

Singular Maps of Surfaces into Hyperbolic 3-Manifolds

by

Eric Katerman

A thesis submitted in partial fulfillment
of the requirements for the
degree of Bachelor of Arts with Honors
in Mathematics

WILLIAMS COLLEGE

Williamstown, Massachusetts

May 14, 2002

I wish to thank Professor Chkhenkeli, my second reader, and Professor Adams, my advisor.

©2002 Eric Katerman

Abstract

We construct singular maps of surfaces into hyperbolic 3-manifolds in order to find upper bounds for meridian length, longitude length, and maximal cusp volume. We also provide ample background and history of hyperbolic geometry and 3-manifold theory for this exposition to be accessible to undergraduate mathematics majors. Generalizations and attempts to strengthen the results herein are also included for completeness.

Contents

1	Introduction	1
2	Definitions	5
2.1	Point-set Topology	5
2.2	Algebraic Topology	6
2.3	3-manifold Theory	8
3	Hyperbolic 2- and 3-Manifold Theory	11
3.1	History	11
3.2	Elements of Hyperbolic Geometry	13
3.3	Universal Covers of Hyperbolic Spaces	14
3.4	Pleated Surfaces	17
3.5	Cusp Diagrams	18
4	Boundary-Compressible Surfaces	20
4.1	Performing Boundary Compressions	22
4.2	Four Cases	27
4.2.1	Meridian to itself	27
4.2.2	One meridian to another	29
4.2.3	Longitude to a meridian	34
4.2.4	Longitude to itself	35
4.3	Summary	36
5	Cusp Size Bounds	37
5.1	Necessary Properties of a Useful Surface	37
5.2	Main Results	38

6	Attempts at Stronger Results	43
6.1	S.C. Boundary Incompressibility of He Surfaces	43
6.1.1	Drilling and Cyclic Covers	44
6.1.2	Induction on Crossing Number	44
6.1.3	Branching Covers	47
6.2	Seifert and Seifert-He Surfaces	53
7	Generalized He Surfaces	55
7.1	Construction	55
8	Conclusions	58
8.1	Summary	58
8.1.1	Main Results	58
8.1.2	Strengthening our Main Results	59
8.1.3	Generalizations	59
8.2	What the Future Holds	60

List of Figures

2.1	The fundamental group is not necessarily commutative. . . .	7
2.2	The two surfaces on the left are embeddings of a sphere and of a torus. A Klein bottle cannot be embedded in \mathbb{R}^3 as represented on the right by a projection of an immersion of a Klein bottle.	9
3.1	A torus with meridian m and longitude l	15
3.2	A properly immersed surface S in a 3-manifold M	16
3.3	Here, S is at least locally a proper map since its boundary is contained entirely within the cusp of M	17
4.1	The four cases to check.	21
4.2	Here, $\alpha \cup \beta$ bound a ∂ -compressing disk D	22
4.3	Form an annulus A by connecting the ends of $N(\alpha)$ to each other by a disk in C which intersects neither l nor m in its interior. Then the number of half-twists in $N(\alpha)$ is the number of half-twists in the A	23
4.4	Here, the boundary of a boundary-compressing disk is shown.	24
4.5	The boundary-compressing disk D with boundary $\partial D = \alpha \cup \beta$ is shown here.	25
4.6	A boundary compressed surface S' in the case that α connects one meridian to another and B has zero half-twists.	26
4.7	The boundary of D in the simple case.	28
4.8	Performing a boundary compression on a surface whose constructed band $B = N(\alpha) \cup A_1 \cup A_2$ has a single half-twist. Notice that γ' , the boundary component which replaces γ_1 and γ_2 , is the product $\gamma_1 \circ \gamma_2$	30

4.9	The dark grey region in the middle is the patch, and the point at the bottom right corner of this patch is a problematic point, but as shown there exists a homeomorphism from a neighborhood around that point to an open disk in \mathbb{R}^2	33
5.1	In this situation, a boundary-compressing disk exists with α as part of its boundary.	38
5.2	The fraction of an ideal triangle's area which intersects the cusp can be no more than $\frac{3}{\pi}$	40
5.3	The lengths of ℓ and m are at least 1.	41
6.1	Crossing c is the one we break open. Note that crossing a might be crossing b rotated some $n\pi + \frac{\pi}{2}$ radians.	45
6.2	Constructing an alternating braid whose components do not cross themselves (although they cross each other).	48
6.3	Creating an alternating projection with strictly more crossings than $\pi(L)$	51
7.1	Here, we cone a projection of the figure-8 knot (K) to the core curve c of a 2-torus \mathbb{T}^2 to produce a surface S	56

Chapter 1

Introduction

Ever since Thurston's work on 3-manifold geometry and topology in the 1970s, the topology and geometry of 3-manifolds have been known to be intrinsically intertwined. Indeed, hyperbolic volume $V(M)$ of a hyperbolic 3-manifold M is an important invariant of non-compact 3-manifolds of finite volume; that is, given two such manifolds M_1 and M_2 which are homeomorphic, $V(M_1) = V(M_2)$. Similarly, the volume $V(C)$ of the maximal cusp C of a knot complement $M = \mathbb{S}^3 \setminus K$ is preserved by homeomorphism and therefore conveys much information about its corresponding manifold, and thus we are interested in finding cusp volume bounds which are easy to calculate so that the process of differentiating between two manifolds may be facilitated. In [Lac00], Lackenby finds an upper bound for the hyperbolic volume $V(M)$ of a knot complement $M = \mathbb{S}^3 \setminus K$ for some hyperbolic alternating knot K admitting a projection with twist number $t \leq c$ where c is the minimal crossing number of K :

$$V(M) \leq 16v_3(t - 1)$$

where $v_3 = 1.01494\dots$ is the hyperbolic volume of an ideal 3-simplex of maximal volume. Since the hexagonal packing of horoballs is the best packing in $\mathbb{C} \times [0, 1]$, we find that the ratio of hyperbolic volume of a fundamental region to the hyperbolic volume of the region of intersection between that fundamental region and the horoballs is $\frac{2v_3}{\sqrt{3}}$, and thus Lackenby's upper bound on hyperbolic volume of a manifold extends to an upper bound on the volume of the maximal cusp of a one-cusped manifold:

$$V(C) \leq \frac{8\sqrt{3}}{v_3}(t - 1)$$

Thus, this is essentially a linear bound on the hyperbolic volume of the maximal cusp $V(C)$ in terms of the twist number of K .

In this paper, we find a bound which improves on Lackenby's bound in many cases; in particular, our bound does better for those K with $t(K)$ close to $c(K)$:

$$V(C) \leq \frac{9c}{2} \left(1 - \frac{1}{c}\right)^2$$

Although ours is essentially a linear bound, it shows that as $c \rightarrow \infty$, $V(C)$ tends toward $4.5c$ rather than $8\sqrt{3}c/v_3 \approx \frac{27c}{2}$, which is what Lackenby's result gives for knots with twist number close to minimal crossing number. In addition, our results hold for all hyperbolic knots, not just hyperbolic alternating knots. To find these bounds, we will use singular maps of disks into the manifold M which have boundary equal to K and are punctured by K at every crossing of a given projection $\pi(K)$. We will show that a surface need only be simple curve boundary-incompressible (which we define in section 2.3), and if the surface does not satisfy this condition, then performing the corresponding boundary compression only improves on the cusp volume bound thereby showing that the bounds given by the original surface hold. We also investigate some techniques which we have explored in order to show that He surfaces are simple curve boundary-incompressible for certain hyperbolic knot complements. This is clearly a stronger result than the one we have found for general properly immersed punctured disks in link complements.

Here we state our main theorem (1.2), which gives us the improvement on Lackenby's result as mentioned above. In our proof of Theorem 5.2 (which we supply in Section 5.2 along with the proofs of the other results that are stated here without proof) we will use Lemma 1.1, which ensures that the area of S' is greater than or equal to the area of the cusp intersecting the surface. We include this lemma here because it is of independent interest since it states that surfaces which are useful in this context need only be s.c. boundary incompressible rather than essential, as surfaces that topologists often concern themselves with tend to be.

Lemma 1.1 *If S is an s.c. boundary-incompressible surface properly immersed in M , then the cusp $N(K)$ of M intersects S only in annuli.*

The following theorem improves Lackenby's result.

Theorem 1.2 *Let K be a knot with c crossings in projection $\pi(K)$, and let $|m|, |\ell|$ denote the lengths of the meridian m of the knot and an ℓ -curve ℓ of the knot. The following bounds hold for K and its maximal cusp C :*

- $|m| \leq 6 - \frac{7}{c}$
- $|\ell| \leq 5c - 6$
- $V(C) \leq \frac{9c}{2} \left(1 - \frac{1}{c}\right)^2$

Theorem 1.3 implies that if a surface of the type that interests us is s.c. boundary compressible, then performing a boundary compression only improves the bounds found in Theorem 1.2.

Theorem 1.3 *Performing s.c. boundary compressions on a properly immersed surface S with one ℓ -boundary component and at least two m -boundary components improves our bounds on meridian length, longitude length, and cusp volume.*

For general motivational purposes, we should mention that topology is the branch of mathematics which investigates what properties of a structured space persist through various deformations of that space. Much of the research in topology has been and will continue to be devoted to categorizing spaces up to certain structure-preserving operations. For example, Poincaré knew that all 2-manifolds which fall into a certain category (called “simply connected”) are equivalent to a sphere, and he posited that the analogous 3-manifold case holds as well, but this conjecture remains open today. Another related example is the fundamental question of knot theory, which is to determine whether two knots are the same. Knot theorists have devoted much energy to attacking this problem, and while many useful invariants have been discovered, no one has written an effective computer program to determine whether two distinct projections represent different knots (see [Ada94a]). This question of categorization and characterization is the heart of topology.

Indeed, this question is the heart of 3-manifold theory, which is the subset of topology that deals with certain 3-dimensional spaces. Incompressible surfaces have been of primary interest in 3-manifold theory since the early 1960’s when Haken developed his theory of normal surfaces (see [Hak62]). Shortly after that, Waldhausen published “On irreducible 3-manifolds which are sufficiently large” [Wal68] in which he proves that, for a large class of 3-manifolds,

each 3-manifold is characterized by its fundamental group. To prove this result, he uses incompressible surfaces as well as Haken's results in [Hak62]. Traditionally, these incompressible surfaces are embedded in M , but recently a few authors (including Agol [Ago00] and Neumann-Coto [NC95]) have found occasion to use singular maps, of which embeddings are a strict subset, to investigate various properties of 3-manifolds. However, many difficulties arise by removing the restriction that these maps be embeddings, which is the primary reason why most authors work with embeddings rather than immersions. We have met with many of these difficulties—including, for example, that the boundary ∂D of a boundary-compressing disk D may itself be immersed, which disallows application of loop theorem—but we were ultimately successful in finding the above bounds using immersed surfaces. We conclude this report with Chapter 8, a recapitulation and a deeper analysis of these concerns, where we have met them with success and where we have not, and what lies ahead.

Chapter 2

Definitions

In this chapter, we define all terms which will be used in this paper. Some background is assumed; for example, we will not define a function, surjectivity, injectivity, continuous maps, or bijections. Although a primary goal of this paper is self-containment, anyone with a background consisting of undergraduate courses in analysis, topology, and algebra should have no trouble filling in these absent definitions.

2.1 Point-set Topology

An n -**manifold** M is a topological space that looks n -dimensional locally, which is to say that M is Hausdorff (there are disjoint neighborhoods around distinct points in M), has a countable basis, and given a point $x \in M$, there exists an open neighborhood of x that is homeomorphic to \mathbb{R}^n . Examples of 3-manifolds include \mathbb{R}^3 , $\mathbb{S}^3 = \mathbb{R}^3 \cup \{\infty\}$, and $\mathbb{S}^2 \times \mathbb{S}^1$, not one of which is homeomorphic to any of the other two. Two spaces S and T are **homeomorphic** if and only if there exists a continuous bijection f with a continuous inverse f^{-1} such that $f(S) = T$ and $f^{-1}(T) = S$. A manifold is **compact** if and only if every open cover $\{U_\alpha\}_{\alpha \in A}$ has a finite subcover $\{U_i\}_{i=1}^n$. Examples of compact 3-manifolds are \mathbb{S}^3 (the one-point compactification of \mathbb{R}^3), $\mathbb{S}^2 \times [0, 1]$, and $\mathbb{S}^1 \times [0, 1] \times [0, 1]$. Although we will be dealing primarily with non-compact manifolds, we include the previous definition for completeness.

A manifold M has boundary if there exists a point $x \in M$ such that any open neighborhood of x is homeomorphic to $\mathbb{R}_+^n = \{(x_1, x_2, \dots, x_n) : x_1 \in \mathbb{R}^+, x_2, \dots, x_n \in \mathbb{R}\}$. Examples of 3-manifolds with boundary are the closed

annulus, $A = \{(x_1, x_2) \in \mathbb{R}^2 \mid 1 \leq x_1^2 + x_2^2 \leq 2\}$ crossed with either $[0, 1]$ or $[0, 1]$. Additionally, we will only consider **connected** manifolds, which means that there exists no disjoint, nonempty open sets $U, V \subset M$ such that $U \cup V = M$. A space X is **path connected** if, given any two points x_0 and x_1 in X , there exists a path with origin x_0 and end x_1 , and X is **locally path connected** if for each point $x \in X$ and each neighborhood U of x , there is a path-connected neighborhood V of x contained in U . In order to allay any possible confusion, we should mention that we have defined path connected and locally path connected in order to define covering spaces below, but we will not need to use these definitions outside of that application.

A **closed** manifold is one which is compact without boundary. An example of a closed manifold is $\mathbb{S}^n = \{(x_1, x_2, \dots, x_n) : \sqrt{\sum_{i=1}^n x_i^2}, x_i \in \mathbb{R}\}$, the n -sphere for some fixed $n \in \mathbb{Z}^+$. Note that while we will be working in non-compact 3-manifolds, the surfaces and curves in those manifolds will sometimes be compact and/or without boundary, so the examples in this and previous paragraphs are important elements of this expository section in order to bring the reader up to speed.

2.2 Algebraic Topology

One tool from algebraic topology that we will develop in some detail is the **fundamental group** of a topological space, which is the group whose elements are homotopy classes of curves in an n -manifold M . First, we define a **closed curve** with **base point** $x_0 \in X$ as a curve $\gamma(t)$ parametrized for $0 \leq t \leq 1$ with $\gamma(0) = \gamma(1) = x_0$. Formally, if γ_0 and γ_1 are closed curves in X , then we say that γ_0 is **homotopic** to γ_1 (written $\gamma_0 \cong \gamma_1$) if there exists a continuous map $F : X \times [0, 1] \rightarrow X$ such that $F(x, 0) = \gamma_0(x)$ and $F(x, 1) = \gamma_1(x)$ for all $x \in X$. The map F is called a **homotopy** from γ_0 to γ_1 . We will say that two closed curves are **homotopically equivalent** if and only if there exists such a map F , and all homotopically equivalent curves make up a single **homotopy class**. It is straightforward to check that homotopy is an equivalence relation.

The fundamental group of M (relative to a point $x_0 \in X$) is concerned with homotopy classes of closed, orientable 1-manifolds in M , and thus in order to investigate the natural group structure which arises we must define a binary operation. It is clear that there exists at least one representative γ from each homotopy class such that $x_0 \in \gamma$. Therefore, we can define the




Figure 2.1: The fundamental group is not necessarily commutative.

binary operation of curve composition (denoted by ‘ \circ ’) between two oriented closed curves $\gamma_1(t)$ and $\gamma_2(t)$, each of which is parametrized on $[0,1]$ (such that $\gamma_1(0) = \gamma_1(1) = \gamma_2(0) = \gamma_2(1) = x_0$) to be

$$\gamma_1 \circ \gamma_2 = \gamma_1(\gamma_2(t)) = \begin{cases} \gamma_2(2t) & 0 \leq t \leq 1/2 \\ \gamma_1(2t) & 1/2 \leq t \leq 1 \end{cases}$$

Clearly this group is closed under curve composition since the composition of two curves with base point x_0 will be another closed curve with the same base point; that is, given γ_1 and γ_2 as above, $\gamma_1 \circ \gamma_2(0) = \gamma_1(0) = x_0 = \gamma_2(1) = \gamma_1 \circ \gamma_2(1)$.

Now we are ready to state the formal definition of the fundamental group of a manifold X relative to some base point $x_0 \in X$. The set of path homotopy classes of loops based at x_0 , with the operation \circ , is called the **fundamental group** of X relative to the base point x_0 . It is denoted by $\pi_1(X, x_0)$, or when choice of base point is immaterial, by $\pi_1(X)$. Notice that this group is not necessarily commutative since, for example, the fundamental group of D (a disk with two holes), which is generated by x and y , is not commutative; namely, $x^2y^2 \neq xyxy$ (see Figure 2.1). Also notice that we do not require representatives from these classes to be embedded, or globally one to one.

2.3 3-manifold Theory

We will take our definition of a metric from analysis, which states that a **metric** on a space X is a function $d : X \rightarrow \mathbb{R}$ which satisfies the following three conditions:

1. $d(x, x) = 0$ for all $x \in X$.
2. $d(x, y) = d(y, x)$ for all $x, y \in X$.
3. $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$.

An example of a metric is the standard Euclidean metric in \mathbb{R}^2 , $d(x, y) = |x - y|$. A bijective map $f : X \rightarrow Y$ is an **isometry** if it preserves distances; that is, f is an isometry if and only if $d_x(a, b) = d_y(f(a), f(b))$ for all $a, b \in X$ where d_x is the metric on X and d_y is the metric on Y . Examples of isometries in the Euclidean plane are translations and rotations. An example of a bijection on the Euclidean plane that is not an isometry is $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $g((x, y)) = (cx, cy)$, which is scaling by some constant factor c .

Hyperbolic 3-manifolds are named as they are because they admit hyperbolic geometric structures, which means that a metric of constant sectional curvature equal to -1 can be put on such manifolds. Constant (sectional) curvature can be defined as satisfying both **homogeneity** (in at least one viewing direction, every point looks the same as every other point) and **isotropy** (from a given point, the space looks the same at any angle). Formally, a geometry satisfies homogeneity if its group of isometries is transitive, and a geometry satisfies isotropy if for any two ordered bases of orthonormal tangent vectors at a point in the space, there is an isometry of the space fixing the point and taking one frame to the other. Thus surfaces which admit a constant curvature metric can be cut into small patches, each of which may be placed isometrically anywhere else on the surface. See Chapter 3 for more on hyperbolic manifolds.

A **proper map** $f : S \rightarrow M$ of a surface S into a manifold M is one which lines up the boundary of the surface with the boundary of the manifold; that is, $\partial f(S) \subset \partial M$ (see Figure 3.3). If there is some such proper map for S , we say that S is **properly mapped into** M . Now suppose that $f : S \rightarrow M$ is an injective continuous map, where S and M are n - and m -manifolds, respectively, where $1 \leq n < m \leq 3$. (Note that this definition holds for general topological spaces, but this restricted definition suffices for




Figure 2.2: The two surfaces on the left are embeddings of a sphere and of a torus. A Klein bottle cannot be embedded in \mathbb{R}^3 as represented on the right by a projection of an immersion of a Klein bottle.

our purposes. See [Mun00]) Let $X \subset M$ be the image set $f(S)$, considered as a subspace of M . Then the function $f' : S \rightarrow X$ obtained by restricting the range of f is bijective. If f' happens to be a homeomorphism of S with X , we say that the map $f : S \rightarrow M$ is an **embedding** of S in M . Informally, an embedding of a curve on a surface or of a surface in a manifold is one with no self-intersections, or, in other words, one which is “globally one to one”. A generalization of an embedding is, then, an **immersion** g under which a curve or surface is “locally one to one” or “locally embedded” (for a comparison between embedded and immersed surfaces, see Figure 2.2). Thus, a **proper embedding** f is an embedding which sends boundary to boundary, and a **proper immersion** g is an immersion which sends boundary to boundary. For a discussion on properly immersed surfaces in hyperbolic 3-manifolds, see section 3.3.

Generally, a **compressing disk** $D \subset M$ for a surface S is such that $D \cap S \supset \partial D$, which does not bound a disk on S . A surface is **incompressible** if there exists no compressing disk (otherwise **compressible**), and if a surface is compressible, we may perform a **compression** by removing a neighborhood $N(\partial D)$ of $\partial D \subset S$ from S . Assuming that $D \cap S = \partial D$, we may glue in two copies (D_1 and D_2) of D on each of the two boundary components in $\partial N(D)$.

Notice that this reduces the genus, or complexity, of the surface, a fact that can be verified by comparing the Euler characteristics of S and S' . We call a surface on which a compression has been performed a **compressed surface**.

A **boundary-compressing disk** is a topological disk $D \subset M$ such that $\partial D = \alpha \cup \beta$ where $\alpha \subset S$, α does not cut a disk off of S , and $\beta \subset \partial M$. (Note that we can assume that ∂D is embedded in S for reasons we will outline in Chapter 5, but we assume neither that $S \cap D = \partial D$ nor S embedded.) We will say that S is **boundary-incompressible** if there exists no such ∂ -compressing disk (otherwise **boundary-compressible**). While $\alpha \subset S$ need not be a simple curve, we will call a surface **simple curve (s.c.) boundary-incompressible** if there does not exist a boundary-compressing disk D with boundary $\partial D = \alpha \cup \beta$ with α a simple curve. One can simplify an s.c. boundary-compressible surface through performing a **boundary compression** by removing a neighborhood of D and gluing in two copies of it back onto S , making the resulting surface, S' , a topologically simpler object. We call S' a **boundary-compressed surface**. A surface which is both incompressible and boundary-incompressible is called **essential**.

Finally, we define covering spaces. Our arguments will occasionally be constructive in nature, and some of these constructions will rely on covering space theory since covering spaces retain some of the good and lose some of the recalcitrant properties of our base spaces. Informally, a **covering space** for (or cover of) an n -manifold M is an n -manifold \widetilde{M} that locally looks just like M and glues up like M but has m times more stuff than M . Formally, let X and \widetilde{X} be path connected, locally path connected spaces, and let $p : \widetilde{X} \rightarrow X$ be continuous. The pair (\widetilde{X}, p) is called a covering space if

1. p is surjective, and
2. for each $x \in X$, there exists an open set U in X containing x such that $p^{-1}(U)$ is a disjoint union of open sets, each of which is mapped homeomorphically onto U by p .

For example, if $M = \mathbb{S}^1 = \{(x, y) \in \mathbb{R}^2 : \sqrt{x^2 + y^2} = 1\}$ and $\widetilde{M} = \mathbb{R}^1$, then (\widetilde{M}, p) is a covering space for M if $p(x) = (\cos x, \sin x)$.

Chapter 3

Hyperbolic 2- and 3-Manifold Theory

In this section we outline some fundamental aspects of hyperbolic 3-manifolds, and we include a brief report on the history of the subject.

3.1 History

Around 300 BC, Euclid wrote *The Elements* in which he lists five axioms and derives consequences from those axioms. The most important of these five axioms is the fifth, which he stated as

That, if a straight line falling on two straight lines make the interior angles on the same side less than two right angles, if produced indefinitely, meet on that side on which are the angles less than two right angles.

Many false proofs attempting to show that this axiom could be derived from the other four were produced through the early 19th century, and while the proofs themselves were incorrect, some important results came out of this work. Indeed, while both of the following statements were known to be equivalent to Euclid's fifth axiom, each were reconsidered since several of those false proofs assumed one or the other (thereby rendering the arguments contained therein circular):

- Given a line and a point not on the line, it is possible to draw exactly one line through the given point parallel to the line.

- The sum of the angles of a triangle is equal to two right angles.

During the early 19th century, both Gauss and Bolyai independently convinced himself that a new, non-Euclidean geometry was possible. That is, it was believed that a geometry which satisfied the first four of Euclid's axioms but not the last existed. Lobachevsky, in his *Geometrical investigations on the theory of parallels* published in 1840, described a non-Euclidean geometry which replaces Euclid's fifth axiom with the following, now known as Lobachevsky's Parallel Postulate:

There exist two lines parallel to a given line through a given point not on the line.

This statement clearly invites the question of whether a geometry exists such that there exist no lines parallel to a given line through a given point not on the line. Riemann published results confirming the positive, and he called the resulting geometry "spherical". In 1868, Beltrami wrote "Essay on the interpretation of non-Euclidean geometry" in which he constructs a model for a 2-dimensional non-Euclidean geometry within a 3-dimensional Euclidean space. This model is known as the *psuedo-sphere*. This discovery quieted all controversy as it provided a concrete model on which Euclid's first four axioms hold but the fifth does not. Using a generalized notion of distance formulated by Cayley in 1859, Klein completed these explorations by showing that there are basically three kinds of 2-dimensional geometries and giving models for each.

Thus, through Klein in the late 19th century, we have the following three 2-dimensional geometries, where a line and a point not on the line are given:

1. **Euclidean** There exists exactly one line through the point parallel to the first line.
2. **Hyperbolic** There exist an infinite number of lines through the point parallel to the first line.
3. **Spherical** No such parallel line through the point exists.

Common models for 2-dimensional versions of the two non-Euclidean geometries are upper-half plane model of hyperbolic 2-space (the shortest lines, or **geodesics**, are vertical half-lines and semi-circles) and the unit sphere model of spherical 2-space (the geodesics are "great circles"). One important note

is that due to the above reformulations of Euclid's fifth axiom, the interior angles of spherical triangles sum to greater than π radians, and those of hyperbolic triangles sum to less than π radians.

Out of all closed, orientable 2-manifolds, all but two types are hyperbolic. That is, orientable surfaces of genus 0 are spherical, those of genus 1 are Euclidean, and the rest are hyperbolic. The genus $g = g(S)$ and the Euler characteristic $\chi = \chi(S)$ of a surface S are related in the following way:

$$\chi = 2 - 2g$$

Therefore, we see that all closed orientable hyperbolic surfaces have negative Euler characteristic. If the Euler characteristic of a hyperbolic surface is calculated using an ideal triangulation such that no edge can be homotoped into the boundary, then we can use this triangulation to calculate the area of the surface because it can be lifted to \mathbb{H}^2 where the area of each ideal triangle is π . In fact, it turns out that the area of a hyperbolic surface S is $\text{Area}(S) = 2\pi|\chi(S)|$, a fact which we will be using to find our bounds below.

3.2 Elements of Hyperbolic Geometry

Any manifold M which admits a metric of constant sectional curvature $K = -x$ for some $x > 0$ can be scaled so that M has constant curvature $K = -1$, and thus all hyperbolic spaces we consider will have $K = -1$. In such a space, all triangles which include their vertices have area strictly less than π , and triangles which do not include their vertices (called **ideal triangles**) have area equal to π . In fact, the set of all ideal triangles (where s is a 2-simplex)

$$\mathbb{T} = \{T \subset \mathbb{H}^2 : T = [s] \setminus \{v_0, v_1, v_2\}, v_i \in \partial\mathbb{H}^2 \text{ for } 0 \leq i \leq 2\}$$

is identical with the set of all hyperbolic triangles with area equal to π . To see this, calculate the area of the hyperbolic triangle T with vertex set $\{(-1, 0), (1, 0), \infty\}$ for $-1, 1 \in \partial\mathbb{H}^2 \setminus \{\infty\}$ by integrating in \mathbb{H}^2

$$\begin{aligned} \int_{-1}^1 \int_{\sqrt{1-x^2}}^{\infty} \frac{1}{y^2} dy dx &= \int_{-1}^1 \left(0 - \left(\frac{1}{\sqrt{1-x^2}} \right) \right) dx \\ &= \sin^{-1}(1) - \sin^{-1}(-1) = \pi \end{aligned}$$

There exists an isometry from T to any other triangle T' with vertices in $\partial\mathbb{H}^2$. Thus the area of any ideal triangle T' is equal to the area of T , or π . This stands in stark contrast to the analogous Euclidean case, where given an area $a > 0$, one can clearly construct a triangle T with vertices at the origin, $(1,0)$, and $(0,2a)$ such that $A(T) = a$.

3.3 Universal Covers of Hyperbolic Spaces

Every closed hyperbolic 2-manifold can be described abstractly as \mathbb{H}^2/Γ where Γ is some torsion-free (all elements have infinite order) discrete group of isometries on \mathbb{H}^2 . These isometries fall into three categories: parabolic, hyperbolic, and elliptic. Parabolic isometries are conjugate to translations of the form $f(z) = z + b$ for some $b \in \mathbb{C}$, and they fix one point in $\partial\mathbb{H}^2$. Hyperbolic isometries are conjugate to those of the form $f(z) = wz$ for some $w \in \mathbb{C}$, and they fix two points on $\partial\mathbb{H}^2$. Elliptic isometries fix an infinite number of points in \mathbb{H}^2 , not just $\partial\mathbb{H}^2$ like the other two types of isometries. Each of these isometries corresponds to elements of

$$\mathrm{PSL}(2, \mathbb{C}) = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} : a, b, c, d \in \mathbb{C}; ad - bc = 1 \right\}$$

the group of 2×2 matrices with complex entries and determinant equal to 1. Specifically, elliptic isometries are elements of $\mathrm{PSL}(2, \mathbb{C})$ with the square of this trace (tr^2) less than 4, hyperbolic isometries are elements with tr^2 greater than 4, and parabolic isometries are elements with tr^2 equal to 4.

All closed hyperbolic 2-manifolds of finite area are (infinitely) covered by \mathbb{H}^2 , and for this reason we call \mathbb{H}^2 the universal cover for hyperbolic 2-manifolds. A useful fact is that all and only hyperbolic surfaces, i.e. surfaces which admit a metric of constant curvature -1, have negative Euler characteristic. Note that a spherical metric of constant curvature 1 can be put on \mathbb{S}^2 and a Euclidean metric of constant zero curvature can be put on the torus whose universal cover is \mathbb{E}^2 .

Analogously, \mathbb{H}^3 is the universal cover of all hyperbolic 3-manifolds. We will be primarily interested in hyperbolic knot complements. We call the knot K hyperbolic if $M = \mathbb{S}^3 \setminus K$ is a hyperbolic 3-manifold, meaning that the universal cover of M is \mathbb{H}^3 . We will refer to $C = N(K)$, a tubular neighborhood of $K \subset \mathbb{S}^3$ as the **cusps** of $\mathbb{S}^3 \setminus K$; unless otherwise noted, we will assume C to be the **maximal cusp**, which is the cusp “blown up” until




Figure 3.1: A torus with meridian m and longitude l .

it first touches itself in M . Notice that C is homeomorphic to $\mathbb{T}^2 \times [0, 1)$, or a solid torus missing its core curve, and thus ∂C is homeomorphic to \mathbb{T}^2 . Here we should note that in a fundamental domain in \mathbb{H}^3 , the lift of C is covered by a finite number of disjoint horoballs, a fact which can be used to produce the details of the argument used to, for example, calculate Lackenby's upper bounds on cusp volume as a consequence of his upper bounds on manifold volume. Additionally, we will sometimes refer to the knot exterior $M_k = \mathbb{S}^3 \setminus \text{int}(C)$ (where $\text{int}(C)$ denotes the interior of C) when we require a compact manifold with boundary.

We are interested in the most simple generating curves on ∂C , the **meridian** and the **longitude** of \mathbb{T}^2 since each is a knot invariant and particular surfaces in $\mathbb{S}^3 \setminus K$ give upper bounds on each, as we discuss below. Neither a meridian nor a longitude can be continuously deformed through \mathbb{T}^2 into the other, a fact which can be checked by homology arguments. We will call a curve $m \in \partial C$ a meridian if it bounds a disk in C but not in ∂C . We will call a curve $l \subset \partial C$ a longitude if it is homotopic to a simple closed curve which not the product of some curve with any number of meridians. (See Figure 3.1 for a picture of a meridian and a longitude on a torus.) Additionally, we will call a curve that is the composition of p meridians and q longitudes a **(p,q)-curve**. In order to facilitate our arguments, we will define an l -curve to be any $(1, q)$ -curve on ∂C .




Figure 3.2: A properly immersed surface S in a 3-manifold M .

Here we revisit the notion of a proper map of a surface into a 3-manifold since we can rephrase it in a way that will aid our constructive arguments below. Informally, a proper map sends the cusps of the surface out the cusps of the 3-manifold (see Figure 3.2, which depicts a 5-cusped 3-manifold M with some surface S properly mapped into it. This is merely a heuristic diagram, but it is helpful because it gives some vague yet concrete picture to something quite abstract.) We say that each cusp C_i of M is the region which looks thin (but is thick in terms of the distance between points) from the ends of the manifold to the circles drawn to give M some dimension, and we note that each $C_i \cong \mathbb{T}^2 \times [0, 1)$, as depicted in Figure 3.2. Therefore, we will call a surface S properly embedded into M if each of its boundary components are embedded in this region. Similarly, we will call a surface S properly immersed into M if S is immersed and its boundary is shot out the cusp of M . This is clarified by Figure 3.3. Note that whenever we have a proper map of a surface S into a manifold M , it will locally look something like Figure 3.3. We also grant ourselves the freedom to push ∂S around within $C \cap S$ if necessary, but in most cases, all we need is that particular boundary components which will be the objects of our focus lie somewhere between $C \cap S$ and the puncture (in the case that the boundary component




Figure 3.3: Here, S is at least locally a proper map since its boundary is contained entirely within the cusp of M .

is a meridian) or between $C \cap S$ and the ℓ -curve boundary (in the case that the boundary component is an ℓ -curve).

3.4 Pleated Surfaces

Define a partition $P = \{P_\alpha\}$ of an n -manifold M^n into (disjoint) path-connected subsets of M^n to be a **foliation** of codimension c if there exists an cover by open sets $U \subset M$ of M^n such that each U comes with a homeomorphism $h : U \rightarrow \mathbb{R}^n$ or $h : U \rightarrow \mathbb{R}_+^n$ and sends each component of $U \cap P_\alpha$ to a hyperplane parallel to the standard hyperplane \mathbb{R}^{n-c} in \mathbb{R}^n . Each P_α is called a **leaf** of the foliation P . A **lamination** is a foliation of a closed subset of M^n . We can attain our goal of calculating Euler characteristic in terms of the number of boundary components of $S_k = S \cap M_k$ by laminating our surface so that we can spin geodesics around the boundary components to simulate an ideal triangulation of S . Since the boundary of properly immersed surfaces are shot out of the cusp of M and therefore are ideal vertices

in some valid ideal triangulation, we will not need to perform this operation in proving our main result, but we will need to perform it in Chapter 7 where we construct surfaces which are not necessarily properly mapped into M . A **pleated surface** is one which has been lifted to \mathbb{H}^3 and each ideal triangle in the ideal triangulation has been straightened. Then the triangles in its ideal triangulation have area π , allowing us to calculate the surface area of a hyperbolic surface S to be $A(S) = 2\pi|\chi(S)|$. Whenever we have an abstract hyperbolic surface S , we will calculate its area by lifting it to \mathbb{H}^3 and pleating it.

3.5 Cusp Diagrams

All hyperbolic knot complements can be decomposed into a finite number of ideal tetrahedra which can be lifted to the universal cover \mathbb{H}^3 by p^{-1} , the inverse of the projection map $p : \mathbb{H}^3 \rightarrow \mathbb{S}^3 \setminus K$. The cusp lifts to a disjoint set of **horoballs**, which are balls tangent to $\partial\mathbb{H}^3$. We will sometimes refer to their boundary spheres, which we call **horospheres**. By convention, we say that a horoball is centered at $x \in \partial\mathbb{H}^3$ if it is tangent to $\partial\mathbb{H}^3$ at x . Centering a horoball at ∞ yields a horizontal plane and the space above it. Doing so allows us to easily picture a fundamental parallelogram region in the boundary plane which, when opposite edges are identified with no twist, corresponds to the boundary of the cusp in $\mathbb{S}^3 \setminus K$. The length of one of these sides is the length of the meridian of C which we denote by $|m|$, and the length of an adjacent side is the length of the longitude of C , which we denote by $|\ell|$.

If we blow up C until it first touches itself (tangentially), then we will call the resulting cusp **maximal**. The maximal cusp horoball diagram may be rearranged through isometries (a subset of Γ) in such a manner that one horoball is centered at ∞ , and we can normalize so that its boundary horosphere is the horizontal plane $z = 1$. Therefore, in a manner similar to how we calculated the (hyperbolic) area of a hyperbolic triangle, we can calculate the maximum (hyperbolic) volume of a hyperbolic chimney by evaluating the following:

$$\int_1^\infty \int_0^{|m| \sin \theta} \int_0^{|\ell| \sin \phi} \frac{1}{z^2} dx dy dz \leq \int_1^\infty \int_0^{|m|} \int_0^{|\ell|} \frac{1}{z^2} dx dy dz$$

$$\begin{aligned}
&= \int_1^\infty \int_0^{|m|} \frac{|\ell|}{z^2} dy dz \\
&= \int_1^\infty \frac{|m||\ell|}{z^2} dz \\
&= \frac{|m||\ell|}{2} (1 - 0) = \frac{|m||\ell|}{2}
\end{aligned}$$

Thus we find that the volume of this chimney is no more than $\frac{1}{2}|m||\ell|$. This chimney volume is called the **maximal cusp volume** for single-cusped 3-manifolds and is always strictly smaller than the volume of the manifold itself because the cusp is a strict subset of its manifold. In fact, the largest ratio of maximal cusp volume to manifold volume, or **cusp density** is realized in the figure-8 knot complement and is

$$\frac{\sqrt{3}}{2v_3} \approx .853$$

where $v_3 = 1.01494\dots$ is the volume of an ideal tetrahedron of maximum volume in \mathbb{H}^3 as described in Chapter 1. One can construct examples of hyperbolic 3-manifolds with arbitrarily small cusp density; for one such construction, see [ACF⁺01].

Chapter 4

Boundary-Compressible Surfaces

In this section we investigate what happens to properly immersed surfaces S with c meridinal punctures and one ℓ -curve boundary component when we perform ∂ -compressions on S . Note that this will allow us to argue that, as long as the process of performing boundary compressions ceases at some point, and as long as the bounds we get from boundary-compressed surfaces are at least as good as their parent surfaces, then the bounds we find from the original surfaces must hold.

Since the pre-image of our original surface is a surface with $c+1$ punctures such that exactly one of the punctures corresponds to an ℓ curve on the cusp and c of the punctures correspond to meridinal curves on the cusp boundary, each possible ∂ -compressing curve naturally falls into one of four categories (see Figure 4.1):

1. $m \leftrightarrow m$ meridian to itself
2. $m_1 \leftrightarrow m_2$ one meridian to another
3. $\ell \leftrightarrow m$ longitude to meridian
4. $\ell \leftrightarrow \ell$ longitude to itself

In each case, we will investigate what the surface looks like locally in a neighborhood $N(\alpha)$ around the boundary-compressing arc $\alpha \subset S$.




Figure 4.1: The four cases to check.

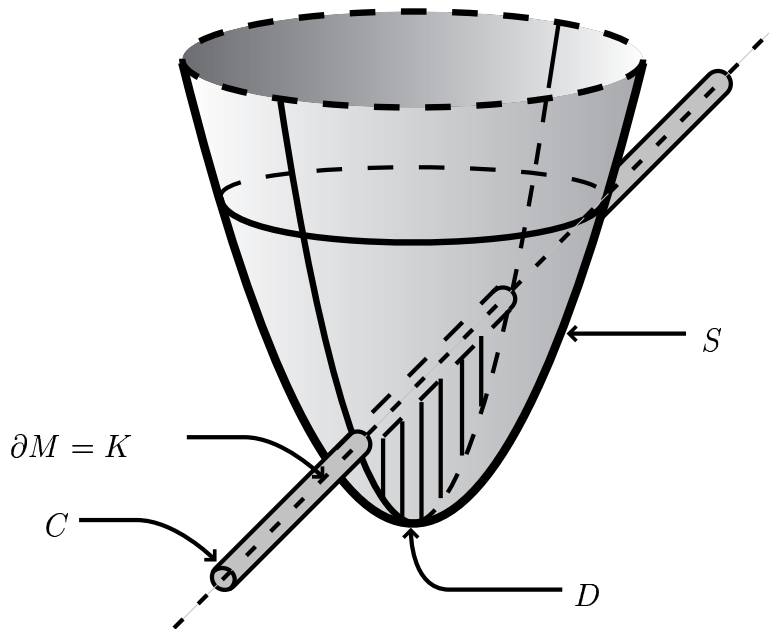


Figure 4.2: Here, $\alpha \cup \beta$ bound a ∂ -compressing disk D .

4.1 Performing Boundary Compressions

Now we describe the precise method of performing a ∂ -compression on particular surfaces. First we will describe the case which is easiest to represent as a figure and because of this probably the most simple to understand; however, it only applies to a small subset of the ∂ -compressions that we are interested in. Let $M = \mathbb{S}^3 \setminus K$ where K is some hyperbolic knot, and let $f : S \rightarrow M$ be a proper embedding of some surface S with two boundary components such that $f(\partial S)$ is two meridians, as in Figure 4.2. Now there exists an embedded ∂ -compressing disk D such that when we remove a neighborhood of it and glue in two copies of it along the new boundary, we have a new, boundary-compressed surface. The resulting surface has only one boundary component, which happens to be a trivial curve in ∂M . This is the simplest case, however, and the non-embedded case is a bit less straight-forward.

We now turn to a definition which is essential to our arguments in Sections 4.2.1, 4.2.2, 4.2.3, and 4.2.4. In each of the below cases we will refer to the neighborhood of α by $N(\alpha) \subset S$ which is a band-like region of S that will

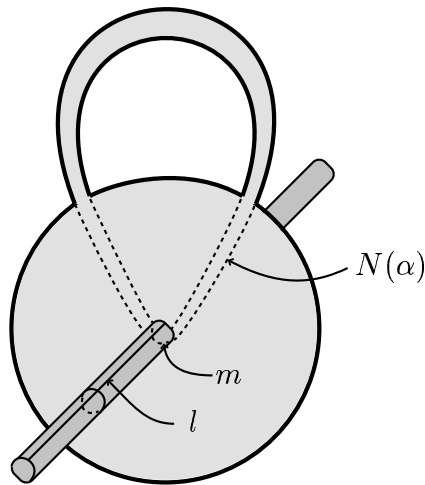


Figure 4.3: Form an annulus A by connecting the ends of $N(\alpha)$ to each other by a disk in C which intersects neither l nor m in its interior. Then the number of half-twists in $N(\alpha)$ is the number of half-twists in the A .

help to determine what the boundary-compression on S around D looks like. Indeed, the surface which results from that boundary compression depends upon the number of half-twists in that band. The definition for “number of half-twists” depends on which of the four cases (see Chapter 4) we are working in, so we will define it here for each of those four cases. Generally speaking, however, we can say that in each case the number of half-twists in $N(\alpha)$

In the case that α is a simple arc connecting a meridian $m \in \partial S$ to itself, connect $N(\alpha)$ to a disk $D \subset C$ such that $\partial D \cap (l \cup m) = \emptyset$, where l is a longitude of C which does not intersect $\partial N(\alpha)$. Figure 4.3 shows the case where $N(\alpha)$ has zero half-twists, and one can see that connecting the two ends of $N(\alpha)$ by a disk in C which does not intersect l or m in its interior will produce an annulus with zero half-twists. Indeed, in general, this procedure will produce an annulus with a well-defined number of half-twists up to \mathbb{Z}_2 .

The second case is when α connects one meridian to another in S . In this case, form an annulus A by connecting the ends of $N(\alpha)$ to each other by a disk in C which does not intersect l , m_1 , or m_2 in its interior. Again, define the number of half-twists in $N(\alpha)$ to be the number of half-twists in A . We see that the number of half-twists in $N(\alpha)$ is again clearly well-defined




Figure 4.4: Here, the boundary of a boundary-compressing disk is shown.




Figure 4.5: The boundary-compressing disk D with boundary $\partial D = \alpha \cup \beta$ is shown here.

up to \mathbb{Z}_2 . Since whether $N(\alpha)$ has an even or odd number of half-twists is the most important factor in determining what the corresponding boundary components look like, we only require that in each of the four cases, the number of half-twists is well-defined up to odd or even, or up to \mathbb{Z}_2 . See figures 4.4, 4.5, and 4.6 for detailed pictures of the case that $N(\alpha)$ has zero half-twists.

The third case is when α connects a meridian m to the ℓ -curve boundary component, ℓ , in S . Again, find a longitude l which does not intersect $\partial N(\alpha)$, and form an annulus A by connecting $N(\alpha)$ to a disk in C whose interior intersects neither l nor m . Define the number of half-twists of $N(\alpha)$ to be the number of half-twists in A .

Finally, in the case where α connects an ℓ -curve boundary component ℓ to itself, find a meridian $m \subset C$ which does not intersect $\partial N(\alpha)$, and form an annulus A by connecting the ends of $N(\alpha)$ with a disk D in C such that $\text{int}(D) \cap (\ell \cup m) = \emptyset$. Once again, define the number of half-twists of $N(\alpha)$




Figure 4.6: A boundary compressed surface S' in the case that α connects one meridian to another and B has zero half-twists.

to be the number of half-twists in A . We see that this is well-defined up to \mathbb{Z}_2 .

4.2 Four Cases

Now, using the boundary-compressing methods described above, we will investigate each of these four cases and show that the resulting boundary-compressed surfaces are topologically simpler objects. Because of this, we will show in Chapter 5 that whether or not properly immersed surfaces boundary compress is immaterial as our bounds on meridian length, longitude length, and cusp volume hold in either case.

4.2.1 Meridian to itself

Here we consider the case where there exists a boundary-compressing disk D with simple boundary $\partial D = \alpha \cup \beta$, where $\alpha \subset S$ connects a meridian $m \subset \partial S$ to itself, β is the connecting arc on ∂M , and α does not cut off a disk on S . Recall the notation used in section 4.1. Let A be the annulus with one boundary component that is $\gamma_1 \subset \partial C \cap S$ and another boundary component that is the closed curve $\gamma_2 \subset \partial S$. (see Figure 4.7). Form $B = A \cup N(\alpha)$.

Lemma 4.1 *Let S be a properly immersed disk with c meridional boundary components and exactly one ℓ -curve boundary component such that S may be boundary-compressed along an arc $\alpha \subset S$ which connects a meridional boundary component to itself. Then performing this boundary compression will result in a surface S' such that $g(S') = 0$, S' has strictly fewer meridional boundary components than S , and S' has exactly one ℓ -curve boundary component.*

Proof. We can rule out a few possibilities from our list of potential ∂ -compressions since a neighborhood of a simple closed curve in S must be an annulus rather than a Möbius band. Thus, since an odd number of half twists in the band implies that there is an Möbius band in S and therefore is not a disk, we can conclude that $N(\alpha)$ must have $n \in 2\mathbb{Z}^+ \cup \{0\}$ half twists.

We are not concerned with the curve of intersection between D and S because it has no effect on the consequences of the theorem. Notice that the condition that $\alpha \cup \beta = \partial D$ is a nontrivial element of $\pi_1(S)$ disallows ∂D to be homotopic to the ℓ -curve boundary component $\ell \subset \partial S$, and thus if




Figure 4.7: The boundary of D in the simple case.

∂D is simple, then performing the boundary compression along it will result in either one or two surfaces, all of which have some number of meridional boundary components, and in the second case one of the two surfaces will also have a ℓ -curve boundary component as well, thereby rendering it a surface from which we can derive the desired bounds. This is because in order to finish the boundary-compression, if at this point we have an embedded (but not properly mapped) surface, we may cap off either of the two pieces of S' and take whichever piece has the ℓ boundary curve. We are assured that this piece will also have at least one meridional puncture because there exists no map from an abstract disk D into some hyperbolic 3-manifold $\mathbb{S}^3 \setminus K$ which is a proper immersion with ∂D mapping to $\ell \approx K$ in \mathbb{S}^3 ; that is, there is no way to bound a disk with a hyperbolic knot in such a way that K does not intersect D .

Generally, since $N(\alpha)$ has an even number of half-twists, the first step of performing a boundary-compression along it will result in two new disks S_1 and S_2 . This is because we can remove a neighborhood of α ($N(\alpha) \subset N(\alpha)$) in such a way that $\partial N(\alpha)$ is parallel in S to $\partial N(\alpha)$. Therefore, since S is a surface of genus 0 and thus each element in $\partial N(\alpha)$ must be capped off by a

disk, then S' is not connected. Again, one of those surfaces has one ℓ -curve boundary component and at least one meridional puncture because if it did not, then ∂D would homotope to $\ell \subset \partial S$, which is a contradiction because a hyperbolic knot cannot bound a disk. Also notice that when $n \in 2\mathbb{Z}^+$, each piece of S' will intersect itself (as well as the other piece), but this does not affect Euler characteristic. Therefore we have shown that since we can choose the piece of S' which retains $\ell \subset \partial S'$, then we can cap it off with a disk $E \subset M$ such that $\partial E = \partial S' \setminus \partial S \subset \partial N(\alpha)$.

□

4.2.2 One meridian to another

This case is slightly more tricky than the last because we cannot get rid of the case where $N(\alpha)$ has an odd number of half-twists so easily. However, as we will see, such cases do not pose much of a problem either since the corresponding boundary compressions lead to the desired consequences.

Lemma 4.2 *Let S be a properly immersed surface with c meridional boundary components and exactly one ℓ -curve boundary component such that S may be boundary-compressed along an arc $\alpha \subset S$ which connects one meridional boundary component, m_1 , to another, m_2 . Then performing this boundary compression will result in a surface S' such that $g(S') = 0$, S' has strictly fewer meridional boundary components than S , and S' has exactly one ℓ -curve boundary component.*

Proof.

First we note that if $N(\alpha)$ has an even number of twists, then arguments analogous to those given in section 4.2.1 apply. For, again, we can ignore the curves of intersection between D and S because they do not affect the calculation of Euler characteristic. This method of performing the boundary compression makes it much more clear that the resulting surface S' does not depend on the intersection curves between D and S . Perhaps the best way to see this is to recall the notation from section 4.1, and let $B = N(\alpha) \cup A_1 \cup A_2$ where $N(\alpha)$ is a neighborhood of the simple arc $\alpha \subset S$ connecting m_1 to m_2 and A_i (for $i = 1, 2$) is the annulus with $m_i \subset \partial S$ and the parallel curve in $\partial C \cap S$ as its boundary curves. Fix A_1 , and rotate A_2 around m_2 n times in the direction opposite the one that B is twisting, where $2n$ is the number of half-twists in B . Now we can continuously push $N(\alpha)$ entirely through




Figure 4.8: Performing a boundary compression on a surface whose constructed band $B = N(\alpha) \cup A_1 \cup A_2$ has a single half-twist. Notice that γ' , the boundary component which replaces γ_1 and γ_2 , is the product $\gamma_1 \circ \gamma_2$.

D in such a way that $\partial N(\alpha)$ never intersects ∂M . Since this is true for any neighborhood $N(\alpha)$ of $\alpha \subset S$, performing that boundary compression is superficially the same as performing one where $N(\alpha)$ has no twists since in each case the operation decreases the number of meridional punctures in S by two and alters the Euler characteristic accordingly.

However, since a band B stretching between two meridians which has an odd number of half-twists exist in many hyperbolic knot complements, we should look at the corresponding boundary compressions in some detail. Note again that the arcs of intersection between D and S are inconsequential since they do not affect the number of boundary components or the Euler characteristic of S' . Even so, since we cannot perform the same trick as

above and push S through D so that $\partial N(\alpha)$ does not intersect ∂M , we must construct a surface which would be the result of a boundary compression along a disk whose boundary lies in $S \cup \partial M$. First, remove a neighborhood of $\alpha \subset S$ and add to C , a $(2, 0)$ -curve γ with singular set equal to n double points, where n is the (odd) number of half twists in $N(\alpha)$. Now connect $\partial N(\alpha)$ to γ with an annulus. (See Figure 4.8 for the case where B has a single half-twist.) This new surface is our boundary compressed surface because the operation we just performed is equivalent to adding two (mutually embedded) disks to the deleted neighborhood of α and then gluing it together along some (p, q) -curve $\gamma \subset \partial M$.

□

While this turns out to be drastically different from the other cases we have already looked at, it surprisingly gives us better bounds than the original surface since S' has one fewer (meridinal) boundary component than S had, and even though the new meridinal boundary component γ is not a simple closed curve on ∂M , the Euler characteristic is the same as it would be if all boundary curves were simple. We will use this fact to prove the lemmas needed for Theorem 5.2.

Next, we present an argument which we originally tried to use to reduce the number of cases to look at from four to three. In particular, we thought that this argument would allow us to collapse the first two cases presented above since, given a boundary-compressing disk D with boundary $\partial D = \alpha \cup \beta$ where α connects one meridian to another, it seems as though we could take a small neighborhood of α in S until it resembles a simple arc α' which connects one of the two boundary components to itself. However the following reduction argument fails because there is no clear way to glue in two copies of a disk in order to create the new ∂ -compressing disk in the case that $N(\alpha)$ has an odd number of half-twists. This is interesting because while seemingly more elegant than our above constructive approach, it ultimately fails because of the nature of non-embedded surfaces.

Here, we argue that from boundary compressions corresponding to the one meridinal boundary component to another type, we can construct a ∂ -compression of the type detailed in section 4.2.1.

Assume that $\alpha \subset S$ is a simple arc connecting two meridinal boundary components of S , m_1 and m_2 . By convention, we will say that m_i are the boundary curves of the intersection of the cusp C with S , and thus $m_i \subset$

$\partial C \cap S$. In addition, assume that $\beta \subset \partial M$ is such that $\alpha \cup \beta$ bounds a ∂ -compressing disk $D \subset M$. Let B be defined as in section 4.1. Now look at a curve γ homotopic in S to the product $m_1 \circ m_2$ in B . First we will show that γ bounds a compressing disk, and then we will argue how we can deform this into a ∂ -compressing disk of the type described above.

Remove a neighborhood $N'(\alpha) \subset N(\alpha)$ from B and put copies α_1, α_2 of α into B at $\partial N'(\alpha)$. Analogously, remove a neighborhood $N(D)$ of the boundary compressing disk D and glue in two copies D_1, D_2 of D so that D_i has boundary $\alpha_i \cup \beta_i$ (where β_i is parallel in C to β). Let $\gamma' \subset (\alpha_1 \cup \alpha_2 \cup m_1 \cup m_2)$ be a closed curve homotopic to γ , and let $A' \subset S$ be the annulus such that $\partial A' = \gamma \cup \gamma'$; clearly this is an annulus since, by construction, $\gamma \cong \gamma'$. Cap A' off by gluing the disk $D_c \subset \partial C$ which has boundary $\partial D_c = \gamma' \cap (m_1 \cup m_2) \cup \beta_1 \cup \beta_2$ and is the obvious subset of ∂C union $D_1 \cup D_2$.

Finally, let $D' = A' \cup D_c$. We claim that D' is a compressing disk. Since $\partial D' = \gamma$ by construction, and $D' \subset M$ is indeed a topological disk in M , we are done. Now notice that we can push a part of γ into, say, m_1 so that part of $\partial D' = \gamma'$ (the deformed γ) is a subset of C , or in other words $\gamma' \cap C \neq \emptyset$. Indeed, we can do this without revoking D' of its topological-disk status, and thus D' is clearly a ∂ -compressing disk of the type outlined in section 4.2.1. Therefore, it seems as though this argument shows that all boundary-compressions where α connects one meridian to another in S give rise to boundary-compressions where α connects a meridian to itself. However, this is not the case when B has some number of half-twists because it is unclear how to remove a neighborhood of D and glue in two copies D_1, D_2 so that the resulting object is a simpler surface.

The question of how to boundary compress along a non-simple arc α is something that we have spent some time on, and although our findings were inconclusive, we should include such work here in order to facilitate future endeavors similar to our own. While what we present in the remainder of this section is not precisely the same as the case presented in the previous paragraph, each case addresses the question of how does one remove a neighborhood of a non-embedded boundary arc of a necessarily non-embedded disk in such a way that the re-gluing of the copies D_1, D_2 are the correct generalization of the typical boundary-compression process which we use throughout this exposition.

Our idea was to remove a neighborhood of α and, wherever the neighborhood was removed twice—that is, wherever α crossed itself—glue in an extra “patch” to make the boundary-compressed surface S' a valid surface.




Figure 4.9: The dark grey region in the middle is the patch, and the point at the bottom right corner of this patch is a problematic point, but as shown there exists a homeomorphism from a neighborhood around that point to an open disk in \mathbb{R}^2 .

(See Figure 4.9 for a close-up look at the inserted patch and why it makes S' a valid surface.) This seems like a reasonable operation not only because it makes S' a valid surface but also because it seems as though it is the correct generalization of the standard “remove a neighborhood, glue in two copies” procedure one follows in the case that ∂D is embedded. Even so, we have decided that this is not the correct generalization of that procedure because boundary-compressions are supposed to result in topologically simpler surfaces whereas we found explicit examples of surfaces with some meridians and an ℓ -curve boundary component which, upon boundary-compressing along some non-embedded arc α in S , did not result in a simpler surface.

In fact, we found examples of disks which after compressing along such an α gave us surfaces S' with $g(S') > 0$, which we determined by calculating the Euler characteristic of the pre-images of those surfaces. Thus, our standard method was the following: pick a disk with some meridians and an ℓ -curve boundary, draw in some non-embedded arc which could make up part of the boundary of a boundary-compressing disk, boundary-compress along that

arc and an arc in the boundary of the manifold using the patching operation described above, look at the pre-image of the resulting space, and calculate the Euler characteristic χ of that pull-back. What we found was that $\chi < 1 - n$ (where n is the number of meridians in S) which indicates that S' is not a disk but instead some surface of higher genus. This is undesirable because the motivation for performing compressions and boundary-compressions is to simplify the surface, and therefore since our generalization does not simplify the surface, it is probably not a good candidate for the general boundary-compressing procedure.

Another general procedure which we considered briefly was to remove a different sort of neighborhood $N'(D)$ from S . Instead of removing $D \times [0, 1]$, we thought about removing a neighborhood of varying width. That is, let $\alpha(t)$ have singular points at $t = \{t_1, \dots, t_n\}$. Construct a “neighborhood” $N'(\alpha)$ as a subset of $N(\alpha)$ such that $N'(\alpha)$ shrinks to a point at each t_i for $1 \leq i \leq n$. Now remove $N'(\alpha)$ from S and glue in two copies D_1, D_2 along $\partial N'(\alpha)$ so that we have added precisely n singular arcs to S in making it S' , a boundary-compressed surface. This seems promising because it does not directly contradict the standard embedded-boundary case, but we have not carried these ideas out to their full fruition.

4.2.3 Longitude to a meridian

Now we investigate the case where $\alpha \subset S$ connects the ℓ -curve boundary component to a meridional puncture in S .

Lemma 4.3 *Let S be a properly immersed surface with c meridional boundary components and exactly one ℓ -curve boundary component such that S may be boundary-compressed along an arc $\alpha \subset S$ which connects the ℓ boundary component to a meridional boundary component. Then performing this boundary compression will result in a properly immersed surface S' such that $g(S') = 0$, S' has strictly fewer meridional boundary components than S , and S' has exactly one ℓ -curve boundary component.*

Proof. This case is marginally different from those we have looked at already. First we consider what performing the boundary-compressing operation looks like in M . Consider the neighborhood $N(\alpha) \subset S$ of $\alpha \subset S$ where $\alpha \cup \beta = \partial D$, the boundary of our boundary-compressing disk, and $\beta \subset \partial M$. We borrow notation from section 4.1 and say that $\gamma_1 = \ell$ and $\gamma_2 = m$ are the two

boundary components of S which α connects, and connect A_1 , the annulus around ℓ , with A_2 , the annulus around m , by $N(\alpha)$ to give us $B = N(\alpha) \cup A_1 \cup A_2$. Since this case is almost identical to that of section 4.2.2, apply the same arguments. The primary difference between the two cases is that performing the boundary compression in this case will result in an unwinding of the ℓ -curve so that the resulting ℓ -curve boundary of S' is the original ℓ -curve composed with m . Also, since m has been subsumed into the resulting ℓ -curve boundary component of S' , S' has at least one fewer meridians than S . Also, this procedure does not add genus to S since we are only gluing in disks along closed curves which are mutually simple (temporary) boundary components of S' , and thus $g(S') = g(S) = 0$.

□

4.2.4 Longitude to itself

Finally, we prove the familiar lemma in the case where α connects ℓ to itself.

Lemma 4.4 *Let S be a properly immersed disk with c meridional boundary components and exactly one ℓ -curve boundary component such that S may be boundary-compressed along an arc $\alpha \subset S$ which connects the ℓ boundary component to itself. Then performing this boundary compression will result in a surface S' such that $g(S) = g(S')$, S' has strictly fewer meridional boundary components than S , and S' has exactly one ℓ -curve boundary component.*

Proof. Begin again by recalling our notation from section 4.1. If B has an even number of half-twists, then removing $N(\alpha)$ from S will result in two surfaces S'_1, S'_2 , where S'_i ($i = 1, 2$) has an ℓ -curve boundary which corresponds to the original ℓ -curve boundary component in S composed with the number of meridians on S'_j for $i \neq j$. Therefore, both S'_1 and S'_2 have an ℓ -curve boundary as well as at least one meridian. Also, this procedure does not add genus to the surface, so $g(S'_1) = g(S'_2) = g(S) = 0$. In the case that B has an odd number of half-twists, S' is a single connected surface with an ℓ -curve boundary which is the original ℓ -curve composed with n meridians, where n is the number of meridians in the original surface S . Therefore, S' is a surface of the desired type, and we are done.

□

4.3 Summary

Thus, in this section, we have shown the following.

Theorem 4.5 *Let S be a properly immersed surface with c meridional boundary components and exactly one ℓ -curve boundary component. Then any boundary compression will result in a surface S' such that $g(S) = g(S')$, S' has strictly fewer meridional boundary components than S , and S' has exactly one ℓ -curve boundary component.*

Proof. Apply lemmas 4.1, 4.2, 4.3, and 4.4.

□

We will use this theorem to find bounds on meridian length, longitude length, and cusp volume in Chapter 5.

Chapter 5

Cusp Size Bounds

In [ACF⁺01], the authors find upper bounds for meridian length, longitude length, and cusp volume for a given cusp C in a particular 3-manifold M that contains certain surfaces. Here we describe those results in greatest generality, which depends on the work done in this paper. After outlining the conditions necessary of a surface in order to obtain bounds on longitude length, meridian length, and cusp volume, we state the results themselves.

5.1 Necessary Properties of a Useful Surface

In order for a surface S to give such bounds, three conditions must be satisfied: first, there must be an ideal triangulation Δ of S such that no edge of Δ can be homotoped into ∂S . In the case that we are dealing with non-ideal boundary components, we simulate an ideal triangulation by finding a geodesic lamination as described in section 2.3. Since edges cannot be homotoped into ∂S , we know that in lift of S along with its triangulation Δ to the universal cover \mathbb{H}^3 is such that edges pass between distinct horoballs, and this gives rise to a pleated surface S' in \mathbb{H}^3 which allows us to calculate the area of the surface S as $\text{Area}(S) = \text{Area}(S') = 2\pi|\chi(S)|$.

Second, we also require that the pleated surface S' intersects the cusp only in annuli so that no two distinct intersection curves of S' and the cusp intersect. This condition is guaranteed when S is s.c. boundary-incompressible which is a consequence of Lemma 5.1.

The third condition is that each surface from which we extract the desired bounds from must have exactly one ℓ -curve boundary curve and at least one




Figure 5.1: In this situation, a boundary-compressing disk exists with α as part of its boundary.

meridional boundary curve. In Chapter 4, we show that since boundary compressions only improve the bounds in which we are interested, we can replace our original He surface with one on which we have performed any outstanding boundary compressions. Thus, requiring that the resulting surfaces have such boundary components ensures that the bounds that we find below hold for any surface that results from performing boundary compressions on the original He surface.

5.2 Main Results

Here we state our main theorem (5.2). In our proof of Theorem 5.2, we will use the following lemma, which ensures that the area of S' is greater than or equal to the area of the cusp intersecting with the surface.

Lemma 5.1 *If S is an s.c. boundary-incompressible surface properly immersed in M , then the cusp $N(K)$ of M intersects S only in annuli.*

Proof. Proof by contrapositive. Assume that $N(K)$ intersects S in regions other than annuli. Then there either exists a simple arc α contained entirely within $S \cap N(K)$ which does not cut off a disk from S but connects one boundary curve to another which also sits in the area of intersection of the surface with the cusp. (See Figure 5.1.) Now, since the cusp is an embedded tubular neighborhood of the knot K in M , we are assured that there exists a homotopy between α and an arc $\beta \subset K$ through $N(K)$. Therefore there exists a disk $D \subset N(K)$ with boundary $\partial D = \alpha \cup \beta$ where $\beta \subset \partial M$. Clearly D is a boundary-compressing disk, which is a contradiction. This ends the proof. □

Theorem 5.2 *Let K be a knot with c crossings in projection $\pi(K)$, and let $|m|, |\ell|$ denote the lengths of the meridian m of the knot and the shortest ℓ -curve on the maximal cusp. The following bounds hold for K and its maximal cusp C :*

- $|m| \leq 6 - \frac{7}{c}$
- $|\ell| \leq 5c - 6$
- $V(C) \leq \frac{9c}{2} \left(1 - \frac{1}{c}\right)^2$

Proof. First, we construct a He surface S in $\mathbb{S}^3 \setminus K$ and notice that S has $c + 1$ boundary components, c corresponding to meridional punctures and 1 corresponding to a longitudinal boundary. By Lemma 5.1, we are assured that the sum of the areas of the regions of intersection between S and $N(K)$ will be less than or equal to the area of the entire surface S . Thus, since the area of each intersection region is equal to the length of its corresponding boundary curve and the area of S is equal to $2\pi\chi(S)$, we have $c|m| + |\ell| \leq 2\pi|\chi(S)|$. Here we use the fact that the ideal packing of pairwise tangent horospheres in \mathbb{H}^2 is such the volume of a manifold corresponding to the quotient space of an appropriate group of isometries is no more than $\frac{\pi}{3}$ times the hyperbolic volume of where those horoballs intersect that region, we find that $c|m| + |\ell| \leq 6|\chi(S)|$. (See Figure 5.2 for the geometric construction used to argue this.)

Now, $|m| \geq 1$ and $|\ell| \geq 1$ because, as seen in Figure 5.3, every horoball corresponding to the lift of the cusp of M will be tangent to another horoball,




Figure 5.2: The fraction of an ideal triangle's area which intersects the cusp can be no more than $\frac{3}{\pi}$.

and since the sides of the fundamental domain are ℓ and m which correspond to parabolic isometries, we can place the centers of copies of horoballs of maximal height at the corners of that fundamental domain. Since these horoballs are disjoint each with radius $1/2$, this implies that $|m| \geq h$ and $|\ell| \geq h$, and thus $|m|/h \geq 1$ and $|\ell|/h \geq 1$. Furthermore, since $h = 1$ by normalization, we find that $|m| \geq 1$ and $|\ell| \geq 1$. Also, since $|\chi(S)| = (c - 1)$ (since S is a $c + 1$ -punctured disk), we can compute first two the desired inequalities without much difficulty:

- $c|m| + |\ell| \leq 6c - 6$

$$\Rightarrow |m| \leq 6 - \frac{7}{c}$$

$$\Rightarrow |\ell| \leq 5c - 6$$

In order to prove the third bound, we fix c , the number of crossings of a projection of K , and maximize the quantity $f(|m|) = |m||\ell|$ as a function of one variable as follows:

$$f(|m|) = |m||\ell| = 6c|m| - 6|m| - c|m|^2$$




Figure 5.3: The lengths of ℓ and m are at least 1.

$$f'(|m|) = 6c - 6 - 2c|m|$$

$$0 = 6c - 6 - 2c|m| \text{ implies that } f \text{ maximized when } |m| = 3 - \frac{3}{c}$$

$$\begin{aligned} f(|m|) &\leq 6c \left(3 - \frac{3}{c} - 6 \left(3 - \frac{3}{c} \right) - c \left(9 - \frac{18}{c} + \frac{9}{c^2} \right) \right) \\ &= 9c - 18 + \frac{9}{c} = 9c \left(1 - \frac{1}{c} \right)^2 \end{aligned}$$

Since $V(C) \leq \frac{1}{2}|m||\ell|$, we have shown that

$$V(C) \leq \frac{9c}{2} \left(1 - \frac{1}{c} \right)^2$$

□

Since it is difficult to show that surfaces with one ℓ boundary curve and n meridional boundary curves are boundary-incompressible, we will show that these bounds are only improved if ∂ -compressions exist. In order to prove this, we will need to make use of the following lemma:

Lemma 5.3 *Let S be a properly immersed surface with c meridional punctures and one ℓ -curve boundary component, and let S' be a properly immersed surface with $c - k$ meridional punctures and one ℓ -curve puncture, where $1 \leq k \leq c - 1$. Then the bounds stated in Theorem 5.2 are improved.*

Proof. We apply Theorem 5.2 to S' and find the following bounds:

- $|m| \leq 6 - \frac{7}{c-k} < 6 - \frac{7}{c}$
- $|\ell| \leq 5(c-k) + 6 < 5c + 6$
- $\frac{|m||\ell|}{2} \leq \frac{9(c-k)}{2} \left(1 - \frac{1}{c-k}\right)^2 \leq \frac{9c}{2} \left(1 - \frac{1}{c}\right)^2$

since $1 \leq c - k \leq c - 1$.

□

Theorem 5.4 *Performing boundary compressions on a surface S with one ℓ -boundary component and at least two m -boundary components improves our bounds on meridian length, longitude length, and cusp volume.*

Proof. First notice that by arguments set forth in Chapter 4, we need only check four cases. By Theorem 4.5, performing any of these four boundary-compressions results in a surface S' which has at least one fewer meridional boundary components and exactly one ℓ boundary component. Now we can apply Lemma 5.3 to finish the proof.

□

In this section we have proved that the bounds which we find for meridian length, longitude length, and cusp volume hold even in the case that our He surface can be boundary-compressed. In the following two sections, we reproduce our attempts to strengthen these results as well as generalize He's construction to make surfaces which give better bounds in terms of toroidal (and higher genus) crossing number.

Chapter 6

Attempts at Stronger Results

Here, we list and describe some of our attempts to show various things relating to other results of this report. We include them for their potentially instructive content, and in fact some of the lemmas and corollaries which arise are of independent interest.

6.1 S.C. Boundary Incompressibility of He Surfaces

We spent some time investigating techniques to show that He surfaces S are boundary incompressible. An outline of the line of reasoning we followed is this:

1. Remove a neighborhood of x so that S sits in $B^3 \setminus K$ rather than $\mathbb{S}^3 \setminus K$. This is helpful because it is easier to take covers of B^3 than it is \mathbb{S}^3 , and it is possible because we can always push α off of x if necessary.
2. Take covers of $B^3 \setminus K$ so that S lifts to S' , for which proving ∂ -incompressibility is more tractable. (See also tower constructions, such as the one Jaco uses to prove the Loop Theorem: [Jac80]). Specifically, find a cover such that S' is the union of mutually embedded surfaces. (See also [ACF⁺01] where the authors use the same fact—that if there exists a boundary compressing-disk downstairs, then there exists one upstairs—to prove that every closed orientable 3-manifold M contains a knot K such that $M \setminus K$ is hyperbolic with cusp volume ≤ 9 .)

3. Use “Menasco machinery” to show that an embedded ∂ -compressing disk $D' \subset S'$ cannot exist. See also [ADFaSP99, Ada94b, Men84, MT91, MT92].

These were the tools with which we began our foray into He surfaces.

6.1.1 Drilling and Cyclic Covers

First, remove a neighborhood of the cone point x as described above. Next drill a number of holes in $M' = B^3 \setminus K$ so the result is $M'' = F \setminus K$ where F is some solid m -handlebody. The idea is then to take some cyclic n -fold cover of M'' so that there exists a projection of $\partial M''$ such that the two strands of K' at each crossing of that projection are on different components of K' (it turns out that K' is a link). By doing this, S' is the union of a set of mutually embedded annuli, and thus one may apply Menasco’s methods to show that a ∂ -compressing disk cannot exist. The problem with this approach is that the number of holes we drill affects the complexity of the problem, and there is no easy way to avoid making the problem really hard.

6.1.2 Induction on Crossing Number

Here, the idea was that if we can show that the He surface of the Hopf link is boundary-incompressible, and that if a ∂ -compressing disk exists in a link L with $c(L) = n$, then a ∂ -compressing disk also exists in the He surface corresponding to L' , the link which results when we break open a crossing in $\pi(L)$ (reduced, alternating) a certain way. We proved that given a reduced alternating projection of a link L with crossing number m , there exists a crossing and a direction to break open that crossing so that the resulting projection is reduced and alternating and corresponds to a link L' with $c(L') = m - 1$.

Lemma 6.1 *Given a reduced alternating projection $\pi(L)$ of an n -crossing alternating link L , there exists a crossing c such that $\pi(L')$, a projection of the link which results when c is replaced by a 0 - or ∞ -tangle in L , is a reduced alternating projection of an $(n - 1)$ -crossing link L' .*

Proof. Locally, a reduced alternating projection of an alternating link L looks like Figure 6.1. Note that the figure is a bit misleading because its north, south, east, and west crossings might correspond to the same strands




Figure 6.1: Crossing c is the one we break open. Note that crossing a might be crossing b rotated some $n\pi + \frac{\pi}{2}$ radians.

of L ; for example, crossing a might be the same as crossing b rotated some $n\pi + \frac{\pi}{2}$ radians. Replace c with a zero tangle rotated $-\frac{\pi}{4}$ radians. Then the north strand is connected to the east strand and the west strand is connected to the south strand. To show that the resulting projection $\pi(L')$ of the resulting link L' is reduced and alternating, it suffices to show that the local part of the projection shown in Figure 6.1 is reduced and alternating. If all crossings are distinct, then by the picture, this is clearly true since it is alternating and there are no outstanding Type I Reidemeister moves to be performed. Also, L' has one fewer crossing than L because this operation is local. If two crossings in Figure 6.1 are the same, then c is a crossing in some twist t (with $n > 1$ crossings) in L , and replacing c either results in a projection which may be reduced by $n - 1$ Type I moves or a projection with a twist t' with $n - 1$ crossings, so we may choose to replace c to achieve the latter. Therefore we have found a way to replace crossing c such that the projection $\pi(L')$ of the link L' is reduced, alternating, and has one fewer crossing than L .

□

The full-blown theorem did not play out because the lemma needed to invoke the inductive hypothesis turned out to be quite difficult to solve using

available means. Below we state this lemma, methods explored in attempting to prove it, as well as the theorem which would result if it were true.

Conjecture 6.2 *If a He surface S for a reduced, alternating projection of a link L with crossing number n is boundary compressible, then there is a link L' which differs from L by one crossing and which admits a reduced, alternating projection $\pi(L')$ for which a He surface S' is boundary-compressible.*

This conjecture is true in cases where $\beta = \partial D \cap \partial \mathbb{S}^3 \setminus L$ of the boundary-compressing disk D in S misses c and $\alpha \cup \beta = \partial D$ is not an ℓ -curve. This is because replacing c with a zero tangle will affect neither β nor α since S' is the same as S except below where we replaced c where we cut along the singular arc γ down to the cone point. However, β does not pass through c implies that α does not intersect γ , so cutting along γ will not affect α , and thus $\alpha \cup \beta$ bound a boundary-compressing disk D in S' . Notice that since we can choose any crossing to apply Lemma 6.1, the above argument holds for any situation where L is a link of more than one components, and in the case that L is a knot, then as long as β does not travel entirely around the knot, then we can find a crossing to which we can apply Lemma 6.1. However, we cannot prove Conjecture 6.2 in this last case. If we could prove this, then we would be able prove the following.

Conjecture 6.3 *If He surfaces for reduced, alternating projections of the Hopf link are boundary incompressible, then He surfaces for reduced, alternating projections of all alternating links are boundary incompressible.*

Proof. Let L be a reduced, alternating projection with n crossings which is boundary compressible. Then we can find a sequence of links $\{L_k\}_{k=1}^{n-2}$ with projections $\pi(L_k)$ which have $n - k$ crossings. This process terminates when $k = n - 2$ because Lemma 6.1 ensures that each $\pi(L_k)$ is reduced and alternating and thus $\pi(L_{n-2})$ must be a projection of the Hopf link. By Conjecture 6.2, a He surface for $\pi(L_k)$ is boundary compressible, but this contradicts our assumption that He surfaces for reduced, alternating projections of the Hopf link are boundary incompressible.

□

Proving that such projections of the Hopf link are boundary incompressible does not seem difficult, but we have not put much energy into it since proving Conjecture 6.2 is difficult. Even so, proving these would imply Conjecture 6.3, which is desirable.

6.1.3 Branching Covers

This attempt showed the most promise, and therefore I will spend a little more time developing the idea behind it here. The motivation for this approach was the failure of the first at the hands of our rampant drilling scheme. We still want to find a cover M' of M such that individual sheets of S' are embedded (although those sheets may intersect each other) so that we can use the Menasco techniques to prove that a ∂ -compressing disk for S' cannot exist in M' . Once we have that, it should be relatively easy to argue for why a disk in the original manifold cannot exist and thus the result would follow. A few things are crucial in making these arguments:

1. The original ∂ -compressing disk in M lifts to a disk in M' .
2. $\partial M'$ has the same alternating property as ∂M has because the Menasco techniques rely on it.
3. There exists a *single* branching axis such that taking an m -fold branched cover (for sufficient m) over it will result in S' having the properties outlined above.
4. It seems difficult to start with a (possibly) immersed disk in M and say that it lifts to an embedded disk in M' . If we know that if there exists an embedding of $\partial D'$, then we can use the Loop Theorem to show that there exists an embedding of D' and go from there, but it is not easy to show that we can find a cover so that such an embedding of the boundary of the disk exists.

It is easy to see *prima facie* that a reasonably chosen cover will preserve the alternating property and thus satisfy our second concern above. Even so, we did run into problems later on after the first concern turned out to be an insuperable difficulty. Before discussing these problems, we review how we dealt with the third concern.

At first, I was confident that given a reduced, alternating projection of a knot K , one could find a path γ which “cut” the projection in such a way that branching over γ would allay our third concern. This turned out not to be the case; however, I did discover that we can find such a γ for alternating knots which admit a reduced, alternating braid representation; we call such a knot an **alternating braid**. This is a consequence of the following lemma.




Figure 6.2: Constructing an alternating braid whose components do not cross themselves (although they cross each other).

Lemma 6.4 *Let L be an alternating braid with braid index b and reduced, alternating projection $\pi(L)$. Then, for some $m \in \mathbb{Z}^+$, there exists a way to compose m copies of L which gives a projection $\pi(L')$ of L' , a link with b components, each of which does not cross itself. In addition, $\pi(L')$ is reduced and alternating.*

Proof. Label the strands of $\pi(L)$ as s_1, s_2, \dots, s_b . Given a strand s_k , there is a number a_k such that when L is composed with itself a_k times, then the strand in the new link corresponding to s_k does not cross itself. Let

$$m = \prod_{k=1}^b a_k$$

and compose copies of $\pi(L)$ to give the link shown in Figure 6.2. The resulting projection $\pi(L')$ with strands $\{s'_k\}$ has the property that each strand s'_k connects to itself and thus $\pi(L')$ has b components, each of which does not cross itself. Additionally, $\pi(L')$ is alternating because composing copies of $\pi(L)$ with itself in the manner outlined above (and displayed in Figure 6.2)

preserves the alternating property since, for example, if the top end of s_1 in Figure 6.2 is adjacent to an undercrossing below it, then the bottom end of s_1 is adjacent to an overcrossing above it, and thus when we connect the bottom end of s_1 to a cut-open copy of L , it will be adjacent to an undercrossing below it. Therefore $\pi(L')$ is alternating. Finally, $\pi(L')$ is reduced because type I moves correspond to extra strands in the braid representation of a link, but since the number of strands in $\pi(L)$ is the same as the number of strands in $\pi(L')$, and since $\pi(L)$ is reduced, $\pi(L')$ must be reduced as well. This completes the proof.

□

With this, we focused on showing that a ∂ -compressing disk cannot exist in the branched cover taken over $B^3 \setminus K$ where K is an alternating braid. We present a rough outline of arguments which seem promising but ultimately fail due to the non-embedded quality of S and the fact that a boundary-compressing disk D may lift to a planar surface with more than one boundary component in M' . Refer to [ADFaSP99, Ada94b, Men84, MT91, MT92] for relevant pictures and details.

Conjecture 6.5 *The He surface of links which have a minimal crossing alternating braid representation are s.c. boundary-incompressible.*

Proof. Let L be such a link, and let $\pi(L)$ be a reduced, alternating projection of its associated braid. By Lemma 6.4, we can compose b with itself m times so that each component of the resulting link L' does not cross itself (although it crosses other components) and its projection $\pi(L')$ is alternating and reduced.

Let $M = S^3 \setminus L$, and let S denote a He surface generated produced from $\pi(L)$. Remove a neighborhood $U(x)$ of the cone point x so what remains is $S \setminus U(x) \subset B^3 \setminus L = M$. Now look at the m -fold branched cover M' of M , where $S \setminus U(x)$ lifts to S' which is a set of individually embedded (but collectively immersed) punctured annuli whose punctures correspond to meridional curves on components of the link other than that whose longitude is the boundary of it. If a ∂ -compressing disk exists in M , then it lifts to a planar surface in M' which has one boundary component $\alpha' \cup \beta'$ and no boundary components corresponding to single meridians of one component of L . This planar surface arises from the fact that we are branching over a particular axis in M , and if that axis intersects the boundary compressing

disk $D \subset M$, then D lifts to something other than a disk. Thus, if we show that no such planar surfaces exist in M' , the proof will be complete.

Let $\gamma' = \alpha' \cup \beta'$, where α' and β' are the lifts of α and β , and assume there exists a planar surface F (as described above) in M' where $\partial F = \alpha' \cup \beta' \cup \bigcup_{i=1}^k (\gamma_i)$, $\alpha \subset F$, α does not cut off a disk from F , $\beta \subset \partial M'$, and each γ_i comes from the intersection of the branching axis and $D \subset M$. Put F in standard position by looking down on it from B_+^3 , travelling around a component of L' , call it L'_1 , in a clockwise manner and pushing F to the right of each bubble it would otherwise pass through; that is, drop F “outside” the loop. Now F will be such that it intersects \mathbb{S}_+^2 either in simple closed curves or simple arcs that have endpoints on β . To show the former do not exist, we can use Menasco’s arguments because $\pi(L')$ is reduced and alternating. That is, since $\pi(L')$ is alternating (Menasco calls this the “alternating property”), all simple closed curve c must intersect the bubbles between \mathbb{S}_-^2 and \mathbb{S}_+^2 in a certain way, and because of this no such curve c can exist.

To show the latter do not exist, assume they do and take an outermost such curve c , and let $\beta^* \subset \beta$ have the same endpoints as c . Now, β^* must touch at least one bubble; otherwise, we could cut F along c and glue in a disk which does not intersect \mathbb{S}_+^2 such that the resulting planar surface F' does not intersect \mathbb{S}_+^2 along c . In fact, β^* must touch an odd number of bubbles because if β touches an even number of bubbles, we can homotope β^* and c back through a bubble near one of the two endpoints of β^* , which reduces the even number of bubbles case to the odd number of bubbles case. In addition, we only have to worry about cases where the number of bubbles is equivalent to 3 (mod 4) because if the number of bubbles is equivalent to 1 (mod 4) then $c \cup \beta^* = a_1 a_2 \dots a_{2n-1}$ for some $n \in \mathbb{Z}^+$. (This is the primary place in the proof that breaks down because we want to say that all sub-surfaces of a disk are disks, which is true, but the lift F of our original boundary-compressing disk D is a planar surface, not necessarily a disk, and thus there might be some boundary in the sub-surface bounded by $c \cup \beta^*$. We continue below under the assumption that F is a disk because we wish to present some arguments that might prove to be useful elsewhere in the future.)

If we are granted that F is a disk, then we can still make some interesting Menasco-type arguments in order to attempt to prove that F cannot exist. The fact that $c \cup \beta^*$ is the product of an odd number of generating letters in $\pi_1(M')$ implies that $c \cup \beta^*$ does not bound a disk (since the most that the word could reduce to is a single non-trivial generating letter). So now




Figure 6.3: Creating an alternating projection with strictly more crossings than $\pi(L)$.

we are faced with the case where the number of bubbles b that β^* crosses is equivalent to 3 (mod 4).

We first investigate the case where $b = 3$, shown in Figure 6.3. Since $c \cup \beta^*$ bound a disk, we can push L through that disk in order to create a new alternating projection of L with strictly more crossings than $\pi(L)$ which is a contradiction since reduced, alternating projections of alternating links are minimal, and that the new projection is alternating contradicts the minimality of the original projection $\pi(L)$. (We have not found arguments for the other 3 (mod 4) cases yet.)

Thus, since no innermost curve c exists, F does not intersect \mathbb{S}_+^2 at all, but since F has part of its boundary on β , this is impossible, and thus F cannot exist. Thus the original boundary-compressing disk D cannot exist, and therefore S is boundary-incompressible.

□

Although the proof of the conjecture itself fails, we have shown the following psuedo-corollary, which is related to the well-known fact in knot theory that two alternating links are splittable if and only if they are obviously

splittable in a reduced, alternating projection. Our result is local rather than global, but only applies to trivial components rather than link components in general.

Corollary 6.6 *Let L be an alternating link with reduced, alternating projection $\pi(L)$ and T be a tangle in $\pi(L)$. Then a trivial component whose projection into $\pi(L)$ is alternating and encircles T (as in Figure 6.3) is non-trivial in $\mathbb{S}^3 \setminus L$.*

As we have noted within the “proof” of Conjecture 6.5, the above argument fails primarily because in general D lifts to a planar surface rather than a disk. Hopefully we will be able to get around this in the future, but we cannot see how to do so now. For some time, we thought that applying the generalized loop theorem (from [Jac80]) would do the trick, but at that point we had not realized that all we require is for the surfaces to be simple curve boundary incompressible (as defined in section 2.3), and therefore we automatically assume that the curves which concern in S are embedded. We include the statement of the general loop theorem for completeness.

Theorem 6.7 *Let M be a 3-manifold and D a compact planar surface with boundary components J_1, \dots, J_k . Let N_1, \dots, N_k be normal subgroups of $\pi_1(M)$. Suppose that $f : D \rightarrow M$ is a map such that $f(\partial D) \subset \partial M$, $[f|J_i] \notin N_i$, $f(J_i)$ is orientation preserving in M for each i , and $f(J_i) \cap f(J_j) = \emptyset$ for $i \neq j$. Then, given regular neighborhoods U_i of $f(J_i)$ in ∂M , there exists a compact planar surface D' with boundary components J'_1, \dots, J'_k and an embedding $g : D' \rightarrow M$ such that for every j (for $1 \leq j \leq k'$), there exists a unique i_j (where $1 \leq i_j \leq k$) so that $g(J'_j) \subset U_{i_j}$ and for some j , $[g|J'_j] \notin N_{i_j}$.*

Indeed, the success of the above arguments depended upon our erroneous assumption that a disk in M lifts to a disk in M' , and the problem with that is if the original disk intersects the branching axis more than once. Since we are not stipulating position of the original disk—only existence—we have no way of avoiding this. Unfortunately, applying even the generalized versions of Dehn’s Lemma and the Loop Theorem become tricky here, and we are not assured that the boundary of the ∂ -compressing disk in M' is embedded. This turns out to be problematic for the alternating property because if the intersection curves of D' with S_+ are immersed, then we cannot apply the arguments that drive Menasco’s techniques. This is when we decided to look at Seifert surfaces and Seifert-He surfaces rather than general He surfaces.

6.2 Seifert and Seifert-He Surfaces

A **Seifert surface** S for a link L is a surface properly embedded into $\mathbb{S}^3 \setminus L$ such that $\partial S = L$. A **minimal genus Seifert surface** is a Seifert surface which is of genus as low or lower than any other Seifert surface for L . It is easy to show that minimal genus Seifert surfaces are both incompressible and boundary-incompressible. Indeed, let S be a minimal genus Seifert surface for a link L , and assume that S is compressible. Performing the compression gives us a surface S' with $\chi(S') > \chi(S)$. However, $\partial S' = \partial S = L$, and therefore S' is a Seifert surface of lower genus than S . This contradicts our assumption that S was a minimal genus Seifert surface, so S must be incompressible. Similarly, assume that S is a minimal genus Seifert surface for a link L , and further assume that S is boundary-compressible. Since S is embedded, performing this boundary-compression results in a surface S' with $\chi(S') > \chi(S)$, a contradiction since, if the boundary of the boundary compressing disk D is $\partial D = \alpha \cup \beta$. Therefore minimal genus Seifert surfaces are both incompressible and boundary-incompressible.

Recalling that our primary motivation for showing that He and Seifert-He surfaces are boundary-incompressible is their ability to give bounds on meridian length, longitude length, and cusp volume, we are also interested in showing that surfaces created using a combination of He and Seifert techniques are boundary-incompressible. A natural starting point is the case where we apply Seifert's algorithm to all but one crossing c of a knot K , and to that last crossing we apply He's algorithm. The result is an immersed surface S with boundary $K \cup m(c)$ where $m(c)$ is the meridian of the understrand in crossing c . We first looked at the simplest possible case: a Hopf link, L . Applying the above algorithm to the Hopf link gives an uncapped projective plane (immersion g of disk D with $\#(S_2(g)) = 2$) with a meridional puncture and a half-twisted band connected to it along L . We tried to prove this to be boundary-incompressible by placing it in a solid 2-handlebody and looking at a particular cover taken along the two branching axes (which correspond to the two handles of the complementary solid 2-torus in \mathbb{S}^3). While this seemed promising, upon further reflection S turned out to be ∂ -compressible for a reason that would render all Seifert-He surfaces of this nature (i.e. where the He algorithm is applied to only one crossing) ∂ -compressible. Thus we relegated this problem to the back-burners because, in general, Seifert-He surfaces are not boundary-incompressible. Unfortunately, we cannot even use the arguments that we used in Chapters 4 and 5 to prove that any bounds

on meridian length, longitude length, and maximal cusp volume that we get from Seifert-He surfaces since, in general, Seifert-He surfaces have non-zero genus and our arguments depend on the fact that He surfaces are topological disks.

Chapter 7

Generalized He Surfaces

In this section, we present preliminary results which come from generalized He surfaces. Using these new surfaces, we hope to prove the following conjecture.

Conjecture 7.1 *Let K be a knot with toroidal crossing number c_T , and let $|m|, |\ell|$ denote the lengths of the meridian m of the knot and the ℓ -curve of the knot. The following bounds hold for K :*

- $|m| \leq 6 - \frac{7}{c_T}$
- $|\ell| \leq 5c_T - 6$
- $V(C) \leq \frac{9c}{2} \left(1 - \frac{1}{c_T}\right)^2$

As we will show below, we can often improve this last bound on account of what sort of curve we are coming to.

7.1 Construction

Let K have some projection $\pi(K)$ onto an unknotted genus n surface $J = \partial H$ in $M = \mathbb{S}^3 \setminus K$, and let γ be the core curve of the handlebody H . Construct a generalized He surface by connecting each of the points on $\pi(K)$ to a point on γ . Notice that the resulting boundary component γ' is some (possibly null) product of γ , say $\gamma' = n\gamma$ for some $n \in \mathbb{Z}^+ \cup \{0\}$. It is possible that γ' is homotopic to a parabolic rather than a hyperbolic element of Γ , the




Figure 7.1: Here, we cone a projection of the figure-8 knot (K) to the core curve c of a 2-torus \mathbb{T}^2 to produce a surface S .

discrete group of isometries which defines M , and in such cases the above bound does not hold because we cannot spin around a parabolic element to mimic an ideal triangulation. This is because spinning around a parabolic element does not send edges in the pleated surface to ideal edges, but rather to longer and longer closed geodesics in M which we can think of as the geodesic arcs below lines on the horosphere centered at ∞ . However, in the case that γ is trivial, we can cap it off and calculate the bounds postulated in Conjecture 7.1. Unfortunately, we have found that this is the only case which we are immediately guaranteed to be able to calculate these bounds because in the cases when γ is nontrivial, we have no way of ensuring that γ is clean even if it is a geodesic. Here, a **clean** geodesic is one which does not intersect the cusp. If γ were a clean geodesic, then we would be sure that the cusp intersects the surface S only in annuli which is analogous to condition 2 described in section 5.1.

Originally, we were hoping that this process generalized to higher genus surface crossing numbers; indeed, we anticipated that we might be able to

project an n -genus link onto an $n - 1$ -genus surface and derive bounds on meridian length, longitude length, and maximal cusp volume in terms of the minimal crossing number on that surface, but these assumptions turned out to be false precisely for the reason mentioned above. Our bounds depend on the surfaces satisfying the condition that the cusp intersect the surface only in annuli because if they do not, we are not justified in saying that the sum of the areas of the sub-regions is less than or equal to the area of the entire surface. If we had better ways of determining whether a given geodesic were clean or not, then we might be able to salvage this line of research because we would then be able to find a projection of a link onto some genus n surface, cone to its core curve c , determine whether the boundary component $\gamma = c^r$ (for some $r \in \mathbb{Z} \cup \{0\}$) is a clean geodesic, and then if it is, calculate bounds for cusp size quantities for M . We plan on revisiting a recent SMALL paper (see [ACF⁺01]) in conjunction with other papers which describe conditions for a manifold having short geodesics in order to determine whether this line of research is worth pursuing further.

Chapter 8

Conclusions

Here we restate the main results and summarize the arguments set forth in this report. We also include speculation on which directions this research may take and how we may improve upon the results.

8.1 Summary

Our ultimate goal is to provide a means of attaining relatively accurate upper and lower bounds for meridian length, longitude length, and maximal cusp volume which one can calculate using some hyperbolic link invariant, such as crossing number or toroidal crossing number. We have been successful in finding upper bounds for these cusp size quantities, and we have done so by utilizing a construction first introduced by He; we call these surfaces He surfaces.

8.1.1 Main Results

First, in Chapter 4, we showed that performing boundary compressions on the surfaces that interest us result in surfaces which give bounds that are better than the original surface. Then, in Chapter 5, we showed that maximal cusp volume is essentially bounded above linearly by crossing number, an invariant for hyperbolic links. We proved that these bounds hold even when the surfaces in question are s.c. boundary incompressible (as defined in Section 2.3). This is true because the boundary-compressing process eventually terminates, so we are guaranteed that a s.c. boundary-compressed surface

results in a surface which gives better bounds than the original He surface, and thus the bounds calculated in terms of crossing number hold.

8.1.2 Strengthening our Main Results

While we have proved what we set out to prove, we would like to prove the stronger result, namely that He surfaces are s.c. boundary incompressible. Conjecture 6.5 implies that, for a large class of links, this is true; however, we were unsuccessful in proving this conjecture due to the unexpected behavior of a boundary-compressing disk D near the axis over which we took a branching cover in order to lift D to a space which is easier to work in. On the way to attempting to prove this, however, we did prove Lemma 6.4, which is of some independent interest since it ensures that given any link L which admits a reduced, alternating braid representation $\pi(L)$, we can find some branched cover of its complement such that the components of $\pi(L')$ do not cross themselves, only each other, and $\pi(L')$ is reduced and alternating. Because of its generality, this result may be useful in constructing examples or counter-examples in the future, especially in the realm of cyclic covers of alternating braids.

8.1.3 Generalizations

Chapter 7 saw our attempts to generalize these methods in order to get better bounds by projecting a link onto a genus n surface, and since given a link L there exists some genus N surface F such that there exists a projection $\pi(L)$ which has no crossings on F , we thought that this approach would prove to be highly informative. However, as we have seen, that construction depends on spinning an edge around γ (the product of the core curve c to which our surface S is coned) in order to simulate an ideal triangulation from which we can calculate Euler characteristic and therefore find the area of the surface, and spinning only works when γ is a geodesic. Indeed, even if γ is a geodesic, we are only ensured that our bounds hold if γ is clean, because otherwise it would intersect the cusp, and thus we could say that the sum of the small areas where the cusp intersects S plus the length of γ is less than or equal to the area of S itself. Notice that this is one of the three necessary conditions we outlined in Section 5.1, and since we are only able to ensure it in special cases that are as of yet more difficult to determine than would make our

attained bounds useful, we hope to find some other aspect or application of this method in order to improve the results of this paper.

8.2 What the Future Holds

Singular maps of surfaces into manifolds are frustrating and recalcitrant at worst, and at other times, they are often only disappointing and unruly. The advantage of using immersed surfaces is their flexibility as they are more general than embedded surfaces and therefore admit a wider range of constructions than embedded surfaces. This advantage does carry along with it detriments which cannot be circumlocuted, however, and because of this trade-off, one has a dilemma to resolve before proving anything interesting with singular maps. Few have been successful in using them to an advantage because proving that they have the properties which topologists usually require surfaces to have is an exceedingly difficult task. Indeed, while much is known about embedded surfaces, relatively little is known about immersed—or, more generally, singular—surfaces, and because of this, any research which is based around them is immediately faced with very tough problems. While it might be too much to hope for, it would be nice to see 3-manifold theory reformulated and rewritten with a focus on immersed and singular surfaces because the machinery which is available to prove things about embedded surfaces simply does not exist for immersed or singular surfaces. This is an arduous task, however, so I do not expect anyone to either tackle it or complete it for quite some time.

A more reasonable task is to generalize and extend the results of this paper or try to fix the above arguments which fail, thus rendering the stated conjectures theorems. Our first attempts to generalize He's construction are outlined in Chapter 7, and although we have discovered errors in our original thinking as well as facts which severely diminish the potential of this procedure, the procedure itself is worth considering in some context. Part of its innovation lies in the fact that the surfaces S that we create from this general procedure (an example of which is shown in Figure 7.1) are not properly immersed surfaces since, in general, one of its boundary components γ will be some (possibly non-trivial) power of c , the core curve of the surface which we are coning to, and clearly this does not necessarily lie within the cusp; according to the definition in Section 2.3, this means that it is not a proper map. In fact, in order to be a useful surface—i.e. one from which

we can derive the desired bounds in terms of genus n crossing number— c must be a clean geodesic. Therefore, since c can be a parabolic element of Γ rather than a geodesic or a dirty or spotted geodesic rather than a clean one at that, we submit ourselves to the nontrivial question of determining when a given closed curve in a manifold is a clean geodesic, which is related to the problem of determining when a manifold has short closed geodesics. For more on these problems, see [ACF⁺01].

We also feel as though there is hope for the other conjectures which we have failed to prove. I believe that some amount of worthwhile progress has been made towards solving each of the Conjectures 6.2, 6.3, 6.5, and 7.1, and hopefully with some more work which comes after having sufficiently broadened my overall mathematical background as well as deepening my understanding of hyperbolic 3-manifolds in particular, these conjectures will be upgraded to theorems. At this point, however, not enough is known about general singular maps to make this a tractable endeavor, and thus we leave them as is until another time when current difficulties no longer exist in the same form.

Bibliography

- [ACF⁺01] Colin Adams, Adam Colestock, Jim Fowler, Danny Gillam, and Eric Katerman, *Cusp size bounds from surfaces in hyperbolic 3-manifolds*, Forthcoming (2001), 1–13.
- [Ada94a] Colin Adams, *The knot book*, WH Freeman, 1994.
- [Ada94b] ———, *Toroidally alternating knots and links*, *Topology* **2** (1994), 353–369.
- [ADFaSP99] Colin Adams, Ryan Dorman, Kerryann Foley, and Jonathan Kravis and Sam Payne, *Alternating graphs*, *J. of Combinatorial Theory* **77** (1999), 96–120.
- [Ago00] Ian Agol, *Bounds on exceptional dehn filling*, *Geometry and Topology* **4** (2000), 431–449.
- [Hak62] Wolfgang Haken, *über das homöomorphieproblem der 3-mannigfaltigkeiten. i.*, *Math. Z.* **80** (1962), 89–120.
- [Jac80] William Jaco, *Lectures on three-manifold topology*, American Mathematical Society, 1980.
- [Lac00] Marc Lackenby, *The volume of hyperbolic alternating link complements*, arXiv pre-print server (2000), 1–21.
- [Men84] William Menasco, *Closed incompressible surfaces in alternating knot and link complements*, *Topology* **23** (1984), 37–44.
- [MT91] William Menasco and Morwen Thistlethwaite, *A geometric proof that alternating knots are non-trivial*, *Math. Proc. Camb. Phil. Soc.* **109** (1991), 425–431.

- [MT92] ———, *Surfaces with boundary in alternating knot exteriors*, J. reine angew. Math. **426** (1992), 47–65.
- [Mun00] James Munkres, *Topology*, Prentice Hall, 2000.
- [NC95] Max Neumann-Coto, *Area-minimizing tori in 3-manifolds*, Topology and its applications **65** (1995), 167–178.
- [Wal68] Friedhelm Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. (2) **87** (1968), 56–88.