

Commensurators of Cusped Hyperbolic Manifolds

Oliver Goodman, Damian Heard, and Craig Hodgson

CONTENTS

1. Introduction
 2. The Commensurability Criterion
 3. Example: Punctured Torus Bundles
 4. Algorithm for Finding Isometries of Tilings
 5. Enumerating Canonical Cell Decompositions I
 6. Enumerating Canonical Cell Decompositions II
 7. Commensurability of Cusps
 8. Experimental Results
 9. Appendix
- Acknowledgments
References

This paper describes a general algorithm for finding the commensurator of a nonarithmetic hyperbolic manifold with cusps and for deciding when two such manifolds are commensurable. The method is based on some elementary observations regarding horosphere packings and canonical cell decompositions. For example, we use this to find the commensurators of all nonarithmetic hyperbolic once-punctured torus bundles over the circle.

For hyperbolic 3-manifolds, the algorithm has been implemented using Goodman's computer program Snap. We use this to determine the commensurability classes of all cusped hyperbolic 3-manifolds triangulated using at most seven ideal tetrahedra, and for the complements of hyperbolic knots and links with up to twelve crossings.

1. INTRODUCTION

Two manifolds or orbifolds M and M' are *commensurable* if they admit a common finite sheeted covering. For hyperbolic n -orbifolds, we can suppose that $M = \mathbb{H}^n/\Gamma$ and $M' = \mathbb{H}^n/\Gamma'$, with Γ and Γ' discrete subgroups of $\text{Isom}(\mathbb{H}^n)$. In this paper, we assume that M and M' are of finite volume and of dimension at least 3. Then, by Mostow–Prasad rigidity, commensurability means that we can conjugate Γ by an isometry g such that $g\Gamma g^{-1}$ and Γ' intersect in a subgroup of finite index in both groups.

Given that the classification of finite-volume hyperbolic manifolds up to homeomorphism appears to be hard, it seems sensible to attempt to subdivide the problem and start with a classification up to commensurability. Looked at in this way, we see a remarkable dichotomy between the *arithmetic* and *nonarithmetic* cases. (See [Maclachlan and Reid 03] for the definition of arithmetic hyperbolic manifolds.)

Define the *commensurator* of Γ to be the group

$$\text{Comm}(\Gamma) = \{g \in \text{Isom}(\mathbb{H}^n) \mid [\Gamma : \Gamma \cap g\Gamma g^{-1}] < \infty\}.$$

Then Γ and Γ' are commensurable if and only if $\text{Comm}(\Gamma)$ and $\text{Comm}(\Gamma')$ are conjugate. Geometrically,

2000 AMS Subject Classification: Primary 57M50;
Secondary 57N10, 57M25

Keywords: Hyperbolic manifolds, commensurator, canonical cell decomposition, hyperbolic links, hyperbolic 3-manifolds

an element of the normalizer of Γ in $\text{Isom}(\mathbb{H}^n)$ represents a *symmetry* (i.e., isometry) of $M = \mathbb{H}^n/\Gamma$. Similarly, an element of the commensurator represents an isometry between finite sheeted covers of M ; this gives a *hidden symmetry* of M if it is not the lift of an isometry of M (see [Neumann and Reid 92a]).

It follows from deep work of Margulis [Margulis 91] (see also [Zimmer 84]) that in dimension ≥ 3 , the commensurator $\text{Comm}(\Gamma)$ is discrete if and only if Γ is not arithmetic. This means that the commensurability class of a nonarithmetic, cofinite-volume, discrete group Γ is particularly simple, consisting only of conjugates of the finite-index subgroups of $\text{Comm}(\Gamma)$. In terms of orbifolds, it means that M and M' are commensurable if and only if they cover a common quotient orbifold.

On the other hand, commensurability classes of arithmetic groups are “big”: we may well have commensurable Γ and Γ' such that the group generated by $g\Gamma g^{-1}$ and Γ' is not discrete for any g .

A hyperbolic n -orbifold is *cusped* if it is noncompact of finite volume. This paper describes a practical algorithm for determining when two cusped hyperbolic n -manifolds cover a common quotient, and for finding a smallest quotient. For nonarithmetic finite-volume *cusped* hyperbolic n -manifolds of dimension $n \geq 3$, this solves the commensurability problem.

Section 2 begins with some elementary observations about horoball packings and canonical cell decompositions of a cusped hyperbolic manifold. This leads to a characterization of the commensurator of a nonarithmetic cusped hyperbolic manifold M as the maximal symmetry group of the tilings of \mathbb{H}^n obtained by lifting canonical cell decompositions of M . In Section 3, we use this to determine the commensurators of nonarithmetic hyperbolic once-punctured torus bundles over the circle.

Section 4 gives an algorithm for finding the isometry group of a tiling of \mathbb{H}^n arising from a cell decomposition of a hyperbolic manifold, and Sections 5 and 6 describe methods for finding all possible canonical cell decompositions for a cusped hyperbolic manifold. Section 7 contains some observations on commensurability of cusps in hyperbolic 3-manifolds that can simplify the search for all canonical cell decompositions.

In three dimensions, each orientable hyperbolic orbifold has the form $M = \mathbb{H}^3/\Gamma$, where Γ is a discrete subgroup of $\text{PSL}(2, \mathbb{C}) = \text{Isom}^+(\mathbb{H}^3)$. The *invariant trace field* $k(\Gamma) \subset \mathbb{C}$ is the field generated by the traces of the elements of $\Gamma^{(2)} = \{\gamma^2 \mid \gamma \in \Gamma\}$ lifted to $\text{SL}(2, \mathbb{C})$. This is a number field if M has finite volume (see [Neumann and Reid 92a, Reid 90, Macbeath 88]). The *invariant quater-*

nion algebra is the $k(\Gamma)$ subalgebra of $M_2(\mathbb{C})$ generated by $\Gamma^{(2)}$. These are useful and computable commensurability invariants (see [Coulson et al. 00, Maclachlan and Reid 03]).

For the *arithmetic* subgroups of $\text{Isom}(\mathbb{H}^3)$, the invariant quaternion algebra is a complete commensurability invariant. In fact, for *cusped* arithmetic hyperbolic 3-orbifolds, the invariant trace field is an imaginary quadratic field, and the quaternion algebra is just the algebra of all 2×2 matrices with entries in the invariant trace field (see [Maclachlan and Reid 03, Theorem 3.3.8]); so the invariant trace field is a complete commensurability invariant. However, most cusped hyperbolic 3-manifolds are nonarithmetic (cf. [Borel 81]), so other methods are needed to determine commensurability.

Damian Heard and Oliver Goodman have implemented the algorithms described in this paper for nonarithmetic hyperbolic 3-manifolds; these are incorporated in the computer program Snap.¹ Using this, we have determined the commensurability classes for all manifolds occurring in the Callahan–Hildebrand–Weeks census [Hildebrand and Weeks 89, Callahan et al. 99] of cusped hyperbolic manifolds with up to seven tetrahedra, and for complements of hyperbolic knots and links up to twelve crossings, supplied by Morwen Thistlethwaite [Hoste and Thistlethwaite 99]. These results are discussed in Section 8, while Section 9 outlines the Dowker–Thistlethwaite notation used to describe links.

This work has uncovered interesting new examples of commensurable knot and link complements (see Examples 2.1 and 2.2), and a new example of a knot with shape field properly contained in the invariant trace field (see Example 7.1). The results have also been used by Button [Button 05] to study fibered and virtually fibered cusped hyperbolic 3-manifolds.

For 1-cusped manifolds we note that “cusp density” (see Section 2) is a very good invariant. We have found only a few examples of incommensurable 1-cusped manifolds that are not distinguished by cusp density (see Example 2.3).

There is also a “dumb” algorithm, based on volume bounds for hyperbolic orbifolds, which works for any (possibly closed) nonarithmetic hyperbolic 3-orbifold, but appears to be quite impractical. If M and M' cover Q with $\text{Vol}(Q) > C$, then the degrees d, d' of the coverings are bounded by $D = \lfloor \text{Vol}(M)/C \rfloor$ and $D' = \lfloor \text{Vol}(M')/C \rfloor$, respectively. Then if M and M' are commensurable, they admit a common covering N of

¹Available online at <http://www.ms.unimelb.edu.au/~snap/>.

degree at most D' over M and at most D over M' . The best current estimate for C for orientable nonarithmetic 3-orbifolds is $0.041\dots$ from recent work of Marshall and Martin [Marshall and Martin 08]. Since $\text{Vol}(M) \approx 2$ is typical, we would have to find all coverings of M' of degree $d' \leq 50$. This means finding all conjugacy classes of transitive representations of $\pi_1(M')$ into S_{50} , a group with about 10^{64} elements!

2. THE COMMENSURABILITY CRITERION

We use the following terminology throughout this paper. A set of disjoint horoballs in \mathbb{H}^n is called a *horoball packing*, and a *cuspidal neighborhood* in a hyperbolic n -orbifold is one that lifts to such a horoball packing.

Lemma 2.1. *The symmetry group of a horoball packing in \mathbb{H}^n is discrete whenever the totally geodesic subspace spanned by their ideal points has dimension at least $n - 1$.*

Proof: Let $\{g_i\}$ be a sequence of symmetries of the packing converging to the identity. Choose horoballs B_1, \dots, B_n whose ideal points span a totally geodesic subspace H of dimension $n - 1$. For i sufficiently large, we can assume that $g_i(B_k) = B_k$ for $k = 1, \dots, n$. But this implies that these g_i fix H pointwise. Since the only such isometries are the identity and reflection in H , the sequence must be eventually constant. \square

Lemma 2.2. *Let $M = \mathbb{H}^n/\Gamma$ be a finite-volume cusped hyperbolic orbifold. The set of parabolic fixed points of Γ spans \mathbb{H}^n .*

Proof: The set of parabolic fixed points is dense in the limit set of Γ , which equals the whole of the sphere at infinity. \square

Lemma 2.3. *Let M, M' be finite-volume cusped hyperbolic orbifolds. Then M and M' cover a common orbifold Q if and only if they admit choices of cuspidal neighborhoods lifting to isometric horoball packings.*

Proof: If M and M' cover Q , choose cuspidal neighborhoods in Q and lift to M and M' . These all lift to the same horoball packing in \mathbb{H}^n , namely the horoball packing determined by our choice of cuspidal neighborhoods in Q . Conversely, both M and M' cover the quotient of \mathbb{H}^n by the group of symmetries of their common horoball packing, which, by Lemmas 2.1 and 2.2, is discrete. \square

We can define the cusp density of a 1-cusped hyperbolic orbifold M as follows. Since M has only one cusp, it has a unique maximal (embedded) cuspidal neighborhood U . The *cusp density* of M is $\text{Vol}(U)/\text{Vol}(M)$. Since the cusp densities of all orbifolds covered by M are the same, it is a commensurability invariant of orbifolds with discrete commensurator.

Choosing a full set of disjoint cuspidal neighborhoods in a noncompact finite-volume hyperbolic n -manifold M determines a “Ford spine.” This is the cell complex given by the set of points in M equidistant from the cuspidal neighborhoods in two or more directions. Cells of dimension $n - k$ contain points equidistant from the cuspidal neighborhoods in $k + 1$ independent directions ($k = 1, \dots, n$). This spine can also be seen intuitively as the “bumping locus” of the cuspidal neighborhoods: blow up the cuspidal neighborhoods until they press against each other and flatten.

Dual to the Ford spine is a decomposition of M into ideal polytopes, generically simplices. The ideal cell dual to a given 0-cell of the Ford spine lifts to the convex hull in \mathbb{H}^n of the set of ideal points determined by the equidistant directions. We call the cell decompositions that arise in this way *canonical*.² For a 1-cusped manifold, the canonical cell decomposition is unique. It is shown in [Akiyoshi 01] that a finite-volume hyperbolic manifold with multiple cusps admits finitely many canonical cell decompositions.

Theorem 2.4. *Cusped hyperbolic n -manifolds M and M' cover a common orbifold if and only if they admit canonical ideal cell decompositions lifting to isometric tilings of \mathbb{H}^n .*

Proof: If M and M' cover Q , choose cuspidal neighborhoods in Q and lift them to M, M' , and \mathbb{H}^n . Constructing the Ford spine and cell decomposition in \mathbb{H}^n clearly yields the lifts of those entities from both M and M' corresponding to our choice of cuspidal neighborhoods.

Conversely, observe that the symmetry group of the common tiling gives an orbifold that is a quotient of both manifolds. \square

Remark 2.5. We can omit the word “canonical” in the above theorem. The proof is unchanged.

The previous theorem gives the following characterization of the commensurator.

²The term is not really ideal, since for manifolds with multiple cusps, there are generally multiple canonical cell decompositions depending on the choice of cuspidal neighborhoods.

Theorem 2.6. *Let $M = \mathbb{H}^n/\Gamma$ be a finite-volume cusped hyperbolic n -manifold with discrete commensurator. Then $\text{Comm}(\Gamma)$ is the maximal symmetry group of the tilings of \mathbb{H}^n obtained by lifting canonical cell decompositions of M ; it contains all such symmetry groups.*

For manifolds with discrete commensurator, we can now define a *truly canonical* ideal cell decomposition as follows. Find $\text{Comm}(\Gamma)$ as in the above theorem. Choose *equal-volume* cusp neighborhoods in $\mathbb{H}^n/\text{Comm}(\Gamma)$. Lift them to M and take the resulting canonical cell decomposition of M . Two such manifolds are commensurable if and only if their truly canonical cell decompositions give isometric tilings of \mathbb{H}^n .

Choosing maximal cusp neighborhoods in $\mathbb{H}^n/\text{Comm}(\Gamma)$ also gives a canonical version of *cuspidensity* for multicusped manifolds.

Theorem 2.6 is the basis for the algorithms described in this paper. Canonical cell decompositions can be computed by the algorithms of Weeks described in [Weeks 93] and implemented in SnapPea.³ In Section 4 below we give an algorithm for finding the isometry groups of the corresponding tilings of \mathbb{H}^n . Combining this with Theorem 2.6 gives an algorithm for finding commensurators of 1-cusped nonarithmetic hyperbolic n -manifolds. In Sections 5 and 6 we extend this algorithm to *multi-cusped* manifolds by describing methods for finding *all* canonical cell decompositions.

For hyperbolic 3-manifolds, these algorithms have been implemented by Heard and Goodman. (These are incorporated in `find commensurator` and related commands in the program Snap.) We conclude this section with some examples discovered during this work.

2.1 Example: A 5-Link Chain and Friends

The five links shown in Figure 1 have commensurable complements, as shown using our computer program. In the first three cases at least it is possible to “see” this commensurability.

The first of these is the 5-link chain C_5 . Thurston [Thurston 78, Chapter 6] explains how to obtain a fundamental region for the hyperbolic k -link chain complements: We span each link of the chain by a disk in the obvious manner. The complement of the union of these five disks is then a solid torus. Once the link is deleted, the disks and their arcs of intersection divide the boundary of the solid torus into ideal squares A, B, C, D, E as shown in Figure 2 with cusps labeled a, b, c, d, e .

³Available online at <http://geometrygames.org/SnapPea/index.html>.




FIGURE 1. Five links with commensurable complements.




FIGURE 2. Cell decomposition on solid torus boundary.

A hyperbolic structure is given by taking two regular pentagonal drums with ideal vertices and adjusting their heights to obtain (ideal) square faces. Glue two drums together as shown in Figure 3, identify the top with the bottom via a $\frac{4\pi}{5}$ rotation, and glue faces as indicated. Edges are then identified in fours: two horizontal with two vertical. It is easy to check that the sum of dihedral angles around each edge is 2π , so this gives a hyperbolic structure, since the angle sum is π at each ideal vertex of a drum.




FIGURE 3. Gluings for the complement of link C_5 .




FIGURE 4. Gluings for the second link complement.

It is clear from the symmetry of the picture that the drums are cells in the canonical cell decomposition obtained by choosing equal-area cusp cross sections. We remark that Neumann–Reid show in [Neumann and Reid 92a, Section 5] that this link complement is nonarithmetic.

Now change the 5-link chain by cutting along the disk shown in Figure 1, applying a half-twist, and regluing to obtain our second link. In the complement, this surgery introduces a half-turn into the gluing between the *A*-faces. See Figure 4. Edges are still identified in fours, two horizontal with two vertical, but there are now only four cusps.

If we repeat the process on the second disk shown in Figure 1, we obtain the third link. Again this corresponds to changing the gluing pattern on our two drums.

Since these link complements are nonarithmetic, the tiling of \mathbb{H}^3 by pentagonal drums covers some canonical cell decomposition of each one. Since their volumes are the same, each one decomposes into two pentagonal drums. It should therefore be possible, in each case, to find five ideal squares meeting at order-4 edges, cutting the complement into one or two solid tori. We leave this as a challenge for the reader.

2.2 Example: Commensurable Knot Complements

Commensurable knot complements seem to be rather rare. Previously known examples include the Rubinstein–Aitchison dodecahedral knots [Aitchison and Rubinstein 92] and examples due to a construction of González-Acuña and Whitten [González-Acuña and Whitten 92] giving knot complements covering other knot complements. For example, the $-2, 3, 7$ pretzel knot has $18/1$ and $19/1$ surgeries giving lens spaces. Taking the universal covers of these lens spaces gives new hyperbolic




FIGURE 5. Projection of knot $9n6$ with $(k, m) = (3, 2)$.




FIGURE 6. Quotient of $9n6$ complement by the 2-fold symmetry.

knots in S^3 whose complements are 18- and 19-fold cyclic covers of the $(-2, 3, 7)$ -pretzel complement.

Our program finds a pair of knots $9n6$ and $12n642$, having nine and twelve crossings respectively, whose complements are commensurable with volumes in the ratio $3 : 4$. Walter Neumann has pointed out that these knots belong to a very pretty family of knots: take a band of k repeats of a trefoil with the ends given m half-twists before putting them together. For example, the knot with $(k, m) = (3, 2)$ is shown in Figure 5.

The half-twists are put in so as to undo some of the crossings of the trefoils (allowing the projection in Figure 5 to be rearranged so as to have nine crossings). The pair of knots found by our program correspond to $(k, m) = (3, 2)$ and $(4, 1)$.

To see that these knots have commensurable complements we find a common quotient orbifold. In each case this is the quotient of the knot complement by its symmetry group; these are dihedral groups of order 12 and order 16 respectively.

The picture of $9n6$ in Figure 5 shows an obvious axis of 2-fold symmetry; on the left of Figure 6 is the quotient, which is the complement of a knot in the orbifold S^3 with singular set an unknot labeled 2. By pulling the knot straight, we see (Figure 6, right) that this is an orbifold whose underlying space is a solid torus with knotted singular locus.




FIGURE 7. View of singular locus from inside solid torus.

If we arrange the singular locus on a torus parallel to the boundary of the solid torus, we see three clasps, three strands in the (vertical) core direction, and a strand with slope 2/1. The view from inside the solid torus looking toward the boundary is shown in Figure 7. (For the knot with $(k, m) = (4, 1)$ we would see four clasps, four strands in the core direction, and a strand with slope 1/1.)

The solid torus with its singular set has three 2-fold symmetries whose axes intersect the solid torus in six arcs, each passing perpendicularly through the core like a skewer, with symmetry group dihedral of order 6. The ends of the arcs are shown as stars in Figure 7.

The quotient orbifold is obtained by taking a slice of the solid torus between two axes and folding closed the top and bottom disks like books. The result is a ball with the axes giving two unknotted arcs of order 2 in the singular set, running out to the boundary (which is now a $(2, 2, 2, 2)$ -pillowcase orbifold). The original singular set gives an arc linking the other two, so that the whole singular locus is an “H” graph labeled with 2’s.

The three diagrams of Figure 8 show what happens to the singular locus in one slice of the solid torus as we fold. We begin with the annulus in the bottom sixth of the previous figure, redrawn after twisting the bottom. This bounds a solid cylinder with the singular locus as shown in the second diagram. Folding along the top and bottom (and expanding the region slightly) gives the final result.

We leave it to the reader to draw similar pictures for the knot with $(k, m) = (4, 1)$ and verify that the result is indeed the same orbifold. Alternatively, Orb⁴ or SnapPea can be used to verify that the appropriate dihedral covers of the final orbifold give the complements of the knots with $(k, m) = (3, 2)$ and $(4, 1)$.




FIGURE 8. Constructing the quotient orbifold for the complement of 9n6.

We remark that Walter Neumann has found an infinite family of new examples of pairs of commensurable knot complements in the 3-sphere; this example is the simplest case.

2.3 Example: Cusp Horoball Pictures

Figure 9 shows the horoball packings of two 1-cusped census manifolds m_{137} and m_{138} as seen from the cusp. Using Snap, we find that the commensurability classes of these two equal-volume manifolds are indistinguish-




FIGURE 9. The maximal horoball packing and canonical cell decomposition as seen from the cusps of manifolds m_{137} and m_{138} .

⁴Available online at <http://www.ms.unimelb.edu.au/~snap/orb.html>.

able by cusp density or invariant trace field. However, their maximal horoball packings and canonical cell decompositions are different. (For example, the edges joining degree-4 vertices in the cusp diagrams are all parallel for m_{137} , but not for m_{138} .)

3. EXAMPLE: PUNCTURED TORUS BUNDLES

The bundles over S^1 with a once-punctured torus as fiber provide an interesting family of 1-cusped hyperbolic 3-manifolds with known canonical triangulations. By analyzing how symmetries of the lifted triangulations of the universal cover appear when viewed from the cusp, we obtain (in all the nonarithmetic cases) strong constraints on what symmetries may be possible. We then show that they all come from symmetries of the manifold. This leads to the following theorem:

Theorem 3.1. *Let $M = \mathbb{H}^3/G$ be an orientable nonarithmetic hyperbolic 3-manifold that is a once-punctured torus bundle over S^1 . Then M has no “hidden symmetries,” i.e., the commensurator of G is the normalizer of G in $\text{Isom}(\mathbb{H}^3)$.*

Let F denote a once-punctured torus, and let $\varphi : F \rightarrow F$ be an orientation-preserving homeomorphism. Let

$$M = M_\varphi = F \times_\varphi S^1 = \frac{F \times [0, 1]}{(x, 0) \sim (\varphi(x), 1)}$$

be the mapping torus of φ . Identifying F with $(\mathbb{R}^2 - \mathbb{Z}^2)/\mathbb{Z}^2$, φ can be represented up to isotopy by an element of $\text{SL}(2, \mathbb{Z})$; since M depends only on the isotopy class of φ , we assume $\varphi \in \text{SL}(2, \mathbb{Z})$. Then M is hyperbolic whenever φ is hyperbolic, i.e., when φ has distinct real eigenvalues; M_φ and $M_{\varphi'}$ are homeomorphic if and only if φ and φ' are conjugate.

Define matrices

$$L = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \quad R = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

For each word w in the symbols L, R define $\varphi_w \in \text{SL}(2, \mathbb{Z})$ as the corresponding matrix product.

Lemma 3.2. *Each $\varphi \in \text{SL}(2, \mathbb{Z})$ is conjugate to $\pm\varphi_w$ for some word w in the symbols L, R . The sign is unique, and w is determined up to cyclic permutations of its letters.*

Let $M = M_\varphi$, where $\varphi = \pm\varphi_w$ is hyperbolic. The so-called *monodromy triangulation* T of M has one tetrahedron for each letter in w and gluings determined by w and




FIGURE 10. The cusp triangulation of $M = M_{\varphi_w}$, where $w = LRRLR$. The triangulation of M has five edges, labeled 0–4, which appear as both edges and vertices of the cusp triangulation.

the sign. It is nicely described in [Floyd and Hatcher 82] and [Guéritaud 06a]. It follows from work of Lackenby [Lackenby 83] that T is the canonical ideal cell decomposition of M . Other proofs of this result have recently been given by Guéritaud [Guéritaud 06a, Guéritaud 06b] and Akiyoshi, Sakuma, Wada, and Yamashita [Akiyoshi et al. 07].

The intersection of T with a (small) torus cross section of the cusp of M , lifted to its universal cover \mathbb{R}^2 , gives the (lifted) cusp triangulation T_0 of M . Note that edges and vertices of T_0 correspond to edges of T seen transversely or end-on respectively. SnapPea provides pictures of these cusp triangulations; see Figure 10 for an example.

We need two things: the first is a combinatorial description of T_0 in terms of w ; the second is an understanding of which edges and vertices of T_0 correspond to the same edges of T in M . Both are outlined briefly here; for detailed explanations we refer the reader to [Floyd and Hatcher 82, Appendix] and [Guéritaud 06a, Sections 3 and 4].

3.1 The Monodromy Triangulation T

The triangulation T is built up in layers by gluing tetrahedra according to the letters of w . We begin with an almost flat ideal tetrahedron projecting onto a punctured torus; the tetrahedron has edges a, b, c, c_- identified as shown in Figure 11.

For each successive letter L or R in the word w , we attach a tetrahedron to the top of the previous tetrahedron, as shown in Figure 12, in the cover $(\mathbb{R}^2 - \mathbb{Z}^2) \times \mathbb{R}$.

After using all the letters of w , the final triangulation of the fiber F differs from the initial triangulation by




FIGURE 11. Initial ideal tetrahedron, projecting to a once-punctured torus.




FIGURE 12. Attaching the next tetrahedron.

the monodromy φ , and we can glue the top and bottom together to obtain an ideal triangulation T of M .

3.2 Combinatorial Description of T_0

Now consider the induced triangulation of a cusp-linking torus (i.e., cusp cross section) in M . Each tetrahedron contributes a chain of four triangles going once around the cusp, as shown in Figure 13.




FIGURE 13. The chain of triangles around a cusp coming from one tetrahedron.




FIGURE 14. Gluing the next chain of triangles.




FIGURE 15. Horizontal strip of triangles for LRRLR.

In the triangulation T_0 of \mathbb{R}^2 this lifts to an infinite chain of triangles forming a (vertical) sawtooth pattern.

Each chain is glued to the next in one of two ways depending on whether the letter is an L or an R (Figure 14).

Before gluing, we adjust slightly the “front” triangles of the first chain and the “back” triangles of the second, so as to create two horizontal edges. After stacking these chains of triangles together we have a decomposition of T_0 into horizontal strips, which can be described combinatorially as follows.

Make a horizontal strip out of triangles corresponding to the letters of w : for each L add a triangle with a side on the bottom edge of the strip and a vertex at the top; vice versa for R ; repeat infinitely in both directions. For example, for $LRRLR$ we have the strip shown in Figure 15.

Fill the plane by reflecting repeatedly in the top and bottom edges of the strip. Then the cusp-linking torus is the quotient of the plane by the group G_0 generated by vertical translation by *four* strips and horizontal translation by one period of the strip (composed with an extra vertical translation by two strips if $\varphi = -\varphi_w$). (See Figure 10 for $w = LRRLR$.)

3.3 Edge/Vertex Correspondence in T_0

To understand which edges of T_0 are identified in T , we define a “direction” on each horizontal strip: right to left on the first strip, left to right on the strips below and above it, and so on, so that adjacent strips have opposite directions. Then we have the following result:




FIGURE 16. Edge identifications in strips of T_0 .




FIGURE 17. Two edge cycles in T_0 .

Lemma 3.3. *Each edge e of T_0 that crosses a strip bounds two triangles, one of which, δ say, lies on the side of e given by the direction on the strip. Then e and the opposite vertex of δ give the same edge of T . Horizontal edges of T_0 correspond to the edges at the opposite vertices of both adjacent triangles.*

This result is illustrated in Figure 16 and will be important in the arguments below.

Proof: The boundary between two tetrahedra in T is a punctured torus consisting of two ideal triangles, homotopic to a fiber of M . In T_0 the boundaries give *edge cycles*, lifts of a cycle of six edges going once around the cusp in the vertical direction as in Figure 13. Figure 17 shows two cycles.

Cycles meeting a given strip form either “ L ’s and Γ ’s” or “backward L ’s and Γ ’s.” This gives the directions: L ’s and Γ ’s read left to right!

We see from Figure 13 that the edges of a cycle are homotopic to three edges of T , say a, b, c , cyclically repeated. The vertices are the same three edges of T arranged so that if an edge is a , its vertices are b and c , and so on (as in Figure 16). Thus for any (forward or backward) L or Γ , the edge at the downstroke (a strip-crossing edge) equals the edge seen at the other end of its horizontal stroke. Further, the edge seen at any horizontal stroke equals the edge at the other end of the downstroke. This proves the lemma. \square




FIGURE 18. A vertex of order ≥ 10 .

3.4 Reduction to T_0

Let G be the image of the holonomy representation of $\pi_1(M)$ in $\text{Isom}(\mathbb{H}^3)$, so that $M = \mathbb{H}^3/G$ as usual. Let x be a fixed point of some maximal parabolic subgroup of G , which we identify with the group of translations G_0 defined above. We regard symmetries and hidden symmetries as isometries of \mathbb{H}^3 . Since M is 1-cusped, we can compose any symmetry or hidden symmetry with an element of G to obtain an isometry that fixes x . Thus if M has a symmetry, it is represented by a symmetry of T_0 ; if it has a nontrivial hidden symmetry, it is represented by a symmetry of T_0 that does not come from a symmetry of M .

Since our combinatorial picture of T_0 is not metrically exact, we know only that each symmetry and hidden symmetry acts as a simplicial homeomorphism (S.H.) of T_0 .

Lemma 3.4. *Simplicial homeomorphisms of T_0 preserve horizontal strips whenever $w \neq (LR)^m$ or $(LLRR)^m$ as a cyclic word, for any $m > 0$.*

Since vertex orders in T_0 are all even, we can define a *straight line* to be a path in the 1-skeleton that enters and leaves each vertex along opposite edges. A *strip* consists of a part of T_0 between two (infinite, disjoint) straight lines such that every interior edge crosses from one side of the strip to the other.

Proof: If an edge of T_0 lies on the edge of a strip, not necessarily horizontal, the opposite vertex of the triangle that crosses that strip must have order at least 6: the straight line going through this vertex has at least the two edges of the triangle on one side.

If there exists an S.H. that is not horizontal-strip-preserving, every vertex of T_0 will lie on a nonhorizontal edge that is an edge of some strip (namely the image of a previously horizontal strip edge).

Suppose there is a vertex of order 10 or more, corresponding to three or more adjacent L ’s or R ’s in w (as in Figure 18).




FIGURE 19. A vertex of order 8.

Here none of the nonhorizontal edges shown can lie on a strip edge because they are all opposite vertices of order 4. So in this case all S.H.'s must preserve horizontal strips.

Suppose there is a vertex of order 8, corresponding to LL or RR in w (as in Figure 19).

Now the only nonhorizontal edges shown that can be strip edges are the vertical ones. Vertex orders on this vertical edge alternate 4, 8, 4, 8. So if an S.H. maps a horizontal strip edge to this vertical one, vertex orders along some horizontal edge must be 4, 8, 4, 8, This determines T_0 and hence w , since $(LLRR)^m$. (The reader can verify that this particular T_0 admits a $\pi/2$ rotation.)

Finally, suppose there is no vertex of order greater than 6. It follows immediately that w is $(LR)^m$ and that T_0 is the tiling of the plane by equilateral triangles. (Again, this admits S.H.'s that are not horizontal-strip-preserving.) \square

Lemma 3.5. *Simplicial homeomorphisms of T_0 coming from symmetries of T preserve the horizontal strip directions whenever $w \neq (LR)^m$ or $(LLRR)^m$.*

Proof: An S.H. of T_0 that comes from a symmetry of T obviously preserves the order of every edge of T . Thus if we label each edge of T_0 with the order of the corresponding edge of T , the labels must be preserved.

The order of an edge of T corresponding to a vertex of T_0 is simply its order as a vertex of T_0 . By Lemma 3.3,




FIGURE 20. Edge orders in T for a vertex of order ≥ 10 .




FIGURE 21. Edge orders in T for a vertex of order 8.

we can use the strip directions to label the corresponding edges of T_0 .

Given a vertex of order 10 or more, we have the situation depicted in Figure 20. The direction on a horizontal strip cannot be reversed, because this would swap an edge of order 4 with one of order 6 or more.

Suppose now that the maximum vertex order is 8. For each vertex of order 8 we have edge orders as in Figure 21.

If there is an S.H. that reverses strip direction, then the other two sides of this diamond figure, wherever it appears, must also be labeled with 4's, as shown in Figure 22.

The 4 on the upper right edge implies that the right-hand vertex of the diamond has order greater than or equal to 8, hence 8. This gives another diamond figure adjacent to the first one. The argument can be repeated, giving the pattern associated with $w = (LLRR)^m$.

If there are no vertices of order greater than 6, then $w = (LR)^m$. \square

Proof of Theorem 3.1: It is shown in [Bowditch et al. 95] that the only arithmetic orientable hyperbolic punctured torus bundles are those for which $w = LR$, LLR (or LRR), or $LLRR$, or powers of these, having invariant trace fields $\mathbb{Q}(\sqrt{-3})$, $\mathbb{Q}(\sqrt{-7})$, and $\mathbb{Q}(\sqrt{-1})$ respectively. Thus, in the nonarithmetic cases, Lemmas 3.4 and 3.5 show that all symmetries of T_0 that come from hidden symmetries of M are represented by simplicial homeomorphisms that preserve the strips and strip directions.

Note that there are two “sister” manifolds for each w depending on whether the monodromy is φ_w or $-\varphi_w$; the




FIGURE 22. Edge orders if an S.H. reverses strip direction.




FIGURE 23. Symmetry of type 1.




FIGURE 24. Symmetry of type 2.




FIGURE 25. Symmetry of type 3.




FIGURE 26. Symmetry of type 4.

triangulations of \mathbb{H}^3 , however, are the same, depending only on w .

It is now easy to see that the possible symmetries of T_0 , modulo the translations in G_0 , are restricted to the types listed below; we show that each, if it occurs at all, comes from an actual symmetry of M :

1. Shifting up or down by two strips. This is realized by the symmetry $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \times 1$ from $F \times_{\pm\varphi_w} S^1$ to itself; see Figure 23. (Note that $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$ is central in $\text{SL}(2, \mathbb{Z})$.)
2. If w is a power, $w = u^m$, then T_0 admits a horizontal translation by $\ell(u)$ triangles, where $\ell(u)$ is the length of the word u . In this case, we have the situation of Figure 24, and $1 \times r_{2\pi/m}$ gives a symmetry of order m (where r_θ denotes rotation of the unit circle S^1 by θ).
For the $-\varphi_w$ case we replace one of the φ_u 's by $-\varphi_u$. After rotating, we also have to apply $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \times 1$ to $F \times [0, 1/m]$.
3. If w is *palindromic*, i.e., w and its reverse w' are the same as cyclic words, then T_0 admits a rotation by π about a point on one of the strip edges. This is realized by $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \times \text{refl}(S^1)$, where $\text{refl}(S^1)$ denotes a reflection of S^1 (Figure 25). Note that conjugation by $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ takes L to L^{-1} and R to R^{-1} .

4. If $\ell(w)$ is even, and rotating w by a half-turn swaps L 's and R 's, then T_0 has a glide reflection mapping a strip to itself, exchanging the top and bottom of the strip. This is realized by $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \times 1$, since conjugation by this matrix swaps L and R . More explicitly,

let $w = uv$, where v is u with L 's and R 's interchanged. The symmetry is the map shown in Figure 26, followed by r_π on the S^1 factor. This symmetry is orientation-reversing.

5. If $\ell(w)$ is even and reversing w swaps L 's and R 's (we might say that w is *antipalindromic*), then T_0 admits a glide reflection with vertical axis, shifting everything up by one strip. This is realized by $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \times \text{refl}(S^1)$, since conjugation by $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ takes L to R^{-1} and R to L^{-1} . This symmetry is orientation-reversing. We leave the picture as an exercise for the reader.

This completes the proof of Theorem 3.1. □

Clearly, if M admits any two of the symmetries 3 to 5, it admits the third, which is a product of the other two. Thus if w is not a power, M may have one of five possible symmetry groups.

4. ALGORITHM FOR FINDING ISOMETRIES OF TILINGS

To apply our commensurability criterion, Theorem 2.4, we need a way to determine when tilings of \mathbb{H}^n arising from ideal cell decompositions are isometric. We pursue this question in a slightly more general setting.

Let M and M' be geometric n -manifolds with universal cover $X^n = \mathbb{E}^n$ or \mathbb{H}^n , each with a given finite decomposition into convex polyhedra. We wish to determine whether the tilings T and T' that cover these two cell decompositions are isometric. (The elements of T , respectively T' , are convex polyhedra in X^n that project to polyhedra in the decomposition of M , respectively M' .)

Necessary and sufficient conditions for the tilings to be isometric are as follows:

1. There is an isometry i of X^n that maps an element of T isometrically onto an element of T' .
2. Whenever i maps $P \in T$ isometrically onto $P' \in T'$, and F is a (codimension-1) face of P , i maps the neighbor of P at F isometrically onto the neighbor of P' at $i(F)$.

Sufficiency follows from the fact that we can proceed from any tile in T to any other by a finite sequence of steps between neighboring tiles.

Let S and S' denote the polyhedra decomposing M and M' respectively.

Let Θ denote the set of all triples (j, p, p') , where $p \in S$, $p' \in S'$, and j is an isometry carrying p onto p' . We can find Θ in a finite number of steps. Condition 1 above is equivalent to Θ being nonempty.

We say that $(j, p, p') \in \Theta$ is *induced* by an isometry i of X^n if there exist $P \in T$ projecting to p and $P' \in T'$ projecting to p' such that i carries P isometrically onto P' and the restriction induces j . Let f be a face of p and let q and q' be the neighbors of p and p' at f and $j(f)$ respectively. The restriction of j to f induces an isometry between certain faces of q and q' . If this extends to (j_q, q, q') , we say that (j, p, p') *extends across* f to (j_q, q, q') . If (j, p, p') is induced by i , then (j_q, q, q') will also be induced by i .

Condition 2 is equivalent to the following: whenever (j, p, p') is induced by i , and f is a face of p , then (j, p, p') extends across f .

Theorem 4.1. *With the above notation, the tilings T and T' are isometric if and only if there exists a nonempty subset I of Θ such that every element of I extends across each of its faces to yield another element of I .*

Proof: If T and T' are isometric with isometry i , simply let I be the set of elements of Θ induced by i .

Conversely, suppose we have $I \subseteq \Theta$ having the stated properties. Choose any element $(j, p, p') \in I$ and any isometry i of X^n that induces it. Clearly condition 1 above is satisfied. For condition 2, note that if i maps P onto P' and this induces $(j, p, p') \in I$, then because this extends across all of its faces, i maps neighbors of P isometrically onto neighbors of P' , and these too induce elements of I . Therefore condition 2 is satisfied with all the induced triples belonging to I . □

An algorithm for finding such a subset $I \subseteq \Theta$ or establishing that none exists is straightforward. If Θ is empty, stop; there is no such subset. Otherwise, choose any element of Θ and put it in I . For each element of I check whether we can extend across all faces. If we can't, remove all elements of I from Θ and start again. If we can, and every element we reach is already in I , stop and output I . If we can extend across faces of all elements of I , and we obtain new elements not yet in I , add those new elements to I and check again.

Each set I found by the algorithm represents an equivalence class of isometries carrying T onto T' : choose an isometry i inducing any element of I ; then isometries i, i' are *equivalent* if $i' = g'ig$, where g is a covering transformation of M and g' is a covering transformation of M' .

To find the symmetry group of a tiling T we apply the algorithm with $T = T', M = M'$, etc. Let Γ denote the group of covering transformations of M . Repeating the algorithm until Θ is exhausted, we obtain finitely many equivalence classes of symmetries. These represent the double cosets of Γ in $\text{Symm}(T)$. Together with Γ they generate $\text{Symm}(T)$.

Remark 4.2. Each set I found by this algorithm also represents a common covering N of M and M' constructed as follows: For each element $r = (j, p, p') \in I$, let p_r denote a copy of p . Let N be the disjoint union of the polyhedra p_r as r varies over I , with the following face identifications: Each face f of p appears as a face f_r of p_r . Whenever r extends across f to $s = (j_q, q, q')$, identify face f_r of p_r with face f_s of q_s .

4.1 Example

Let M be the Euclidean torus obtained by gluing a 1×3 rectangle along opposite pairs of edges. Let M' be




FIGURE 27. Isometric tilings of the plane arising from Euclidean tori M, M' .




FIGURE 28. A common covering of M, M' .




FIGURE 29. Another common covering of M, M' .

another such torus obtained by gluing a 2×1 rectangle. Subdivide each into unit squares as shown in Figure 27, so that S , the subdivision of M , equals $\{A, B, C\}$, while $S' = \{P, Q\}$. Both S and S' lift to tilings of the plane by unit squares.

Then Θ consists of 48 elements, since there are eight ways of isometrically mapping one unit square onto another and six combinations of squares to be mapped. Suppose we start with (t_{AP}, A, P) , where t_{AP} denotes the translation carrying A onto P in Figure 27. We can extend (t_{AP}, A, P) across the marked edge (whose image in M' is also shown) to obtain (t_{BP}, B, P) . Continuing until we have extended across every available edge, we obtain six elements of Θ , which we can abbreviate to $\{t_{AP}, t_{BP}, t_{CP}, t_{AQ}, t_{BQ}, t_{CQ}\}$. These give rise to the common covering of M and M' depicted in Figure 28.

Suppose instead we begin with a clockwise rotation r_{AP} through 90° carrying A onto P . Then we obtain a different common cover, shown in Figure 29.

Suppose finally we subdivide one of the squares in S' into two triangles as in Figure 30. Now the tilings are clearly not isometric.




FIGURE 30. Two nonisometric tilings of the plane.

The algorithm might start with t_{AP} , but when it tries to extend this across the marked edge it will find that the induced mapping from an edge of A to an edge of Q does not extend to any isometry from A to Q .

4.2 Combinatorial Construction of Coverings

Mostow–Prasad rigidity often allows us to work in a purely combinatorial setting; this motivates the following. The algorithm described above can be viewed as a method for constructing a common cover of two manifolds that admit suitable subdivisions into polyhedra.

The algorithm works equally well for topological spaces M and M' admitting n -dimensional polyhedral decompositions (i.e., obtained by gluing n -dimensional Euclidean or finite-volume hyperbolic convex polyhedra pairwise along their $(n - 1)$ -dimensional faces such that open faces of dimension $q \leq n$ embed). We then let each element $(j, p, p') \in \Theta$ represent a combinatorial isomorphism between p and p' . (We can barycentrically subdivide the decompositions of M and M' in order to realize the j 's as piecewise linear homeomorphisms in a canonical way.)

If the algorithm is successful, it constructs, in general, a branched covering of M and M' , branched over the codimension-2 skeletons of M and M' .

When M and M' are PL-manifolds, we would like to construct an unbranched cover that is also a manifold. It is sufficient to check that the cover we construct does not branch over any codimension-2 cell of M or M' . For if this is the case we proceed by induction on q to show that our cover is not branched over any codimension- q cell of M (or M').

By construction, there is no branching over cells of codimension 0 and 1, and we assume that there is no branching over cells of codimension 2. So suppose $q \geq 3$ and there is no branching over cells of codimension $\leq q - 1$. Then the link of a codimension- q cell in the cover is a connected unbranched covering of the link of a codimension- q cell of M . But since the latter is a $(q - 1)$ -sphere, hence simply connected, since $q \geq 3$, the covering is a homeomorphism.

In the case of most interest to us, namely that M and M' are cusped hyperbolic 3-manifolds, Mostow–Prasad rigidity ensures that if N topologically covers both M and M' , then the induced hyperbolic structures on N are isometric.

Frequently, an ideal cell decomposition of a hyperbolic 3-manifold consists entirely of ideal tetrahedra. If this is the case for M and M' , then our algorithm will always find a common cover, branching over the edge sets

of the two manifolds (elements of Θ will always extend over their faces because all tetrahedra are combinatorially equivalent). This will indicate that the two manifolds are commensurable only if the covering is in fact unbranched.

When an element of Θ induces a mapping of an edge e of (the ideal cell decomposition of) M onto an edge e' of M' , the resulting cover will be unbranched if and only if the order of e (the number of tetrahedra that glue around it) is equal to the order of e' . For example, this has the following corollary.

Theorem 4.3. *If a finite-volume hyperbolic 3-manifold M admits a decomposition into ideal tetrahedra such that the order of every edge is 6, then M is commensurable with the complement of the figure-eight knot.*

5. ENUMERATING CANONICAL CELL DECOMPOSITIONS I

For 1-cusped hyperbolic manifolds with discrete commensurator we now have all the ingredients of an effective method for testing commensurability. For multicusped manifolds we need a way to search through the (finite) set of all canonical ideal cell decompositions. In this section and the next we describe two alternative approaches to this problem, while in Section 7 we show how the search can be restricted for greater efficiency.

The first approach is very simple-minded: if we can bound the degree with which M covers its commensurator quotient, we can enumerate all cusp cross sections that could possibly cover equal-area cross sections in the quotient. Such degree bounds can be obtained from estimates on the minimum volume of cusped nonarithmetic hyperbolic 3-manifolds [Meyerhoff 85, Adams 92, Neumann and Reid 92b].

The second approach is geometric and may be of more theoretical interest, since it truly finds *all* the canonical cell decompositions. In fact, it associates with a given c -cusped M , a convex polytope in \mathbb{R}^c whose k -dimensional faces, for $0 \leq k < c$, are in one-to-one correspondence with canonical cell decompositions of M .

Let M be a hyperbolic orbifold with (horoball) cusp neighborhoods C_1, \dots, C_m and let N be a degree- d quotient of M with corresponding cusp neighborhoods c_1, \dots, c_n . Let $\pi : M \rightarrow N$ be the covering projection and let $c_{j(i)} = \pi(C_i)$. If we choose horospherical cross sections in N with “area” (i.e., codimension-1 volume) equal to 1, then the area of ∂C_i is equal to the degree with which C_i covers $c_{j(i)}$. The sum of the areas of the C_i covering c_j will be d for each c_j .

Thus, in order to find a possible degree- d quotient of M via the canonical cell decomposition, we can enumerate the possible integer area vectors as follows. For each n , $1 \leq n \leq m$, and each partition of $\{1, \dots, m\}$ into n nonempty subsets I_1, \dots, I_n , enumerate all area vectors (a_1, \dots, a_m) such that each a_i is a positive integer and $\sum_{i \in I_j} a_i = d$ for $j = 1, \dots, n$.

Example: if $m = 3$ we have partitions

$$\begin{aligned} & \{\{1, 2, 3\}\}, \quad \{\{1, 2\}, \{3\}\}, \quad \{\{1, 3\}, \{2\}\}, \\ & \{\{1\}, \{2, 3\}\}, \quad \{\{1\}, \{2\}, \{3\}\}. \end{aligned}$$

The first partition admits $\frac{1}{2}(d-1)(d-2)$ area vectors, the next three admit $d-1$ each, and the last, just one.

We can enumerate area vectors corresponding to possible quotients of degree d as follows. Any quotient of M of degree d has $n \leq m$ cusps and has a corresponding area vector (a_1, \dots, a_m) with sum $\sum_{i=1}^m a_i = nd$. So if we fix an integer $D \geq m$ and enumerate all positive-integer area vectors summing to at most D , we will find all quotients of degree d with $n \leq m$ cusps such that $nd \leq D$. In particular, this will include all quotients with at most $m-1$ cusps, provided that $d \leq D/(m-1)$. Since the canonical cell decomposition is determined by the *ratio* of the areas, any m -cusped quotient is found using the area vector $(1, \dots, 1)$ with $\sum_i a_i = m \leq D$. So we will find *all* canonical cell decompositions arising from quotients of degree $d \leq D/(m-1)$.

6. ENUMERATING CANONICAL CELL DECOMPOSITIONS II

Let M be a hyperbolic n -manifold with $c > 0$ cusps. Then the set of possible choices of (not necessarily disjoint) horospherical cross sections dual to the cusps of M is parameterized by the vector of their areas (i.e., $(n-1)$ -dimensional volumes) in $\mathbb{R}_{>0}^c = \{(v_1, \dots, v_c) \in \mathbb{R}^c \mid v_1 > 0, \dots, v_c > 0\}$. In fact, it turns out to be more convenient to parameterize by the $(n-1)$ th root of area, a quantity we will refer to as *size*. Multiplying a size vector $v \in \mathbb{R}_{>0}^c$ by a constant $\lambda > 0$ has the effect of shifting the corresponding horospherical cross sections a distance $\log(\lambda)$ down (i.e., away from) the cusps.

For each $v \in \mathbb{R}_{>0}^c$ we obtain a Ford spine. It is clear from the definition that if we choose a set of disjoint horospherical cross sections and then shift them all up or down by the same amount, we get the same Ford spine. Thus the set of possible Ford spines is parameterized by the set of rays in $\mathbb{R}_{>0}^c$, or equivalently, by points in the open $(c-1)$ -dimensional simplex $\mathcal{S} = \{(v_1, \dots, v_c) \in$

$\mathbb{R}_{>0}^c \mid v_1 + \dots + v_c = 1$. For a 1-cusped manifold the Ford spine is unique.

Dual to each Ford spine is a *canonical cell decomposition* $D(v)$. As we vary $v \in \mathcal{S}$, the decomposition changes only when the combinatorics of the spine change. To better understand this dependence we now review an alternative approach to defining $D(v)$, namely the original one in [Epstein and Penner 88].

We work in Minkowski space $\mathbb{E}^{n,1}$ with the inner product $*$ defined by

$$x * y = x_1y_1 + \dots + x_ny_n - x_{n+1}y_{n+1}$$

for $x = (x_1, \dots, x_{n+1})$, $y = (y_1, \dots, y_{n+1})$ in \mathbb{R}^{n+1} . Then hyperbolic space \mathbb{H}^n is the upper sheet of the hyperboloid $x * x = -1$, and the horospheres in \mathbb{H}^n are represented by the intersections of hyperplanes, having lightlike normal vectors, with \mathbb{H}^n . Each such hyperplane H has a unique Minkowski normal \mathbf{n} such that $x \in H$ if and only if $x * \mathbf{n} = -1$.

Let Γ denote the group of covering transformations of \mathbb{H}^n over M . Each size vector v gives rise to a Γ -invariant set of horospheres in \mathbb{H}^n . The resulting set of normals in Minkowski space is invariant under the action of the group Γ . The convex hull of this set of points, which we shall refer to as the *Epstein–Penner convex hull*, intersects every ray based at the origin passing through a point in the upper sheet of the hyperboloid. The boundary of this convex set is a union of closed convex n -dimensional polytopes having coplanar lightlike vertices. Epstein and Penner [Epstein and Penner 88] show that these project to a locally finite, Γ -invariant set of ideal polyhedra in \mathbb{H}^n , which in turn projects to a finite set $D(v)$ of ideal hyperbolic polyhedra in M .

Starting with a given ideal hyperbolic cell decomposition D of M and a size vector v , we proceed next to describe necessary and sufficient conditions for D to be the canonical cell decomposition $D(v)$, as in [Weeks 93] and [Sakuma and Weeks 95a].

Let C be a cell of D . Then v determines a horospherical cross section to each ideal vertex of C . We lift C and this choice of horospheres to \mathbb{H}^n . In Minkowski space, this gives a convex (Euclidean) n -dimensional polytope whose vertices are the hyperplane normals for these horospheres. Whenever C is not a simplex, it is necessary to add the condition that these vertices be coplanar. If this is satisfied for all the nonsimplicial cells of D , we can lift each cell to an n -dimensional polytope in Minkowski space with vertices corresponding to the choice of horospheres determined by the size vector v . If C and C' are

neighboring cells in D , it is necessary that the angle between neighboring lifts into Minkowski space be convex upward. Together these conditions are also sufficient to imply $D = D(v)$.

These conditions can be expressed as a set of linear equations and linear inequalities on the entries of the size vector v .

Proposition 6.1. *Let D be an ideal hyperbolic cell decomposition of a cusped hyperbolic n -manifold M with c cusps. Then there exist matrices L_D and F_D with c columns such that for $v \in \mathbb{R}_{>0}^c$, $D(v) = D$ if and only if $L_D v = 0$ and $F_D v > 0$.*

Note 6.2. Here v is written as a column vector, and the condition $F_D v > 0$ means that each entry of the vector $F_D v$ is positive.

Proof: Let v_i denote the entry of v corresponding to the i th cusp of M . Let \mathbf{n}_j be the vertex representative for the j th vertex of C lifted to Minkowski space for the choice of horospheres given by $v = (1, \dots, 1)$. Then for an arbitrary size vector v , the corresponding representative is $\mathbf{n}_j/v_{c(j)}$, where $c(j)$ is the cusp of the j th vertex of C .

The coplanarity condition on a nonsimplicial cell C gives a set of linear equations satisfied by v , one for each vertex of C in excess of $n + 1$, as follows. Let N_v be a Euclidean normal to the hyperplane containing $\{\mathbf{n}_0/v_{c(0)}, \dots, \mathbf{n}_n/v_{c(n)}\}$ such that $(\mathbf{n}_j/v_{c(j)}) \cdot N_v = 1$ for $j = 0, \dots, n$, where \cdot denotes the Euclidean dot product. Writing M_C for the inverse of the matrix with rows \mathbf{n}_j , we obtain $N_v = M_C(v_{c(0)}, \dots, v_{c(n)})^t$, which is a linear function of v . For $j > n$, $\mathbf{n}_j/v_{c(j)}$ belongs to this hyperplane if and only if $\mathbf{n}_j \cdot N_v - v_{c(j)} = 0$, which is linear in v . The full set of constraints for C gives a matrix equation $L_C v = 0$.

The convexity condition at an $(n-1)$ -cell f of D , being the common face of n -cells C and C' , can be expressed as follows. Let N_v , as above, be the defining normal for the hyperplane containing the lift of C determined by v . Let $\mathbf{n}'_k/v_{c(k)}$ be a vertex of an adjacent lift of C' , not in the lift of f . This vertex lies above the hyperplane if and only if $(\mathbf{n}'_k/v_{c(k)}) \cdot N_v > 1$, or equivalently, $\mathbf{n}'_k \cdot N_v - v_{c(k)} > 0$. We refer to the left-hand side of this inequality as the *tilt* at f of v and express the condition as $F_f v > 0$, where F_f is a suitable row vector. (Note that the sign of our tilt function is opposite to that of [Weeks 93] and [Sakuma and Weeks 95b].)

Finally, concatenate the matrices L_C into a matrix L_D and the rows F_f into a matrix F_D . □

Let \mathcal{P}_D denote the set of $v \in \mathbb{R}_{>0}^c$ such that $L_D v = 0$ and $F_D v > 0$. We call this the *parameter cell* of D , since it contains all cusp size parameters v such that $D(v) = D$. Each $v \in \mathbb{R}_{>0}^c$ belongs to a parameter cell, namely $\mathcal{P}(v) = \mathcal{P}_{D(v)}$. The parameter cell \mathcal{P}_D is nonempty if and only if D is a canonical cell decomposition.

It is shown in [Akiyoshi 01] that the number of canonical cell decompositions is finite. Therefore $\mathbb{R}_{>0}^c$ is a union of finitely many parameter cells.

Proposition 6.3. *Each $v \in \mathbb{R}_{>0}^c$ can be perturbed to obtain a nearby vector v' such that $\mathcal{P}(v')$ has dimension c and $\mathcal{P}(v)$ is a face of $\mathcal{P}(v')$ (or equals $\mathcal{P}(v)$ if this has dimension c).*

Proof: Let $D = D(v)$. If L_D is zero (or empty) then $\mathcal{P}(v)$ is an open subset, hence a c -dimensional cell, and we just set $v' = v$.

Otherwise, perturb v such that it leaves the linear subspace determined by $L_D v = 0$. For a small perturbation, the Epstein–Penner convex hull changes as follows: no dihedral angle between adjacent n -faces goes to π , but some nonsimplicial n -faces may be subdivided if their vertices become noncoplanar.

It follows that L_D may lose rows and F_D may gain rows. Let D' be the new decomposition. Then $L_{D'} \neq L_D$ because $L_{D'} v' = 0$, while $L_D v' \neq 0$. Repeat until $L_{D'}$ is zero (or empty). Then $\mathcal{P}(v')$ has dimension c .

Now v belongs to a face of $\mathcal{P}_{D'}$, since it satisfies $F_f v > 0$ for each face f common to D and D' , and $F_{f'} v = 0$ for each face f' of D' not in D . The former condition amounts to $F_D v > 0$. We have to show that the latter is equivalent to $L_D v = 0$. But that is equivalent to the coplanarity of the lifted vertices of each n -cell of D . Such a cell may be subdivided by new faces f' in D' . Then the vertices will be coplanar at v if and only if the tilt at each subdividing face is zero, i.e., if and only if $F_{f'} v = 0$ for the subdividing faces. \square

The above proposition implies that each face of a parameter cell is another parameter cell; the decomposition corresponding to a parameter cell is a refinement of the decompositions corresponding to its faces.

Remark 6.4. It is tempting to suppose that the canonical cell decomposition of M corresponding to a c -dimensional parameter cell must consist entirely of ideal simplices, but this need not be the case. In general, we can have cell decompositions with nonsimplicial cells such that L_D is a

zero matrix. (For example, this occurs for the Borromean rings complement; see Section 6.1 below.)

We now have the following algorithm for finding all canonical cell decompositions. First we find all the c -dimensional parameter cells:

1. Choose an arbitrary $v \in \mathbb{R}_{>0}^c$.
2. Perturb v if necessary, as in the proof of Proposition 6.3, so that $\mathcal{P}(v)$ has dimension c , and add it to our list of cells.
3. If the closure of the cells we have found so far does not contain the whole of $\mathbb{R}_{>0}^c$, choose a new v not in the closure of any cell found so far and repeat step 2.

By the finiteness result quoted above, this algorithm eventually terminates. We can then enumerate all canonical cell decompositions by enumerating the faces of all dimensions of the cells $\mathcal{P}(v)$.

While the computational geometry involved in implementing the above algorithm is certainly possible, it is not particularly nice. We explain a refinement that gives a little more insight and an algorithm that is easier to implement.

For a decomposition D of M , let Σ_D denote the row vector obtained by adding together the rows of F_D . We define the *tilt polytope* of M to be the set of $v \in \mathbb{R}_{>0}^c$ such that $\Sigma_D \cdot v < 1$ for all canonical cell decompositions D of M .

Proposition 6.5. *The tilt polytope T of M is bounded. The parameter cells \mathcal{P}_D of M are the cones over the origin of those faces of T that are not contained in $\partial\mathbb{R}_{>0}^c$.*

Proof: We show that the closure of a c -dimensional parameter cell $\overline{\mathcal{P}}_D$ has bounded intersection with T . Let v be a unit vector in $\overline{\mathcal{P}}_D$. Then since v is not contained in every face of \mathcal{P}_D , we have $\Sigma_D \cdot v > 0$. The length of any multiple of v contained in T is bounded by $1/(\Sigma_D \cdot v)$. Since this is continuous in v , and the set of such v is compact, $\overline{\mathcal{P}}_D \cap T$ is bounded. Since T is a union of finitely many such sets, it is bounded.

Let us write H_D for the half-space $\{x \in \mathbb{R}^c \mid \Sigma_D \cdot x < 1\}$. Then T is the intersection of all the H_D 's with $\mathbb{R}_{>0}^c$. We will show that if v belongs to a c -dimensional parameter cell \mathcal{P}_D , and $\mathcal{P}_{D'}$ is any other parameter cell, then the ray generated by v leaves H_D before it leaves $H_{D'}$.

It will then follow that a ray in \mathcal{P}_D penetrates the (nonempty) face of T generated by H_D . Since a ray not in \mathcal{P}_D belongs to the closure of some other parameter

cell, it does not leave H_D first and therefore does not pass through the same face of T . Since cones on the lower-dimensional faces of a $(c-1)$ -dimensional face of T are the faces of a c -dimensional parameter cell, the result then follows from Proposition 6.3.

It remains to show that H_D cuts off any ray in \mathcal{P}_D closer to the origin than $H_{D'}$, for all parameter cells $\mathcal{P}_{D'} \neq \mathcal{P}_D$. Equivalently, for $v \in \mathcal{P}_D$, $\Sigma_D \cdot v > \Sigma_{D'} \cdot v$.

Firstly, let \mathcal{P}_D and $\mathcal{P}_{D'}$ be any two parameter cells such that $\mathcal{P}_{D'}$ is a face of \mathcal{P}_D , and let v belong to \mathcal{P}_D . The rows of $F_{D'}$ are a proper subset of the rows of F_D , and since $F_f \cdot v > 0$ for each row, $\Sigma_D \cdot v > \Sigma_{D'} \cdot v$. If instead \mathcal{P}_D is a face of $\mathcal{P}_{D'}$, then the rows of $F_{D'}$ omitted from F_D are precisely those for which $F_f \cdot v$ vanishes. Therefore in that case, $\Sigma_D \cdot v = \Sigma_{D'} \cdot v$.

Next, let \mathcal{P}_D be a c -dimensional parameter cell, and let $\mathcal{P}_{D'}$ be arbitrary, with $D' \neq D$. Choose $v \in \mathcal{P}_D$ and $v' \in \mathcal{P}_{D'}$. Let $\mathcal{P}_{D_1}, \dots, \mathcal{P}_{D_m}$ be the parameter cells through which the straight line $v_t := (1-t)v + tv'$ passes for $0 \leq t \leq 1$ (so that $D_1 = D$ and $D_m = D'$). For v_t in \mathcal{P}_{D_i} , $\Sigma_{D_i} \cdot v_t \geq \Sigma_{D_{i+1}} \cdot v_t$, while for $v_t \in \mathcal{P}_{D_{i+1}}$, $\Sigma_{D_i} \cdot v_t \leq \Sigma_{D_{i+1}} \cdot v_t$. Since the difference between these terms is (affine) linear in t , the former inequality must hold for all lesser values of t , in particular when $v_t = v$. Note also that the first such inequality is strict, namely, $\Sigma_{D_1} \cdot v > \Sigma_{D_2} \cdot v$. It follows that $\Sigma_D \cdot v > \Sigma_{D'} \cdot v$ for arbitrary D' . \square

Let T_0 be a polytope resulting from the intersection of $\mathbb{R}_{>0}^c$ with some of the half-spaces H_D defined in the above proof. If $T_0 \supsetneq T$, some face $A = \overline{T_0} \cap \partial H_D$ of T_0 will contain a point $v \in \mathbb{R}_{>0}^c$ not in \overline{T} , and thus not in $\overline{T} \cap \partial H_D$, nor in the cone on this, $\overline{\mathcal{P}_D}$. Therefore $L_D v \neq 0$ or $F_f v < 0$ for some row F_f of F_D . This gives a test for when T_0 properly contains T ; when satisfied, it yields a new half-space $H_{D(v)}$ whose intersection with T_0 is strictly smaller. After a finite number of intersections we arrive at $T_0 = T$. See Figure 31.

The computational geometry involved in the above is relatively straightforward. By using homogeneous coordinates we can treat an unbounded region, such as $\mathbb{R}_{>0}^c$, as a polytope with some vertices “at infinity.” The face A of T_0 , as defined above, is the convex hull of those vertices v of T_0 satisfying $\Sigma_D \cdot v = 1$. If any of these satisfy $L_D v \neq 0$ or $F_f v < 0$ for some row F_f of F_D , we conclude that $T_0 \neq T$. If such v lies in $\partial \mathbb{R}_{>0}^c$, we perturb it a little to bring it inside $\mathbb{R}_{>0}^c$ before determining a new half-space $H_{D(v)}$.




FIGURE 31. T_0 is a partially computed tilt polytope. The face of T_0 containing $v \in \mathcal{P}_D$ has a vertex w not in $\overline{\mathcal{P}_D}$. Therefore $D(w)$ gives another face of T .




FIGURE 32. The Borromean rings.

6.1 Example: The Borromean Rings Complement

Let M be the complement of the Borromean rings in S^3 (shown in Figure 32). Then M is an arithmetic hyperbolic 3-manifold with three cusps. It may be realized by gluing two regular ideal hyperbolic octahedra in the pattern shown in Figure 33 (see [Thurston 78]). Letters indicate face identifications; cusps are numbered 0, 1, 2.

Equal-area cusp cross sections give rise to symmetrically placed ideal vertex cross sections in the two octahedra. It follows that the corresponding ideal cell decomposition consists of precisely these two octahedra. We will first compute its parameter cell.

Let us position the tiling of \mathbb{H}^3 by ideal octahedra such that one of them, call it C , has a lift with vertices at $\{(\pm 1, 0, 0), (0, \pm 1, 0), (0, 0, \pm 1)\}$ in the pro-




FIGURE 33. Gluings for the Borromean rings complement.




FIGURE 34. Initial cell decomposition of M into ideal octahedra and the corresponding projective parameter cell.




FIGURE 35. Decomposition into ideal tetrahedra and corresponding parameter cell.

jective ball model. When the cusp cross sections all have equal area and contain the center of the octahedron, their Minkowski normals are $\mathbf{n}_0 = (1, 0, 0, 1)$, $\mathbf{n}_1 = (0, 1, 0, 1)$, $\mathbf{n}_2 = (0, 0, 1, 1)$, $\mathbf{n}_3 = (-1, 0, 0, 1)$, $\mathbf{n}_4 = (0, -1, 0, 1)$, $\mathbf{n}_5 = (0, 0, -1, 1)$. A neighboring lift of the other octahedron, call it C' , has vertex representatives $\{\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_2, \mathbf{n}'_3, \mathbf{n}'_4, \mathbf{n}'_5\}$, where \mathbf{n}'_i is the Minkowski metric reflection of \mathbf{n}_i in the hyperplane spanned by $\{\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_2\}$. Then $\mathbf{n}'_3 = (1, 2, 2, 3)$, $\mathbf{n}'_4 = (2, 1, 2, 3)$, $\mathbf{n}'_5 = (2, 2, 1, 3)$. (See the left-hand side of Figure 34.)

For size vector $v = (v_0, v_1, v_2)^t$, the resulting ideal vertex representatives become \mathbf{n}_i/v_i , \mathbf{n}_{i+3}/v_i , and \mathbf{n}'_{i+3}/v_i for $i = 0, 1, 2$. The hyperplane containing \mathbf{n}_0/v_0 , \mathbf{n}_1/v_1 , \mathbf{n}_2/v_2 , and \mathbf{n}_3/v_0 is $\{\mathbf{x} \mid \mathbf{x} \cdot N_v = 1\}$, where N_v is calculated as in the proof of Proposition 6.1, giving

$$N_v = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & -1 \\ -1 & 2 & 0 & -1 \\ -1 & 0 & 2 & -1 \\ 1 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} v_0 \\ v_1 \\ v_2 \\ v_0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} v.$$

Then \mathbf{n}_4/v_1 and \mathbf{n}_5/v_2 lie in this hyperplane if and only if $\mathbf{n}_4 \cdot N_v - v_1 = 0 = (2, -2, 0)v$ and $\mathbf{n}_5 \cdot N_v - v_2 =$

$0 = (2, 0, -2)v$. Hence $L_C = \begin{pmatrix} 2 & -2 & 0 \\ 2 & 0 & -2 \end{pmatrix}$. By symmetry, $L_{C'}$ is the same. Next we compute the tilt of f_0 , the face between C and C' . This will be positive if \mathbf{n}'_3/v_0 lies above the hyperplane defined by N_v , i.e., if $\mathbf{n}'_3 \cdot N_v - v_0 = (-2, 2, 2)v > 0$. Thus $F_{f_0} = (-2, 2, 2)$, and the tilt is positive at $v = (1, 1, 1)^t$. By symmetry, the other faces of the octahedra also have positive tilt. The parameter cell of this decomposition $\{v \mid v_0 = v_1 = v_2\}$ is, projectively, a point (as shown in the right half of Figure 34).

Suppose next we increase v_0 slightly. The faces of C still have positive tilt. We claim that C is subdivided into four tetrahedra containing the edge $\mathbf{n}_0, \mathbf{n}_3$, as shown on the left of Figure 35.

Let f_1 be the triangle with vertices $\{\mathbf{n}_0, \mathbf{n}_2, \mathbf{n}_3\}$. Then f_1 has positive tilt if and only if $\mathbf{n}_4 \cdot N_v - v_1 = (2, -2, 0)v = F_{f_1}v > 0$. Letting f_2 be the triangle with vertices $\{\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_3\}$, we obtain $F_{f_2} = (2, 0, -2)$. By symmetry, the tilts of the other two faces shown in Figure 35 are the same. Four more triangles, having the same tilts, subdivide C' . Therefore M is subdivided into eight simplices in the parameter cell shown on the right of Figure 35.

Increasing v_0 , eventually $F_{f_0}v = 0$, so that the faces of C and C' vanish from the decomposition. The four simplices around the edge $\mathbf{n}_1, \mathbf{n}_2$ become an octahedron. Further increasing v_0 so that $F_{f_0}v < 0$, we claim




FIGURE 36. Decomposition into square-based pyramids and corresponding parameter cell.




FIGURE 37. All parameter cells for canonical decompositions of the Borromean rings complement.

that this octahedron splits at the square dual to the edge $\mathbf{n}_1, \mathbf{n}_2$ into two square-based pyramids with vertices $\{\mathbf{n}_0, \mathbf{n}'_0, \mathbf{n}_3, \mathbf{n}_1\}$ and $\{\mathbf{n}_0, \mathbf{n}'_0, \mathbf{n}_3, \mathbf{n}_2\}$, as shown on the left of Figure 36.

Let N'_v define the hyperplane containing $\{\mathbf{n}_0/v_0, \mathbf{n}_1/v_1, \mathbf{n}_3/v_0, \mathbf{n}'_3/v_0\}$. Then

$$N'_v = \begin{pmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix} v.$$

Let C_1 be the square-based pyramid with vertices $\{\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_3, \mathbf{n}'_3, \mathbf{n}'_0\}$, where $\mathbf{n}'_0 = (-1, 2, 2, 3)$ is the Minkowski reflection of \mathbf{n}'_3 in the plane containing $\mathbf{n}_1, \mathbf{n}_2, \mathbf{n}_4, \mathbf{n}_5$. We find that $L_{C_1} = (0, 0, 0)$. Let f_3 be the triangle with vertices $\{\mathbf{n}_0, \mathbf{n}_3, \mathbf{n}'_3\}$. Then $F_{f_3} = (1, -1, -1)$ and $F_{f_3}v > 0$. Therefore M is divided into four square-based pyramids in the parameter cell shown on the right of Figure 36, as mentioned in Remark 6.4.

By symmetry, the full set of parameter cells is projectively as shown in Figure 37. On the three line segments containing the center point, each of the original octahedra is subdivided into two square-based pyramids.

7. COMMENSURABILITY OF CUSPS

The number of canonical cell decompositions that can arise for a multicusped manifold can be quite large, plac-

ing practical limits on the usefulness of our methods. We show next how it is often possible to greatly reduce the complexity of computing the commensurator of a manifold.

A horospherical cross section of a cusp in an orientable hyperbolic 3-manifold is a Euclidean torus, well defined up to similarity. We can position and scale a fundamental parallelogram in \mathbb{C} such that one vertex lies at the origin and the two adjacent edges end at 1 and a point z in the upper half-plane. Such a z is called a *cusps shape parameter*. Alternative choices of fundamental parallelogram yield parameters differing by the action of $SL_2\mathbb{Z}$ by Möbius transformations. By choosing a suitable fundamental domain for the action of $SL_2\mathbb{Z}$ on the upper half-plane, we can make a canonical choice of shape parameter for each cusp.

If one cusp covers another, there is an induced covering of Euclidean tori. It follows that their cusp shape parameters are related by the action of an element of $GL_2\mathbb{Q}$. Let us call cusp shapes *commensurable* if they are so related. As we shall see shortly, it is not hard to determine when two cusp shapes are commensurable.

Suppose we are trying to determine the (discrete) commensurator quotient Q of M . If M has cusps of incommensurable shape, these necessarily cover distinct cusps of Q . Therefore any assignment of cusp neighborhoods in Q will yield a Ford spine and tiling whose symmetry group is the whole commensurator (the group of covering transformations of Q).

It follows that we can start by choosing arbitrary horospheres in each of a set of representatives for the commensurability classes of cusps of M . Then as we vary our choices of horosphere in the remaining cusps we will be sure to find a tiling whose symmetry group is the commensurator of M . The easiest case is that in which no two cusps of M are commensurable. Then any choice of horospheres at all will do.

The harder case is that in which M has multiple commensurable cusps. If the symmetry group of M is non-

trivial, it may act by nontrivially permuting some of these cusps. Conceptually we should first divide M by its symmetry group and then find the commensurator of this quotient.

For practical purposes it is hard to work with nonmanifold quotients. Instead, whenever two cusps are related by a symmetry of M , we choose symmetrically equivalent cusp neighborhoods. Let c be the number of orbits under the action of $\text{Symm}(M)$ on the cusps. Let d be the number of distinct commensurability classes of cusp shape. Then the parameter space of relative horosphere positions we need to search, in order to find a tiling whose symmetry group equals the commensurator, has dimension $c - d$.

Our algorithm is really only a slight modification of the algorithm of Weeks [Weeks 93] for finding the symmetries of a cusped hyperbolic manifold $M = \mathbb{H}^3/\Gamma$. In order to find the symmetries, we need only consider the canonical cell decomposition arising from a choice of cusp neighborhoods such that all boundary tori have equal area. Then the symmetries of M are the symmetries of the lifted tiling that normalize Γ . Equivalently, they are the symmetries of the tiling for which the covering described in Remark 4.2 has degree one.

Returning to the question of when two cusp shapes are commensurable, we note first that cusp shapes of M belong to the invariant trace field of M . But if k is any number field, and α, α' are irrational elements of k , they are related by an element of $\text{GL}_2 \mathbb{Q}$ if and only if

$$(c\alpha + d)\alpha' = a\alpha + b \quad (7-1)$$

is soluble for $a, b, c, d \in \mathbb{Q}$ such that

$$ad - bc \neq 0. \quad (7-2)$$

We can replace (7-2) with the condition that a, b, c, d be not all zero, since for $\alpha, \alpha' \notin \mathbb{Q}$, (7-2) follows automatically from (7-1) and the fact that $\{1, \alpha\}$ are linearly independent over \mathbb{Q} . Regarding k as a finite-dimensional vector space over \mathbb{Q} , we see that (7-1) has nontrivial solutions if and only if $\{1, \alpha, \alpha', \alpha\alpha'\}$ is linearly dependent over \mathbb{Q} . In particular, if $[k : \mathbb{Q}] < 4$, then all irrationals are commensurable in this sense. We thank Ian Agol for pointing out this condition.

7.1 Example: A Knot with Cusp Field Not Equal to Invariant Trace Field

An interesting example uncovered during this work is the complement of the knot 12n706 shown in Figure 38. This has one torus cusp with shape parameter $z = 6i$ generating a cusp field $\mathbb{Q}(i)$ that is strictly contained in its




FIGURE 38. Knot 12n706.

invariant trace field $\mathbb{Q}(i, \sqrt{3})$. This answers a question of Neumann–Reid [Neumann and Reid 92a], who asked whether the figure-eight knot and the two dodecahedral knots of Aitchison–Rubinstein [Aitchison and Rubinstein 92] are the only such examples.⁵

8. EXPERIMENTAL RESULTS

We have implemented the algorithms described here and used them to compute the commensurability classes of all 4929 manifolds in the Hildebrand–Weeks census of cusped hyperbolic 3-manifolds [Callahan et al. 99] and all 7969 complements of the 8614 hyperbolic knots and links up to 12 crossings. After replacing nonorientable manifolds (in the 5-census) by their double covers and removing duplicates, we obtained a total of 12783 orientable manifolds. Tables of all the results are available at <http://www.ms.unimelb.edu.au/~snap/>.

In the process it was necessary to identify the arithmetic manifolds and arrange them into their own separate commensurability classes. Criteria for arithmeticity and for commensurability of arithmetic manifolds are given in [Coulson et al. 00] and [Maclachlan and Reid 03]. In fact, a cusped manifold is arithmetic if and only if its invariant trace field is imaginary quadratic and it has integer traces. Cusped arithmetic manifolds are commensurable if and only if they have the same invariant trace field. They are therefore classified by the discriminant d of the invariant trace field $\mathbb{Q}(\sqrt{d})$.

There were 142 arithmetic manifolds in six commensurability classes. Table 1 lists some of the arithmetic manifolds found with discriminant -3 , -4 , or -7 , and all with discriminant -8 , -11 , or -15 .

⁵Alan Reid informs us that Nathan Dunfield has found another example: the 15-crossing knot 15n132539.

d	Manifolds
-3	$m000, m002, m003, m004 = 4a1, m025, m203 = 6a5 = 11n318, s118, s961, 8a39, 10a280, 10n130, 10n143, 10n155, 11n539, 12a3285, \dots$
-4	$m001, m124 = 8n10 = 10n139, m125, m126, m127, m128, m129, s859, v1858, 8n8 = 9n34 = 10n112, 9n36, 10a242, 10n74, 11n545, \dots$
-7	$m009, m010, s772, s773, s774, s775, s776 = 6a8 = 8n7, 8a25, 8a37, 10n73 = 12n968, 10n113, 11n498, 12n3068, 12n3078, 12n3093, \dots$
-8	$v2787, v2788, v2789, 9a73, 9a74, 12a3292, 12a3296, 12n2625 = 12n2630, 12n2972, 12n3088, 12n3098 = 12n3099$
-11	$12a2126, 12a2961, 12a3039, 12a3230, 12a3295$
-15	$12a3169, 12a3273, 12a3284, 12a3300, 12a3307, 12a3308$

TABLE 1. Selected arithmetic manifolds.

Manifold	s.g.	c.d.	c.vol, c.cl	Field	nc, cnc	Cusp Density
10a297	20	40	0.365076519	4, 1025(2)	5, 1	0.642619992
10a291	4	40	0.365076519	4, 1025(2)	4, 1	0.642619992
10a277	4	40	0.365076519	4, 1025(2)	3, 1	0.642619992
12n2492	4	40	0.365076519	4, 1025(2)	3, 1	0.642619992
12n2899	4	40	0.365076519	4, 1025(2)	4, 1	0.642619992
12n1189	2	2	7.175483613, 0	7, -76154488(-2)	2, 2	0.589830477
12n1190	2	2	7.175483613, 1	7, -76154488(-2)	2, 2	0.589830477
12n1481	2	2	7.175483613, 2	7, -76154488(-2)	2, 2	0.589830477
12n2348	2	2	7.175483613, 3	7, -76154488(-2)	3, 3	0.644747497
12n2580	2	2	7.175483613, 4	7, -76154488(-2)	3, 3	0.631898787
m045	4	4	0.818967911	3, -107(-2)	1, 1	0.608307263
m046	4	4	0.818967911	3, -107(-2)	1, 1	0.608307263
v3379	8	8	0.818967911	3, -107(-2)	2, 1	0.608307263
v3383	4	8	0.818967911	3, -107(-2)	3, 1	0.608307263
v3384	8	8	0.818967911	3, -107(-2)	2, 1	0.608307263
12a1743	8	16	0.818967911	3, -107(-2)	2, 1	0.608307263
v3376	4	4	1.637935822, 0	3, -107(-2)	2, 1	0.690189995
v3377	4	4	1.637935822, 0	3, -107(2)	1, 1	0.690189995
v3378	4	4	1.637935822, 0	3, -107(2)	1, 1	0.690189995
12a2937	8	8	1.637935822, 1	6, -1225043(1)	3, 2	0.608307263
9a94	4	24	0.575553268	4, 144(1)	3, 2	0.844133714

TABLE 2. Selected nonarithmetic manifolds.

Degree	Discriminant	Signature	Minimum Polynomial
4	1025	0, 2	$x^4 - x^3 + 3x^2 - 2x + 4$
7	-76154488	1, 3	$x^7 - 2x^6 + 3x^5 - 5x^4 + x^3 - 8x^2 - 2x - 4$
3	-107	1, 1	$x^3 - x^2 + 3x - 2$
6	-1225043	0, 3	$x^6 - 2x^5 - 2x^3 + 30x^2 - 52x + 29$
4	144	0, 2	$x^4 - x^2 + 1$

TABLE 3. Fields in Table 2.

The naming of manifolds in the tables is as follows. Manifolds whose names begin with m , s , or v belong to the 5-, 6-, or 7-tetrahedra census of cusped manifolds respectively. The rest are knot and link complements in the form $\langle \text{number of crossings} \rangle \langle \text{alternating or nonalternating} \rangle \langle \text{index in table} \rangle$. The link tables were provided by Morwen Thistlethwaite and are included with current versions of Snap and Tube.

The remaining 12641 manifolds are nonarithmetic, falling into 11709 commensurability classes. A few of them are shown in Table 2. Column headings are as follows: s.g. is the order of the symmetry group; c.d. is the degree of the manifold over its commensurator quotient; c.vol is the volume of the commensurator quotient, optionally followed by c.cl, commensurability class numbered from 0, when incommensurable manifolds are listed

name	Dowker code	name	Dowker code
4a1	dadbcd	12a2937	lcbeechjklaiiefgbd
6a5	fbccdefacb	12a2961	lcbfdcfahibekdlgj
6a8	fcbbbceafb	12a3039	lcbhbcfaikbjelgdh
8a25	hbbfcfegahdb	12a3169	lccfcdfhcjakblieg
8a37	hcbcccfghadeb	12a3230	ldbccdcfagiebkdllhj
8a39	hdbbbbceagbhdf	12a3273	ldbdbbceagbidkflhj
8n7	hcbcccDFgHEaB	12a3284	ldbfbbbeaibkjldgfh
8n8	hcbcccDFgHEab	12a3285	ldbfbbcfaikbjlehgd
8n10	hdbbbbceagBHDF	12a3292	ldccccdegjhkflacib
9a73	ibcdfhfagbice	12a3295	ldccccdgjakhbflice
9a74	ibcdfhfaiabecg	12a3296	ldccccdgjhkbelcfia
9a94	icbecdhfiabeg	12a3300	lebbdbbceagbidkflhj
9n34	icbdccceagBHDF	12a3307	lebcbbcfaikbljdheg
9n36	icbecdHfiabeG	12a3308	lfbbbbbbceagbidkflhj
10a242	jbeefghjiaecbd	12n968	lbbjcDFhIJLAbKEG
10a277	jcbfbceagbidjfh	12n1189	lbbjceaHbKiDjgLF
10a280	jcccddeghjibcfa	12n1190	lbbjceaHbKiDjgLF
10a291	jdbbbceagbidjfh	12n1481	lbbjchaEGJDbkFli
10a297	jebbbbceagbidjfh	12n1848	lbcidfIagIcKBHJ
10n73	jbbhcEFihGJAdB	12n2348	lcbbhceahbkiDjgLF
10n74	jbbhcfaHGIdJE	12n2492	lcbdfceagbiDJKfLH
10n112	jcbcecdFgIHabJE	12n2580	lcbeceahEGJDbkFli
10n113	jcbcecFaHJIBDGE	12n2625	lcbfdcfaiJbKLEDHG
10n130	jcbfbceagbiDjfh	12n2630	lcbfdcfaiJbKLHDEG
10n139	jdbbbceagbiDjfh	12n2899	ldbbbfceagbiDjgLF
10n143	jdbbbceagbiDjfh	12n2972	ldbccdcfaIJbKLEDHG
10n155	jebbbbceagbiDjfh	12n3068	ldbdbbceagbiDkFLHJ
11n318	kbchdEfcHiJKaBG	12n3078	ldbfbbbeaIbKJLDGFH
11n498	kcbcfcfahJbIkeGD	12n3088	ldccccdeGacJBk1Fih
11n539	kcbcedceahbIJDGKF	12n3093	ldccccdeGHijKlAFBc
11n545	kcbfccdIfJabKGEH	12n3098	ldccccdeGHijKlAFBc
12a1743	lbbjchfjialkedbg	12n3099	ldccccdeGHijKlAFBc
12a2126	lbcidgjlhafkaie		

TABLE 4. Dowker–Thistlethwaite codes of all knots and links mentioned in this paper.

with the same commensurator volume. Thus manifolds are commensurable if and only if they have the same entry in this column. The invariant trace field is described by its degree, discriminant, and a number specifying which root of the minimum polynomial generates it (with sign corresponding to choice of complex conjugate); nc, cnc gives the number of cusps in the manifold and the number of cusps in the commensurator quotient; cusp density is computed using equal-area cusps in the commensurator quotient.

The first group of manifolds have the smallest commensurator volume found (among nonarithmetic mani-

folds) and are the link complements that appeared in Section 2.1. They all have “hidden symmetries,” i.e., commensurabilities not arising from the symmetry group of the manifold. The total number of nonarithmetic manifolds having hidden symmetries is 148. The next group of manifolds shows that incommensurable manifolds are not always distinguished by cusp density. The third group of manifolds includes manifolds whose classes are distinguished by invariant trace field but not by cusp density, and manifolds with the same commensurator quotient volume but different invariant trace fields. One should not get the impression that cusp density is a poor invari-

ant: in fact, among the 11278 cusp densities found, only 417 grouped together incommensurable manifolds. The final line gives data for the nonarithmetic manifold with highest cusp density found. (The maximum possible cusp density is $0.853276\dots$, which occurs for the figure-eight knot complement.) More details of the fields occurring above are listed in Table 3.

9. APPENDIX

The indexing system used for knots and links, here and in Snap, may still be subject to change. This is due to the difficulty of determining whether two nonhyperbolic links are equivalent and the consequent possibility that duplicates will later be discovered and removed from the tables.

For this reason we provide Table 4, giving the Dowker–Thistlethwaite codes of all the knots and links that we refer to. Since this code is not well known for links, we describe here how this works.

Firstly, there is the trivial matter of passing between the alphabetic codes used by Snap and their numerical forms. The Dowker code for a link with n crossings and k components is a permutation of the even integers $2, \dots, 2n$, with possible sign changes, bracketed into k subsequences, e.g., $(6, -8) (-10, 14, -12, -16, -2, 4)$. To express this alphabetically, we encode n and k as the first two letters using $a = 1, b = 2$, etc. Then follow k letters giving the lengths of the bracketed subsequences. Finally, there are n letters giving the sequence of even integers using $a = 2, b = 4$, etc. and $A = -2, B = -4$, etc. The alphabetic code for the above example is thus `hbbfcDEgFHAb`.

To go from a link diagram to its Dowker code, proceed as follows. Traverse each component, numbering the crossings, starting with 1 on an overcrossing. When the first component is done, continue with consecutive numbers on the next component. Each crossing will receive two numbers. We can number the crossings in such a way that every crossing gets one even and one odd number. (This follows easily from the fact that we can two-color the plane containing a link diagram.) Negate any even number that labels an overcrossing (for an alternating link there will not be any). For each odd number $1, 3, \dots, 2n - 1$ write down the corresponding even number: this gives a sequence of n even numbers. A component with $2j$ crossings will have j odd numbers on it, so there will be j corresponding numbers for it in the code; bracket together the numbers for each component. See, for example, Figure 39.




FIGURE 39. Complete the numbering indicated on the above two-component link to obtain the Dowker–Thistlethwaite code $(6, -8) (-10, 14, -12, -16, -2, 4)$.

The reverse procedure, going from a code to a link diagram, is a little bit trickier, but essentially the same as the procedure for knots, described in [Adams 94]. We draw the first component as a knot, with extra as-yet-unconnected crossings on it. When adding further components, we will encounter crossings with other link components.

ACKNOWLEDGMENTS

We thank Ian Agol for pointing out a simplification to our method of determining the commensurability of Euclidean tori, Gaven Martin for information on current volume bounds, and Walter Neumann for several interesting discussions on this work. We also thank Alan Reid, Genevieve Walsh, and the referee for their helpful comments on the paper.

This work was partially supported by grants from the Australian Research Council.

REFERENCES

- [Adams 92] C. Adams. “Noncompact Hyperbolic 3-Orbifolds of Small Volume.” In *Topology ’90 (Columbus, OH, 1990)*, pp. 1–15, Ohio State Univ. Math. Res. Inst. Publ. 1. Berlin: De Gruyter, 1992.
- [Adams 94] C. Adams. *The Knot Book*. New York: W. H. Freeman and Company, 1994.
- [Aitchison and Rubinstein 92] I. R. Aitchison and J. H. Rubinstein. “Combinatorial Cubings, Cusps and the Dodecahedral Knots.” In *Topology ’90 (Columbus, OH, 1990)*, pp. 273–310, Ohio State Univ. Math. Res. Inst. Publ. 1. Berlin: De Gruyter, 1992.
- [Akiyoshi 01] H. Akiyoshi. “Finiteness of Polyhedral Decompositions of Cusped Hyperbolic Manifolds Obtained by the Epstein–Penner’s Method.” *Proc. Amer. Math. Soc.* 129:8 (2001), 2431–2439.
- [Akiyoshi et al. 07] H. Akiyoshi, M. Sakuma, M. Wada, and Y. Yamashita. *Punctured Torus Groups and 2-Bridge Knot Groups (I)*, Lecture Notes in Mathematics, 1909. New York: Springer, 2007.
- [Borel 81] A. Borel. “Commensurability Classes and Volumes of Hyperbolic 3-Manifolds.” *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)* 8:1 (1981), 1–33.

- [Bowditch et al. 95] B. H. Bowditch, C. Maclachlan, and A. W. Reid. “Arithmetic Hyperbolic Surface Bundles.” *Math. Ann.* 302:1 (1995), 31–60.
- [Button 05] J. O. Button. “Fibred and Virtually Fibred Hyperbolic 3-Manifolds in the Censuses.” *Experiment. Math.* 14:2 (2005), 231–255.
- [Callahan et al. 99] P. J. Callahan, M. V. Hildebrand, and J. R. Weeks. “A Census of Cusped Hyperbolic 3-Manifolds” (with microfiche supplement). *Math. Comp.* 68:225 (1999), 321–332.
- [Coulson et al. 00] D. Coulson, O. Goodman, C. Hodgson, and W. Neumann. “Computing Arithmetic Invariants of 3-Manifolds.” *Experiment. Math.* 9 (2000), 127–152.
- [Epstein and Penner 88] D. B. A. Epstein and R. C. Penner. “Euclidean Decompositions of Noncompact Hyperbolic Manifolds.” *J. Differential Geom.* 27:1 (1988), 67–80.
- [Floyd and Hatcher 82] W. Floyd and A. Hatcher. “Incompressible Surfaces in Punctured-Torus Bundles.” *Topology Appl.* 13 (1982), 263–282.
- [González-Acuña and Whitten 92] F. González-Acuña and W. C. Whitten. “Imbeddings of Three-Manifold Groups.” *Mem. Amer. Math. Soc.* 99:474 (1992).
- [Guéritaud 06a] F. Guéritaud (with an appendix by D. Futer). “On Canonical Triangulations of Once-Punctured Torus Bundles and Two-Bridge Link Complements.” *Geometry and Topology* 10 (2006), 1239–1284.
- [Guéritaud 06b] F. Guéritaud. “Géométrie hyperbolique effective et triangulations idéals canoniques en dimension 3. PhD thesis, Univ. Paris-sud, Orsay, 2006.
- [Hildebrand and Weeks 89] M. Hildebrand and J. Weeks. “A Computer Generated Census of Cusped Hyperbolic 3-Manifolds.” In *Computers and Mathematics (Cambridge, MA, 1989)*, pp. 53–59. New York: Springer, 1989.
- [Hoste and Thistlethwaite 99] J. Hoste and M. Thistlethwaite. “Knotscape.” Available online (<http://www.math.utk.edu/~morwen/knotscape.html>), 1999.
- [Lackenby 83] M. Lackenby. “The Canonical Decomposition of Once-Punctured Torus Bundles.” *Comment. Math. Helv.* 78 (2003) 363–384.
- [Macbeath 88] A. M. Macbeath. “Commensurability of Co-compact Three-Dimensional Hyperbolic Groups.” *Duke Math. J.* 50:4 (1983), 1245–1253. Erratum: *Duke Math. J.* 56:1 (1988), 219.
- [Maclachlan and Reid 03] C. Maclachlan and A. Reid. *The Arithmetic of Hyperbolic 3-Manifolds*. New York: Springer-Verlag, 2003.
- [Margulis 91] G. A. Margulis. *Discrete Subgroups of Semisimple Lie Groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete, 17. Berlin: Springer-Verlag, 1991.
- [Marshall and Martin 08] T. H. Marshall and G. J. Martin. “Minimal Co-volume Hyperbolic Lattices II. Preprint, 2008.
- [Meyerhoff 85] R. Meyerhoff. “The Cusped Hyperbolic 3-Orbifold of Minimal Volume.” *Bull. Amer. Math. Soc.* 13 (1985), 154–156.
- [Neumann and Reid 92a] W. D. Neumann and A. W. Reid. “Arithmetic of Hyperbolic Manifolds.” In *Topology '90 (Columbus, OH, 1990)*, pp. 273–310, Ohio State Univ. Math. Res. Inst. Publ. 1. Berlin: De Gruyter, 1992.
- [Neumann and Reid 92b] W. D. Neumann and A. W. Reid. “Notes on Adams’ Small Volume Orbifolds.” In *Topology '90 (Columbus, OH, 1990)*, pp. 311–314, Ohio State Univ. Math. Res. Inst. Publ. 1. Berlin: De Gruyter, 1992.
- [Reid 90] A. W. Reid. “A Note on Trace-Fields of Kleinian Groups.” *Bull. London Math. Soc.* 22:4 (1990), 349–352.
- [Sakuma and Weeks 95a] M. Sakuma and J. Weeks. “Examples of Canonical Decompositions of Hyperbolic Link Complements.” *Japan. J. Math. (N.S.)* 21:2 (1995), 393–439.
- [Sakuma and Weeks 95b] M. Sakuma and J. Weeks. “The Generalized Tilt Formula.” *Geom. Dedicata* 55:2 (1995), 115–123.
- [Thurston 78] W. P. Thurston. *The Geometry and Topology of Three-Manifolds*. Princeton: Princeton University Math. Dept., 1978. Also available online (<http://msri.org/publications/books/gt3m>).
- [Weeks 93] J. R. Weeks. “Convex Hulls and Isometries of Cusped Hyperbolic 3-Manifolds.” *Topology Appl.* 52:2 (1993), 127–149.
- [Zimmer 84] R. Zimmer. *Ergodic Theory and Semi-simple Lie Groups*. Boston: Birkhäuser, 1984.

Oliver Goodman, Department of Mathematics and Statistics, University of Melbourne, Parkville, Victoria 3010, Australia (oag@optusnet.com.au)

Damian Heard, Department of Mathematics and Statistics, University of Melbourne, Parkville, Victoria 3010, Australia (damian.heard@gmail.com)

Craig Hodgson, Department of Mathematics and Statistics, University of Melbourne, Parkville, Victoria 3010, Australia (cdh@ms.unimelb.edu.au)

Received January 29, 2007; accepted October 1, 2007.