DEHN SURGERY FROM A HYPERBOLIC PERSPECTIVE

MARC LACKENBY

Ever since Thurston's work on hyperbolic 3-manifolds [27, 28], the two fields of hyperbolic geometry and Dehn surgery have been very closely related. On the one hand, the 'thick-thin' decomposition of hyperbolic 3-manifolds due to Margulis [15] demonstrates the fundamental role that Dehn filling plays in the subject. In the other direction, hyperbolic techniques have proved to be an extremely powerful way of solving some of the hardest questions about Dehn surgery.

These notes are based on lectures given at ICERM in July 2019. Their aim is to provide a quick introduction to the use of hyperbolic methods in the theory of Dehn surgery.

1. Hyperbolic structures

The theory of hyperbolic 3-manifolds is a vast subject. We can only give a very rapid sketch of the parts of the field that are most relevant to Dehn surgery. We refer the reader to [4, 25, 27] for more details.

Hyperbolic 3-manifolds are, by definition, Riemannian manifolds locally modelled on hyperbolic 3-space $\mathbb{H}^3$. Recall that $\mathbb{H}^3$ is defined to be either of the following Riemannian manifolds:

(1) *upper half-space* $\{(x, y, z) \in \mathbb{R}^3 : z > 0\}$ with Riemannian metric ds^2_{Eucl}/z^2; here ds^2_{Eucl} is the normal Euclidean metric;
(2) *the Poincaré ball model*, which is the open unit ball in $\mathbb{R}^3$ with Riemannian metric at a point p given by

$$ds^2_{\text{Eucl}} \left(\frac{2}{1 - d_{\text{Eucl}}(p, 0)^2} \right)^2;$$

here d_{Eucl} is Euclidean distance and 0 is the origin.

These two Riemannian manifolds are isometric.

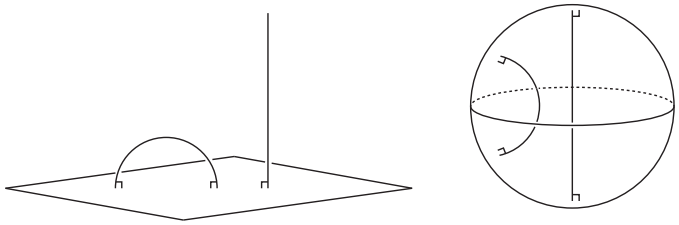

FIGURE 1. Two models of hyperbolic 3-space, with some geodesics. Left: Upper half-space. Right: The Poincaré ball model.

We mention three simple facts about the geometry of hyperbolic space:

(1) Angles in both models of hyperbolic space are the same as in Euclidean space. This is because at each point, the hyperbolic metric is just a multiple of the Euclidean metric.
(2) All points and all directions in hyperbolic space 'look the same'. More specifically, given any two points p_1 and p_2 and any orthogonal transformation between their tangent spaces, there is a hyperbolic isometry taking p_1 to p_2 and that acts on the tangent spaces in the specified way. This is proved by a direct construction. For example, in the special case where p_1 and p_2 are the origin in the Poincaré ball model, the required isometry is just an orthogonal transformation of Euclidean space.
(3) In both models, geodesics are Euclidean lines and circles that hit the boundary orthogonally. This is proved as follows. Geodesics are uniquely determined by a starting point p and tangent vector v at that point. So, one starts with any specific geodesic that is a line or circle hitting the boundary orthogonally and one uses the isometry established in (2) to take this geodesic to the one starting at p and with tangent vector v. The isometries constructed in (2) preserve Euclidean lines and circles hitting the boundary orthogonally and so the geodesic through p with tangent vector v is also of this form.

Definition 1.1. A compact 3-manifold M has a (complete) *hyperbolic structure* if it satisfies any of the following equivalent conditions:

(1) $M - \partial M$ has a complete Riemannian metric where all sectional curvatures are -1;
(2) $M - \partial M$ has a complete Riemannian metric that is locally isometric to $\mathbb{H}^3$;
(3) $M - \partial M$ is the quotient of $\mathbb{H}^3$ by a discrete group of isometries acting freely.

The manifold M is then a (complete) *hyperbolic manifold*.

The equivalence of the three definitions is a fundamental fact [4, 25]. In the proof, one shows that, with any of these definitions of a complete hyperbolic 3-manifold M, the universal cover of $M - \partial M$, with its induced Riemannian metric, is isometric to $\mathbb{H}^3$.

From now onwards, we will drop the word 'complete' and so whenever we refer to a 'hyperbolic structure' or a 'hyperbolic manifold', we are referring to the above definition.

The following result is the central part of Thurston's Geometrisation Conjecture [28]. It was proved by Perelman [21, 23, 22] although Thurston had already proved it in the important case when M is Haken.

Theorem 1.2 (Perelman 2002-03)**.** *Let M be a compact orientable 3-manifold other than $S^1 \times D^2$ or $T^2 \times I$. Then M has a hyperbolic structure if and only if M is irreducible, atoroidal and not Seifert fibred.*

We recall the definitions of the above terms.

Definition 1.3. Let M be a compact orientable 3-manifold.

(1) We say that M is *irreducible* if any properly embedded 2-sphere bounds a 3-ball.
(2) We say that M is *atoroidal* if any properly embedded torus is compressible or parallel to a component of ∂M.
(3) We say that M is *Seifert fibred* if it admits a foliation by circles.

This definition of Seifert fibred is very short but not immediately useful. It was actually proved by Epstein [10] that a manifold that is Seifert fibred in the above sense in fact admits a foliation by circles, where each circle has a nice local model. This was Seifert's original definition [26].

Theorem 1.2 is useful in the particular case where M is the exterior of a knot in the 3-sphere. The theorem in this case was proved by Thurston [28].

Theorem 1.4 (Thurston 1982)**.** *Let M be the exterior of a knot K in the 3-sphere. Then M has a hyperbolic structure if and only if K is neither a satellite knot nor a torus knot.*

We will focus on hyperbolic manifolds with finite volume. There is a very nice characterisation of such manifolds [28, 4].

Theorem 1.5. *A compact orientable hyperbolic 3-manifold M has finite volume if and only if the following all hold:*

(1) ∂M is empty or tori;
(2) M is neither $S^1 \times D^2$ nor $T^2 \times I$.

A crucial result is the following [20, 24].

Theorem 1.6 (Mostow-Prasad rigidity 1968, 1973)**.** *Any finite-volume hyperbolic structure on a manifold with dimension at least* 3 *is unique up to isometry.*

Example 1.7. One of the most important examples of a hyperbolic structure is that on the figure-eight knot complement, which we now describe.

We start with two tetrahedra and we glue their faces in pairs according to the recipe shown in Figure 2. It is not hard to see that every point in this space, apart from the image of the vertices, has a neighbourhood homeomorphic to an open ball in $\mathbb{R}^3$. However, the eight vertices are all identified to a single point and a regular neighbourhood of this point is not homeomorphic to an open ball in $\mathbb{R}^3$. But if we remove this point, the resulting space is an open 3-manifold. It is a remarkable fact that this manifold is in fact homeomorphic to the complement of the figure-eight knot. Assuming this, we can impose a complete finite-volume hyperbolic structure on the figure-eight knot complement as follows.

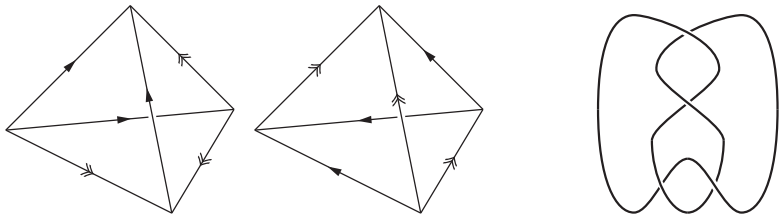

FIGURE 2. Instructions on how to glue two tetrahedra: each face in one tetrahedron is glued to a face in the other tetrahedron in the only way that is compatible with the edge identifications. Removing the point that is the image of the vertices gives the complement of the figure-eight knot (shown on the right)

Start with two copies of $\mathbb{H}^3$ with its Poincaré ball model. On each sphere at infinity, consider four points that are the vertices of a regular Euclidean tetrahedron

centred at the origin. We now form an 'ideal tetrahedron' with these four vertices at infinity. This is like a hyperbolic polyhedron, in that its edges are geodesics and its faces are totally geodesic, but its vertices are at infinity (see Figure 4). So, topologically, an ideal tetrahedron is a tetrahedron with its vertices removed. Now glue these two ideal tetrahedra using hyperbolic isometries, again using the recipe specified in Figure 2. One can check that the resulting manifold is hyperbolic by finding a neighbourhood of each point that isometric to a ball in $\mathbb{H}^3$. The proof that this Riemannian manifold is complete requires a little more work [27].

2. Cusps

When M is a finite-volume orientable hyperbolic 3-manifold, each end of $M - \partial M$ has a neighbourhood of the form $T^2 \times [1, \infty)$. This has a specific geometry, which we investigate in this section.

Definition 2.1. The subset $\{(x, y, z) : z \geq 1\}$ of upper half-space is known as a *horoball*. More generally, the image of this subset under any hyperbolic isometry is known as a horoball. In the upper half-space model, these are either subsets of the form $\{(x, y, z) : z \geq c\}$ for some $c > 0$ or Euclidean balls intersecting the boundary plane at a single point. In the Poincaré ball model, they are Euclidean balls intersecting the boundary sphere at a single point. The boundary of a horoball is a *horosphere*.

Each end of $M - \partial M$ has the following geometric model. It is obtained from the horoball $\{(x, y, z) : z \geq 1\}$, by taking the quotient modulo a discrete group of Euclidean translations isomorphic to $\mathbb{Z} \times \mathbb{Z}$ that preserve the z co-ordinate. (See Figure 3.) This is known as a *cusp* or *horoball neighbourhood* of the end.

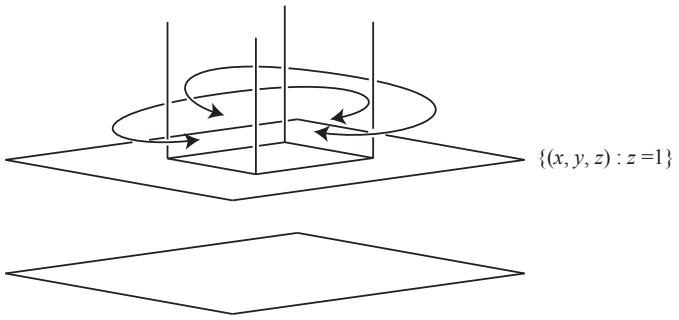

Figure 3. A fundamental domain in upper-half space for a cusp of a hyperbolic 3-manifold. Its four vertical faces are glued in pairs to form a copy of $T^2 \times [1, \infty)$.

Given a horoball neighbourhood of an end of $M - \partial M$, its inverse image under the covering map $\mathbb{H}^3 \to M - \partial M$ is a union of disjoint horoballs. By performing an isometry to $\mathbb{H}^3$, we may assume that one of these horoballs is $\{(x, y, z) : z \geq 1\}$. The image of $\{(x, y, z) : z = \text{a constant } c\}$ in M is an immersed torus that inherits a Euclidean metric. For $c \geq 1$, this torus is embedded in M. As we decrease c below 1, it initially remains embedded, until at some value $c = c_{\min}$, it becomes non-embedded. The image of $\{(x, y, z) : z \geq c_{\min}\}$ is a *maximal cusp*. The boundary of a maximal cusp is an immersed torus, which we call a *maximal cusp torus*.

When ∂M is a single torus, this maximal cusp is uniquely determined. This is because there is only the single parameter c that can be varied and so there is a unique infimal value of c for the interval of z co-ordinates where the torus is embedded. When ∂M consists of more than one torus, then again each end has a well-defined maximal cusp, but these need not be disjoint from each other. They might need to be shrunk to make them disjoint.

Example 2.2. In Example 1.7, we described the hyperbolic structure on the figure-eight knot complement that was obtained by gluing two ideal tetrahedra together. Place a horoball about each of the ideal vertices of each ideal tetrahedron. This is a Euclidean ball in the Poincaré ball model; we choose these balls to be disjoint and all to have the same Euclidean size. The intersection of these horoballs with the ideal tetrahedra patches together to form a cusp. The boundary of the cusp consists of eight Euclidean triangles that are glued together as shown in Figure 5, forming a Euclidean torus.

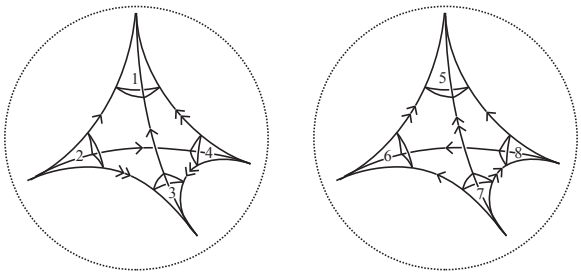

FIGURE 4. A cusp for the figure-eight knot complement.

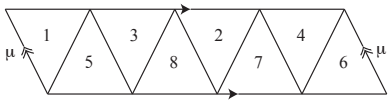

FIGURE 5. The arrangement of Euclidean equilateral triangles that glue together to form the toral boundary of the cusp

3. The thick-thin decomposition

Definition 3.1. Let M be a hyperbolic 3-manifold and let $\epsilon > 0$. The *ϵ-thin* and *ϵ-thick* parts of M are as follows:

$$M_{(0,\epsilon]} = \{x \in M - \partial M : \text{there is a geodesic loop based at } x \text{ with length } \leq \epsilon\},$$
$$M_{[\epsilon,\infty)} = \mathrm{cl}(M - (M_{(0,\epsilon]} \cup \partial M)).$$

Theorem 3.2 (Margulis)**.** *There is a universal $\epsilon > 0$ such that, for any finite-volume hyperbolic 3-manifold, each component of $M_{(0,\epsilon]}$ is one of*

(1) a cusp;
(2) a solid torus regular neighbourhood of a closed geodesic with length less than ϵ;
(3) a closed geodesic of length ϵ.

In fact, Margulis's result [15] is considerably more general than this, as there is a version of the theorem that applies to any discrete subgroup of any semi-simple Lie group. But in the case of a finite-volume hyperbolic 3-manifold M, this theorem implies that, topologically, M is obtained from its ϵ-thick part by attaching solid tori, or in other words, by Dehn filling.

An $\epsilon > 0$ satisfying the conclusion of Theorem 3.2 for all finite-volume hyperbolic 3-manifolds is known as a *Margulis constant*.

The thick part of a finite-volume hyperbolic 3-manifold also has some useful structure.

Theorem 3.3 (Jørgensen)**.** *Let ϵ be a Margulis constant. Then there is a universal constant $k > 0$ such that for any finite-volume hyperbolic 3-manifold M, its ϵ-thick part may be triangulated using at most k Vol(M) tetrahedra.*

Proof outline. The proof of this is sketched in [27] and is carefully presented in [16]. A brief summary is as follows, but where we focus on the case where $M_{(0,\epsilon]}$ is empty for simplicity.

One picks a maximal collection P of points in $M_{[\epsilon,\infty)}$ such that no two of these points are closer than $\epsilon/4$. The open balls of radius $\epsilon/8$ about P are therefore disjoint, and hence we obtain the inequality

$$|P| \operatorname{Vol}(B(\epsilon/8)) \leq \operatorname{Vol}(M),$$

where $B(\epsilon/8)$ is a ball in $\mathbb{H}^3$ of radius $\epsilon/8$. By the maximality of P, every point in $M_{[\epsilon,\infty)}$ lies within $\epsilon/4$ of some point in P. Moreover, the number of points in P that lie within $\epsilon/2$ of a specific point of P is bounded above by $\operatorname{Vol}(B(5\epsilon/8))/\operatorname{Vol}(B(\epsilon/8))$. This is because the open ball of radius $\epsilon/2$ about this point lifts to a ball in $\mathbb{H}^3$, because it lies in $M_{[\epsilon,\infty)}$. The inverse image of P in this ball is well-spaced: the balls of radius $\epsilon/8$ about these points are all disjoint. These balls of radius $\epsilon/8$ all lie within $B(5\epsilon/8)$, and so we obtain the required bound.

This set of points P has an associated Voronoi domain. This is a cell structure, where the interior of each 3-cell consists of the set of points in $M_{[\epsilon,\infty)}$ that are closer to a specific point of P than to any other point of P. We can subdivide this cell structure into a triangulation, by subdividing each 2-cell into triangles (without introducing any new vertices) and then subdividing each 3-cell into tetrahedra by coning from the central vertex. This forms a triangulation of $M_{[\epsilon,\infty)}$.

One can easily find an upper bound on the number of tetrahedra in this triangulation that is a linear function of $|P|$, as follows. Each tetrahedron has one vertex that is a point in P and its opposite face is a subset of a 2-cell of the Voronoi domain. Thus, to obtain the required upper bound on the number of tetrahedra, it suffices to find a universal upper bound on the number of triangles in the boundary 2-sphere of each 3-cell. This 3-cell is centred at a point v in P. These triangles form a triangulation of the 2-sphere. The vertices of this triangulation are 0-cells in the Voronoi domain, which are isolated points that are equidistant from at least four points of P, one of which is v. So, a vertex in the 2-sphere is specified by choosing 3 other points in P, each of which is at most $\epsilon/2$ from v. Thus, there is a universal upper bound for the number of vertices in the 2-sphere and hence for the number of triangles. Thus, the number of tetrahedra in this triangulation of $M_{[\epsilon,\infty)}$ is at most $k \operatorname{Vol}(M)$ for some universal constant $k > 0$. □

Another useful observation is the following.

Theorem 3.4. *Let M be a finite-volume orientable hyperbolic 3-manifold. Let $\epsilon > 0$ be a Margulis constant. Then $M_{[\epsilon,\infty)}$ also admits a finite-volume hyperbolic structure.*

Proof. It is fairly straightforward to verify that $M_{[\epsilon,\infty)}$ satisfies the hypotheses of Theorem 1.2. For example, suppose that it were reducible with a reducing sphere S^2. Then S^2 lies in M, which is irreducible, and it therefore bounds a ball in $M - \partial M$. No component of $M_{(0,\epsilon]}$ can lie in this ball. This is because any cusp is non-compact, and because no closed geodesic lies within a 3-ball. Hence, the 2-sphere bounds a ball in $M_{[\epsilon,\infty)}$, contradicting the assumption that it is a reducing sphere. The verification that $M_{[\epsilon,\infty)}$ is atoroidal is similar but more complicated. The reason that $M_{[\epsilon,\infty)}$ is not Seifert fibred follows from the fact that any Dehn filling of a Seifert fibred space is either also Seifert fibred or reducible. Since M is hyperbolic and obtained by Dehn filling $M_{[\epsilon,\infty)}$, it could not be the case that $M_{[\epsilon,\infty)}$ is Seifert fibred. □

4. Hyperbolic Dehn surgery

We start with some terminology.

Definition 4.1. Let M be a 3-manifold, and let $s_1, \dots, s_n$ be slopes on distinct toral components of ∂M. Then $M(s_1, \dots, s_n)$ denotes the manifold obtained by Dehn filling along $s_1, \dots, s_n$.

Thurston's work on Dehn surgery is the motivation for much of the work in this article. His 'hyperbolic Dehn surgery theorem' [28, 27] is as follows.

Theorem 4.2 (Thurston 1979)**.** *Let M be a finite-volume orientable hyperbolic 3-manifold. Then there is a finite set $\mathcal{E}$ of slopes on ∂M such that $M(s_1, \dots, s_n)$ admits a finite-volume hyperbolic structure provided each $s_i \notin \mathcal{E}$. Moreover, if $\mathcal{E}$ is chosen appropriately, the cores of the attached solid tori form a union of disjoint geodesics in $M(s_1, \dots, s_n)$ with arbitrarily small length and these are the shortest geodesics in the manifold.*

We are therefore led to the following definition.

Definition 4.3. A slope s on ∂M is *exceptional* if $M(s)$ is not hyperbolic.

Thurston examined the case of the figure-eight knot exterior, and showed that it has exactly 10 exceptional slopes [27]. This led Gordon [14] to ask the following.

Question 4.4. Does a toral boundary component of a finite-volume orientable hyperbolic manifold have at most 10 exceptional slopes? Is the figure-eight knot exterior the unique such manifold with exactly 10 exceptional slopes?

The author and Meyerhoff [18] answered the first of these questions.

Theorem 4.5 (L-Meyerhoff 2015)**.** *Let M be a finite-volume orientable hyperbolic 3-manifold with a single toral boundary component. Then M has at most 10 exceptional slopes.*

We still do not know that the figure-eight knot is the unique worst case. However, we do know that there are only finitely many manifolds with 10 exceptional slopes. In fact, the following result was proved by Agol [3].

Theorem 4.6 (Agol 2010)**.** *Only finitely many finite-volume orientable hyperbolic 3-manifolds with a single toral boundary component have 9 or more exceptional slopes.*

Agol's proof provides an algorithm to compute this finite list of manifolds with 9 or more exceptional slopes, but it is far from practical.

5. Slope distance

Definition 5.1. The *distance* $\Delta(s_1, s_2)$ between slopes s_1 and s_2 on a torus T is the minimal number of intersection points between curves with these slopes.

If we fix a basis for $H_1(T)$ and represent slopes s_1 and s_2 as $\pm(p_1, q_1)$ and $\pm(p_2, q_2)$ with respect to this basis, then we have the formula

$$\Delta(s_1, s_2) = \left|\det \begin{pmatrix} p_1 & p_2 \\ q_1 & q_2 \end{pmatrix}\right|.$$

The exceptional slopes for the figure-eight knot are (using the standard meridian and longitude basis) as follows:

$$\begin{array}{ccccc} (1,0), & (-4,1), & (-3,1), & (-2,1), & (-1,1), \\ (0,1), & (1,1), & (2,1), & (3,1), & (4,1). \end{array}$$

The biggest distance between any two of these slopes is 8, which is the distance between $(-4, 1)$ and $(4, 1)$. Gordon [14] conjectured that the maximal distance between any two exceptional slopes on the boundary of a finite-volume orientable hyperbolic manifold is always at most 8. This was also proved by the author and Meyerhoff [18].

Theorem 5.2 (L-Meyerhoff 2015)**.** *The distance between any two exceptional slopes on a one-cusped finite-volume orientable hyperbolic 3-manifold is always at most* 8*.*

Gordon conjectured that the figure-eight knot exterior and another hyperbolic manifold (the 'figure-eight knot sister') are the unique cases where the distance 8 between exceptional slopes is realised. We still do not know whether this is true.

Once we have an upper bound on the distance between exceptional slopes, we have an upper bound on the number of such slopes. A particularly nice argument for this was given by Agol [2].

Lemma 5.3 (Agol 2000)**.** *Let $\mathcal{S}$ be a collection of slopes on a torus T, where any two of these have distance at most some integer Δ. Let p be any prime greater than Δ. Then $|\mathcal{S}| \leq p + 1$.*

Proof. Fix a basis for $H_1(T)$, so that each slope corresponds to a pair (a, b) of integers, up to sign. Each slope $\pm(a, b)$ determines a point in the projective line $\mathbb{PF}_p^1$ over the field $\mathbb{F}_p$ via

$$\pm(a, b) \mapsto [a : b] \mod p.$$

For distinct (a_1, b_1) and (a_2, b_2) in $\mathcal{S}$, we know that

$$0 < \left|\det \begin{pmatrix} a_1 & a_2 \\ b_1 & b_2 \end{pmatrix}\right| < p,$$

and so these map to distinct points in $\mathbb{PF}_p^1$. So, $|\mathcal{S}| \leq |\mathbb{PF}_p^1| = p + 1$. □

This lemma is remarkably accurate. For example, the following table provides, for each integer Δ between 0 and 10, the maximal size of a collection of slopes with the property that any two slopes in the collection have distance at most Δ. We can see that Agol's bound is sharp in each of these cases except $\Delta = 0, 7$.

Maximal intersection number Δ	0	1	2	3	4	5	6	7	8	9	10
Maximal number of slopes	1	3	4	6	6	8	8	10	12	12	12

Agol [3] used this lemma and the following theorem to deduce Theorem 4.6.

Theorem 5.4 (Agol 2010)**.** *For all but finitely many finite-volume orientable hyperbolic 3-manifolds with a single toral boundary component, the distance between any two exceptional slopes is at most* 5*.*

6. Slope length

Definition 6.1. Fix a horoball neighbourhood N of ∂M. Let s be a slope on ∂M. The *length* $L(s)$ of this slope, with respect to N, is the length of any Euclidean geodesic representative of s on ∂N. When we do not refer to N, we take it to be the maximal cusp.

Example 6.2. The slope μ in Figure 5 (which is actually the meridian slope of the figure-eight knot) has length 1. To see this, we take a maximal horoball neighbourhood of the end, as described in Example 2.2. This was obtained by taking four identically-sized horoballs in the two copies of $\mathbb{H}^3$, forming their intersection with the two ideal tetrahedra and then gluing them together. When the horoball neighbourhood is maximal, these four horoballs in each copy of $\mathbb{H}^3$ are tangent. We now consider a face of these ideal tetrahedra containing a geodesic representative of μ. Realise this ideal triangle in the upper half-space model, as shown in Figure 6, with one vertex at the point at infinity, another vertex at the origin and the final vertex at $(1, 0, 0)$. This face intersects the boundary of the maximal cusp in three arcs, two of which are parts of Euclidean circles and the remaining one is $\{(x, y, z) : 0 \leq x \leq 1,\ y = 0,\ z = 1\}$. Since this arc lies in $\{z = 1\}$, its hyperbolic length is equal to its Euclidean length, which is 1. Thus, we see that the length of each side of each Euclidean triangle that forms the boundary of the maximal cusp is 1.

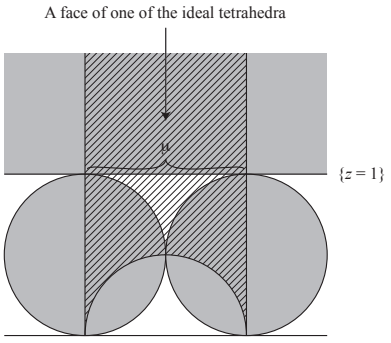

Figure 6. The length of the meridian slope on the figure-eight knot is 1.

The following lemma, due to Meyerhoff [19], gives that 1 is a universal lower bound on slope length.

Lemma 6.3. *For any finite-volume orientable hyperbolic 3-manifold M, the length of each slope on the boundary of a maximal cusp is at least* 1.

Proof. The universal cover of $M - \partial M$ is $\mathbb{H}^3$. The inverse image of the maximal cusp is a union of horoballs in $\mathbb{H}^3$. Use the upper half space model for $\mathbb{H}^3$ and by performing an isometry if necessary, we may arrange that one of the horoballs, denoted B_∞, is $\{(x, y, z) : z \geq 1\}$. Since the cusp is maximal, some covering translate B_1 of B_∞ is tangent to it. In the maximal cusp torus, pick a geodesic representative for s through the point of tangency and lift it to a geodesic arc $\tilde{s}$ in ∂B_∞ starting at $B_\infty \cap B_1$. This has the same length as s. There is a covering transformation taking the start of $\tilde{s}$ to its end. This preserves B_∞, but it takes B_1 to another horoball B_2. (See Figure 7.) Since B_1 and B_2 both have Euclidean diameter 1 and their interiors are disjoint, the length of $\tilde{s}$ is at least 1. □

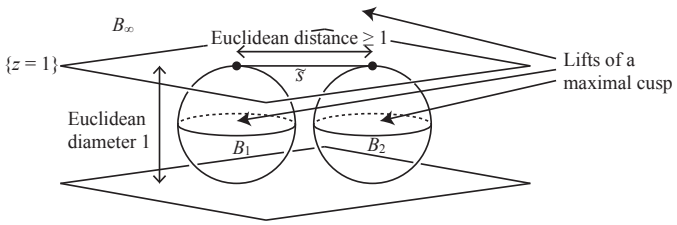

FIGURE 7. The length of each slope on the boundary of a maximal cusp is at least 1

One of the most important results about slope length is the following theorem, known as the 2π-theorem [5].

Theorem 6.4 (Gromov-Thurston)**.** *Let M be a compact orientable finite-volume hyperbolic 3-manifold. Let N be a horoball neighbourhood of ∂M. Let $s_1, \ldots, s_n$ be slopes on distinct components of ∂M. Suppose that $L(s_i) > 2\pi$ with respect to N for each i. Then $M(s_1, \ldots, s_n)$ has a complete negatively curved Riemannian metric. Hence, by Theorem 1.2, $M(s_1, \ldots, s_n)$ also admits a complete finite-volume hyperbolic structure.*

The proof is surprisingly direct. One removes N from $M - \partial M$ and then one attaches a collection of solid tori, as specified by the filling. Each solid torus is assigned a Riemannian metric that is negatively curved and that, near the boundary, agrees with that of N near ∂N. This metric is highly symmetric: it is invariant under the translational and rotational symmetry of the solid torus. Hence, it is specified by a relatively small amount of data: just three real-valued functions. These functions can be chosen more or less explicitly, so that all the sectional curvatures are negative, provided that the length of each s_i is more than 2π.

The explicit nature of the metric actually provides more information. In particular, if the lengths of all the s_i tend to infinity, then the sectional curvatures can be arranged to be arbitrarily close to -1 and the volumes of the solid tori tend to the volume of the removed cusps. One can compare the volume of this metric with

the volume of a hyperbolic metric on $M(s_1, \dots, s_n)$, using Gromov norm [27]. We obtain the following result [8, 12]

Theorem 6.5 (Futer-Kalfagianni-Purcell 2008, Cooper-L 1998, Thurston 1980)**.** *If $L(s_i) \geq \ell_{\min} > 2\pi$ for each i, then the hyperbolic volume of $M(s_1, \dots, s_n)$ is bounded as follows:*

$$\left(1 - \left(\frac{2\pi}{\ell_{\min}}\right)^2\right)^{3/2} \mathrm{Vol}(M) \leq \mathrm{Vol}(M(s_1, \dots, s_n)) < \mathrm{Vol}(M).$$

An upper bound on the volume of M is important, due to the following result, that relates volume and surgery.

Theorem 6.6 (Jørgensen)**.** *Let $V > 0$. Then there is a finite list of finite-volume orientable hyperbolic 3-manifolds $\{M_1, \dots, M_n\}$ such that any orientable hyperbolic 3-manifold with volume at most V is obtained by Dehn filling some M_i.*

Proof. We saw in Theorem 3.2 that when M is a finite-volume orientable hyperbolic 3-manifold, it is obtained by Dehn filling $M_{[\epsilon,\infty)}$ where ϵ is a Margulis constant. Theorem 3.4 states that $M_{[\epsilon,\infty)}$ admits a finite-volume hyperbolic structure. We also saw in Theorem 3.3 that $M_{[\epsilon,\infty)}$ admits a triangulation where the number of tetrahedra is bounded above by a linear function of $\mathrm{Vol}(M)$. Hence, in our case, there is an upper bound on the number of tetrahedra, and so there is a finite list $\{M_1, \dots, M_n\}$ of possibilities for $M_{[\epsilon,\infty)}$. □

This theorem implies that if we have a sequence of orientable hyperbolic 3-manifolds M_i with bounded volume, then we may pass to a subsequence where each M_i is obtained by Dehn filling some fixed finite-volume hyperbolic 3-manifold. In fact, we can deduce a little more from the proof.

Proposition 6.7. *Let M_i be a sequence of orientable hyperbolic 3-manifolds with bounded volume. Then, there is a finite-volume orientable hyperbolic 3-manifold M, a subsequence $M_{i(j)}$ and slopes $\sigma_1^{i(j)}, \dots, \sigma_k^{i(j)}$ on distinct components of ∂M such that $M_{i(j)}$ is homeomorphic to $M(\sigma_1^{i(j)}, \dots, \sigma_k^{i(j)})$. Moreover, these slopes are pairwise distinct.*

Proof. Only the final claim remains to be proved. As argued above, we may pass to a subsequence $M_{i(j)}$ where the thick part of $M_{i(j)}$ is some fixed manifold M'. We may also assume that the number of components of $\partial M_{i(j)}$ is constant. Thus, $M_{i(j)}$ is homeomorphic to $M'(\sigma_1^{i(j)}, \dots, \sigma_l^{i(j)})$ for slopes $\sigma_1^{i(j)}, \dots, \sigma_l^{i(j)}$ on $\partial M'$. If $\sigma_1^{i(j)}$ has a constant subsequence, pass to this. Otherwise pass to a subsequence where the slopes $\sigma_1^{i(j)}$ are pairwise distinct. Repeat with $\sigma_2^{i(j)}$, and so on. If any slope $\sigma_m^{i(j)}$ is constant, then perform this Dehn filling. Thus, we reattach some components of thin part of $M_{i(j)}$. Let M be the resulting manifold. Then $M_{i(j)}$ is $M(\sigma_1^{i(j)}, \dots, \sigma_k^{i(j)})$ but now the filling slopes are pairwise distinct. Note that M is obtained from $M_{i(j)}$ by removing some components of its thin part. As in the proof of Theorem 3.4, we can again verify that this satisfies the topological conditions for the existence of a finite-volume hyperbolic structure. □

7. Surgical finiteness

We saw in Theorem 6.4 that when we Dehn fill a hyperbolic 3-manifold along a slope with length more than 2π, the resulting manifold M is hyperbolic. It is natural to wonder whether a given manifold M can be obtained in more than one way by this procedure. The following theorem [8] is a finiteness result that showcases some of the geometric methods that we have discussed.

Theorem 7.1 (Cooper-L 1998)**.** *Let M be a compact orientable 3-manifold and let $\epsilon > 0$. Then there are only finitely many finite-volume hyperbolic 3-manifolds X and slopes $s_1, \dots, s_n$ on ∂X with $L(s_i) \geq 2\pi + \epsilon$ for each i, with respect to some horoball neighbourhood of the cusps, and with $X(s_1, \dots, s_n)$ homeomorphic to M.*

Corollary 7.2. *A 3-manifold can be obtained by p/q surgery on a hyperbolic knot K in S^3 with $|q| > 11$ for only finitely many K and p/q.*

Proof. We will see in Section 8 that a slope p/q on the exterior of a hyperbolic knot in S^3 satisfies $L(p/q) > (0.558)|q|$. Hence if $|q| \geq 12$, then $L(p/q) > 6.96 > 2\pi + 0.41$. Now apply Theorem 7.1. □

Proof of Theorem 7.1. Suppose that, on the contrary, M is homeomorphic to $X_i(s_1^i, \dots, s_{n(i)}^i)$ for infinitely many finite-volume hyperbolic 3-manifolds X_i and slopes $s_1^i, \dots, s_{n(i)}^i$ with $L(s_j^i) \geq 2\pi + \epsilon$. By Theorem 6.4, M has a hyperbolic structure. By Theorem 6.5, the volume of X_i is at most a constant times the volume of M, where the constant depends only on ϵ. Thus, the X_i have bounded volume. So, by Proposition 6.7, we may pass to a subsequence where every X_i is obtained by Dehn filling some fixed hyperbolic manifold Y. Moreover, the filling slopes are pairwise distinct. So, M is homeomorphic to $Y(\sigma_1^i, \dots, \sigma_m^i)$, where $\sigma_1^i, \dots, \sigma_m^i$ are slopes on distinct components of ∂Y. These slopes all have length more than 2π with respect to some horoball neighbourhood N^i of the cusps. The reason for this is that the slopes $s_1^i, \dots, s_{n(i)}^i$ were assumed to have length at least $2\pi + \epsilon$ with respect to some horoball neighbourhood of ∂X_i. This horoball neighbourhood is closely approximated by a horoball neighbourhood of the relevant cusps of Y, and so the lengths of the corresponding slopes on ∂Y are also more than 2π. The extra slopes on ∂Y that are filled to form X_i are pairwise distinct. Hence, their lengths are all eventually more than 2π as well. We may pass to a subsequence and re-order the slopes so that $\sigma_1^i, \dots, \sigma_k^i$ are independent of i and $\sigma_{k+1}^i, \dots, \sigma_m^i$ are pairwise distinct. Note that $k < m$, as otherwise X_i and the filling slopes $s_1^i, \dots, s_{n(i)}^i$ are independent of i. Let $Z = Y(\sigma_1^i, \dots, \sigma_k^i)$, which is a fixed hyperbolic manifold, by the 2π theorem (Theorem 6.4). Thus, M is homeomorphic to $Z(\sigma_{k+1}^i, \dots, \sigma_m^i)$. So, by hyperbolic Dehn surgery (Theorem 4.2), the manifold $Z(\sigma_{k+1}^i, \dots, \sigma_m^i)$ is hyperbolic for all sufficiently large i. Moreover, as i tends to infinity, the length of the shortest geodesic in $Z(\sigma_{k+1}^i, \dots, \sigma_m^i)$ tends to zero. Hence, this length is eventually less than the length of the shortest geodesic in M. This is a contradiction. □

8. Slope length and distance

In this section, we establish the following simple but important relationship between slope length, slope distance and the area of a cusp torus.

Lemma 8.1. *Let s_1 and s_2 be slopes on a Euclidean torus T. Then*

$$\Delta(s_1, s_2) \leq \frac{L(s_1)L(s_2)}{\text{Area}(T)}.$$

Proof. Pick an orientation on s_1, making it a primitive element of $H_1(T)$ and extend it to a basis $\{s_1, s_3\}$ for $H_1(T)$. We can visualise these slopes by lifting to the universal cover of T, which is the Euclidean plane $\mathbb{R}^2$. The inverse image of some basepoint in T is a lattice in $\mathbb{R}^2$. We may assume that one of the lattice points is the origin O. If we lift s_1, s_2, s_3 to paths starting at O, the endpoints of these paths are lattice points, which we shall also label s_1, s_2, s_3. (See Figure 8.)

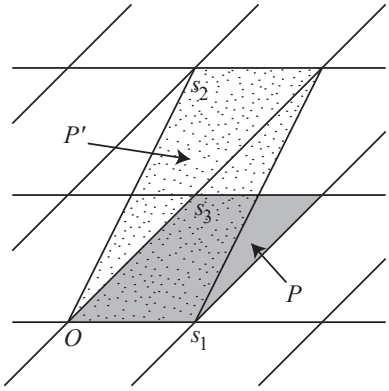

FIGURE 8. The picture that leads to the inequality relating slope distance, slope length and cusp area

Let P be the parallelogram with corners O, s_1, s_3 and s_1+s_3. Since $\{s_1, s_3\}$ forms a basis for $H_1(T)$, this parallelogram P forms a fundamental domain and hence has area equal to the area of T. There are integers p and q such that $s_2 = ps_1 + qs_3$. Taking intersection numbers with s_1, we obtain

$$\langle s_1, s_2\rangle = p\langle s_1, s_1\rangle + q\langle s_1, s_3\rangle = \pm q.$$

Let P' be the parallelogram with sides s_1 and s_2. Then its area satisfies

$$\text{Area}(P') = |q|\text{Area}(P) = \Delta(s_1, s_2)\text{Area}(P).$$

But the area of P' is at most the product of two adjacent sides. We deduce that $L(s_1)L(s_2) \geq \Delta(s_1, s_2)\text{Area}(T)$, as required. □

This inequality is used to establish the lower bound on $L(p/q)$ that appeared in the proof of Corollary 7.2. The meridian slope ∞ on the exterior of a knot in the 3-sphere has length at most 6. An upper bound of 2π is a consequence of Theorem 6.4, since the 3-sphere does not admit a hyperbolic structure. But as we will see in Theorem 10.1, this upper bound on the length of the meridian slope can be improved to 6. We will see in Theorem 9.3 that the boundary T of a maximal cusp has area at least 3.35. Hence, by Lemma 8.1,

$$L(p/q) \geq \frac{\Delta(\infty, p/q)\text{Area}(T)}{L(\infty)} \geq \frac{|q|3.35}{6} \geq (0.558)|q|.$$

9. The area of the boundary of the maximal cusp

The inequality in Lemma 8.1 is a key method for producing an upper bound on the distance between exceptional slopes. The area of T is the denominator. Thus, if one can get good lower bounds on the area of T, we can obtain good upper bounds on the distance between exceptional slopes. Indeed, much of the main progress in the hyperbolic approach to Dehn surgery has been made by finding increasingly good lower bounds on the area of a maximal cusp torus. The following sequence of theorems provides, in chronological order, some of the known lower bounds on cusp torus area [5, 1, 7, 13].

Theorem 9.1 (Thurston 1980)**.** *The area of a maximal cusp torus is at least* $\sqrt{3}/2$*.*

Theorem 9.2 (Adams 1987)**.** *The area of a maximal cusp torus is at least* $\sqrt{3}$*.*

Theorem 9.3 (Cao-Meyerhoff 2001)**.** *The area of a maximal cusp torus is at least* 3.35*.*

Theorem 9.4 (Gabai-Meyerhoff-Milley 2009)**.** *The area of a maximal cusp torus is at least* 3.7 *unless* M *is a member of one of finitely many explicit families.*

We will explain some of the ideas behind these theorems later, and will make Theorem 9.4 a little more precise. But first we give the proof of Thurston's result, which is quite elementary.

Proof of Theorem 9.1. The boundary ∂N of a maximal cusp lifts to a horosphere in $\mathbb{H}^3$, which is isometric to a Euclidean plane. If we pick a basepoint in ∂N, its inverse image in the plane is a lattice. Lemma 6.3 gives that the length of each slope is at least 1 and hence these lattice points are all at least distance 1 from each other. Thue's theorem gives that the densest possible disc packing in the plane is the hexagonal lattice. When the distance between adjacent lattice points of a hexagonal lattice is exactly 1, it has co-area $\sqrt{3}/2$. Hence, the area of ∂N is at least $\sqrt{3}/2$. □

These lower bounds on area give the following upper bounds on the distance between exceptional slopes. Then, applying Lemma 5.3 with a suitably chosen prime, we also get the following upper bounds on the number of exceptional slopes.

Area(∂N)	Upper bound on $\Delta(s_1, s_2)$ using 2π-theorem	Prime p	Upper bound on number of exceptional slopes
$\sqrt{3}/2$	45	47	48
$\sqrt{3}$	22	23	24
3.35	11	13	14
3.7	10	11	12

Table 1. Lower bounds on cusp torus area and the resulting upper bounds on the number of exceptional slopes

We can see that the final upper bound on the number of exceptional slopes is 12, which is not far off the desired bound of 10. To get to this bound, we will use an improved version of the 2π-theorem. In addition, we will use not the just the statement of Theorem 9.4 but also the underlying techniques developed by Gabai, Meyerhoff and Milley.

10. The 6-theorem

We saw in the previous section that good lower bounds on cusp torus area lead to good upper bounds on the distance between exceptional slopes. There is another way of improving the utility of Lemma 8.1, which is to reduce the upper bound on the length of an exceptional slope. The following result [17, 2] improves the 2π-theorem, by reducing the critical slope length from 2π to 6. It is known as the 6-theorem.

Theorem 10.1 (Agol, L 2000)**.** *Let M be a compact orientable finite-volume hyperbolic 3-manifold. Let N be a horoball neighbourhood of ∂M. Let $s_1, \dots, s_n$ be slopes on distinct components of ∂M. Suppose that, for each slope s_i, $L(s_i) > 6$ with respect to N. Then $M(s_1, \dots, s_n)$ is irreducible and atoroidal and has infinite, word hyperbolic fundamental group. Hence, by Theorem 1.2, $M(s_1, \dots, s_n)$ is hyperbolic.*

Proof outline. We will focus on the conclusion that $M(s_1, \dots, s_n)$ is irreducible, because this illustrates the main ideas.

Let $L \subset M(s_1, \dots, s_n)$ be the cores of the attached solid tori. Suppose that $M(s_1, \dots, s_n)$ is reducible, and pick a reducing sphere S^2 in $M(s_1, \dots, s_n)$ that is transverse to L and that intersects L as few times as possible. Let $F \subset M$ be $S^2 - \text{int}(N(L))$.

Then F is essential in M. In other words, it is incompressible and boundary-incompressible. For suppose that it were compressible. Then this would also give a compression disc for S^2, which would then compress S^2 to two 2-spheres. At least one of these would have to be a reducing sphere. Since the disc was a compression disc for the planar surface F, it has to separate the components of ∂F. Hence, each of these spheres intersects L fewer times than the original S^2. This contradicts our minimality assumption.

The proof that F is boundary-incompressible also uses the fact that $|S^2 \cap L|$ is minimal. Consider a potential boundary-compression disc D. This intersects ∂M in an arc, which lies in an annulus A that is the closure of a component of $\partial M - \partial F$. If this arc is inessential in A, then D can be easily be modified to form a compression disc for F, which we have shown to be impossible above. On the other hand, if the arc is essential in A, then it joins distinct components ∂F, which correspond to points of $S^2 \cap L$ with opposite sign. This disc can then be used to perform an isotopy of S^2 that reduces $|S^2 \cap L|$ by two, as shown in Figure 9. This is again a contradiction.

We note that $|\partial F| \geq 3$, since otherwise F is a disc or an annulus. But any properly embedded annulus in a finite-volume hyperbolic 3-manifold is boundary-parallel. Hence, it is boundary-compressible, which we have shown not to be the case. Furthermore, if F were a disc, then it would form a compression disc for ∂M. But each cusp of a hyperbolic manifold M is π_1-injective.

As $|\partial F| \geq 3$, F admits an ideal triangulation, which is just an expression of $F - \partial F$ as a union of triangles glued along their edges and with their vertices all removed.

We homotope each edge of the ideal triangulation to a geodesic. The only obstruction to performing this is if the edge is homotopic into a cusp, but this implies that the surface is boundary-compressible or compressible, which we have shown not to be the case. We may then homotope each ideal triangle so that it is totally geodesic. The result is a (possibly non-embedded) surface $\overline{F}$ that is known as

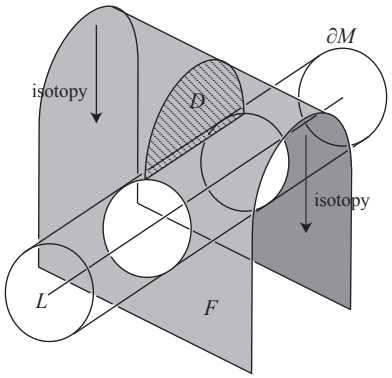

FIGURE 9. A boundary-compression disc D for F joining distinct components of ∂F leads to an isotopy that reduces $|S^2 \cap L|$.

pleated. Since it is a union of hyperbolic ideal triangles glued along their boundary geodesics, $\overline{F}$ inherits a hyperbolic metric. So, applying Gauss-Bonnet and using the fact that F is a planar surface,

$$\text{Area}(\overline{F}) = -2\pi\chi(\overline{F}) = 2\pi(|\partial F| - 2) < 2\pi|\partial F|.$$

We will show that each component of ∂F contributes at least 2π to the area of $\overline{F}$. This will lead to a contradiction.

Consider any component of N, which is of the form $T^2 \times [1, \infty)$. For any $t \geq 1$, let T_t be the torus $T^2 \times \{t\}$. We will examine how $\overline{F}$ intersects these tori. Consider any component X of $\overline{F} \cap N$. Now N is of the form $T^2 \times [1, \infty)$ and far out towards $T^2 \times \{\infty\}$, X looks like a collection of half-open annuli, each one corresponding to a component of ∂F. In fact, X can contain at most one of these half-open annuli, since otherwise the topology of $T^2 \times [1, \infty)$ could be used to build a boundary-compression disc for F. Suppose that X has exactly one such half-open annulus, corresponding to a single component of ∂F. Then X forms a homology between this component of ∂F and the curves $X \cap \partial N$. Hence, for homological reasons, the intersection between X and T_t contains a curve of slope s_i, for any $t \geq 1$, where s_i is the filling slope on the relevant component of ∂M. Hence, the length of $X \cap T_t$ is at least $L(s_i)/t$.

We wish to find a lower bound for the area of X. The co-area formula gives such a lower bound, in terms of the length of curves $X \cap T_t$. Specifically,

$$\text{Area}(X) \geq \int_{t=1}^{\infty} \frac{\text{Length}(X \cap T_t)}{t}\, dt \geq \int_{t=1}^{\infty} \frac{L(s_i)}{t^2}\, dt = L(s_i).$$

Thus, we deduce that each component of ∂F contributes at least $L(s_i)$ to the area of $\overline{F}$. This is a contradiction if $L(s_i) \geq 2\pi$.

Of course, we would like to find a contradiction if $L(s_i) > 6$. To achieve this, we use the parts of $\overline{F}$ in $M - N$. Associated to N, there is a spine S for M made up of totally geodesic cells. This is defined as follows:

$$S = \{x \in M - \partial M : x \text{ does not have a unique closest point in } N\}.$$

In the case of the figure-eight knot, the spine S is shown in Figure 10. To see that S is a spine for M, we will exhibit a retraction r of $M - S$ onto N. Each point

x in $M - S$ has, by definition, a unique closest point in N. We define $r(x)$ to be this point. It is not hard to see that r is a homotopy equivalence and hence each component of $M - S$ is homotopy equivalent to $T^2 \times [1, \infty)$. In fact, this implies that it is homeomorphic to $T^2 \times [1, \infty)$. Hence, again we deduce that each component X of $\overline{F} \cap (M - S)$ is incident to at most one component of ∂F. We will show that, when such a component X is incident to ∂F, it has area at least 2π, which will give us our required contradiction.

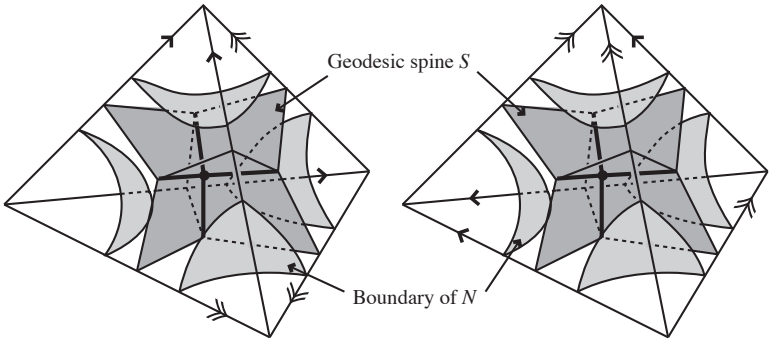

FIGURE 10. The geodesic spine for the figure-eight knot complement.

Let $\tilde{S}$ be the inverse image of S in $\mathbb{H}^3$. Here, we are working with the upper half-space model for $\mathbb{H}^3$ and have arranged for a component of the inverse image of N to be $B_\infty = \{(x, y, z) : z \geq 1\}$. Let $\tilde{E}$ be the component of $\mathbb{H}^3 - \tilde{S}$ containing B_∞. Let $\tilde{E}_t$ be $\tilde{E} \cap \{(x, y, z) : z = t\}$. For $t \geq 1$, this is a horosphere, but as t decreases below 1, discs are removed the horosphere, where these discs lie below $\tilde{S}$. (These discs are not necessarily disjoint.)

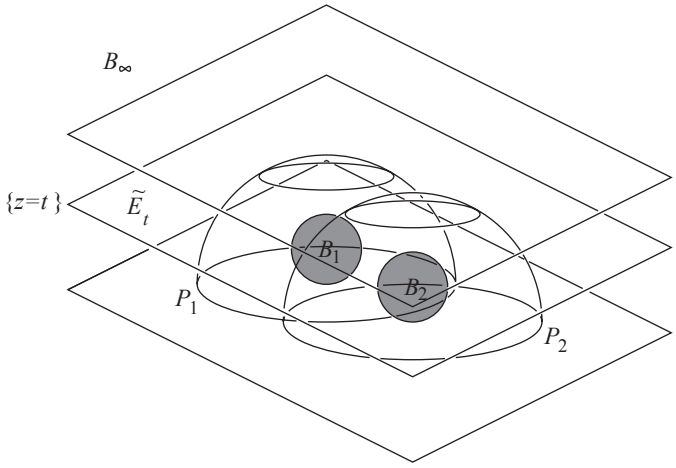

FIGURE 11. The set $\tilde{E}_t$ is obtained from a horosphere by removing (not necessarily disjoint) discs.

We claim that these discs are closest to each other in the case of the figure-eight knot and that these discs are largest in the case of the figure-eight knot. To see this,

consider two such discs that are removed from $\{(x, y, z) : z = t\}$ to form $\tilde{E}_t$. These discs lie below two totally geodesic planes P_1 and P_2 that are equidistant between B_∞ and a horoball B_1, and between B_∞ and a horoball B_2. We now perform a sequence of modifications to B_1 and B_2. Each modification will only make the discs bigger and will bring them closer. At the end of these modifications, we will reach the arrangement of horoballs as shown in Figures 6 and 12. These modifications are: slide B_1 and B_2, without changing their size, until they are tangent; scale them both by the same Euclidean factor until one (B_2, say) is tangent to B_∞; then enlarge B_1, keeping it tangent to B_2, until it also just touches B_∞. The resulting arrangement of horoballs is the same as the horoballs in the case of the figure-eight knot, which proves the claim.

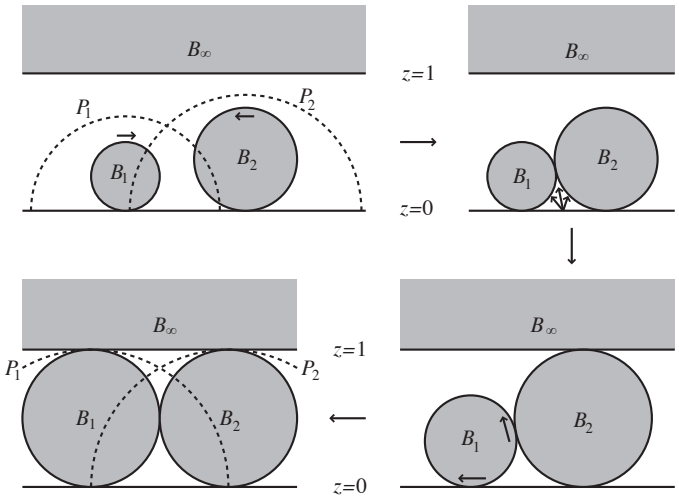

FIGURE 12. Moving horoballs into the figure-eight knot configuration

Since the construction of $\tilde{E}_t$ is invariant under the covering transformations preserving B_∞, we obtain a well-defined surface E_t in $M - \partial M$, which is the image of $\tilde{E}_t$ under the covering map. These sets E_t fill up the entirety of $M - S$.

Consider any component X of $\overline{F} \cap (M - S)$ that is incident to a component of ∂F. The co-area formula again gives a lower bound on the area of X in terms of the lengths of $X \cap E_t$ as t varies:

$$\text{Area}(X) \geq \int \frac{\text{Length}(X \cap E_t)}{t}\, dt.$$

So, we need to find a lower bound on the length of $X \cap E_t$. Now, E_t is a torus T_t with some discs removed. The 1-manifold $X \cap E_t$ can be extended to a 1-manifold in T_t by inserting arcs into these discs, so that the resulting curves are homologous to the slope s_i. Hence, the resulting curves have length at least $L(s_i)/t$. We care about the part of these curves lying in E_t. It is possible to show that the length of $X \cap E_t$ is minimised when the discs removed from T_t are as large as possible and when the curves are a concatenation of arcs, each one running between discs that are closest. In other words, these lengths are minimised in the case where M is the figure-eight knot and X is a totally geodesic surface that is a union of pieces as in

Figure 13. So,

$$\frac{\text{Length}(X \cap E_t)}{L(s_i)/t} \geq f(t),$$

where $f(t)$ is this ratio in the case of the figure-eight knot. So,

$$\text{Area}(X) \geq \int \frac{\text{Length}(X \cap E_t)}{t}\, dt \geq \int \frac{L(s_i)}{t^2} f(t)\, dt = L(s_i) \int_t \frac{f(t)}{t^2}\, dt.$$

This final integral is easily computed: it is the ratio of the area of X to the slope length in Figure 13. In other words, it is the ratio of $(\pi/3)$ to 1. Thus, we deduce that

$$\text{Area}(X) \geq (\pi/3) L(s_i).$$

This is a contradiction if $L(s_i) > 6$ for each i.

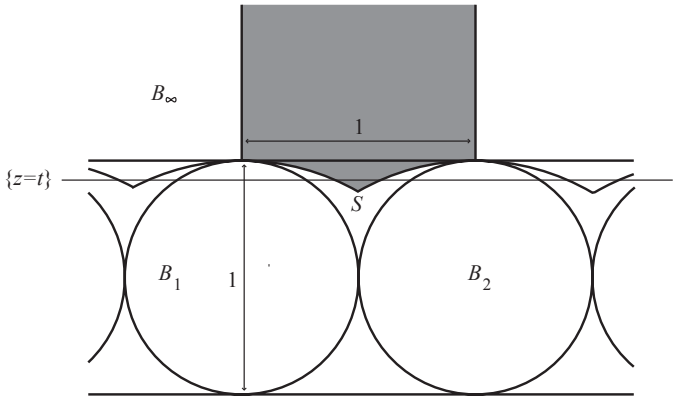

FIGURE 13. The grey region has area $\pi/3$. Its intersection with $\{(x, y, z) : z = 1\}$ has length 1.

The above argument only established the irreducibility of $M(s_1, \ldots, s_n)$, but the argument can be readily adapted to establish the remaining topological conclusions. To show that $M(s_1, \ldots, s_n)$ is atoroidal, one considers a punctured torus in M rather than a punctured sphere. To show that $\pi_1(M(s_1, \ldots, s_n))$ has infinite fundamental group, one actually shows that each of the cores of the attached solid tori has infinite order. Otherwise, there is an immersed disc in $M(s_1, \ldots, s_n)$ with boundary mapping to a multiple of the core curve. This restricts to a planar surface in M. One boundary component is not a meridian, but the remaining boundary components all have slopes in the set $\{s_1, \ldots, s_n\}$. The above argument was robust enough that we can ignore the area contribution from the boundary component that does not have slope in the set $\{s_1, \ldots, s_n\}$.

The final claim in the theorem, which is technically the hardest part to prove, is that $\pi_1(M(s_1, \ldots, s_n))$ is word hyperbolic. In other words, loops in $M(s_1, \ldots, s_n)$ satisfy a linear isoperimetric inequality. More precisely, any homotopically trivial loop ℓ in $M(s_1, \ldots, s_n)$ bounds a disc with area at most a constant times the length of ℓ. Here, area is measured using any Riemannian metric on $M(s_1, \ldots, s_n)$. It is convenient to use a metric that is obtained from the hyperbolic metric on $M - \partial M$ by removing cusps and attaching some Riemannian solid tori.

We perturb ℓ a little so that misses the core curves in $M(s_1, \ldots, s_n)$ and therefore corresponds to a loop in M, which we will also call ℓ. Since ℓ is homotopically

trivial in $M(s_1, \ldots, s_n)$, it bounds a disc D in $M(s_1, \ldots, s_n)$. We suppose that D intersects L as few times as possible. Its restriction to M is a planar surface F. One of its boundary components is ℓ and the others have slopes in $\{s_1, \ldots, s_n\}$. The aim is to find a linear upper bound on the area of F in terms of the length of ℓ. Then, filling in the Riemannian solid tori to form $M(s_1, \ldots, s_n)$, we will obtain a disc in $M(s_1, \ldots, s_n)$ with area that is linearly bounded above by the length of ℓ.

The first thing to do is homotope ℓ to a geodesic $\overline{\ell}$ in M. This homotopy is realised by a mapped-in annulus. Its area can be bounded above linearly in terms of the length of ℓ. We now homotope F to a pleated surface $\overline{F}$. Unlike in the above argument, it is no longer the case that each component of $\partial F \cap \partial M$ contributes at least 2π to the area of $\overline{F}$. We would able to deduce this conclusion if each component of ∂F gave rise to a half-open annulus X running from infinity in a cusp to S. But this cannot be the case for every component of ∂F, as this would imply that $\overline{F}$ had too much area. Instead, X must contain parts of $\overline{\ell}$ in its boundary. Thus, we deduce that for each component of ∂F, we get a definite contribution to the length of $\overline{\ell}$. So, the length of ℓ is at least a constant times $|\partial F|$. Thus, the area of $\overline{F}$ is bounded above by a constant times the length of ℓ as required. □

11. The Epstein-Penner decomposition and the Adams horoballs

In order to improve Theorem 10.1, it will be necessary to have a better picture of the arrangement of horoballs in $\mathbb{H}^3$. Henceforth, we will fix a horoball neighbourhood N of ∂M. We will choose the covering map $\mathbb{H}^3 \to M - \partial M$ so that one component B_∞ of the inverse image of N is $\{(x, y, z) : z \geq 1\}$ in the upper half-space model.

Consider any other component B_1 of the inverse image of N. The group of covering transformations of the cusp can be realised as a group of Euclidean translations of $\mathbb{H}^3$ that preserve B_∞. This group acts on the inverse image of N, and so sends B_1 to a lattice of horoballs in $\mathbb{H}^3$, all of which have the same Euclidean diameter. But it is a key observation of Adams [1] that there is always another horoball in the inverse image of N that has the same Euclidean size as B_1 but that is not in this lattice.

One way to understand this is to consider the spine S defined in the proof of Theorem 10.1. Dualise S to form a union of ideal polyhedra. Thus, dual to each totally geodesic face, we have a bi-infinite geodesic. Dual to an edge of S, we have a totally geodesic face. Dual to a vertex of S, we have an ideal polyhedron. Thus, we have expressed $M - \partial M$ as a union of ideal polyhedra glued isometrically along their faces. This is known as the *Epstein-Penner decomposition* of M [11]. In the case of the figure-eight knot complement, the Epstein-Penner decomposition is the ideal triangulation given in Example 1.7.

The Epstein-Penner decomposition depends only on a choice of horoball neighbourhood N of the cusps. When ∂M is a single torus, then S is independent of the choice of N, and hence so too is the Epstein-Penner decomposition. When $|\partial M| > 1$, then a canonical choice of N can be made. The Epstein-Penner decomposition is a very important construction; for example it encodes all the symmetries of M.

Suppose now that N is a maximal horoball neighbourhood of an end of $M - \partial M$. Then it cannot be expanded any further, and so there is a horoball B_1 in the inverse image of N that is tangent to B_∞. Between B_∞ and B_1, there is a totally geodesic face of S and dual to this face is an edge of the Epstein-Penner decomposition. Pick an orientation on this edge. Then this edge intersects the cusp torus at *two* points,

one point entering N and one point leaving it. Thus, incident to B_∞, there are *two* full-sized horoballs that do not differ by the covering transformations of B_∞. These are called the *Adams horoballs*. (See Figure 14.)

If we consider their vertical projection onto ∂B_∞ and thence onto ∂N, we see two discs with disjoint interiors. So, the argument establishing Theorem 9.1 can be improved. Thue's theorem gives that the densest possible packing of discs of diameter 1 in the plane is given by the hexagonal packing. But we now have an arrangement of discs in ∂B_∞ that projects to two discs in ∂N with disjoint interiors and so the area of ∂N is at least $\sqrt{3}$. This proves Theorem 9.2.

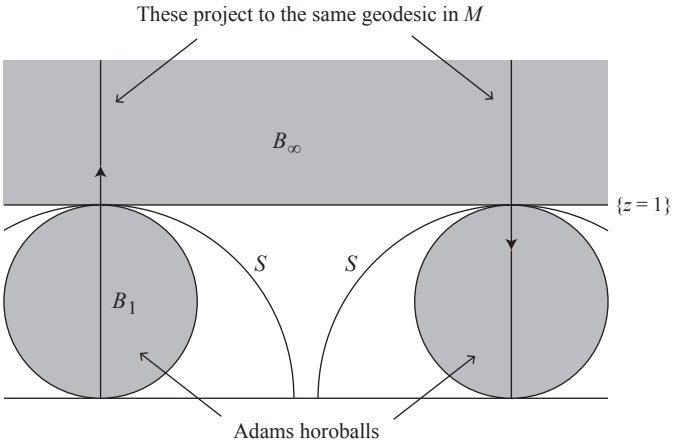

FIGURE 14. The Adams horoballs are two full-sized horoballs that do not differ by a covering transformation preserving B_∞

12. Cusp area due to Cao and Meyerhoff

The above argument of Adams used the two Adams horoballs, both of which have Euclidean diameter 1. But how close can these horoballs be? If we could argue that they are not too close, then we would obtain an improvement to Adams's lower bound on the area of the maximal cusp torus.

So consider two Adams horoballs B_1 and B_2 that are closest to each other. We have two bi-infinite geodesics γ_1 and γ_2 running from B_1 to B_∞, and from B_2 to B_∞. These two geodesics map to the same geodesic in M (but with opposite orientations). We also have a geodesic γ_3 joining the points at infinity of B_1 and B_2. Cao and Meyerhoff showed that γ_3 maps to a different geodesic in M. In particular, the intersection points between γ_3 and $\partial B_1 \cup \partial B_2$ map to points in ∂N that are different from the images of $\gamma_1 \cap \partial B_\infty$ and $\gamma_2 \cap \partial B_\infty$. Thus, about these points in ∂N, we actually get discs in ∂N arising from *four* horoballs.

If B_1 and B_2 were to touch, then γ_3 would have zero length outside of N, and so these four horoballs would all have maximal size. Hence, in this case, the area of ∂N would be at least $2\sqrt{3}$. On the other hand, if B_1 and B_2 were to be some positive distance apart, then we would obtain an improvement to our application of Thue's theorem. Cao and Meyerhoff [7] formalised this into a computer-assisted analysis, thereby proving Theorem 9.3.

13. Orthoclasses

It is obvious from the above discussion that it is important to understand how close two horoballs in the inverse image of N can be. We could perform a covering transformation taking one of these to B_∞. So, equivalently, we need to understand how close a horoball B can be to B_∞. The *orthodistance* of B is just the hyperbolic distance between B and B_∞. If a horoball has orthodistance d, then a simple calculation in hyperbolic geometry gives that its Euclidean diameter is e^{-d}.

Joining B to B_∞, there is a geodesic (that is not necessarily an edge of the Epstein-Penner decomposition). But if we pick an orientation on this geodesic, then we see that its image in M intersects ∂N in two points, one point where the geodesic goes perpendicularly into N and one point where the geodesic leaves N. Thus, we see that, just as in the case of the Adams horoballs, they come in pairs. The formal definition is as follows.

Definition 13.1. Let B_1 and B_2 be two horoballs in the inverse image of N. We say that they are in the same *orthoclass* if either they differ by a covering transformation preserving B_∞ or there is a covering transformation g such that $g(B_1) = B_\infty$ and $g(B_\infty) = B_2$.

Thus, we see that when two horoballs are in the same orthoclass, they have the same orthodistance to B_∞ and the same Euclidean diameter.

We order the orthodistances of the orthoclasses into increasing order: $0 = o(1) \leq o(2) \dots$. We let $e_n = e^{o(n)/2}$. Thus, the Euclidean diameter of the horoballs in this orthoclass is e_n^{-2}.

14. Improving the 6-theorem

A particularly important quantity is e_2. This corresponds to the second orthoclass. Thus, e_2^{-2} is the diameter of the second-largest horoballs, or more specifically, the largest horoballs that are not the Adams horoballs. We have the following extension of the 6-theorem [18].

Theorem 14.1 (L-Meyerhoff 2015)**.** *Let M be a compact orientable 3-manifold, with boundary a torus and with interior admitting a complete finite-volume hyperbolic structure. Let s be a slope on ∂M with length at least*

$$\frac{\pi e_2}{\arcsin(e_2/2)}$$

if $e_2 \leq \sqrt{2}$, and length at least

$$\frac{2\pi e_2}{2\arcsin(\sqrt{1-e_2^{-2}}) + e_2^2 - 2\sqrt{e_2^2-1}}$$

if $e_2 > \sqrt{2}$. Then, $M(s)$ is hyperbolic.

This improves Theorem 10.1 in the following sense. We denote by $L(e_2)$ the critical slope length that is given in the above result. A graph of $L(e_2)$ is shown in Figure 15. When $e_2 = 1$, Theorem 14.1 gives the same critical slope length as the 6-theorem. But as e_2 increases, the critical slope length decreases, tending to zero.

The idea behind the proof is as follows. When $e_2 = 1$, it gives no improvement over the 6-theorem. But for larger values of e_2, we have the following two effects:

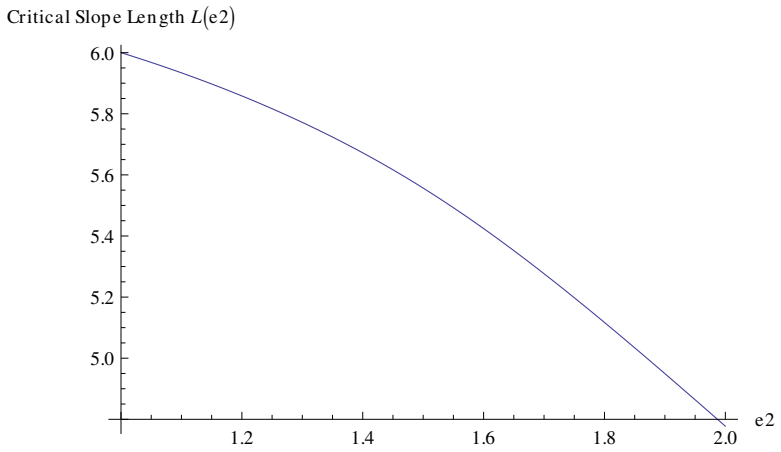

FIGURE 15.

(1) the Euclidean distance between the centres of the Adams horoballs is at least $e_2 > 1$;
(2) apart from the two Adams horoballs and their translates under the group of covering transformations that preserve B_∞, all other horoballs in the inverse image of N have Euclidean diameter at most $e_2^{-2} < 1$.

These two observations allow us to get better estimates than in the proof of the 6-theorem. Specifically, in that proof, we considered a planar surface F properly embedded in M, with boundary curves having slope s. We homotoped F to a pleated surface $\overline{F}$. We considered the spine S and the submanifold $M - S$ which is topologically just $T^2 \times [1, \infty)$. This is a union of the surfaces E_t, each of which is the torus T_t minus some (not necessarily disjoint) discs. The key part of the proof was to find a lower bound on the area of each component X of $\overline{F} - S$ that runs up to infinity in the cusp. The point is that if we could show that the area of each such component X is more than 2π, then we would obtain the desired contradiction to the existence of F. Using the co-area formula, we obtain that this area is least

$$\text{Area}(X) \geq \int \frac{\text{Length}(X \cap E_t)}{t}\, dt.$$

In our setting, we can find better lower bounds on the length of $X \cap E_t$ than in the proof of the 6-theorem. We obtain that

$$\frac{\text{Length}(X \cap E_t)}{L(s)/t} \geq \begin{cases} 1 & \text{if } t \geq 1; \\ \max\left\{0, 1 - \frac{2}{e_2}\sqrt{1-t^2}\right\} & \text{if } \min\{1/\sqrt{2}, 1/e_2\} \leq t \leq 1; \\ 0 & \text{otherwise.} \end{cases}$$

Geometrically, this is the ratio between the length of $X \cap E_t$ and the slope length in the case shown in Figure 16, at least when $e_2 \leq \sqrt{2}$. This arrangement in Figure 16 is where two full-sized horoballs are exactly e_2 apart (which is as close as they can be).

Applying the co-area formula gives a lower bound of 2π on the area of X, in the situation where the length of the slope s is at least the quantity in the theorem. This proves Theorem 14.1.

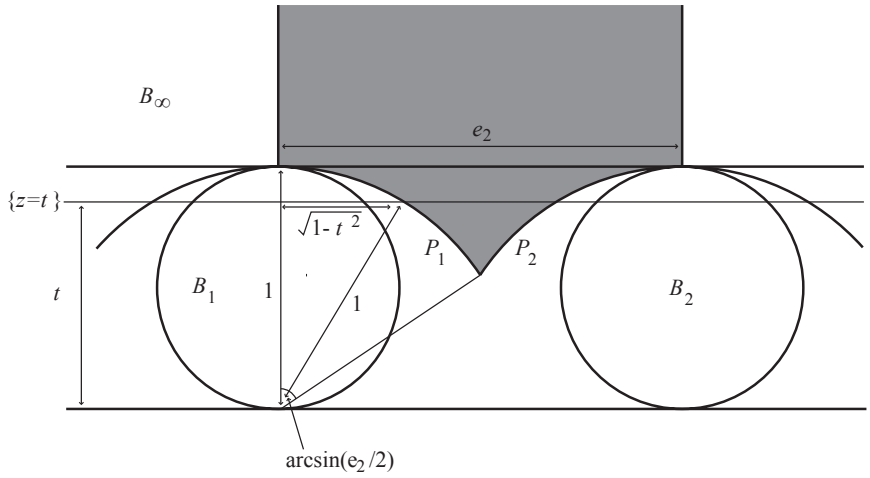

FIGURE 16. Two full-sized horoballs B_1 and B_2 that are close as they can be. Equidistant between B_1 and B_∞ is a totally geodesic plane P_1, and P_2 is defined similarly. The shaded surface above P_1 and P_2 has area $2\arcsin(e_2/2)$ and contributes e_2 to slope length.

What good does this do us? Well, we already know that when $e_2 = 1$, the second-largest horoballs have diameter 1 and so the area of ∂N is at least $2\sqrt{3}$. This constitutes an improvement to Theorem 9.3. On other hand, when $e_2 > 1$, then we get an improvement to the 6-theorem.

But to fully exploit these observations, we need to use the following theory developed by Gabai, Meyerhoff and Milley [13].

15. Triples and Mom manifolds

Definition 15.1. A (p, q, r)-*triple* is a triple of horoballs $\{B_1, B_2, B_3\}$ in the inverse image of N such that

(1) B_1 and B_2 lie in the rth orthoclass;
(2) B_1 and B_3 lie in the qth orthoclass;
(3) B_2 and B_3 lie in the pth orthoclass.

One useful way to think of triples is when considering the Epstein-Penner decomposition. Suppose that a face of this is an ideal triangle. Then the three geodesics in its boundary give rise to three orthoclasses, $o(p)$, $o(r)$ and $o(q)$, say and hence a (p, q, r)-triple. The key idea of Gabai, Meyerhoff and Milley is that when there are enough different triples, then this actually forces the Epstein-Penner decomposition to be of a certain form, enough to be able to determine a lot of information about the topology of M.

A fairly elementary but important fact about triples is as follows.

Lemma 15.2. *There are no (n, n, n)-triples for any positive integer n.*

We have already seen this in the case $n = 1$. This is because a $(1, 1, 1)$-triple would give rise to three horoballs, one of which is B_∞ say, and the other two are full-sized horoballs B_1 and B_2. The three geodesics joining the points at infinity

of these horoballs would then project to the same geodesic in M. But we already mentioned in Section 12 that this could not be the case.

Definition 15.3. A *geometric Mom-n structure* is a collection of n triples

$$(p_1, q_1, r_1), \ldots, (p_n, q_n, r_n),$$

no two of which are equivalent under the action of the action of $\pi_1(M)$ and such that $p_1, q_1, r_1, \ldots, p_n, q_n, r_n$ all lie in the same n-element subset of $\mathbb{Z}$.

The rationale for this definition is as follows. We saw above that the edges of the Epstein-Penner decomposition determine orthoclasses, and when the faces of the Epstein-Penner decomposition are triangles, then these faces correspond to triples. Conversely, each orthoclass determines a geodesic running from the cusp to the cusp. Each triple determines a mapped-in triangle attached to them. Gabai, Meyerhoff and Milley show that, under certain geometric hypotheses, these geodesics and triangles can actually be used to form a cell complex embedded in M. More specifically, one starts with the horoball neighbourhood N, one attaches to it 1-cells that run along the geodesics in the orthoclasses and 2-cells corresponding to the triples. Since in a geometric Mom-n structure, there are as many triples as orthoclasses, this space has Euler characteristic zero. Hence, when we thicken it up, it forms a 3-manifold M' with zero Euler characteristic embedded within M. As is true of any compact 3-manifold, $\chi(\partial M') = 2\chi(M')$. Hence, as long as we know that M' has no 2-sphere boundary components, we deduce that $\partial M'$ is a collection of tori. Gabai-Meyerhoff-Milley show that, again under certain geometric hypotheses, each component of $\partial M'$ either bounds a solid torus disjoint from M' or is parallel to a component of ∂M. In other words, M is obtained by Dehn filling M'. But M' is built in a very specific way, and so they can prove the following result.

Theorem 15.4 (Gabai-Meyerhoff-Milley 2009)**.** *Let M be a finite-volume orientable hyperbolic 3-manifold. If M contains a geometric Mom-2 involving the $o(1)$ and $o(n)$ orthoclasses, then either $e_n \geq 1.5152$ or M is obtained by Dehn filling one of the manifolds m125, m129 or m203.*

Here, the manifold notation is from the census [6]. They also prove a result about Mom-3 manifolds. To state it, we need the following terminology. A geometric Mom-3 is *torus-friendly* if it does not possess exactly two triples of type (p, q, r) for any set of distinct positive indices p, q and r.

Theorem 15.5 (Gabai-Meyerhoff-Milley 2009)**.** *Let M be a finite-volume orientable hyperbolic 3-manifold. If M contains a torus-friendly geometric Mom-3 involving the orthoclasses $o(1)$, $o(2)$ and $o(3)$, then either $e_3 \geq 1.5152$ or M is obtained by Dehn filling one of the manifolds m412, s596, s647, s774, s776, s780, s785, s898 or s959.*

16. The lower bound on cusp area due to Gabai-Meyerhoff-Milley

The lower bound on cusp torus area in Theorem 9.4 was proved using Mom technology, together with a computer assisted proof. In fact, we can now describe the families of manifolds mentioned in the theorem's statement. They are just the manifolds obtained by Dehn filling one of m412, s596, s647, s774, s776, s780, s785, s898 or s959. So, if we assume that M is not of this form, then we deduce that either $e_3 \geq 1.5152$ or M does not contain a torus-friendly geometric Mom-3.

Corresponding to the first three orthoclasses, there are three geodesics in M, that intersect the cusp torus ∂N at 6 points. The idea is to place discs in the cusp torus of certain sizes at these 6 points. The exact sizes of these discs are a little complicated. For example, at each of the two points of intersection between ∂N and the $o(2)$ geodesic, a disc of radius $(e_4/e_2) - (e_4/2)$ is placed. These sizes are chosen so that if there is an overlap between two of these discs, then we deduce the existence of a triple involving some of the first three orthoclasses. Thus, assuming that there is no geometric Mom-2 or torus-friendly geometric Mom-3, then the set of overlaps between these discs is restricted. Hence, one can get an explicit lower bound on the area of ∂N, in terms of e_2, e_3 and e_4, by adding up the area of these discs and then subtracting the area of possible overlaps.

The computer-assisted part of the proof is an analysis of the set of possible values of e_2, e_3 and e_4. The first step is to reduce to a bounded set of values. By definition, e_2, e_3 and e_4 are all at least 1. It is also not too difficult to show that if they are large then the area of ∂N is large. For example, we may place discs of radius $e_2/2$ about the two $o(1)$ points in ∂N. These discs cannot overlap and so we deduce that the area of ∂N is at least $\sqrt{3}(e_2)^2$, by Thue's theorem. So if $e_2 \geq 1.462$, then the area of ∂N is more than 3.7, as required.

Once we have restricted to a bounded set of possible values of e_2, e_3 and e_4 to consider, we divide this set up into little boxes. For each box, we have an explicit lower bound on the area of ∂N. One can think each box as specifying values of e_2, e_3 and e_4, but with small error bars. The lower bound on the area of ∂N is a function of e_2, e_3 and e_4 and so one gets a lower bound on the area, with explicit error, and hence an absolute lower bound. Since a computer is being used, the errors arising from floating point arithmetic must also be taken into account. But they can too can be bounded explicitly.

By running over the parameter space of possible values of e_2, e_3 and e_4, an explicit lower bound of 3.7 for the area of the maximal cusp torus is established.

17. The maximal number of exceptional Dehn surgeries

The proof of Theorems 4.5 and 5.2 builds on the above analysis. One of the key points is that when e_2 is close to 1, then the area lower bounds established by Gabai, Meyerhoff and Milley are good enough to prove Theorems 4.5 and 5.2. As e_2 increases, the area lower bounds due to Gabai, Meyerhoff and Milley become less good, but then the improvement to the 6-theorem given by Theorem 14.1 becomes useful.

In fact, the parameter space that is used in [18] is 6-dimensional. Three of the parameters are e_2, e_3 and e_4, as above. The other three are used to specify the size and shape of the cusp torus. The advantage of this approach is that, having an explicit cusp shape, we can compute directly the number of slopes less than the critical slope length provided by Theorem 14.1.

For each point in the parameter space, we obtain a lower bound on the cusp area. If this is definitely more than the area as specified by the cusp shape parameters, then we know that this cannot correspond to a manifold. We can also exclude some other points in the parameter space, using various tests. For example, one of these tests determines whether there is a possible location for the two Adams horoballs that does not force the existence of a geometric Mom-2.

Thus, the computer-assisted part of the proof is somewhat similar to that used by Gabai, Meyerhoff and Milley. One firsts bounds the parameters to a compact set. Then one divides this set into little boxes. For each box, we determine whether it can be excluded by one of the above tests. If a box cannot be excluded, then the set of slopes with length at most the bound in Theorem 14.1 is computed. For each box, the computer finds that the maximal number of slopes is 10 and the maximal distance between any two such slopes is 8, thereby nearly completing the proof of Theorems 4.5 and 5.2.

All that remains is to consider the manifolds M that are obtained by Dehn filling one of m412, s596, s647, s774, s776, s780, s785, s898 or s959. For any given manifold M, it is not hard to determine its exceptional Dehn fillings. The filling slopes with length more than 2π are not exceptional, and so all that one needs to do is to enumerate all the slopes with length at most 2π, perform the Dehn filling and determine whether the resulting manifold is hyperbolic. This is not particularly hard to do in practice. Of course, there are actually infinitely many 3-manifolds M that are obtained by Dehn filling one of m412, s596, s647, s774, s776, s780, s785, s898 or s959. But it is possible to determine the exceptional fillings on any such manifold, essentially by filling one boundary torus of the relevant census manifold at a time. We refer the reader to [18] for more details.

18. Exercises

(1) If you have not already done so, download Snappy [9]. Familiarise yourself with it, so that you can form a picture of the cusp neighbourhood of the knots 4_1 and 6_1. This gives a view of upper half-space from infinity. The inverse image of the cusp is a collection of horoballs, including $\{(x, y, z) : z \geq 1\}$. A horoball is *full-sized* if its Euclidean diameter is 1. Up to the action of the fundamental group of the cusp, how many full-sized horoballs does 4_1 have? What about 6_1?

You will not need Snappy to answer the remaining questions, but it is a very useful tool.

(2) What is the area of the maximal cusp torus for the figure-eight knot complement?

(3) Determine the length of the second shortest slope on the figure-eight knot exterior.

(4) Let p/q and p'/q' be slopes on the boundary of the figure-eight knot exterior. Show that there is a homeomorphism of the figure-eight knot exterior taking p/q to p'/q' if and only if $p'/q' = \pm p/q$. (A version of Mostow-Prasad rigidity implies that any homeomorphism of a finite-volume hyperbolic 3-manifold to itself is homotopic to an isometry. You may assume this.)

(5) (a) Let s be a Euclidean geodesic on the boundary of a cusp $N = T^2 \times [1, \infty)$. Let $s \times [1, \infty)$ be the vertical annulus over it. Compute the area of $s \times [1, \infty)$ in terms of the length of s.
(b) What is the relationship between the volume of N and the area of ∂N?

(6) Let $\mathcal{S}$ be a collection of distinct slopes on the torus.
(a) Show that if $\Delta(s, s') \leq 1$ for each $s, s' \in \mathcal{S}$, then $|\mathcal{S}| \leq 3$.
(b) Show that if $\Delta(s, s') \leq 2$ for each $s, s' \in \mathcal{S}$, then $|\mathcal{S}| \leq 4$.

(7) Show that for any finite-volume orientable hyperbolic 3-manifold M, there is a choice of horoball neighbourhood of ∂M where all the cusps have the

same volume. Show that any two such horoball neighbourhoods determine the same geodesic spine and hence the same Epstein-Penner decomposition.

(8) (*Harder*) Complete the proof of Theorem 3.4 by showing that $M_{[\epsilon,\infty)}$ is atoroidal, where M is a finite-volume orientable hyperbolic 3-manifold and ϵ is a Margulis constant.

(9) (*Harder*) Let M be the exterior of the figure-eight knot. Show that provided p/q avoids a finite set of slopes, $M(p/q)$ is homeomorphic to $M(p'/q')$ if and only if $p'/q' = \pm p/q$.

References

1. Colin C. Adams, *The noncompact hyperbolic* 3*-manifold of minimal volume*, Proc. Amer. Math. Soc. **100** (1987), no. 4, 601–606. [14, 20]
2. Ian Agol, *Bounds on exceptional Dehn filling*, Geom. Topol. **4** (2000), 431–449. [8, 15]
3. ———, *Bounds on exceptional Dehn filling II*, Geom. Topol. **14** (2010), no. 4, 1921–1940. [7, 9]
4. Riccardo Benedetti and Carlo Petronio, *Lectures on hyperbolic geometry*, Universitext, Springer-Verlag, Berlin, 1992. [1, 2, 3]
5. Steven A. Bleiler and Craig D. Hodgson, *Spherical space forms and Dehn filling*, Topology **35** (1996), no. 3, 809–833. [10, 14]
6. Patrick J. Callahan, Martin V. Hildebrand, and Jeffrey R. Weeks, *A census of cusped hyperbolic* 3*-manifolds*, Math. Comp. **68** (1999), no. 225, 321–332, With microfiche supplement. [25]
7. Chun Cao and G. Robert Meyerhoff, *The orientable cusped hyperbolic* 3*-manifolds of minimum volume*, Invent. Math. **146** (2001), no. 3, 451–478. [14, 21]
8. Daryl Cooper and Marc Lackenby, *Dehn surgery and negatively curved* 3*-manifolds*, J. Differential Geom. **50** (1998), no. 3, 591–624. [11, 12]
9. Marc Culler, Nathan M. Dunfield, Matthias Goerner, and Jeffrey R. Weeks, *SnapPy, a computer program for studying the geometry and topology of* 3*-manifolds*, Available at `http://snappy.computop.org`. [27]
10. D. B. A. Epstein, *Periodic flows on three-manifolds*, Ann. of Math. (2) **95** (1972), 66–82. [3]
11. D. B. A. Epstein and R. C. Penner, *Euclidean decompositions of noncompact hyperbolic manifolds*, J. Differential Geom. **27** (1988), no. 1, 67–80. [20]
12. David Futer, Efstratia Kalfagianni, and Jessica S. Purcell, *Dehn filling, volume, and the Jones polynomial*, J. Differential Geom. **78** (2008), no. 3, 429–464. [11]
13. David Gabai, Robert Meyerhoff, and Peter Milley, *Minimum volume cusped hyperbolic three-manifolds*, J. Amer. Math. Soc. **22** (2009), no. 4, 1157–1215. [14, 24]
14. C. McA. Gordon, *Dehn filling: a survey*, Knot theory (Warsaw, 1995), Banach Center Publ., vol. 42, Polish Acad. Sci. Inst. Math., Warsaw, 1998, pp. 129–144. [7, 8]
15. D. A. Každan and G. A. Margulis, *A proof of Selberg's hypothesis*, Mat. Sb. (N.S.) **75 (117)** (1968), 163–168. [1, 6]
16. Tsuyoshi Kobayashi and Yo'av Rieck, *A linear bound on the tetrahedral number of manifolds of bounded volume (after Jørgensen and Thurston)*, Topology and geometry in dimension three, Contemp. Math., vol. 560, Amer. Math. Soc., Providence, RI, 2011, pp. 27–42. [6]
17. Marc Lackenby, *Word hyperbolic Dehn surgery*, Invent. Math. **140** (2000), no. 2, 243–282. [15]
18. Marc Lackenby and Robert Meyerhoff, *The maximal number of exceptional Dehn surgeries*, Invent. Math. **191** (2013), no. 2, 341–382. [7, 8, 22, 26, 27]
19. Robert Meyerhoff, *A lower bound for the volume of hyperbolic* 3*-manifolds*, Canad. J. Math. **39** (1987), no. 5, 1038–1056. [10]
20. G. D. Mostow, *Quasi-conformal mappings in n-space and the rigidity of hyperbolic space forms*, Inst. Hautes Études Sci. Publ. Math. (1968), no. 34, 53–104. [3]
21. Grigori Perelman, *The entropy formula for the Ricci flow and its geometric applications*, arxiv:math.DG/0211159, 2002. [2]
22. ———, *Finite extinction time for the solutions to the Ricci flow on certain three-manifolds*, arxiv:math.DG/0307245, 2003. [2]
23. ———, *Ricci flow with surgery on three-manifolds*, arxiv:math.DG/0303109, 2003. [2]
24. Gopal Prasad, *Strong rigidity of* $\mathbf{Q}$*-rank* 1 *lattices*, Invent. Math. **21** (1973), 255–286. [3]

25. John G. Ratcliffe, *Foundations of hyperbolic manifolds*, second ed., Graduate Texts in Mathematics, vol. 149, Springer, New York, 2006. [1, 2]
26. Herbert Seifert and William Threlfall, *Seifert and Threlfall: a textbook of topology*, Pure and Applied Mathematics, vol. 89, Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1980, Translated from the German edition of 1934 by Michael A. Goldman, With a preface by Joan S. Birman, With "Topology of 3-dimensional fibered spaces" by Seifert, Translated from the German by Wolfgang Heil. [3]
27. William P. Thurston, *The geometry and topology of three-manifolds*, Princeton Univ. Math. Dept. Notes, 1979. [1, 4, 6, 7, 11]
28. ———, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 3, 357–381. [1, 2, 3, 7]

Mathematical Institute, University of Oxford,
Woodstock Road, Oxford OX2 6GG, United Kingdom