.

It is immediate that, $\alpha_i \cap \alpha_j = \beta_i \cap \beta_j = 1$ for $i \neq j$. When α_i and α_j are not partners, then $\alpha_i \cap \beta_j = 1$. This calculation makes it clear that we may form many maximal complete 1-systems: For each $i = 1, ..., g$, choose one of the two pairs of partners from $\{\alpha_{2i-1}, \alpha_{2i}\}$ and $\{\beta_{2i-1}, \beta_{2i}\}$. Together with the δ_0 curve, this forms $2g + 1$ curves that pairwise intersect once. By Theorem 1.2.1 this is maximal.

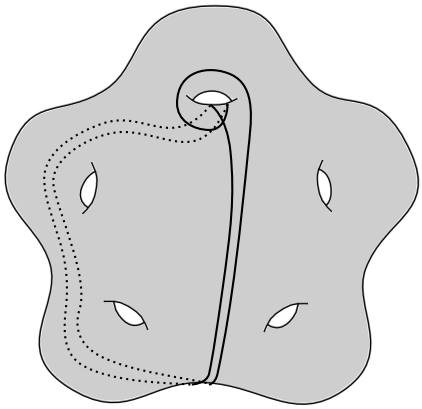

Figure 31. The curves $\alpha_1, \alpha_2 \in A$

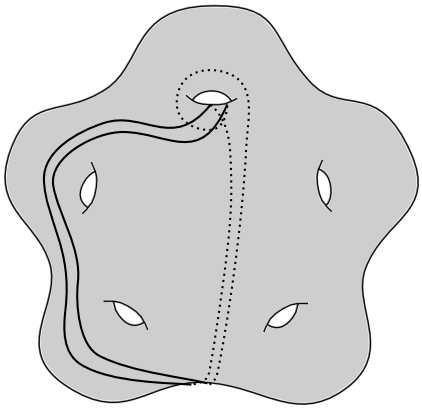

Figure 32. The curves $\beta_1, \beta_2 \in B$

More precisely, for $\epsilon \in \{1, -1\}^g$, let $A(\epsilon) = \{\alpha_{2i-1}, \alpha_{2i} | \epsilon_i = 1\}$ and $B(\epsilon) = \{\beta_{2i-1}, \beta_{2i} | \epsilon_i = -1\}$, and let $\Gamma(\epsilon) = \{\delta_0\} \cup A(\epsilon) \cup B(\epsilon)$. Several examples in genus 5 and 7 are shown in Figure 35, Figure 36, Figure 37, and Figure 38.

Lemma 3.5.1. *For each $\epsilon \in \{-1, 1\}^g$, the collection of curves $\Gamma(\epsilon)$ forms a maximal complete 1-system.*

There is a natural action of $MCG^*(S)$ (see §2.1) on systems of conjugacy classes of curves on S. The following proposition, whose proof is postponed to §3.6, implies that the maximal complete 1-systems from Lemma 3.5.1 represent many distinct orbits.

Proposition 3.5.2. *If there exists $\phi \in MCG^*(S)$ and $\epsilon, \epsilon' \in \{-1, 1\}^g$ such that $\phi \cdot \Gamma(\epsilon) = \Gamma(\epsilon')$, then either $|A(\epsilon)| = |A(\epsilon')|$ or $|A(\epsilon)| = |B(\epsilon')|$. Hence $\sum \epsilon_i = \pm \sum \epsilon'_i$.*

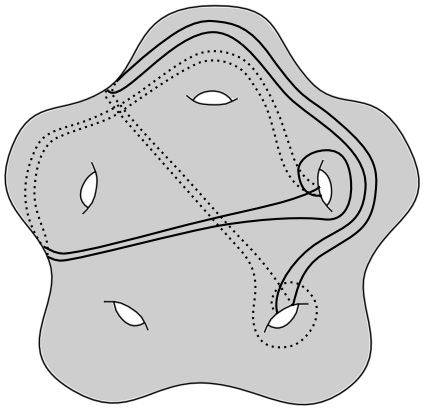

Figure 33. The curves $\alpha_3, \alpha_4 \in A$ and $\beta_5, \beta_6 \in B$

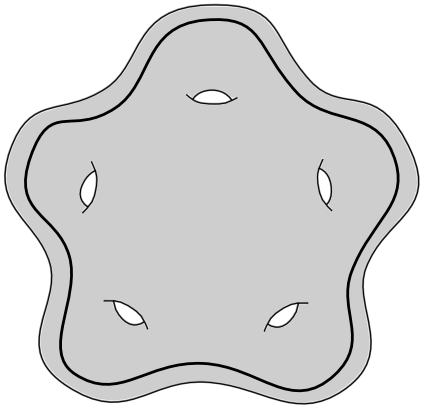

Figure 34. The curve δ_0

Using Proposition 3.5.2 to distinguish mapping class group orbits, it would suffice to work with

$$\Gamma(k) := \Gamma(\underbrace{1, .., 1}_{k}, \underbrace{-1, ..., -1}_{g-k}). \tag{3.1}$$

However, we expect that the existence of a mapping class relating $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ is in fact a much more restrictive condition on ϵ and ϵ'. To formulate this precisely, note that there is an action of the dihedral group D_g on $\{1, -1\}^g$ given by letting the generators act by a g-cycle and a reversal of the list, respectively. Letting $\mathbb{Z}/2\mathbb{Z}$ act by taking ϵ to $-\epsilon$, we have an action of $\mathbb{Z}/2\mathbb{Z} \oplus D_g$ on $\{1, -1\}^g$.

Conjecture. *If $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ are in the same $MCG^*(S)$-orbit, then ϵ and ϵ' are in the same $(\mathbb{Z}/2\mathbb{Z} \oplus D_g)$-orbit in $\{1, -1\}^g$.*

We will also require the simple observation:

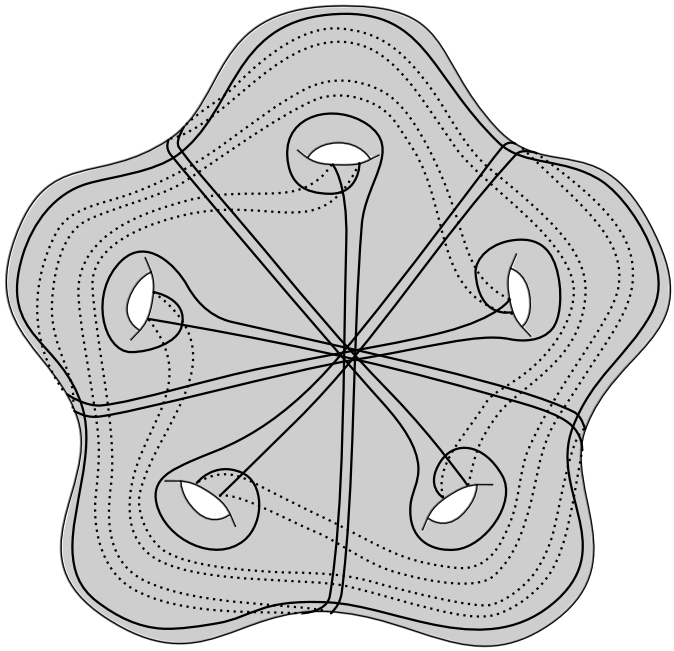

Figure 35. The system of curves $\Gamma(1,1,1,1,1)$

Lemma 3.5.3. *A maximal complete* 1*-system is filling.*

Proof. Suppose not. Then there is a simple closed curve α disjoint from the curves in our 1-system. Cut open S along α, and cap off the two resulting boundary components created with disks. Note that the resulting surface may be disconnected. In any case, the set of $2g+1$ curves obtained forms a maximal complete 1-system of curves on a surface of genus $g' \leq g-1$, contradicting Theorem 1.2.1. □

Let $N(g)$ indicate the number of $MCG^*(S)$-orbits among maximal complete 1-systems. Proposition 3.5.2 provides a lower bound for $N(g)$, and we will use a simple argument involving the dual square complex to derive an upper bound. We obtain:

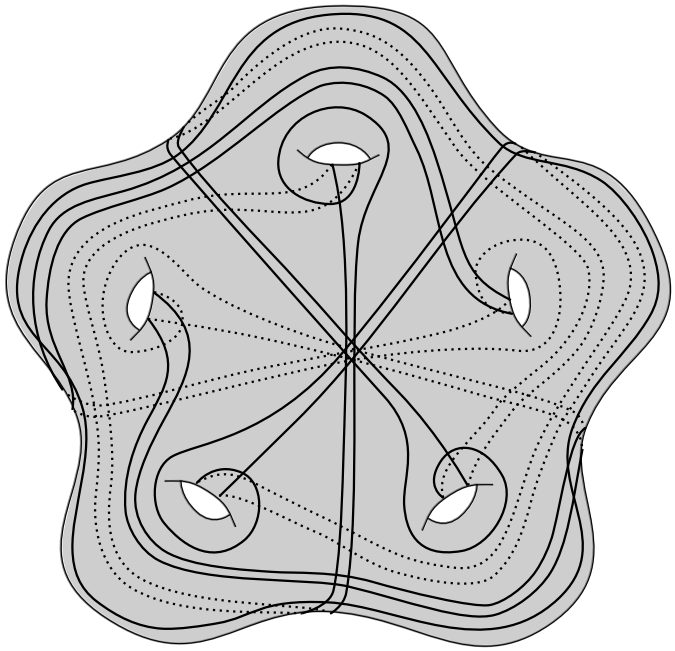

Figure 36. The system of curves $\Gamma(1,-1,1,1,-1)$

Theorem 3.5.4. *For all g, we have the upper bound*

$$N(g) \leq (4g^2+2g)!$$

and, for g odd, we have the lower bound

$$N(g) \geq \frac{g-1}{2}.$$

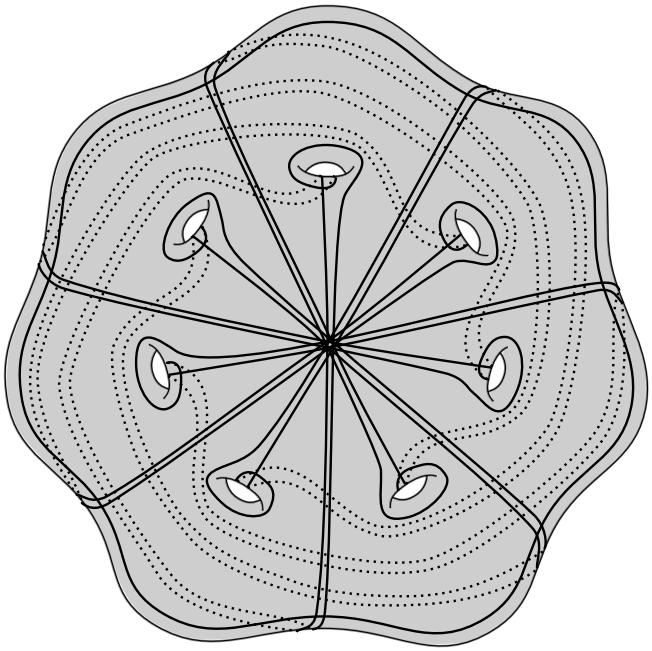

Figure 37. The system of curves $\Gamma(1,1,1,1,1,1,1)$

Proof. Recall the notation $\Gamma(k)$ as in (Equation 3.1). For g odd, the lower bound follows from Proposition 3.5.2: If $\phi \cdot \Gamma(k) = \Gamma(l)$, then $2k - g = \pm(2l - g)$, or equivalently $k = l$ or $k + l = g$. Thus $\{\Gamma(1), \Gamma(2), ..., \Gamma(n)\}$ form $n = (g-1)/2$ distinct mapping class group orbits.

Towards the upper bound, consider the set of isomorphism classes of square complexes $\mathcal{S}(\lambda)$ that are dual to realizations λ of maximal complete 1-systems. By Lemma 3.5.3 and Proposition 3.1.5, an isomorphism of square complexes $\mathcal{S}(\lambda) \cong \mathcal{S}(\lambda')$ yields a homeomorphism of S taking λ to λ'. Thus there is a well-defined map from the set of isomorphism classes of square complexes dual to maximal complete 1-systems to the set of $MCG^*(S)$-orbits of maximal complete 1-systems. Since every finite collection of curves has a realization, this map is surjective. Hence

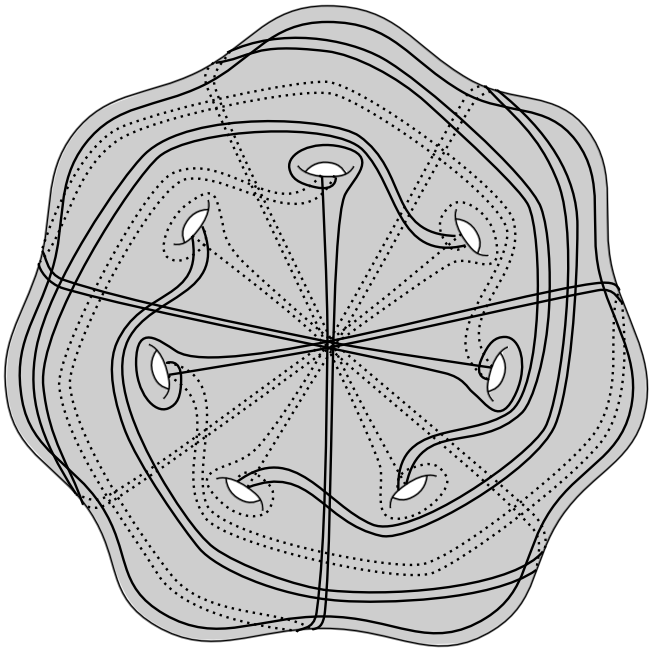

Figure 38. The system of curves $\Gamma(1,-1,1,-1,-1,1,-1)$

an upper bound for the number of possible square complexes dual to a realization of a maximal complete 1-system produces an upper bound for $N(g)$. The following is a schematic of our approach:

$$\left\{ MCG^*(S)\cdot\Lambda \;\middle|\; \begin{array}{l}\Lambda \text{ is a maximal}\\ \text{complete 1-system}\end{array} \right\} \twoheadleftarrow \left\{ \begin{array}{l}\text{isomorphism}\\ \text{class of } \mathcal{S}(\lambda)\end{array} \;\middle|\; \begin{array}{l}\lambda \text{ is a realization}\\ \text{of a maximal complete}\\ \text{1-system } \Lambda\end{array} \right\}$$

For each realization λ of a maximal complete 1-system, each of the curves in λ passes through exactly $2g$ squares of $\mathcal{S}(\lambda)$. We may thus view $\mathcal{S}(\lambda)$ as the quotient of the disjoint

union of $2g+1$ annuli, each of which is built from $2g$ squares, where the quotient map identifies squares in pairs. There are at most $\binom{2g(2g+1)}{2,\dots,2}$ pairings, and each pair of matched squares has two possible identifications, giving at most

$$2^{g(2g+1)}\binom{2g(2g+1)}{2,\dots,2} = 2^{g(2g+1)} \cdot \frac{(2g(2g+1))!}{2^{g(2g+1)}} = (2g(2g+1))!$$

square complexes $\mathcal{S}(\lambda)$. □

This concludes the proof of Theorem B.

3.6 Restricting mapping class group orbits via the dual cube complex

This section is devoted to the following strengthening of Proposition 3.5.2:

Proposition 3.6.1. *If $\phi \in MCG^*(S)$ satisfies $\phi \cdot \Gamma(\epsilon) = \Gamma(\epsilon')$, then*

- *$\phi(\delta_0) = \delta_0$, and*
- *Either $\phi \cdot (A(\epsilon), B(\epsilon)) = (A(\epsilon'), B(\epsilon'))$ or $\phi \cdot (A(\epsilon), B(\epsilon)) = (B(\epsilon'), A(\epsilon'))$.*

In other words, either ϕ preserves the sets of up and down curves, or it exchanges them. The proof of this proposition will follow from a coarse picture of the cube complex $\mathcal{C}(\Gamma(\epsilon))$.

Lemma 3.6.2. *In the complex $\mathcal{C}(A \cup B \cup \{\delta_0\})$, the triples that form 3-cubes are the following:*

1. *$\{\alpha_i, \alpha_j, \alpha_k\}$ or $\{\beta_i, \beta_j, \beta_k\}$, for distinct $i, j, k \in \{1, ..., 2g\}$.*
2. *$\{\alpha_{2i-1}, \alpha_{2i}, \delta_0\}$ or $\{\beta_{2i-1}, \beta_{2i}, \delta_0\}$, for $i \in \{1, ..., g\}$.*
3. *$\{\alpha_{2j-1}, \alpha_{2j}, \beta_i\}$ or $\{\beta_{2j-1}, \beta_{2j}, \alpha_i\}$, for $i \in \{1, ..., 2g\}$ and $j \in \{1, ..., g\}$.*

4. $\{\alpha_i, \beta_j, \delta_0\}$ *for* $i, j \in \{1, ..., 2g\}$.

Proof. Using Lemmas 3.4.1 and 3.4.2, we determine whether a triple of curves forms a 3-cube by choosing a realization of the curves, and observing whether there is a triangular component of the complement. For each of the curves, we fix choices of realizations as in Figure 31, Figure 32, and Figure 34.

If δ_0 is one of the three curves, we arrange the possible ways to choose the other two curves according to whether the curves are chosen as 'up' or 'down' (i.e. from A or B): If both of the other curves are up, then there is a triangle in the complement of the trio if and only if the other two curves were partners. If one of the curves is up and one is down, there is such a triangle. The other cases are similar.

On the other hand, if δ_0 is not one of the three curves: If all of the curves are up, there is such a triangle. If two of the curves are up and one is down, there is a triangle in their complement if and only if the two up curves are partners. The other cases are similar. □

We proceed with an examination of hyperplanes of maximal cubes in the cases for $\epsilon \in \{1, -1\}^g$ where $|A(\epsilon)|, |B(\epsilon)| > 1$.

Lemma 3.6.3. *When* $|A(\epsilon)|, |B(\epsilon)| > 1$*, the sets of hyperplanes of maximal cubes of* $\mathcal{C}(\Gamma(\epsilon))$ *correspond to one of the following lists of curves:*

1. *The* $2|\epsilon^{-1}(1)|$ *curves* $A(\epsilon)$.

2. *The* $2|\epsilon^{-1}(-1)|$ *curves* $B(\epsilon)$.

3. *The 5 curves* $\{\alpha_{2i-1}, \alpha_{2i}, \beta_{2j-1}, \beta_{2j}, \delta_0\}$*, for* $i, j \in \{1, ..., g\}$ *such that* $\alpha_{2i} \in A(\epsilon)$ *and* $\beta_{2j} \in B(\epsilon)$*.*

Proof. Using Corollary 3.4.6, in order to check whether a subset of curves from $\Gamma(\epsilon)$ forms an n-cube, it is enough to check whether every triple forms a 3-cube. By Lemma 3.6.2 we have a complete list of such 3-cubes.

The curves $A(\epsilon)$ and $B(\epsilon)$ form cubes of dimensions $2|\epsilon^{-1}(1)|$ and $2|\epsilon^{-1}(-1)|$, respectively, by Lemma 3.6.2 and Corollary 3.4.6. If one adds a down curve β_i to $A(\epsilon)$, then a pair of up curves that are not partners will not form a 3-cube with this down curve β_i, by Lemma 3.6.2. If one adds δ_0 to $A(\epsilon)$, then again a pair of up curves that are not partners will not form a 3-cube with δ_0. The analogous statements hold for $B(\epsilon)$. By Corollary 3.4.6, the cubes of dimension $2|\epsilon^{-1}(1)|$ and $2|\epsilon^{-1}(-1)|$ containing these sets of hyperplanes, respectively, must be maximal. The same analysis shows that a maximal cube containing δ_0 must contain a pair of partner up curves and a pair of partner down curves. □

The cases in which either of $|A(\epsilon)|$ or $|B(\epsilon)|$ are less than or equal to 1 are quite similar, so we list the relevant information without proof.

Lemma 3.6.4. *When* $|B(\epsilon)| = 0$ *(resp.* $|A(\epsilon)| = 0$*), the sets of hyperplanes of maximal cubes of* $\mathcal{C}(\Gamma(\epsilon))$ *correspond to one of the following lists of curves:*

1. *The* $2g$ *curves* $A(\epsilon)$ *(resp.* $B(\epsilon)$*).*

2. *The 3 curves* $\{\alpha_{2i-1}, \alpha_{2i}, \delta_0\}$ *(resp.* $\{\beta_{2i-1}, \beta_{2i}, \delta_0\}$*), for* $i \in \{1, ..., g\}$*.*

When $|B(\epsilon)| = 1$ *(resp.* $|A(\epsilon)| = 1$*), the sets of hyperplanes of maximal cubes of* $\mathcal{C}(\Gamma(\epsilon))$ *correspond to one of the following lists of curves:*

1. *The* $2g-2$ *curves* $A(\epsilon)$ *(resp.* $B(\epsilon)$*).*
2. *The 5 curves* $\{\alpha_{2i-1}, \alpha_{2i}, \beta_{2j-1}, \beta_{2j}, \delta_0\}$*, for* $i, j \in \{1, ..., g\}$ *such that* $\alpha_{2i} \in A(\epsilon)$ *and* $\beta_{2j} \in B(\epsilon)$*.* □

Proof of Proposition 3.6.1. The simple observation we exploit is that cube complex isomorphisms must send maximal cubes to maximal cubes.

Suppose $|A(\epsilon)|, |B(\epsilon)| > 1$. By Lemma 3.6.3 and Lemma 3.6.4, the maximal cubes that the hyperplane corresponding to δ_0 passes through are all 5-dimensional, while the maximal cubes that the hyperplane corresponding to α_i (resp. β_i) passes through, for any $i \in \{1, ..., 2g\}$, include one of two even-dimensional maximal cubes corresponding to the curves $A(\epsilon)$ and $B(\epsilon)$. Thus any isomorphism of cube complexes $\Phi : \mathcal{C}(\Gamma(\epsilon)) \cong \mathcal{C}(\Gamma(\epsilon'))$ must take the hyperplane corresponding to δ_0 to itself. By Theorem 3.3.1, the induced mapping class fixes δ_0. Similarly, the even-dimensional maximal cubes whose hyperplanes correspond to $A(\epsilon)$ and $B(\epsilon)$) must be sent to the pair of even-dimensional maximal cubes whose hyperplanes correspond to $A(\epsilon')$ and $B(\epsilon')$. The remaining cases are similar. □

We conclude that if $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ are $MCG^*(S)$-equivalent, then the pair of numbers $|A(\epsilon)|$ and $|B(\epsilon)|$ is equal to $|A(\epsilon')|$ and $|B(\epsilon')|$, and Proposition 3.5.2 follows.

3.7 Appendix: Finite-dimensionality of the cube complex dual to a system of curves

Given a finite collection of homotopy classes of curves Λ on S, we describe why $\mathcal{C}(\Lambda)$ is finite-dimensional, with finitely many maximal cubes. This material is well known, but we include it for completeness. See (Sageev, 1995), (Sageev, 2012, p.21), (Chatterji and Niblo, 2005) for details.

Let Λ indicate a finite collection of homotopy classes of closed curves on S. Choose an identification of the universal cover $\widetilde{S}$ with $\mathbb{H}^2$, and enumerate the geodesic representatives of the curves from Λ as $\{\gamma_1, \gamma_2, ...\}$.

We identify a complete geodesic in $\mathbb{H}^2$ with a pair of distinct points $p, q \in S^1$. We say that an ordered sequence of pairwise intersecting complete geodesics $\left(\{p_1^1, p_1^2\}, ..., \{p_k^1, p_k^2\}\right)$ is *spoked* if the list $\left(p_1^{\delta_1}, ..., p_k^{\delta_k}\right)$, for some choices of $\delta_1, ..., \delta_k \in \{1, 2\}$, appears uninterrupted in the cyclic ordering of $\{p_j^i\}$ induced by the counter-clockwise orientation of S^1.

Lemma 3.7.1. *Let $(\{p_1, q_1\}, ...\{p_n, q_n\})$ be n spoked geodesics in $\mathbb{H}^2$. Let θ_i be the counter-clockwise angle from $\{p_i, q_i\}$ to $\{p_{i+1}, q_{i+1}\}$ for $i = 1, ..., n-1$, and θ_n the counter-clockwise angle from $\{p_n, q_n\}$ to $\{p_1, q_1\}$. Then $\sum_{i=1}^{n} \theta_i \leq \pi$.*

Proof. Induction. The base case $n = 2$ is obvious.

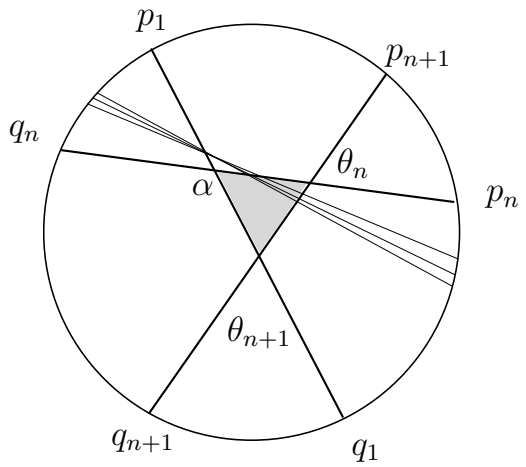

Figure 39. The inductive step for Lemma 3.7.1.

For the inductive step, consider Figure 39. By hyperbolic geometry, $\alpha \geq \theta_n + \theta_{n+1}$. Using the inductive step to estimate the angles with $\{p_{n+1}, q_{n+1}\}$ thrown out, we see

$$\begin{aligned} \pi &\geq \alpha + \sum_{i=1}^{n-1} \theta_i \\ &\geq \theta_n + \theta_{n+1} + \sum_{i=1}^{n-1} \theta_i. \end{aligned}$$

□

Lemma 3.7.2. *There is a bound, depending on S and Λ, on the maximal size of a subset of $\widetilde{\Lambda}$ that pairwise has nonempty intersection.*

Proof. There is a bound to the maximal size of a spoked subset of the lifts of these geodesics: Since there are finitely many intersections, there is a uniform positive lower bound to the angles at intersection points, contradicting Lemma 3.7.1. □

Lemma 3.7.3. *For each $k > 0$, there are finitely many $\pi_1 S$-orbits of sets of k pairwise-intersecting elements of $\widetilde{\Lambda}$.*

Proof. We proceed by induction: The case $k = 1$ is clear, since there are finitely many curves in Λ.

For the inductive step, assume there were infinitely many $\pi_1 S$-orbits of sets of $k+1$ pairwise-intersecting elements of $\widetilde{\Lambda}$. Since there are only finitely many $\pi_1 S$-orbits of sets of k such elements, we would find infinitely many $\pi_1 S$-orbits of $k+1$ such elements that all share the same k elements $\{\gamma_1, ..., \gamma_k\}$. In this case, we would find infinitely many elements from $\widetilde{\Lambda}$ intersecting a compact set whose boundary consists of finitely many arcs of the γ_i, contradicting the fact that there are only finitely many intersection points of elements of Λ. □

By Lemma 3.7.2, $\widetilde{\mathcal{C}(\Lambda)}$ is finite-dimensional. By Lemma 3.7.3, the quotient of $\widetilde{\mathcal{C}(\Lambda)}$ by the action of $\pi_1 S$ consists of finitely many maximal cubes.

APPENDICES

Appendix A

PERMISSION LETTER FOR INCLUSION OF JOINT WORK

April 21, 2014

To whom it may concern,

I grant my permission for the inclusion of the collaborative work with Jonah Gaster that is the subject of Chapter 3.

Sincerely,

Tarik Aougab

CITED LITERATURE

Ahlfors, L. V.: Lectures on quasiconformal mappings, volume 38 of University Lecture Series. Providence, RI, American Mathematical Society, second edition, 2006.

Aougab, T.: Large collections of curves pairwise intersecting exactly once. Preprint, 10 2012.

Bers, L.: Simultaneous uniformization. Bull. Amer. Math. Soc., 66:94–97, 1960.

Bonahon, F.: Bouts des variétés hyperboliques de dimension 3. Ann. of Math. (2), 124(1):71–158, 1986.

Bonahon, F.: Variations of the boundary geometry of 3-dimensional hyperbolic convex cores. J. Differential Geom., 50(1):1–24, 1998.

Bridson, M. R. and Haefliger, A.: Metric spaces of non-positive curvature, volume 319 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Berlin, Springer-Verlag, 1999.

Canary, R. D., Epstein, D. B. A., and Green, P. L.: Notes on notes of Thurston. In Fundamentals of hyperbolic geometry: selected expositions, volume 328 of London Math. Soc. Lecture Note Ser., pages 1–115. Cambridge, Cambridge Univ. Press, 2006.

Canary, R. D. and McCullough, D.: Homotopy equivalences of 3-manifolds and deformation theory of Kleinian groups. Mem. Amer. Math. Soc., 172(812):xii+218, 2004.

Casson, A. J. and Bleiler, S. A.: Automorphisms of surfaces after Nielsen and Thurston, volume 9 of London Mathematical Society Student Texts. Cambridge, Cambridge University Press, 1988.

Chatterji, I. and Niblo, G.: From wall spaces to CAT(0) cube complexes. Internat. J. Algebra Comput., 15(5-6):875–885, 2005.

Chesebro, E. and DeBlois, J.: Algebraic invariants, mutation, and commensurability of link complements. To appear in Pacific Journal of Mathematics, 2012.

Choi, Y.-E. and Series, C.: Lengths are coordinates for convex structures. J. Differential Geom., 73(1):75–117, 2006.

Culler, M.: Lifting representations to covering groups. Adv. in Math., 59(1):64–70, 1986.

de Graaf, M. and Schrijver, A.: Making curves minimally crossing by reidemeister moves. Journal of Combinatorial Theory, Series B, 70(1):134–156, 5 1997.

Dumas, D.: Complex projective structures. In Handbook of Teichmüller theory. Vol. II, volume 13 of IRMA Lect. Math. Theor. Phys., pages 455–508. Eur. Math. Soc., Zürich, 2009.

Dumas, D.: Skinning maps are finite-to-one. Preprint, 2012.

Dumas, D. and Kent, IV, R. P.: Experiments with skinning maps. In Preparation.

Dumas, D. and Kent, IV, R. P.: Slicing, skinning, and grafting. Amer. J. Math., 131(5):1419–1429, 2009.

Dumas, D. and Wolf, M.: Projective structures, grafting and measured laminations. Geom. Topol., 12(1):351–386, 2008.

Farb, B. and Margalit, D.: A primer on mapping class groups, volume 49 of Princeton Mathematical Series. Princeton, NJ, Princeton University Press, 2012.

Fathi, A., Laudenbach, F., and Poénaru, V.: Travaux de Thurston sur les surfaces, volume 66 of Asterérisque. Société Mathématique de France, 1979. Séminaire Orsay.

Gardiner, F. P.: Measured foliations and the minimal norm property for quadratic differentials. Acta Math., 152(1-2):57–76, 1984.

Gardiner, F. P. and Masur, H.: Extremal length geometry of Teichmüller space. Complex Variables Theory Appl., 16(2-3):209–237, 1991.

Goldman, W. M.: The symplectic nature of fundamental groups of surfaces. Adv. in Math., 54(2):200–225, 1984.

Goldman, W. M.: Invariant functions on Lie groups and Hamiltonian flows of surface group representations. Invent. Math., 85(2):263–302, 1986.

Goldman, W. M.: Trace coordinates on Fricke spaces of some simple hyperbolic surfaces. In Handbook of Teichmüller theory. Vol. II, volume 13 of IRMA Lect. Math. Theor. Phys., pages 611–684. Eur. Math. Soc., Zürich, 2009.

Hejhal, D. A.: Monodromy groups and linearly polymorphic functions. In Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973), pages 247–261. Ann. of Math. Studies, No. 79. Princeton, N.J., Princeton Univ. Press, 1974.

Hubbard, J. H.: The monodromy of projective structures. In Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), volume 97 of Ann. of Math. Stud., pages 257–275. Princeton, N.J., Princeton Univ. Press, 1981.

Kamishima, Y. and Tan, S. P.: Deformation spaces on geometric structures. In Aspects of low-dimensional manifolds, volume 20 of Adv. Stud. Pure Math., pages 263–299. Tokyo, Kinokuniya, 1992.

Kapovich, M.: Hyperbolic manifolds and discrete groups, volume 183 of Progress in Mathematics. Boston, MA, Birkhäuser Boston Inc., 2001.

Keen, L. and Series, C.: Continuity of convex hull boundaries. Pacific J. Math., 168(1):183–206, 1995.

Kent, IV, R. P.: Skinning maps. Duke Math. J., 151(2):279–336, 2010.

Lehto, O.: Univalent functions and Teichmüller spaces, volume 109 of Graduate Texts in Mathematics. New York, Springer-Verlag, 1987.

Luo, F.: Grothendieck's reconstruction principle and 2-dimensional topology and geometry. In Handbook of Teichmüller theory. Vol. II, volume 13 of IRMA Lect. Math. Theor. Phys., pages 733–765. Eur. Math. Soc., Zürich, 2009.

Malestein, J., Rivin, I., and Theran, L.: Topological designs. Geom. Dedicata, 168(1):221–233, 2014.

Marden, A.: The geometry of finitely generated kleinian groups. Ann. of Math. (2), 99:383–462, 1974.

Maskit, B.: Kleinian groups, volume 287 of Grundlehren der Mathematischen Wissenschaften. Berlin, Springer-Verlag, 1988.

McMullen, C.: Iteration on Teichmüller space. Invent. Math., 99(2):425–454, 1990.

McMullen, C.: Complex earthquakes and Teichmüller theory. J. Amer. Math. Soc., 11(2):283–320, 1998.

Otal, J.-P.: Courants géodesiques et produits libres. Doctoral dissertation, Université Paris-Sud, Orsay, 1988.

Otal, J.-P.: Thurston's hyperbolization of Haken manifolds. In Surveys in differential geometry, Vol. III (Cambridge, MA, 1996), pages 77–194. Int. Press, Boston, MA, 1998.

Przytycki, P.: Arcs intersecting at most once. Preprint, 02 2014.

Sageev, M.: Ends of group pairs and non-positively curved cube complexes. Proc. London Math. Soc. (3), 71(3):585–617, 1995.

Sageev, M.: $cat(0)$ cube complexes and groups. PCMI Lecture Notes, to appear, 2012.

Series, C.: On Kerckhoff minima and pleating loci for quasi-Fuchsian groups. Geom. Dedicata, 88(1-3):211–237, 2001.

Thurston, W. P.: Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds. Ann. of Math. (2), 124(2):203–246, 1986.

VITA

Education

· M. S., Mathematics, University of Illinois at Chicago, 2008-2010

· B. S., Mathematics, University of Wisconsin-Madison, 1999-2002

Research

Papers

A family of non-injective skinning maps with critical points

· http://arxiv.org/abs/1212.6210, To appear in *Transactions of the AMS*.

Invited Presentations

"A non-injective skinning map with a critical point"

· Joint Mathematics Meetings, AMS Special Session on Deformation Spaces of Geometric Structures on Low-Dimensional Manifolds, Baltimore, MD, January 2014

· SUNY Stonybrook, Geometry Seminar, Stonybrook, NY, December 2013

· CUNY Graduate Center, Hyperbolic Geometry Seminar, New York, NY, December 2013

· Temple University, Geometry-Topology Seminar, Philadelphia, PA, December 2013

· University of Michigan at Ann Arbor, Geometry Seminar, Ann Arbor, MI, October 2013

- Boston College, AMS Sectional Meeting, Chestnut Hill, MA, April 2013
- University of Wisconsin at Madison, Geometry Seminar, Madison, WI, November 2012
- Geometric Structures and Representation Varieties (GEAR) Junior Retreat, University of Illinois at Champaign-Urbana, Champaign, IL, July 2012
- Moduli spaces and dynamical systems (poster session), ICERM, Providence, RI, April 2012
- University of Illinois at Chicago, Graduate Student Colloquium, Chicago, IL, April 2012
- Institut Henri Poincaré, Graduate Student Seminar, Paris, FR, March 2012

"The measurable Riemann mapping theorem and its consequences"

- Geometry, Topology, and Dynamics of Character Varieties, National University of Singapore, Singapore, July 2010

Teaching

- Instructor, University of Illinois at Chicago, Multivariable Calculus, Fall 2013
- Grader, University of Illinois at Chicago, Complex Analysis, Spring 2013
- Teaching Assistant, GEAR Junior Retreat, A short course on hyperbolic 3-manifolds, Summer 2012
- Grader, University of Illinois at Chicago, Differential Geometry, Spring 2010

· Teaching Assistant, University of Illinois at Chicago, Calculus I and II, Business Calculus, Precalculus, 2008-2010