Volume preserving surgeries on hyperbolic 3-manifolds

Peter Rudzis
Advisor: Dr. Rolland Trapp

August 19, 2016

Abstract

In this paper, we investigate two types of surgeries, the parallel surgery and the series surgery, on finite-volume hyperbolic 3-manifolds. In [2] Harnois and Trapp show that these surgeries preserve hyperbolicity and volume when performed on fully augmented links. Using methods similar to Adams in [1], we prove that the parallel and series surgeries preserve hyperbolicity and volume when performed on any finite-volume hyperbolic 3-manifold.

1 Introduction

A **knot** is a continuous embedding of the circle $\mathbb{S}^1$ in the the sphere $\mathbb{S}^3$, while a **link** is a continuous embedding of multiple circles in $\mathbb{S}^3$. Each embedded circle in a link is called a **component** of the link. Since knots can be regarded as links of one component, in this paper we will let the word "link" refer to a knot or link, unless stated otherwise. Knot theory, the study of knots and links, formalizes our intuition about the tangled stuff we encounter in everyday life.

The central objective of knot theory is to distinguish links. Two links K and K' are said to be **equivalent** if there is an orientation preserving homeomorphism $h : \mathbb{S}^3 \to \mathbb{S}^3$ such that $h(K) = K'$. Intuitively, this is to say we can simply bend and stretch K to obtain K'. This equivalence relation partitions the set of links into equivalence classes. Distinguishing two links thus means showing that they belong to different (or the same) equivalence class.

The main tools for distinguishing links are link invariants. If $\mathcal{L}$ is the set of links, then a **link invariant** is simply a function $f : \mathcal{L} \to A$ such that links in the same equivalence class belong to the same level set of f. Said differently, a link invariant is some property of links which equivalent links share in common.

Central to this paper is a link invariant known as **hyperbolic volume**. Often the complement of a link admits the structure of a hyperbolic 3-manifold, and such a link is called a **hyperbolic link**. An important property of hyperbolic 3-manifolds is that hyperbolic volume is a topological invariant. Thus any two equivalent hyperbolic links will have the same hyperbolic volume, and this means that hyperbolic volume is a link invariant.

In [2] Harnois and Trapp investigate two types of surgeries on hyperbolic links. They show that if a link belongs to a special class of links known as fully augmented links, then these surgeries preserve hyperbolicity and volume. The purpose of this paper is to generalize this result and show that we can perform the surgery on any finite-volume hyperbolic 3-manifold and obtain another finite-volume hyperbolic 3-manifold with the same volume as the original.

2 Preliminaries

While this paper is knot theory-inspired, the major concepts and tools utilized come from hyperbolic geometry. This section introduces some of the ideas from hyperbolic geometry which will be needed.

2.1 Hyperbolic 3-manifolds

Euclidean geometry is founded on a set of simple, intuitively appealing axioms, along with the famous the parallel postulate:

Given a line and a point not on the line, there is exactly one line though the point parallel to the given line.

Hyperbolic geometry is what we get when we replace the parallel postulate with:

Given a line and a point not on the line, there is more than one line though the point parallel to the given line.

Among the consequences of this fact is that, while in Euclidean geometry the sum of the angles of a triangle is π, in hyperbolic geometry the sum of the angles of a triangle is less than π.

Hyperbolic space can be difficult to visualize, and thus we use models to understand it. These models take hyperbolic points, lines, and planes and associate with them corresponding Euclidean structures. The upper half-space models are especially useful. The **upper half-plane model** associates points in $\mathbb{H}^2$ with points in the half-space $\{(x, y) \in \mathbb{R}^2 : y > 0\}$ and associates hyperbolic lines with rays and arcs of circles intersecting the x-axis at right angles. The **upper half-space model** for 3-space associates points of $\mathbb{H}^3$ with points of the half-space $\{(x, y, z) \in \mathbb{R}^3 : z > 0\}$, hyperbolic lines with rays and arcs of circles intersecting the x, y-plane at right angles, and hyperbolic planes with hemispheres and half-planes intersecting the xy-plane at right angles.

Note that these models skew distances, in the sense that the arc length along an arc of a circle associated with hyperbolic line segment is not in general the hyperbolic length. Hyperbolic length, area, and volume must be computed in a different way. However the upper half-space models are conformal. This means that the the Euclidean angle between the tangents of two intersecting arcs, at the point of intersection, is the hyperbolic angle between the two corresponding hyperbolic lines.

Definition 2.1. *A **hyperbolic manifold** is a complete Riemannian manifold with constant sectional curvature -1.*

Intuitively a hyperbolic manifold is just a space which locally looks like hyperbolic space. Geometry and topology interact closely in the study of hyperbolic 3-manifolds, a fact best manifested by Mostow's Rigidity theorem.

Theorem 1 (Mostow Rigidity, special case)**.** *Given two homotopic, hyperbolic 3-manifolds M and N, there is a unique isometry between them.*

In particular, one of consequences of this theorem is that two finite-volume hyperbolic 3-manifolds which are homotopic will always have the same hyperbolic volume.

2.2 Covering spaces

Definition 2.2 (Covering space)**.** *Let $p : E \to B$ be a continuous surjective map. We say that an open set $U \subset B$ is **evenly covered** by p if $p^{-1}(U)$ is a disjoint union of open sets $V_i \subset E$, such that the restriction of p to any one V_i is a homeomorphism. If every point of B is evenly covered by p, then p is called the **covering map**, and E is called the **covering space** of B.*

Two covering spaces $p : E \to B$ and $p' : E' \to B$ are said to be **equivalent** if there is a homeomorphism $h : E \to E'$ such that $p = p' \circ h$. The homeomorphism h is called an **equivalence of covering spaces**. The set of all equivalences of covering spaces for E forms a group, known as the **group of covering transformations**. The group of covering transformations is related in a special way to the fundamental group of the base space. Rather than stating this relationship in full generality, we state the relationship for when $E = \mathbb{H}^3$.

Proposition 2.1. *There is a bijective map $\phi : \pi(B, x_0) \to \Gamma$, where $\pi(B, x_0)$ denote the fundamental group of B, and Γ is the group of covering transformations for B covered by $\mathbb{H}^3$.*

For the purposes of this paper, we are mainly interested in the case where $E = \mathbb{H}^3$. In this case, $\mathbb{H}^3$ is always the **universal cover** for the spaces which it covers. A space E is a universal cover for a space B if, given any covering map $p' : E' \to B$, there is a covering map $q : E \to E'$ such that $p \circ q$ is a covering map for B.

Example 2.1 (Punctured-torus)**.** *//TODO*

Example 2.2 (Whitehead link)**.** *//TODO*

2.3 Gluing Polyhedra

Reversing the process of finding a fundamental region for a manifold, if we start out with a finite-family $\mathcal{P}$ of disjoint hyperbolic polyhedra in $\mathbb{H}^3$, we can identify the faces of the polyhedra to obtain a manifold M. Ratcliffe in [4] describes precisely when the resulting manifold is hyperbolic.

Definition 2.3 (G-sided pairing)**.** *Let G be a group of isometries of $\mathbb{H}^3$. Let $\mathcal{P}$ be a finite family of polyhedra, and let $\mathcal{S}$ be the set of their sides. A G-**side pairing** for $\mathcal{P}$ is a set $\Phi \subset G$, $\Phi = \{g_S : S \in \mathcal{S}\}$, indexed by $\mathcal{S}$ and satisfying for each $S \in \mathcal{S}$*

1. *for some $S' \in \mathcal{S}$, $g_S(S') = S$;*
2. *the isometries g_S and $g_{S'}$ are related by $g_{S'} = g_S^{-1}$;*
3. *if S is a side P in $\mathcal{P}$ and S' is a side of P' in $\mathcal{P}$, then $P \cap g_S(P') = S$.*

Since (1) and (2) imply each side S' is uniquely determined by S, there is an equivalence relation generated on the set $\Pi = \cup_{P\in\mathcal{P}} P$, and the equivalence classes are called **cycles** of Φ. Ratcliffe defines the **solid angle** subtended by a polyhedron $P \in \mathcal{P}$ at a point in $x \in P$ to be the number

$$\omega(P,x) = 4\pi \frac{\mathrm{Vol}(P \cap B(x,r))}{\mathrm{Vol}(B(x,r))}, \tag{1}$$

where radius r is chosen to be less than the distance from x to the nearest side of P which does not contain it. If $[x] = \{x_1, ..., x_m\}$ is a finite cycle of Φ, and P_i is the polyhedra containing x_i, for $i = 0, ..., 1$, then the **solid angle sum** of $[x]$ is defined to be the number

$$\omega[x] = \sum_1^m \omega(x_i, P_i). \tag{2}$$

Definition 2.4 (Proper G-side pairing)**.** *A G-side pairing Φ for a finite family of polyhedra $\mathcal{P}$ is said to be **proper** if every cycle of Φ is finite and has a solid angle sum 4π.*

Ratcliffe proves the following theorem, of which we state a special case.

Proposition 2.2 (4,Theorem 10.1.2, special case)**.** *Let G be a group of isometries of $\mathbb{H}^3$, and let M be a space obtained by gluing together a finite family $\mathcal{P}$ of disjoint convex polyhedra in $\mathbb{H}^3$ by a proper G-side pairing Φ. Then M is a 3-manifold with an $(\mathbb{H}^3, G)$-structure such that the natural injection of P° into M is an $(\mathbb{H}^3, G)$-map for each P in $\mathcal{P}$.*

In fact, Proposition 2.2 may be extended to hold true for a family of polyhedra which are not necessarily convex.

Proposition 2.3. *Proposition 2.2 still holds true if we remove the convexity requirement on the polyhedra in $\mathcal{P}$.*

Proof. Let G be a group of isometries, and suppose M is the space obtained by gluing together a finite-family $\mathcal{P}$ of disjoint, not necessarily convex polyhedra in $\mathbb{H}^3$ by a proper G-side pairing Φ. Choose a polyhedron $P \in \mathcal{P}$, and partition P into two polyhedra $P' \cup P''$. Translate P' and P'' with isometries of G so that P', P'', and all the polyhedra in $\mathcal{P}\backslash\{P\}$ are disjoint (but we still denote their translates by P' and P''). Let $\tilde{\mathcal{P}}$ denote the resulting family of polyhedra, and let $\tilde{\mathcal{S}}$ be the set of their sides. For each $\tilde{S} \in \tilde{\mathcal{S}}$, let $h_{\tilde{S}}$ be the translation in G which takes a side of $\mathcal{S}$ to $\tilde{S}$ when P' and P'' are translated. Let $k_{\tilde{S}} = h_{\tilde{S}} g_S h_{\tilde{S}'}^{-1}$, where $\tilde{S}'$ is the side in $\tilde{\mathcal{S}}$ which S' maps to when P' and P'' are translated. Let $\tilde{\Phi} = \{k_{\tilde{S}} : \tilde{S} \in \tilde{\mathcal{S}}\}$.

Lemma 2.1. *The set of isometries $\tilde{\Phi}$ is a proper G-side pairing for $\tilde{\mathcal{P}}$*

Proof. First we check that $\tilde{\Phi}$ is a G-side pairing. Condition (1) is immediate. (2) follows by considering the inverse of $k_{\tilde{S}} = h_{\tilde{S}} g_S h_{\tilde{S}'}^{-1}$. To prove that (3), let $\tilde{S}$ be a side of $\tilde{P} \in \tilde{\mathcal{P}}$ and $\tilde{S}'$ be the side of $\tilde{P}' \in \tilde{\mathcal{P}}$, such that $\tilde{S} = k_{\tilde{S}}\tilde{S}'$. Let $S = h_{\tilde{S}}^{-1}\tilde{S}$ and $S' = h_{\tilde{S}'}^{-1}\tilde{S}'$, and let $P = h_{\tilde{S}}^{-1}\tilde{P}$ and $P' = h_{\tilde{S}'}^{-1}\tilde{P}'$. Then $g_S S' = S$, and so $P \cap g_S(P') = S$. Then $h_{\tilde{S}}(h_{\tilde{S}}^{-1}\tilde{P}) \cap h_{\tilde{S}}(g_{S'}(h_{\tilde{S}'}^{-1}\tilde{P}')) = h_{\tilde{S}'}(S') \Rightarrow \tilde{P} \cap k_{\tilde{S}}(\tilde{P}') = \tilde{S}$, as required.

Moreover $\tilde{\Phi}$ is proper. By construction of our family of polyhedra $\tilde{\mathcal{P}}$, the face identifications from $\tilde{\Phi}$ may be obtained by first translating P' and P'' back to obtain the old family of polyhedra $\mathcal{P}$ and identifying the rest of the sides by the isometries in Φ, a proper G-side pairing. Let $[\tilde{x}]$ be a cycle of $\tilde{\Phi}$, and let $[x]$ be the set of points to which the points of $[\tilde{x}]$ map to when we translate P' and P'' back to the old family of polyhedra. Observe that $[x]$ is a cycle of Φ. Suppose $\tilde{y}' \in P'$ and $\tilde{y}'' \in P''$ are two points in $[\tilde{x}]$ which are identified when P' and P'' are translated back, and let y be the point in $[x]$ to which $\tilde{y}$ and $\tilde{y}'$ are mapped. Let P be the polyhedron in $\mathcal{P}$ containing y, and observe that $\omega(P, y) = \omega(P', \tilde{y}') + \omega(P'', \tilde{y}'')$. From this we can see that $\omega[x] = \omega[\tilde{x}]$. But the solid angle sum with respect of Φ at any point is 4π since Φ is proper. Therefore $\tilde{\Phi}$ is also proper. □

Since some of the face identifications from $\tilde{\Phi}$ may be realized by translating P' and P'' back to obtain the original family $\mathcal{P}$, and the rest of the face identifications may be obtained from the isometries in Φ, it is clear that both $\tilde{\Phi}$ and Φ give us the same manifold M.

Suppose instead of just partitioning one of the polyhedra in $\mathcal{P}$, we partition every polydron $P \in \mathcal{P}$, into say n_P components, where n_P is a positive integer depending on P. Translating these components via isometries in G, we can obtain a new collection of disjoint polyhedra $\mathcal{P}'$. Using our result above, it is a matter of induction to show there exists a proper G-side pairing Φ' for $\mathcal{P}'$ which gives us the manifold M.

It is a well known-fact that any polyhedron may be partitioned into finitely many convex polyhedra. In fact, there are algorithms which find in the minimal convex partition of polyhedra. See for example [] Thus there is some finite family of disjoint convex polyhedra $\mathcal{P}_{con}$ and a proper G-side pairing Φ_{con} for $\mathcal{P}_{con}$ which gives us the manifold M. Proposition 2.2 then tells us that M is hyperbolic. □

2.4 Surgeries on tangles with consecutive crossing disks

Doing a parallel surgery or series surgery are the processes in which we modify a link by replacing a parallel or series tangle with a different tangle, as indicated in Figures 1 and 2. In [2] Harnois and Trapp show that if we start out with a fully augmented link which contains the parallel (resp. series) tangle, then a parallel (resp. series) surgery preserves hyperbolicity and hyperbolic volume. The question which we investigate in this paper is whether this result can be extended to parallel and series surgeries on any hyperbolic 3-manifold.

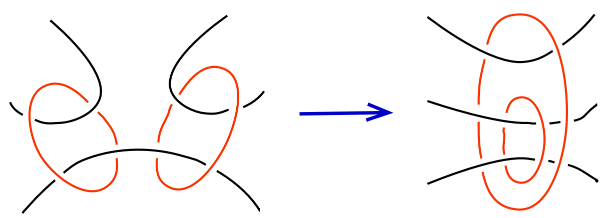

Figure 1: Parallel surgery

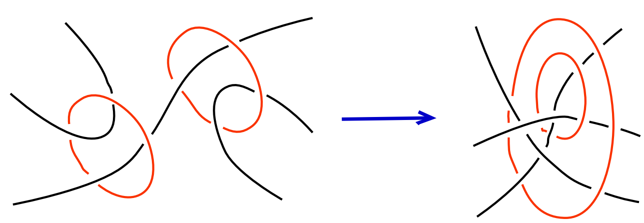

Figure 2: Series surgery

An important fact is that the parallel and series surgeries can be obtained simply by cutting open the manifold along crossing disks and regluing the resulting faces in a different way. In Figure 3 this process is illustrated. The **crossing disks** are the twice-punctured disks colored green. It is easy to see how an analogous construction works for the series surgery.

Note that a twice-punctured disk is a thrice-punctured sphere. Adams work with surgeries on manifolds with embedded thrice-punctured spheres is a good place to start for approaching our question. A key fact which Adams proves is that

Proposition 2.4. *Any thrice-punctured sphere, incompressibly embedded in a hyperbolic 3-manifold, is isotopic to a totally geodesic thrice-punctured sphere.*

This means that an appropriately embedded thrice-punctured sphere will lift to a plane in the covering space $\mathbb{H}^3$.

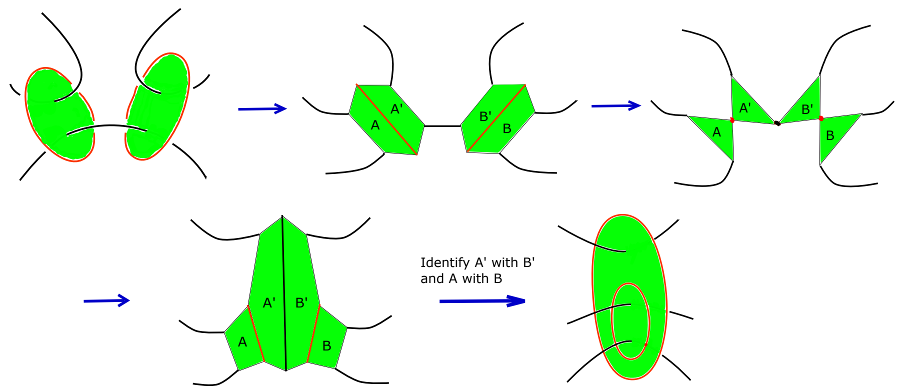

Figure 3: Parallel surgery obtained by surgery on crossing disks

Let M be a finite volume hyperbolic 3-manifold containing in it either a parallel or series tangle complement. Thus M has embedded in it two incompressible, thrice-punctured spheres S_0 and S_1 which have a common puncture. Let $M' = M - (N(S_0) \cup N(S_1))$, where $N(S_0)$ and $N(S_1)$ are regular neighborhoods of S_0 and S_1. Thus the boundary of M' consists of two thrice-punctured spheres $S_0^0, S_0^1 \subset \partial N(S_0)$ and two thrice-punctured spheres $S_1^0, S_1^1 \subset \partial N(S_1)$. Let $\mu_0 : S_0^0 \to S_0^1$ and $\mu_1 : S_1^0 \to S_1^1$ be the pair of identifying isometries which give us back M. Let $\lambda_0 : S_0^0 \to S_1^0$ and $\lambda_1 : S_0^1 \to S_1^1$ be a new pair of identifying isometries, which give us a manifold N. The main result is that

Theorem 2. *If M is a finite-volume hyperbolic 3-manifold, then so is N, and $Vol(M) = Vol(N)$.*

Note that the parallel and series surgeries are obtained precisely in this way, by cutting along embedded thrice-punctured spheres in the tangles. So this result would imply that the hyperbolicity and volume-preserving properties of the parallel and series surgeries extend to finite-volume hyperbolic manifolds in general.

Our strategy for proving this theorem is to induce the surgery in the manifold M just by changing the gluing instructions for the sides of a fundamental region. This is the same type of argument which Adams gives in [1]. However, additional constraints are present since, unlike in Adams situation, we must do surgery to two thrice-punctured spheres, rather than just one, in the same manifold. Specifically, like Adams, we construct a fundamental domain Ω' such that S_0 and S_1 lift to the boundary of Ω'. Regluing sides of Ω' then corresponds to cutting open S_0 and S_1 and regluing the faces. The bulk of our argument goes into justifying our construction of Ω'.

3 Constructing a nice fundamental region

Adams shows in [1] that an embedded thrice-punctured sphere in hyperbolic 3-manifold may be isotoped to be totally geodesic. Thus we may assume S_0 and S_1 lift to planes in $\mathbb{H}^3$. Let N be the cusp neighborhood shared in common by S_0 and S_1, and let $B(N)$ be a horoball which N lifts to in $\mathbb{H}^3$. Since N will always intersect S_0 and S_1, there is a pair of planes P_0 and P_1 which always intersect $B(N)$, where $p(P_0) = S_0$ and $p(P_1) = S_1$. Note that P_0 and P_1 are parallel since S_0 and S_1 are disjoint. Moreover, as we shrink N, the points of $B(N)$ will approach an ideal point, and hence the planes P_0 and P_1 will share an ideal point in common.

Let b be any point in the tangle in M and consider the fundamental group $\pi_1(M)$ based at b. Each subgroup corresponding to the cusps in the boundary of S_i shall be denoted $\pi_1(S_i)$, $i = 0, 1$. The fundamental group of each S_i is denoted in the same way, and it is easy to see that both groups are isomorphic. Recall that there is a monomorphism $\phi : \pi_1(M) \to \mathrm{Isom}^+(\mathbb{H}^3)$. Let Γ denote the group of covering transformations $\phi\pi_1(M)$, with its subgroups $\Gamma_i = \phi\pi_1(S_i)$, $i = 0, 1$. Note that the group Γ_i preserves plane P_i.

Let x_0 be a point in P_0, and let $x_1 = Rx_0$ where R is the reflection taking P_0 to P_1. Let $\Omega_0 = \{x \in \mathbb{H}^3 : d(x, x_0) \le d(x, T_0 x_0) \text{ for all } T_0 \in \Gamma_0\}$, and let $\Omega_1 = R(\Omega_0)$, where R is the reflection which takes P_0 to P_1. Let $\Omega_{planes} = \{x \in \mathbb{H}^3 : d(x, P_0 \cup P_1) \le d(x, T(P_0 \cup P_1)) \text{ for all } T \in \Gamma\}$. Define

$$\Omega = \Omega_0 \cap \Omega_1 \cap \Omega_{planes}. \tag{3}$$

A fundamental region for the action of Γ on $\mathbb{H}^3$ is said to be **locally finite** if every point of $\mathbb{H}^3$ has a neighborhood which intersects only finitely many Γ-images of Ω. We claim

Theorem 3. *The set Ω is a fundamental region for Γ, which is locally finite and the union of two convex polyhedra. Moreover $\Omega \cap P_0$ and $\Omega \cap P_1$ are ideal squares and fundamental polygons for Γ_0 acting on P_0 and Γ_1 acting on P_1 respectively, and the boundary of Ω intersects P_0 and P_1 at right angles.*

Proof. For $T \in \Gamma$ and $i = 0, 1$, define $\Omega_i(T) = \{x : d(x, x_i) \le d(x, Tx_i)\}$, and note that $\Omega_i = \cap_{T_i \in \Gamma_i} \Omega_i(T_i)$. Define $D_{i,j}(T) = \{x : d(x, P_i) \le d(x, TP_j)\}$, and for $i, j \in \{0, 1\}$ let $D_{i,j} = \cap_{T \in \Gamma} D_{i,j}(T)$. Next define

$$C_0 = \Omega_0 \cap \Omega_1 \cap D_{0,0} \cap D_{0,1} \tag{4}$$

and

$$C_1 = \Omega_0 \cap \Omega_1 \cap D_{1,0} \cap D_{1,1}. \tag{5}$$

The fact that Ω is the union of two convex sets is useful, and we will prove it immediately. We will later show that these two convex sets are polyhedra. We have

Lemma 3.1. *The sets C_0 and C_1 are convex, and $\Omega = C_0 \cup C_1$.*

Proof. In the first place, observe that

$$\Omega = (\Omega_0 \cap \Omega_1) \cap [(D_{0,0} \cap D_{0,1}) \cup (D_{1,0} \cap D_{1,1})] = C_0 \cup C_1. \tag{6}$$

Next note that each $\Omega_i(T)$ is simply a half-space, since its boundary is given by the equation of a plane $d(x, x_i) = d(x, Tx_i)$. The boundary of each $D_{i,j}(T)$ is also simply a plane. It is the plane of reflection for the reflection taking P_i to TP_j. Thus $D_{i,j}(T)$ is also a half-space. Thus each Ω_i and $\Omega_{i,j}$ is an intersection of half-spaces and hence a convex set. Then for the same reason C_0 and C_1 are convex sets. □

To set up for the proof that Ω is a fundamental region, we need to describe the geometry of Ω more explicitly. Let $F_0 := \Omega_0 \cap P_0 = \{x \in P_0 : d(x, x_0) \leq d(x, T_0x_0) \text{ for all } T_0 \in \Gamma_0\}$. The set F_0 is an example of a Dirichlet domain for the action of Γ_0 on P_0. Ratcliffe shows that this is a fundamental domain for the action of Γ_0 on P_0 [4, theorem 6.6.13]. This same type of fundamental domain is studied by Adams in [1]. It is a ideal hyperbolic square.

Note that Ω_0 is bounded by planes which intersect P_0 orthogonally. This follows from the fact that each point $x \in \partial\Omega_0$ satisfies equation of a plane $d(x, x_0) = d(x, T_0x_0)$ for some $T_0 \in \Omega_0$. We get orthogonality because both x_0 and T_0x_0 are contained in P_0 for every $T_0 \in \Gamma_0$. In particular, $\partial\Omega_0$ consists of the four planes which contain the edges of the ideal square F_0 and intersect P_0 orthogonally.

These facts about Ω_0 carry over to Ω_1, since Ω_1 is simply the reflection of Ω_0. Specifically, the set $F_1 := \Omega_1 \cap P_1$ is an ideal square, and the set Ω_1 is the region bounded by the four planes which contain the edges of F_1 and intersect P_1 orthogonally. Moreover F_1 is a fundamental polygon for Γ_1 acting on P_1. Indeed, the reflection R taking P_0 to P_1 projects to an isometry $p \circ R$ taking S_0 to S_1, where $p : \mathbb{H}^3 \to M$ is the covering map. In particular, $\pi_1(S_1) = (p \circ R)_*\pi_1(S_0)$, and so $\Gamma_1 = R\Gamma_0R$. Since $\Omega_0 \cap P_0$ is a fundamental domain for Γ_0, it follows that $\Omega_1 \cap P_1 = R(\Omega_0 \cap P_0)$ is a fundamental domain for Γ_1.

The sets Ω_0 and Ω_1 have two common boundary planes. To see this, let x_c be the common ideal point of P_0 and P_1. Let Q be one of the two planes in $\partial\Omega_0$ which has x_c as an ideal point. Since R fixes x_c and Q is orthogonal to P_0, it follows that Q is the unique plane containing line $Q \cap P_0$ which is preserved by R. Since R takes edges of F_0 to edges of F_1, it follows that $Q \cap P_1$ is an edge of P_1, and since x_c is an ideal point of P_1, Q is also orthogonal to P_1. Thus Q is also in $\partial\Omega_1$.

Let H_c and H'_c the two half-spaces exterior to Ω_0 and Ω_1, bounded respectively by each of the two planes in $\partial\Omega_0 \cap \partial\Omega_1$. Let H_i and H'_i denote half-spaces exterior to Ω_i, bounded by the pair of planes in $\partial\Omega_i - \partial\Omega_{(i+1) \bmod 2}, i = 0, 1$. Note that $H_2 = RH_1$ and $H'_2 = RH'_1$. Some possible geometric realizations of these sets in the upper half-space model are illustrated in Figures 4,5, and 6. Note that by choosing the common ideal point x_c of P_0 and P_1 to be 0 and choosing the vertices opposite x_c in F_0 and F_1 to be -1 and 1 respectively, we can ensure that P_0 and P_1 are hemispheres of equal Euclidean radius. Choosing the points in this way, by reflective symmetry of the squares F_0 and F_1, it is clear that H_1, H_1', H_2, and H_2' are half-balls of equal Euclidean radius, and H_c and H_c' are half-balls of equal Euclidean radius.

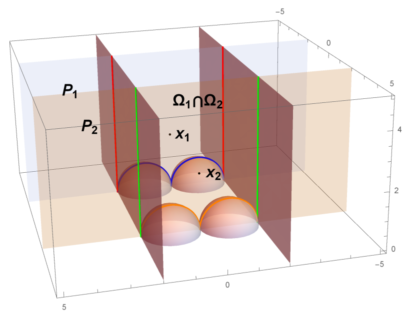

Figure 4: Boundary of $\Omega_0 \cap \Omega_1$ with shared ideal point at ∞

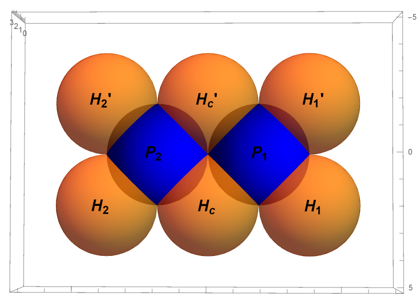

Figure 5: Boundary of $\Omega_0 \cap \Omega_1$ with shared ideal point at 0

Let P_R be the plane of reflection of R, and let K_0 and K_1 be the half-spaces containing P_0 and P_1 respectively. We make the key observation that $\Omega_1 \backslash \Omega_0 \subset K_0$ and $\Omega_0 \backslash \Omega_1 \subset K_1$. Indeed, we see that $\Omega_i = H_c^c \cap H_c'^c \cap H_i^c \cap H_i'^c, i = 0, 1$. Thus $\Omega_1 \backslash \Omega_0 = (H_1^c \cap H_1'^c) \backslash (H_0^c \cap H_0'^c) = (H_0 \cup H_0') \backslash (H_1 \cup H_1')$. Choose $y \in (H_0 \cup H_0') \cap K_2$, and let b_i and b_i' denote the ideal points at the centers of half-balls H_i and H_i', $i = 0, 1$. It is easy to see that the Euclidean distance of y from b_1 and b_1' is less than the Euclidean distance from either b_0 or b_0'. Since H_0, H_0', H_1, H_1' are all of equal radius, it follows that $y \in H_1 \cup H_1'$. This shows that $(H_0 \cup H_0') \backslash (H_1 \cup H_1')$ contains no points of K_1. Thus $\Omega_1 \backslash \Omega_0 \subset K_0$, and $\Omega_0 \backslash \Omega_1 = R(\Omega_1 \backslash \Omega_0) \subset R(K_0) = K_1$.

Observe that $\Omega_0 \cap P_0 \subset \Omega_1 \cap P_0$ and $\Omega_0 \cap P_0 \subset \Omega_1 \cap P_0$. Also note that $P_0, P_1 \subset \Omega_{planes}$. From these two facts it follows that $F_0 = P_0 \cap \Omega$ and $F_1 = P_1 \cap \Omega$.

This proves part of this theorem since F_0 is a fundamental domain for Γ_0 acting on P_0 and F_1 is a fundamental domain for Γ_1 acting on P_1.

Lemma 3.2. *The set Ω is a fundamental region for Γ acting on $\mathbb{H}^3$.*

Proof. A **fundamental set** for a group Γ acting $\mathbb{H}^3$, is subset of $\mathbb{H}^3$ which contains exactly one point of every orbit Γu. By [4,theorem 6.6.11] of Ratcliffe, to prove that Ω is a fundamental region, it is sufficient to find a fundamental set F such that $\Omega^\circ \subset F \subset \Omega$.

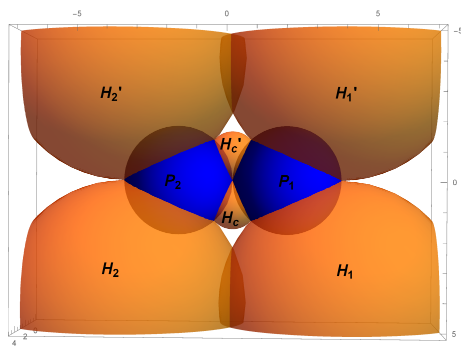

Figure 6: Boundary of $\Omega_0 \cap \Omega_1$ with shared ideal point at 0

We would like to construct our fundamental set F as follows: For each orbit Γu, choose a $y \in \Gamma u \cap \Omega_0 \cap \Omega_1$ such that $d(y, P_0 \cup P_1) \leq d(y, T(P_0 \cap P_1))$ for all $T \in \Gamma$. Let F denote the set of all chosen points. Clearly this is a valid construction of a fundamental set if (i) $\Gamma u \cap \Omega_0 \cap \Omega_1$ is non-empty, and (ii) the points of Γu do attain a minimum distance from P_0 and from P_1.

To verify (i), suppose for a contradiction that, for some $u \in \mathbb{H}^3$, $\Gamma u \cap \Omega_0 \cap \Omega_1 = \emptyset$. Note that Ω_0 and Ω_1 are Dirichlet domains for the action of Γ_0 on P_0 and the action of Γ_1 on P_1 respectively. Then by [4, theorem 6.6.11] Ω_0 and Ω_1 are fundamental domains for the action of Γ_0 on P_0 and the action of Γ_1 on P_1 respectively. Thus $\Gamma u \cap \Omega_0$ and $\Gamma u \cap \Omega_1$ are non-empty. Let u_0 be a point in $\Gamma u \cap \Omega_0$, let $p(i) = i$ mod 2, and for $i > 0$ choose $u_i \in \Gamma u_{i-1} \cap \Omega_{p(i)}$. In this way, we generate a sequence of points $\{u_i\}$ in Γu.

In the upper half-space model, choose the ideal point x_c common to P_0 and P_1 to be ∞, and choose the vertices opposite x_c in ideal squares F_0 and F_1 to be -1 and 1 respectively. See Figure 4.

As before, let K_0 and K_1 denote the half-spaces on each side of the plane of reflection of R. As we observed before, $\Omega_{p(i-1)} \backslash \Omega_{p(i)} \subset K_{p(i)}$. Since by our assumptions $u_{i-1} \in \Omega_{p(i-1)} \backslash \Omega_{p(i)}$ it follows $u_{i-1} \in K_{p(i)}$. Then $d_E(\text{proj}(u_{i-1}), P_{p(i)}) < d_E(\text{proj}(u_i), P_{p(i)})$, where d_E denotes the Euclidean distance, and $\text{proj}(x)$ denotes the projection of x onto the ideal plane. Let $\Pi_j(x)$ denote the Euclidean half-plane determined by the point x and the ideal line at the base of P_j, $j = 0, 1$. Now $d(u_{i-1}, P_{p(i)}) = d(u_i, P_{p(i)})$, because $\Gamma_{p(i)}$ preserves $P_{p(i)}$ and distance. Thus $\Pi_{p(i)}(u_{i-1})$ and $\Pi_{p(i)}(u_i)$ make the same angle with $P_{p(i)}$, because this angle uniquely determines the distance of each point from $P_{p(i)}$. See [5, Proposition 4.1.1]. Since $d_E(\text{proj}(u_{i-1}), P_{p(i)}) < d_E(\text{proj}(u_i), P_{p(i)})$, it is just a matter of similar triangles to show that $d_E(\text{proj}(u_{i-1}), u_{i-1}) < d_E(\text{proj}(u_i), u_i)$. Thus the distance of u_i from the ideal plane is monotonically increasing. Note that this implies there are infinitely many distinct points in the sequence, since there are no repeats. On the other hand, since each u_i belongs to either $\Omega_1 - \Omega_2$ or to $\Omega_2 - \Omega_1$, the sequence is restricted to the four hemispheres $H_1 \cup H_1' \cup H_2 \cup H_2'$ which are of finite radius by our choices for x_c and the vertices of F_0 and F_1. Thus the sequence $\{u_i\}$ is restricted to a compact set, the region inside the four hemispheres and above the horizontal Euclidean plane which contains u_0. Thus the sequence has an accumulation point in $\mathbb{H}^3$. But Γ is a discrete group, and so no orbit of Γ can have an accumulation point, a contradiction.

(ii) Suppose that for every $v \in \Gamma u$ there is a point, $v' \in \Gamma u$ such that $d(v, P_0 \cup P_1) > d(v', P_0 \cup P_1)$. Then there is a sequence $\{v_k\}$ of points in Γu such that $d(v_k, P_0 \cap P_1)$ is monotonically decreasing. By choosing a subsequence and possibly reindexing we can assume that $d(v_k, P_0)$ is monotonically decreasing. Let $D_{i,j}$ for $i, j = 0, 1$ be defined as before. Evidently $D_{0,0}$ contains no points of Γu. A useful fact, which shall be used again later in this proof, is that $\Omega_i \cap D_{i,i}$ is a fundamental domain for the action of Γ on $\mathbb{H}^3$, $i = 0, 1$. Indeed, it is precisely the fundamental domain constructed by Adams in [1], about the point x_i. But this implies $D_{0,0}$ contains a point of Γu, since $\Omega_0 \cap D_{0,0}$ contains a point of Γu, a contradiction.

Therefore F is a fundamental set for the action of Γ on $\mathbb{H}^3$. Suppose $x \in F$. Then $x \in \Omega_0 \cap \Omega_1$. It is clear that x has been chosen so that $x \in \Omega_{planes}$, because a point point in $\Gamma u \cap \Omega_0 \cap \Omega_1$ which minimizes the distance from $P_0 \cup P_1$ also minimizes the distance of all points in Γu. Indeed, if a point $p \in \Gamma x$ minimizes the distance from P_i, then $T_i p \in \Omega_0 \cap \Omega_1$ for some covering transformation $T_i \in \Gamma_i$, which preserves the distance from P_i. Thus $F \subset \Omega$. Now suppose $y \in \Omega^\circ$. Then certainly $y \in \Omega_0 \cap \Omega_1$. Moreover if y is in the orbit Γu, since $d(y, P_0 \cup P_1) < d(y, P_0 \cap P_1)$, y will be the point (uniquely) chosen in $\Gamma u \cap \Omega_0 \cap \Omega_1$ in the construction of F. Thus $y \in F$. Thus $\Omega^\circ \subset F \subset \Omega$. □

Next we show that

Lemma 3.3. *Ω is a locally finite.*

Proof. Observe that $C_i \subset \Omega_i \cap D_{i,i}$ for $i = 0, 1$, and so $\Omega \subset (\Omega_0 \cap D_{0,0}) \cup (\Omega_1 \cap D_{1,1})$. Suppose for a contradiction that there is a ball $B \subset \mathbb{H}^3$ which intersects infinitely many Γ-images of Ω. Let $T_i\Omega$ be an infinite sequence of these Γ-images. Then for every T_i, B intersects either $T_i(\Omega_0 \cap D_{0,0})$ or $T_i(\Omega_1 \cap D_{1,1})$. Thus B intersects Γ-images of either $\Omega_0 \cap D_{0,0}$ or $\Omega_1 \cap D_{1,1}$ infinitely many times. But the fundamental domain constructed by Adams in [1] is locally finite, a contradiction. □

Note that the result here is stronger than we need, since we have shown that every ball intersects finitely many Γ-images of Ω, wheres local finiteness only requires some ball to intersect finitely many Γ-images. The next two lemmas complete the argument. Our proofs are based on the proof of Theorem 6.7.1 in [4].

Lemma 3.4. *The sets C_0 are C_1 are convex polyhedra.*

Proof. Let $\mathcal{S}_0$ and $\mathcal{S}_1$ denote the sets of sides of C_0 and C_1. Recall a convex polyhedron is a convex set whose set sides is locally finite. Choose a point $x \in C_0$. If $x \in C_0^\circ$ of $x \in C_0^c$, then of course some ball centered at x intersects no sides of C_0, and so we can assume that $x \in \partial C_0$. Since Ω is locally finite, there is a ball B centered at x which intersects only finitely many Γ-images of Ω, say $T_1\Omega, ..., T_m\Omega$. Decreasing the radius of B as needed, we may assume that every $T_i\Omega$ contains x. Since $\Omega = C_0 \cup C_1$, it follows that for each T_i either $C_0 \cap T_iC_0$ or $C_0 \cap T_iC_1$ is a nonempty convex subset of ∂C_0. Then by [4, theorem 6.2.6(1)], for each T_i there is a side S_i of C_0 which contains either $C_0 \cap T_iC_0$ or $C_0 \cap T_iC_1$.

If $y \in B \cap \partial C_0$, then for some T_i either $y \in C_0 \cap T_iC_0$ or $y \in C_0 \cap T_iC_1$. Indeed, since $y \in \partial C_0$, y must be contained in C_0 and some Γ-image TC_0 or TC_1. And T is one of the transformations T_i because B only intersects sets $T_1\Omega, ..., T_m\Omega$. Thus

$$B \cap \partial C_0 \subset \cup_{i=1}^m (C_0 \cap T_iC_0) \cup (C_0 \cap T_iC_1) \tag{7}$$

Thus

$$B \cap C_0 \subset S_1 \cup \cdots \cup S_m. \tag{8}$$

If S is a side of D, which B intersects, then B intersects S° since $S = \bar{S^\circ}$. Then one of the S_i intersects S°. By [4, Theorem 6.2.6(3)], this is only possible if $S = S_i$. Therefore B intersects finitely many sides of C_0, so $\mathcal{S}_0$ is locally finite. We conclude that C_0 is a convex polyhedron. A similar argument shows that C_1 is a convex polyhedron. □

This concludes the argument that Ω is a fundamental region which is the union of two convex polyhedra.

□

We will now use the fundamental region Ω constructed above to obtain another fundamental region for M which is more amenable to our needs. Cut open Ω along the planes P_0 and P_1 to obtain three components Ω_0, Ω_1, and Ω_2. Note that Ω_0, Ω_1, and Ω_2 are again polyhedra. Choose $T_0, T_1, T_2 \in \Gamma$ such that $\Omega_0' = T_0\Omega_0$, $\Omega_1' = T_1\Omega_1$, and $\Omega_2' = T_2\Omega_2$ are disjoint. Define $\Omega' = \Omega_0' \cup \Omega_1' \cup \Omega_2'$, and observe that Ω' must also be a fundamental region for the action of Γ on $\mathbb{H}^3$.

Moreover, we see that Ω' intersects $p^{-1}(S_0)$ and $p^{-1}(S_1)$ only in its boundary, at two copies of F_0 and two copies of F_1. Denote the two copies of F_i in $\partial\Omega'$ by F_i^0 and F_i^1, $i = 0, 1$. Let $f_i : F_i^0 \to F_i^1$, $i = 0, 1$, be the identifying isometries which give us back Ω. Let $g_j : F_0^j \to F_1^j$, $j = 0, 1$, be another pair identifying isometries. Note that g_0 and g_1 induce the parallel surgery in the manifold M.

Let $\mathcal{P}$ denote the family of disjoint polyhedra which compose Ω, and let $\mathcal{S}$ be the set of their sides. Let $\Phi = \{T_S : S \in \mathcal{S}\}$ be the proper Γ-side pairing which gives us the manifold M. Similarly, let $\tilde{\mathcal{P}}$ be the family of polyhedra which compose Ω' and let $\tilde{\mathcal{S}}$ be the set of their sides. For every side $\tilde{S}$ of $\tilde{\Omega_k}$ other than $F_{i,j}$ for $i, j = 0, 1$, there is a corresponding side $S = T_k^{-1}\tilde{S} \subset \partial\Omega$. Write $h_{\tilde{S}} = T_k$; thus $h_{\tilde{S}}$ is the isometry which takes a side S of Ω to its corresponding side $\tilde{S}$ in $\tilde{\Omega}$. In addition, define

$$R_{\tilde{S}} = \begin{cases} h_{\tilde{S}} T_S h_{\tilde{S}'}^{-1} & S \in \mathcal{S} \\ g_j & \tilde{S} = F_0^j \\ g_j^{-1} & \tilde{S} = F_1^j \end{cases} \tag{9}$$

where $\tilde{S}'$ is the side in $\tilde{\mathcal{S}}$ corresponding to S'. Let $\tilde{\Phi} = \{R_{\tilde{S}} : \tilde{S} \in \tilde{\mathcal{S}}\}$. Observe that the isometries in $\tilde{\Phi}$ have been chosen to be precisely those which induce the parallel (resp. series) surgery. Since face reidentifications does not change the volume, we know from Proposition 2.2 that once it is demonstrated that $\tilde{\Phi}$ is a proper G-side pairing for $\tilde{\mathcal{P}}$ we are done.

Proof of Theorem 2. First we show that $\tilde{\Phi}$ is a G-side pairing for $\tilde{\mathcal{P}}$. Considering Definition 2.3 we see immediately that (1) is satisfied. (2) follows by in view of the form of the inverse of $R_{\tilde{S}}$. To prove (3), if $\tilde{S} = F_i^j$ for some $i, j \in \{0, 1\}$, note that (3) is satisfied because all the F_i^j are isometric squares, and we have defined each g_i to be an isometry between them. If $\tilde{S}$ is not equal to any F_i^j, let $\tilde{P} \in \tilde{\mathcal{P}}$ be the polyhedron containing $\tilde{S}$, and let $\tilde{S}'$ be the side of $\tilde{P}' \in \tilde{\mathcal{P}}$, such that $\tilde{S} = R_{\tilde{S}}\tilde{S}'$. Let $S = h_{\tilde{S}}^{-1}\tilde{S}$ and $S' = h_{\tilde{S}'}^{-1}\tilde{S}'$, and let $P = h_{\tilde{S}}^{-1}\tilde{P}$ and $P' = h_{\tilde{S}'}^{-1}\tilde{P}'$. Then $T_S S' = S$, and so $P \cap T_S(P') = S$. Then $h_{\tilde{S}}(h_{\tilde{S}}^{-1}\tilde{P}) \cap h_{\tilde{S}}(T_{S'}(h_{\tilde{S}'}^{-1}\tilde{P}')) = h_{\tilde{S}'}(S') \Rightarrow \tilde{P} \cap R_{\tilde{S}}(\tilde{P}') = \tilde{S}$. Thus $\tilde{\Phi}$ is a G-side pairing.

Next we show that $\tilde{\Phi}$ is proper. Each side identification, other than the identifications for the F_i^j, is given by first translating the polyhedra which make up Ω' back to Ω, and then identifying the sides by the isometries in Φ. Thus since Ω has a proper G-side pairing, every point in $\Omega \backslash \cup F_i^j$ has a solid angle of 4π. If a point $\tilde{x}$ is at the interior of one of the F_i^j, then the cycle consists of two points $\{\tilde{x}, \tilde{x}'\}$, and the other point $\tilde{x}'$ is in the interior of another F_i^j. Thus since all the $F_{i,j}$ are flat, we get a solid angle of 2π from each of the two points. These give us a solid angle sum of 4π. Finally if $\tilde{x}$ is in the edge of one of the F_i^j, note that $\tilde{x}$ cannot be a vertex of F_i^j since the vertices are ideal points. By our construction, there is a dihedral angle of $\pi/2$ at each of the edges of F_i^j. At each edge, this F_i^j will be adjacent to one of its Γ-images, lying in the same plane as itself. Thus when the square faces are identified, each edge will meet three other edges, all with a dihedral angle of $\pi/2$. Then the solid angle of $\tilde{x}$ at each edge is π, and so the solid angle sum is 4π, as required.

□

4 Open Questions

A natural class generalization of the result proven in this paper would proving volume-preserving properties for tangle surgeries on tangles containing n crossing disks.

- If for example we imposed the condition that all the crossing disks in a tangle shared exactly one common cusp, then a geometric situation, quite similar to the one in this paper, would be imposed on manifold. In particular the crossing disks would lift to parallel planes, all of which share a single ideal point. It seems likely that the methods used in this paper could be extended to this situation as well.

- If we imposed a slightly less strong condition on shared cusps, say the disks share cusps pairwise, then the geometric situation would be more complex. Any two planes would share an ideal point, but there would be many more conditions which would need to be satisfied to get a nice fundamental domain, using methods similar to those in this paper.

- If we allow some or all of our n crossing disks to have no cusp in common, fewer geometric conditions can be determined. It seems likely that this investigation would require different methods than those used in this paper.

5 Acknowledgments

I would like to thank Dr. Rolland Trapp for his advice and guidance. Without his constant support and superb insight, this project would not have been possible. Another big thank you goes to Dr. Corey Dunn for his encouragement and his wizardly direction of the REU program with Dr. Trapp. This research project was funded both by NSF grant DMS-1461286 and by California State University, San Bernardino.

References

1. Adams, Colin. *Thrice-punctured spheres in hyperbolic 3-manifolds.* Trans. Amer. Math. Soc. 287(2) 645-656. Feb. 1985.

2. Chazelle, Bernard. *Convex Partitions of Polyhedra: A Lower Bound and Worst-case Optimal Algorithm.* Society for Industrial and Applied Mathematics. 13(3) 488-507, 1984.

3. Harnois, John and Rolland, Trapp. *Hyperbolic Tangle Surgeries on Augmented Links.* preprint.

4. Munkres, James R. *Topology, 2nd ed.* Upper Saddle River, NJ: Pearson Education Inc., 2000.

5. Ratcliffe, John G. *Foundations of Hyperbolic Manifolds.* New York: Springer, 2006.

6. Stahl, Saul. *A Gateway to Modern Geometry: The Poincaré Half-Plane, 2nd ed.* Sudbury, MA: Jones and Bartlett Pub.,2008.