

Infinitely many virtual geometric triangulations

David Futer¹ | Emily Hamilton² | Neil R. Hoffman³

¹Department of Mathematics, Temple University, Philadelphia, Pennsylvania, USA

²Department of Mathematics, California Polytechnic State University, San Luis Obispo, California, USA

³Department of Mathematics, Oklahoma State University, Stillwater, Oklahoma, USA

Correspondence

David Futer, Department of Mathematics, Temple University, Philadelphia, PA 19122, USA.

Email: dfuter@temple.edu

Funding information

NSF, Grant/Award Number:

DMS-1907708; Simons Foundation,

Grant/Award Number: #524123

Abstract

We prove that every cusped hyperbolic 3-manifold has a finite cover admitting infinitely many geometric ideal triangulations. Furthermore, every long Dehn filling of one cusp in this cover admits infinitely many geometric ideal triangulations. This cover is constructed in several stages, using results about separability of peripheral subgroups and their double cosets, in addition to a new conjugacy separability theorem that may be of independent interest. The infinite sequence of geometric triangulations is supported in a geometric submanifold associated to one cusp, and can be organized into an infinite trivalent tree of Pachner moves.

MSC 2020

57K32, 20F65 (primary), 20E26, 57M10, 57R05 (secondary)

1 | INTRODUCTION

A hyperbolic 3-manifold M is called *cusped* if it is noncompact and has finite volume. Every cusped 3-manifold M admits a topological ideal triangulation: that is, a decomposition into finitely many tetrahedra whose vertices have been removed, with faces identified in pairs by affine maps. A *geometric ideal triangulation* is a stronger notion, where each tetrahedron is isometric to the convex hull of 4 non-coplanar points in \mathbb{H}^3 , and where the tetrahedra are glued by isometry to give the complete hyperbolic metric on M . See Definition 2.2 for precise details. The focus of this paper is on geometric triangulations.

The presence of a geometric triangulation makes the geometry of M much more accessible to both practical and theoretical study. On the practical side, geometric ideal triangulations are central to the workings of the computer program SnapPy [7] that computes hyperbolic structures

and rigorously verifies their geometric properties [20]. On the theoretical side, Thurston's original proof of the hyperbolic Dehn filling theorem implicitly assumed that the 3-manifold at hand admits a geometric triangulation [29]. Similarly, Neumann and Zagier's work on volume assumes that the complement of some closed geodesic in M admits a geometric triangulation [23].

Despite the importance of geometric triangulations, the first part of the following conjecture has been open for multiple decades.

Conjecture 1.1. *Let M be a (finite volume) cusped hyperbolic 3-manifold. Then*

- (1) (Folklore) M admits at least one geometric ideal triangulation.
- (2) M admits infinitely many geometric ideal triangulations.

In the 1980s, Conjecture 1.1.(1) was widely believed to follow from the work of Epstein and Penner [10]. More precisely, the community believed that a geometric ideal polyhedral decomposition of M can always be subdivided to give a geometric ideal triangulation. It took time to realize that a naive refinement of the Epstein–Penner cell decomposition does not suffice; see the discussion of coning in Section 2 for a description of some of the challenges. To our knowledge, the first record of Conjecture 1.1.(1) in the literature is by Petronio [25, Conjecture 2.3], in 2000. See also Petronio and Porti for a useful account of the history [26].

By contrast, Conjecture 1.1.(2) is new. We propose this tantalizing strengthening of the original conjecture because searching for infinite and flexible sequences of geometric triangulations might provide a pathway to finding at least one. Indeed, our main result can be interpreted as a proof of concept that such a pathway exists in the context of finite covers and Dehn filling.

Passing to covers makes both parts of Conjecture 1.1 more amenable. Toward Part (1) of the Conjecture, Luo, Schleimer and Tillmann showed that every cusped hyperbolic manifold M has a finite cover that supports a geometric triangulation [22]. We recall their proof strategy in Section 2, and incorporate several of their ideas in the proof of our theorems. Our main result, in the direction of Conjecture 1.1.(2), is the following.

Theorem 1.2. *Let M be a cusped hyperbolic 3-manifold and $A \subset M$ a horocusp. Then there is a finite cover $\widehat{M} \rightarrow M$, such that A lifts to a cusp $\widehat{A} \subset \widehat{M}$, with the following properties.*

- \widehat{M} admits infinitely many geometric ideal triangulations.
- For every sufficiently long slope s on $\partial\widehat{A}$, the Dehn filling $\widehat{M}(s)$ admits infinitely many geometric ideal triangulations.

Part of the interest of Theorem 1.2 comes from the fact that direct constructions of geometric ideal triangulations are only known in special classes of manifolds. For example, Guéritaud proved that certain well-studied triangulations of hyperbolic once-punctured torus bundles are geometric [13]. Futer extended Guéritaud's method to hyperbolic 2-bridge link complements [13, Appendix]. Guéritaud and Schleimer proved that if M is a generic multi-cusped hyperbolic manifold, then long Dehn fillings of M will admit geometric triangulations [15]. Ham and Purcell found geometric ideal triangulations of highly twisted link complements, by adapting Guéritaud and Schleimer's construction to some especially nice triangulations of fully augmented links [16].

There have also been attacks on Conjecture 1.1.(1) that have attempted to subdivide a geometric polyhedral decomposition into geometric ideal tetrahedra. Wada, Yamashita, and Yoshida [30], building on work of Yoshida [32], described a sufficient condition on the dual 1-skeleton of a polyhedral decomposition to make such a subdivision possible. Sirotkina proved that a

subdivision is always possible if each 3-cell has at most six faces [28]. Goerner proved that a subdivision is always possible if each 3-cell is a (not necessarily regular) ideal dodecahedron [12]. Champanerkar, Kofman, and Purcell have constructed interesting examples of link complements admitting a decomposition into regular ideal bipyramids, which can then be subdivided into geometric ideal tetrahedra [5, Theorem 3.5].

To our knowledge, there is only one prior paper constructing infinitely many geometric triangulations on the same hyperbolic manifold. Dadd and Duan showed that the figure-8 knot complement, which decomposes into two regular ideal tetrahedra, supports infinitely many geometric triangulations [9]. Their proof strategy is very delicate, in that it does not extend to the figure-8 sister manifold, which also decomposes into two regular ideal tetrahedra.

Given a cusped manifold M , the *topological Pachner graph* of M is the graph whose vertices are isotopy classes of (topological) ideal triangulations, with edges corresponding to 2–3 moves and their inverses. (See Definition 3.5 for the definition of a 2–3 move, and Figure 3 for an illustration.) The *geometric Pachner graph* of M is the induced subgraph whose vertices are geometric ideal triangulations. The infinitely many geometric triangulations found by Dadd and Duan [9] are organized in the form of an infinite ray in a single component of the geometric Pachner graph of the figure-8 knot complement. In a generic situation, the infinitely many geometric triangulations constructed in Theorem 1.2 contain an even greater amount of structure.

Theorem 1.3. *Let M be a cusped hyperbolic 3-manifold containing a non-rectangular cusp. Then there exists a finite cover $\hat{M} \rightarrow M$ such that the geometric Pachner graph of \hat{M} contains a subgraph homeomorphic to an infinite trivalent tree.*

The hypothesis on a non-rectangular cusp can be explained as follows. As we describe in Section 2, every non-compact end of M has the form $A \cong T \times [0, \infty)$, where T is a torus endowed with a Euclidean metric that is well defined up to similarity. We say that A is *rectangular* if the Euclidean metric on T admits a rectangular fundamental domain, and *non-rectangular* otherwise. By the work of Nimershiem [24], the Euclidean structures on cusp tori of hyperbolic 3-manifolds form a dense subset of the moduli space of $\mathcal{M}(T^2)$. Since rectangular tori represent a codimension-one slice of $\mathcal{M}(T^2)$, one can say that a generic cusped 3-manifold satisfies the hypotheses of Theorem 1.3.

The infinite trivalent tree mentioned in Theorem 1.3 can be identified with the dual 1-skeleton of the Farey graph. See Definition 3.9 and Figure 4 for a review of the Farey graph; in brief, its vertices correspond to *slopes*, or simple closed curves on a torus, and to rational numbers in \mathbb{RP}^1 . The branches of the trivalent tree of Theorem 1.3 limit to every point of \mathbb{RP}^1 . In particular, the infinite sequence of geometric triangulations that we will construct can be chosen to approach any rational or irrational foliation on a cusp torus of \hat{M} . Manifolds with rectangular cusps satisfy a slightly weaker version of Theorem 1.3; see Remark 6.12 for details.

1.1 | Proof strategy

Next, we outline the main ideas in the proofs of Theorems 1.2 and 1.3. Both proofs use the same initial setup and general strategy. Since having a non-rectangular cusp simplifies the argument considerably, Theorem 1.3 will be proved first.

Let M be a cusped hyperbolic 3-manifold. We will obtain geometric triangulations by subdividing the canonical (Epstein–Penner) polyhedral decompositions of covers of M . Section 2 reviews

the Epstein–Penner construction [10], emphasizing the way in which the canonical polyhedral decomposition \mathcal{P} depends on the choice of neighborhoods of the cusps. That section also reviews the process of subdivision via *coning* and lays out a sufficient condition (involving an order on the cusps) that ensures \mathcal{P} can be subdivided into geometric ideal tetrahedra. See Lemma 2.9, which is essentially due to Luo, Schleimer, and Tillmann [22], for details.

In Section 3, we describe a particular feature of the canonical polyhedral decomposition \mathcal{P} that occurs in the “generic” scenario when a manifold M has multiple cusps, one cusp A is chosen to be sufficiently small, and there is a unique shortest path from A to the other cusps. In this situation, Guéritaud and Schleimer [15] show that the canonical polyhedral decomposition \mathcal{P} has only one or two cells poking into this cusp A . These cells fit together to form a submanifold called a *drilled ananas* (see Definition 3.3). In Lemma 3.6, we show that a drilled ananas admits an infinite sequence of geometric ideal triangulations. When A is a non-rectangular cusp, these triangulations are arranged in a trivalent tree of 2–3 moves, as described in Theorem 1.3.

To build covers of M satisfying the above-mentioned conditions, we will need to *separate* certain subgroups and subsets of $\pi_1(M)$ from group elements that cause undesired coincidences. Section 4 reviews several key definitions and results about separability that are needed for our purposes. The strongest result that is needed for the proof of Theorem 1.3 is Theorem 4.4, due to Hamilton, Wilton, and Zalesskii [18], which provides separability of double cosets of peripheral subgroups.

With this background in hand, we can begin to construct covers. Assuming that M has a non-rectangular cusp, Section 5 produces a sequence of finite covers $\hat{M} \rightarrow \dot{M} \rightarrow M$, with increasingly strong properties. In particular, \dot{M} contains a drilled ananas, while \hat{M} has a polyhedral decomposition $\hat{\mathcal{P}}$ that can be subdivided via coning. It will follow that \hat{M} admits an infinite trivalent tree of geometric ideal triangulations, establishing Theorem 1.3.

1.2 | New separability tools

To prove Theorem 1.2, which handles hyperbolic manifolds with rectangular cusps and provides an additional conclusion about Dehn fillings, we need stronger separability tools than what was previously available in the literature. The following new result may be of independent interest. In the theorem statement, a *peripheral* subgroup of $\Gamma = \pi_1(M)$ is a subgroup coming from the inclusion of a cusp.

Theorem 1.4 (Conjugacy separation of peripheral cosets). *Let $M = \mathbb{H}^3/\Gamma$ be a cusped hyperbolic 3-manifold. Let H and K be (maximal) peripheral subgroups of Γ corresponding to distinct cusps of M . Let $g \in \Gamma$ be an element such that K is disjoint from every conjugate of gH . Then there is a homomorphism $\varphi : \Gamma \rightarrow G$, where G is a finite group, such that $\varphi(K)$ is disjoint from every conjugate of $\varphi(gH)$.*

Theorem 1.4 has the following topological interpretation. A maximal peripheral subgroup $H \subset \Gamma$ is the stabilizer of a horoball $\bar{B} \subset \mathbb{H}^3$. Given $g \in \Gamma \setminus H$, the coset gH is the set of all elements of Γ that move \bar{B} to $g\bar{B}$. Connecting these two horoballs is a geodesic arc $\tilde{\beta}$ that projects to an arc $\beta \subset M$. We wish to find a finite cover $\hat{M} \rightarrow M$ where the cusp corresponding to K lifts, and where every preimage of β connects distinct cusps. Theorem 1.4 provides such a cover, corresponding to the subgroup $\hat{\Gamma} = \varphi^{-1} \circ \varphi(K)$ that contains K but excludes every conjugate of gH .

Several precursors of Theorem 1.4 appear in the recent literature on 3-manifold groups. Given a peripheral subgroup K and a single element $g \in \Gamma$ that is disjoint from every conjugate of K , it

is straightforward to find a finite quotient that witnesses this disjointness [18, Lemma 4.5]. Given non-conjugate subgroups H and K , Chagas and Zalesskii find a finite quotient of Γ where their images are not conjugate [4]. Given a pair of non-conjugate peripheral subgroups H and K , Wilton and Zalesskii use an argument of Hamilton to construct a finite quotient $\varphi: \Gamma \rightarrow G$, such that non-trivial elements of $\varphi(H)$ and $\varphi(K)$ always lie in distinct conjugacy classes [31, Lemma 4.6].

It is worth recalling the proof of the last result. First, take a hyperbolic Dehn filling $M(s)$ corresponding to a quotient $\Gamma \rightarrow \Gamma(s)$, where K stays parabolic but non-trivial elements of H become loxodromic. In particular, the quotient of H is represented by loxodromic matrices of trace not equal to 2. Then, take a congruence quotient of the matrix group $\Gamma(s)$ in a matrix group over a finite ring (see Definition 4.6), where the traces of these loxodromic matrices can still be distinguished from 2. Our contribution to this narrative is that we achieve even stronger separability for non-conjugate parabolic subgroups H and K , separating the image of K from the image of every conjugate of gH .

The proof of Theorem 1.4 appears in Section 4, and uses a similar two-step method: first construct an appropriate Dehn filling, and then analyze the congruence quotients related to the Dehn filling. As part of this analysis, we apply tools from algebraic number theory, including a theorem of Hamilton [17, Corollary 2.5] (restated below as Proposition 4.9), to control the traces of an entire coset gH .

Using the separability Theorem 1.4, we prove Theorem 1.2 in Section 6. If \mathring{M} is a cover of M containing a drilled ananas \mathring{N} , as above, we use the topological interpretation of Theorem 1.4 to construct two additional covers $\widetilde{M} \rightarrow \overline{M} \rightarrow \mathring{M}$ where the ananas \mathring{N} lifts but most edges of the polyhedral decomposition $\widetilde{\mathcal{P}}$ connect distinct cusps. In particular, \widetilde{M} has a drilled ananas \widetilde{N} and a polyhedral decomposition $\widetilde{\mathcal{P}}$ that can be subdivided into ideal tetrahedra via coning, which implies infinitely many geometric triangulations. Then, we build a cover $\widehat{M} \rightarrow \widetilde{M}$ where the drilled ananas \widetilde{N} has two distinct lifts. One of these lifts supports infinitely many geometric triangulations, while the other gets filled to obtain the Dehn filling conclusion of the theorem. The opening paragraphs of Section 6 outline this construction in much greater detail.

2 | TRIANGULATIONS AND POLYHEDRAL DECOMPOSITIONS

This section reviews some standard definitions about hyperbolic manifolds and their polyhedral decompositions and triangulations. Then, it proves Lemma 2.9 and Corollary 2.10, which will be our main ways to obtain a geometric triangulation from a polyhedral decomposition.

For the remainder of this paper, the symbol M denotes a cusped, orientable hyperbolic 3-manifold. (Since Theorems 1.2 and 1.3 construct finite covers, no generality is lost in assuming that M is orientable.) We will use \widetilde{M} to denote the universal cover of M , which is isometric to \mathbb{H}^3 . Other decorations, such as \widehat{M} and \overline{M} , denote finite-sheeted covers of M .

Definition 2.1. Let M be a cusped hyperbolic manifold, and let $\widehat{f}: \widehat{M} \rightarrow M$ be a finite cover. Let $A \subset M$ be an embedded submanifold. We say that A lifts to \widehat{M} if the inclusion map $\iota: A \hookrightarrow M$ lifts to an inclusion $\widehat{\iota}: A \hookrightarrow \widehat{M}$. In this case, the image $\widehat{A} = \widehat{\iota}(A)$ is called a *lift* of A . The lift \widehat{A} forms only one component of $\widehat{f}^{-1}(A)$, and covers A with degree one.

We remark that a lift is distinct from a path-lift. If $\gamma \subset M$ is a (parametrized) closed curve based at x , then γ always has a path-lift $\widehat{\gamma}$ starting at any preimage $\widehat{x} \in \widehat{f}^{-1}(x)$. This path-lift is only a lift if it returns to \widehat{x} .

Definition 2.2. A *geometric ideal polyhedron* P is the convex hull in \mathbb{H}^3 of $n \geq 4$ non-coplanar points in $\partial\mathbb{H}^3$. If $n = 4$, the polyhedron is called a *geometric ideal tetrahedron*, and its isometry class is determined by the cross-ratio of its 4 vertices. The polyhedron P and its boundary ∂P inherit an orientation from the embedding $P \hookrightarrow \mathbb{H}^3$.

An ideal polyhedron P is called an *ideal pyramid* if P contains an ideal vertex v (called the *apex*) and a unique face F not incident to v (called the *base*). It follows that every edge of P either belongs to ∂F (in which case it is called a *base edge*) or connects v to a vertex of F (in which case it is called a *lateral edge*). Every pyramid is either an ideal tetrahedron, or has a unique choice of apex and base.

A *geometric ideal polyhedral decomposition* \mathcal{P} is a decomposition of M into geometric ideal polyhedra, glued together by orientation-reversing isometries along their boundary faces. The cusps of M are therefore in bijection with the equivalence classes of ideal vertices in \mathcal{P} . If all the cells are ideal tetrahedra, the decomposition \mathcal{P} is called a *geometric ideal triangulation*, and denoted as \mathcal{T} . The preimage of \mathcal{P} in a cover $\tilde{M} \rightarrow M$ is denoted as $\tilde{\mathcal{P}}$, and similarly for other decorations.

Convention 2.3. All triangulations and polyhedral decompositions described below are presumed to be geometric, unless specified otherwise. While there is a rich theory of topological ideal triangulations of 3-manifolds, sometimes endowed with extra data, our focus in this paper is on geometry.

Definition 2.4. A (*closed*) *horocusp* A is the quotient of a closed horoball in \mathbb{H}^3 by a discrete group G of parabolic isometries, where $G \cong \mathbb{Z} \times \mathbb{Z}$. Topologically, A is homeomorphic to $T^2 \times [0, \infty)$ and ∂A is isometric to a flat torus. The interior of A is called an *open horocusp*.

If $M = \mathbb{H}^3/\Gamma$ is a finite-volume hyperbolic 3-manifold, an *open horocusp in M* is an embedded noncompact end that is isometric to an open horocusp. A (*closed*) *horocusp in M* is the closure of an open horocusp in M . In particular, a horocusp $A \subset M$ is homeomorphic to $T^2 \times [0, \infty)$ with a finite number of points of tangency on $T^2 \times \{0\}$ identified in pairs.

A *horocusp collection* in M is a union of closed horocusps A_1, \dots, A_n containing all the noncompact ends of M , such that the interiors of the A_i are pairwise disjoint.

For a hyperbolic 3-manifold $M = \mathbb{H}^3/\Gamma$, we typically work with $\tilde{M} = \mathbb{H}^3$ in the upper half-space model. The preimage of a horocusp collection in M is a collection of (closed) horoballs in \mathbb{H}^3 with disjoint interiors, called a *packing*. When we mention a horoball \tilde{A} in \mathbb{H}^3 in this context, we implicitly assume that \tilde{A} is one of the horoballs in the packing, meaning that \tilde{A} covers one of the specified horocusps of M . We further conjugate Γ in $\text{Isom}(\mathbb{H}^3) \cong \text{PSL}(2, \mathbb{C})$ so that ∞ is a parabolic fixed point of Γ , which means that a horoball \tilde{A} about ∞ occurs in the packing. All other horoballs in the packing are tangent to points of \mathbb{C} . The packing horoballs with largest Euclidean diameter (equivalently, the horoballs closest to \tilde{A}) are called *full-sized*.

Definition 2.5. Let M be a cusped hyperbolic 3-manifold with horocusp collection A_1, \dots, A_n . An *orthogeodesic* is an immersed geodesic segment γ that begins at ∂A_i and ends at ∂A_j , such that γ is orthogonal to ∂A_i and ∂A_j at the respective endpoints. The case $A_i = A_j$ is permitted. If A_i is tangent to A_j , then a point of tangency is considered an orthogeodesic of length 0. We note that an orthogeodesic is necessarily the shortest path in its homotopy class.

In a similar fashion, an orthogeodesic in \mathbb{H}^3 is the shortest path between a pair of disjoint horoballs \tilde{A}, \tilde{A}' . This path is necessarily a geodesic segment that is orthogonal to $\partial\tilde{A}$ and $\partial\tilde{A}'$.




FIGURE 1 The construction of a canonical polyhedral decomposition in a cusped hyperbolic surface. We have a horoball packing of \mathbb{H}^2 (black) and the universal cover $\tilde{\Sigma}$ of the cut locus Σ (red). Dual to $\tilde{\Sigma}$ is a canonical triangulation $\tilde{\mathcal{P}}$ (blue). Every vertex $\tilde{v} \in \tilde{\Sigma}$ is the center of a ball (gray) tangent to a maximal collection of horoballs that contain the vertices of the cell of $\tilde{\mathcal{P}}$ dual to \tilde{v} .

A collection of horocusp in a hyperbolic manifold M determines a canonical decomposition of M into polyhedra, as follows.

Definition 2.6. Let M be a cusped hyperbolic 3-manifold, endowed with a horocusp collection A_1, \dots, A_n . The *Ford–Voronoi domain* $\mathcal{F} \subset M$ consists of all points of M that have a unique shortest path to the union of the A_i . The complement $\Sigma = M \setminus \mathcal{F}$, called the *cut locus*, is a 2-dimensional cell complex consisting of finitely many totally geodesic polygons. The combinatorial dual of Σ is denoted as \mathcal{P} and called the *canonical polyhedral decomposition* determined by (M, A_1, \dots, A_n) . This polyhedral decomposition has one geodesic edge for each polygonal face of Σ , one totally geodesic 2-cell for each edge of Σ , and one 3-cell for each vertex of Σ . The edges of \mathcal{P} are bi-infinite extensions of orthogeodesics between the cusps. See Figure 1 for a 2-dimensional example.

The top-dimensional cells of \mathcal{P} can be characterized as follows. By construction, every 3-cell $P \subset \mathcal{P}$ is dual to a vertex $v \in \Sigma$. There is a metric ball D centered at v , which is tangent to some number of horocusp (corresponding to the ideal vertices of P), and disjoint from their interiors. Furthermore, the collection of cusps tangent to D is maximal with respect to inclusion.

In the context of closed surfaces, the construction of the canonical decomposition \mathcal{P} dates back to the work of Voronoi and Delaunay in the early 20th century. Epstein and Penner [10] gave a characterization of \mathcal{P} using convexity in the hyperboloid model of \mathbb{H}^3 . As a consequence, \mathcal{P} is sometimes called the Delaunay or Epstein–Penner decomposition of M .

The canonical polyhedral decomposition \mathcal{P} determined by a choice of horocusp is always geometric. Thus every cusped hyperbolic 3-manifold admits a geometric polyhedral decomposition. Furthermore, one may attempt to subdivide the polyhedra of \mathcal{P} into tetrahedra by *coning*.

Definition 2.7. Let P be a (geometric) ideal polyhedron, and let v be an ideal vertex of P . The *coning of P from v* is the decomposition of P into (geometric) ideal pyramids whose apex is v and whose bases are the polygonal faces of P not incident to v . If P is an ideal pyramid and w is a vertex of the base of P , the coning of P from w results in ideal tetrahedra, because every face of P not incident to w is an ideal triangle.

Definition 2.8. Let \mathcal{P} be a (geometric) ideal polyhedral decomposition of M , and let V denote the set of cusps of M . Consider a strict partial order $<$ on V . Observe that $<$ imposes a (strict) partial order on the vertices of any polyhedron $P \subset \mathcal{P}$, because vertices of P map to cusps of M . The *coning of P* induced by $<$ is the following subdivision: if P has a unique $<$ -minimal vertex v , then P is coned from v ; otherwise, P is not subdivided at all.

The *iterated coning of P* induced by $<$ is the following two-step procedure. First, cone P from its unique $<$ -smallest vertex (if such a vertex exists), which either leaves P unchanged or decomposes it into pyramids. Second, cone each pyramid from the unique $<$ -smallest vertex of its base (if such a vertex exists).

Lemma 2.9. Let M be a cusped hyperbolic 3-manifold, and $<$ a strict partial order on the set of cusps of M . Let \mathcal{P} be a (geometric) ideal decomposition of M , with the property that every polyhedron $P \subset \mathcal{P}$ has a unique $<$ -minimal vertex v_P . Then the iterated coning of \mathcal{P} induced by $<$ produces a well-defined subdivision of \mathcal{P} into geometric ideal pyramids.

Furthermore, if $<$ gives a total order of the vertices of every polyhedron, then the iterated coning of \mathcal{P} produces a geometric ideal triangulation.

Proof. Suppose that P and P' are polyhedra of \mathcal{P} that are identified along a face F . We need to check that the iterated coning of P and P' induces the same subdivision of the face F . We show this by considering two cases.

Case 1: F does not have a unique $<$ -minimal vertex. In this case, we claim that F will not be subdivided at all. For, polyhedron P will be coned from its unique minimal vertex v_P , which is not contained in F by hypothesis. This produces a collection of pyramids, with F a base of one of the pyramids. Since F does not have a unique minimal vertex, the second stage of iterated coning does not subdivide F at all. An identical argument applies to P' , proving the claim.

Case 2: F has a unique $<$ -minimal vertex w . In this case, we claim that F will be subdivided by coning from w . If $w = v_P$ is minimal in all of P , then P will be subdivided into pyramids by coning from w , hence F will be also. Otherwise, if $w \neq v_P$, then P will be subdivided into pyramids by coning from v_P , and F will be the base of one of these pyramids. At the second stage of the iterated coning, the pyramid in P containing F will be coned from the minimal vertex of F , namely, w . An identical argument applies to P' , proving the claim.

Finally, observe that if $<$ gives a total order of the vertices of each cell, then we must be in Case 2: every face F has a minimal vertex. Thus every face is subdivided into triangles, and every pyramid is subdivided into geometric ideal tetrahedra. \square

The “furthermore” statement of Lemma 2.9 was previously observed by Luo, Schleimer, and Tillmann [22, Lemma 7]. They also used the separability of peripheral subgroups (Proposition 4.3) to show that every cusped 3-manifold M with a geometric polyhedral decomposition \mathcal{P} has a finite cover \hat{M} such that any order on the cusps of \hat{M} imposes a total order on the vertices of each polyhedron of $\hat{\mathcal{P}}$. Compare Lemma 5.4 below. Consequently, \hat{M} has a geometric ideal triangulation.

In fact, a total order is not necessary to produce a geometric triangulation.

Corollary 2.10. Let M be a cusped hyperbolic 3-manifold, and $<$ a strict partial order on the set of cusps of M . Let \mathcal{P} be a (geometric) ideal decomposition of M , with the property that every polyhedron $P \subset \mathcal{P}$ has a unique $<$ -minimal vertex v_P . Let \mathcal{P}' be the pyramidal refinement of \mathcal{P} guaranteed by

Lemma 2.9. Then every choice of diagonals in the non-triangular faces of \mathcal{P}' leads to a decomposition of \mathcal{P}' into geometric ideal tetrahedra.

Proof. Following Lemma 2.9, let \mathcal{P}' be the subdivision into ideal pyramids coming from the iterated coning of \mathcal{P} induced by \prec . Then every non-triangular 2-cell $F \subset \mathcal{P}'$ must be the base of exactly two pyramids. Thus the non-tetrahedral pyramids of \mathcal{P}' are glued in pairs, with each pair forming a bipyramid that is joined to other cells along ideal triangles only. This means that we have complete freedom to choose diagonals of every non-triangular 2-cell F , subdividing the two pyramids adjacent to F into tetrahedra, without impacting the choices anywhere else in the manifold. Every such choice produces a subdivision of \mathcal{P}' into geometric ideal tetrahedra. \square

3 | AN INFINITE TREE OF TRIANGULATIONS

In the last section, we described a construction of Luo, Schleimer, and Tillmann for decomposing an ideal polyhedral decomposition \mathcal{P} into a geometric ideal triangulation. Having one geometric triangulation is clearly a prerequisite to having infinitely many. In this section, we describe a particular geometric feature called a *drilled ananas* (see Definition 3.3) that admits an infinite sequence of geometric triangulations. By embedding a drilled ananas inside a triangulation of a cusped manifold M , we obtain an infinite sequence of ideal triangulations of M . See Lemma 3.6 for the construction of an infinite sequence of triangulations, and Proposition 3.10 for a more refined description of an infinite trivalent tree of triangulations.

Consider a polyhedral decomposition \mathcal{P} of M . Let $A \subset M$ be a horocusp, chosen small enough that for every polyhedron $P \subset \mathcal{P}$, the intersection $P \cap A$ consists of neighborhoods of ideal vertices. Then \mathcal{P} induces a decomposition of the torus ∂A into Euclidean polygons, which truncate the ideal vertices of polyhedra of \mathcal{P} . We call this decomposition the *cellulation* of ∂A , and denote it $C(A)$. If \mathcal{P} is the canonical polyhedral decomposition (determined by some choice of cusps), then $C(A)$ satisfies the *Delaunay condition*: the vertices of every polygon can be inscribed on a circle, where the interior of the circle does not contain any other vertices.

In the following proposition, A is a horocusp of M . Let A^t be a sub-horocusp of A such that $d(\partial A, \partial A^t) = t$. A particular feature occurs when t becomes sufficiently large.

Proposition 3.1 (Guéritaud–Schleimer [15]). *Let M be a cusped hyperbolic 3-manifold, endowed with a choice of horocusps A, B_1, \dots, B_n for $n \geq 1$. Assume that an orthogeodesic α from A to B_1 is the unique shortest path from A to $\cup_{j=1}^n B_j$. Then, for every sufficiently small sub-horocusp $A^t \subset A$, the canonical decomposition \mathcal{P} determined by A^t, B_1, \dots, B_n contains a unique edge from A^t to $\cup_{j=1}^n B_j$. This edge is the bi-infinite extension of α .*

Furthermore, there are one or two 3-cells of \mathcal{P} that meet the cusp A^t . Each such 3-cell is an ideal pyramid with an apex at A^t and all lateral edges identified to α . If A is a rectangular cusp, then the single 3-cell meeting A^t is a rectangular pyramid and the induced cellulation of ∂A^t is a rectangle. Otherwise, if A is a non-rectangular cusp, then the two 3-cells meeting A^t are isometric ideal tetrahedra, and the induced cellulation of ∂A^t consists of two isometric, acute triangles.

This result is due to Guéritaud and Schleimer [15, Section 4.1], and appears in the form of a discussion with the explicit hypothesis that the cusp torus ∂A is not rectangular. For completeness, we reproduce an expanded version of their proof.

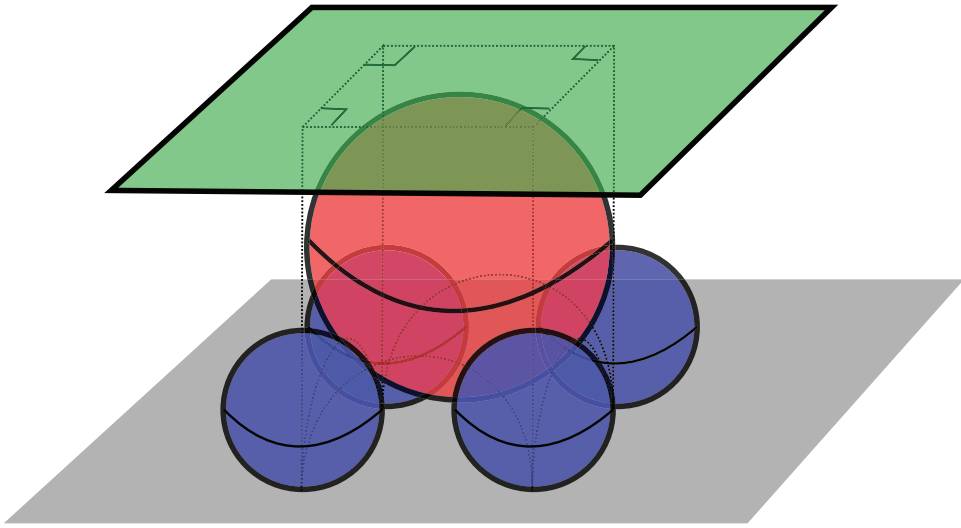


FIGURE 2 A ball D (shown in red) resting on four full-sized horoballs (blue) and tangent to a horoball about ∞ (green). In the setting of Proposition 3.1, all of the full-sized horoballs are in the same orbit of $K \cong \mathbb{Z}^2$. It follows that the ball D can be tangent to four full-sized horoballs at once if and only if K has a rectangular fundamental domain, as in this figure. The center of D corresponds to an ideal rectangular pyramid (dotted).

Proof. Set $M = \mathbb{H}^3/\Gamma$, where we view \mathbb{H}^3 in the upper half-space model. As described in Section 2, we conjugate Γ in $\text{Isom}(\mathbb{H}^3)$ so that a horoball \tilde{A} about ∞ covers the horocusp A . Then $K = \text{Stab}_\Gamma(\infty) \cong \mathbb{Z}^2$ can be identified with $\pi_1(A) \subset \Gamma$. Every preimage of B_j is a horoball tangent to a point of \mathbb{C} .

By hypothesis, there is a unique shortest orthogeodesic α from A to $\cup_{j=1}^n B_j$, which leads from A to B_1 . Extend α to be a bi-infinite geodesic. After shrinking A by an appropriate distance, we may further assume that α is the unique shortest orthogeodesic from A to $A \cup (\cup_{j=1}^n B_j)$. Consequently, there is a choice of horoball \tilde{B}_1 covering B_1 , such that all of the full-sized horoballs tangent to points of \mathbb{C} are in the K -orbit of \tilde{B}_1 .

Consider a ball $D \subset \mathbb{H}^3$ that rests on the collection of horoballs tangent to points of \mathbb{C} (see Figure 2). If the Euclidean diameter of D is sufficiently large, then D will only touch the full-sized horoballs. Furthermore, if we shrink A by a sufficiently large distance $t \geq 0$, producing a sub-horocusp A^t whose preimage horoball $\tilde{A}^t \subset \tilde{A}$ is at sufficient Euclidean height, then D will also be disjoint from \tilde{A}^t . Inflating D to a maximal (hyperbolic) radius produces a ball D^+ that is tangent to \tilde{A}^t and some number of horoballs from the orbit $K \cdot \tilde{B}_1$, and is disjoint from all other horoballs. Observe that D^+ is tangent to either 3 or 4 horoballs in $K \cdot \tilde{B}_1$, and that the case of 4 occurs precisely when $\partial A^t = \partial \tilde{A}^t/K$ has a rectangular fundamental domain.

Now, consider the cut locus Σ^t corresponding to the horocusp collection A^t, B_1, \dots, B_n . Let $\tilde{\Sigma}^t$ be the preimage of Σ^t in \mathbb{H}^3 , and consider the polyhedral decomposition $\mathcal{P} = \mathcal{P}^t$ dual to Σ^t . By Definition 2.6, every vertex $v \in \tilde{\Sigma}^t$ is equidistant from some collection of horoballs (which is maximal with respect to inclusion), and conversely every point equidistant from a maximal collection of horoballs is a vertex of $\tilde{\Sigma}^t$. In particular, the hyperbolic center of the ball D^+ constructed in the previous paragraph must be a vertex $w \in \tilde{\Sigma}^t$. By Definition 2.6, this vertex w is dual to a polyhedron P_w in the canonical decomposition \mathcal{P} , whose ideal vertices lie in the horoballs tangent to D^+ . Recall that D^+ is tangent to \tilde{A}^t and either 3 or 4 full-sized horoballs in $K \cdot \tilde{B}_1$.

If there are 3 full-sized horoballs tangent to D^+ , then P_w is an ideal tetrahedron. Furthermore, up to the action of $K = \text{Stab}_T(\infty)$, there must be exactly one other ideal tetrahedron in \mathcal{P} that intersects \tilde{A}^t . In this case, the induced cusp cellulation of ∂A^t consists of two isometric triangles. These triangles must be acute: otherwise, the circle that circumscribes the 3 vertices of an obtuse triangle would contain the fourth vertex of the parallelogram in its interior, violating the Delaunay condition. Since the two triangles of ∂A^t are isometric, the two tetrahedra meeting A^t are also isometric. Furthermore, each of the two tetrahedra has 3 edges identified to α , corresponding to the fact that the cusp cellulation has a single vertex at $\alpha \cap \partial A^t$.

If there are 4 full-sized horoballs tangent to D^+ , as in Figure 2, then the polyhedron P_w is an ideal rectangular-based pyramid. In this case, P_w is the only cell of \mathcal{P} meeting A^t , and the induced cusp cellulation of A^t is a single rectangle with a vertex at $\alpha \cap \partial A^t$. Consequently, all lateral edges of P_w are identified to α . \square

Remark 3.2. Akiyoshi [1] proved that the canonical polyhedral decomposition \mathcal{P}^t determined by the horocusp collection A^t, B_1, \dots, B_n must stabilize as $t \rightarrow \infty$. Thus, for sufficiently large t , there is a stable geometric decomposition $\mathcal{P} = \mathcal{P}^t$ independent of t . This is our motivation for dropping the superscript t .

In the polyhedral decomposition \mathcal{P} that occurs in Proposition 3.1, the cells that enter the special cusp A^t fit together to form a geometric object that we call a *drilled ananas*.


Definition 3.3. A *drilled ananas* is a 3-manifold N homeomorphic to $T^2 \times [0, \infty) \setminus \{x\}$, where $x \in T^2 \times \{0\}$, and endowed with a complete hyperbolic metric with the following properties. For some $y > 0$, the non-compact end $T^2 \times [y, \infty) \subset N$ is isometric to a horocusp, such that each cross-section $T^2 \times \{y'\}$ for $y' > y$ is a horotorus. The boundary $\partial N = T^2 \times \{0\} \setminus \{x\}$ is made up of two totally geodesic ideal triangles, with vertices at x . These two triangles are glued by isometry along their edges to form a standard two-triangle triangulation of a once-punctured torus, with shearing and bending allowed along the edges of this boundary triangulation. If the base of the drilled ananas is composed of a single totally geodesic ideal rectangle, then we make a choice of diagonal to decompose N into two ideal tetrahedra.

The horocusp $T^2 \times [y, \infty) \subset N$ is called the *cusp* of N , while a regular neighborhood of $x \in T^2 \times \{0\}$ is called the *thorn* of N .


Here are a few notes on terminology and past usage. The term *thorn* was coined by Baker and Cooper [2]. A drilled ananas is a special case of a *topological ideal polyhedron* in the work of Guéritaud [14], and a slightly less special case of an *ideal torihedron* in the work of Champanerkar, Kofman, and Purcell [5, Definition 2.1]. Both of these generalizations capture the idea of placing a hyperbolic structure on a 3-manifold endowed with a polyhedral graph on its boundary. Guéritaud coined the term *ananas* (French for *pineapple*) to describe a topological polyhedron with the topology of a solid torus. (Compare the definition of a filled ananas immediately above Claim 6.10.) The object N in Definition 3.3 can be obtained by removing the core of a solid torus, hence *drilled ananas*.

With the above definition, the conclusion of Proposition 3.1 can be rephrased as follows.

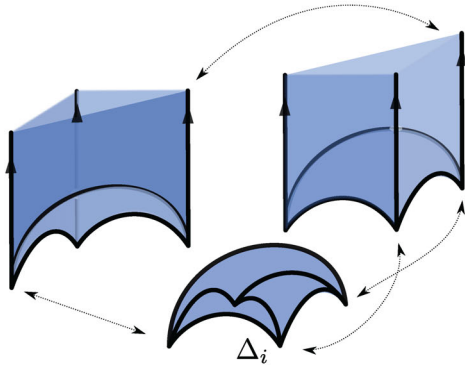
Corollary 3.4. *Let M be a cusped hyperbolic manifold satisfying the hypotheses of Proposition 3.1, and let $\mathcal{P} = \mathcal{P}^t$ be the polyhedral decomposition produced by that proposition. Let $N \subset M$ be the*




(I) Gluing two ideal tetrahedra to form a drilled ananas N_i .



(II) In a 2–3 move, the two tetrahedra glued along a vertical face, shown in part I, are replaced by three tetrahedra that share an edge dual to that face (dashed). The result appears in part III.



(III) After a 2–3 move, the original ananas N_i decomposes into three ideal tetrahedra.



(IV) The new ananas N_{i+1} is the union of the two tetrahedra in N_i which are incident to ∞ . The next 2–3 move will be along a face that is exterior to the fundamental domain shown here.

FIGURE 3 The inductive step of Lemma 3.6

submanifold obtained by gluing together all cells of \mathcal{P} that have an ideal vertex in cusp A , along their shared faces. Then N is a drilled ananas composed of two acute ideal tetrahedra or one ideal rectangular pyramid. Furthermore, N is convex, with an angle less than π at every edge of $\mathcal{P} \cap \partial N$.

Proof. The conclusion that N is a drilled ananas is immediate from Proposition 3.1 and Definition 3.3. It remains to check that N is convex.

By Proposition 3.1, each cell of \mathcal{P} comprising N is an ideal pyramid with a base along ∂N . If N contains a single rectangular pyramid, we subdivide it into two isometric ideal tetrahedra T, T' by choosing a diagonal along ∂N . Otherwise, N already consists of two isometric, acute-angled ideal tetrahedra T, T' . The three lateral faces of T are glued to the three lateral faces of T' , and all lateral edges of T, T' are identified to the single geodesic α that connects the thorn of N to the cusp of N .

Let $\theta_1, \theta_2, \theta_3$ be the dihedral angles of T at the edges identified to α . Since T and T' are necessarily isometric, these are also the dihedral angles of T' at the edges identified to α . See Figure 3(I). Then the three internal angles along the edges of ∂N are $2\theta_1, 2\theta_2, 2\theta_3$.

If T and T' are cells of \mathcal{P} , then Proposition 3.1 says that all of their angles are acute. Thus $2\theta_i < \pi$ for every i , hence ∂N is locally convex at every edge on its boundary. Otherwise, if T and T' were created by subdividing an ideal pyramid, the cusp $A \subset N$ is rectangular; hence, the Euclidean triangles truncating the tips of T and T' have a right angle θ_1 and two acute angles θ_2, θ_3 . Consequently, the internal angles along ∂N are $2\theta_1 = \pi$ and $2\theta_2, 2\theta_3 < \pi$. In either case, ∂N is locally convex, hence N is convex. \square

A key feature of a drilled ananas is that it is made up of two tetrahedra glued along three faces. Each of these three faces supports a local move, called a 2–3 move.

Definition 3.5. Let \mathcal{T}_0 be a topological ideal triangulation of a 3-manifold N , possibly with boundary. Let T and T' be distinct tetrahedra in \mathcal{T}_0 that are glued together along a face F , and observe that $T \cup_F T'$ is a bipyramid. A *topological 2–3 move* replaces $T \cup T'$ with three tetrahedra glued together along a central edge dual to F , while all other tetrahedra of \mathcal{T}_0 remain the same. The resulting triangulation is denoted as \mathcal{T}_1 . We say that a 2–3 move is *geometric* if \mathcal{T}_0 and \mathcal{T}_1 are both geometric triangulations. Assuming \mathcal{T}_0 is geometric, a 2–3 move will be geometric whenever the bipyramid $T \cup_F T'$ is strictly convex.

In Figure 3, tetrahedra T and T' are shown in panel (I), with face F a darker shade of blue. The dual edge to F is dashed in panel (II), and the resulting three tetrahedra are shown in panel (III).

Lemma 3.6. *A drilled ananas N admits an infinite sequence of geometric triangulations, connected by geometric 2–3 moves.*

The inductive construction that proves the lemma is illustrated in Figure 3.

Proof. Let N be a drilled ananas. By definition, ∂N is subdivided into two totally geodesic ideal triangles. Consequently, N itself can be subdivided into two geometric ideal tetrahedra T, T' by coning those triangles to the non-compact end at ∞ . As in Proposition 3.1, we think of each of T, T' as a triangular pyramid with a base along ∂N . The three lateral faces of T are glued to the three lateral faces of T' , and all lateral edges of T, T' are identified to a single geodesic α .

As in Corollary 3.4, let $\theta_1, \theta_2, \theta_3$ be the dihedral angles of T and T' at the edges identified to α . See Figure 3(I). Since these angles are positive and $\theta_1 + \theta_2 + \theta_3 = \pi$, two of the three angles must be strictly less than $\pi/2$. Then the three internal angles along the edges of ∂N are $2\theta_1, 2\theta_2, 2\theta_3$. At least two of these angles (say, $2\theta_1$ and $2\theta_2$) are strictly less than π .

We will prove the lemma by induction. Set $N_0 = N$. The key inductive claim is as follows.

Claim 3.7. Let N_i be a drilled ananas with a two-tetrahedron geometric triangulation. Let $\beta_i \subset \partial N_i$ be a boundary edge such that the internal angle at β_i is less than π . Then performing a diagonal exchange on β_i results in a geometric ideal tetrahedron Δ_i and a new sub-ananas $N_{i+1} \subset N_i$, where $N_i = N_{i+1} \cup \Delta_i$ and N_{i+1} again has a two-tetrahedron geometric triangulation.

The proof of the claim is almost immediate. Removing β_i from a triangulation of ∂N_i results in a quadrilateral. Let β'_i be the opposite diagonal of this quadrilateral. Observe that the geodesic representative of β'_i lies strictly inside N_i , because N_i is locally convex at β_i . The join of β_i, β'_i is a (geometric) ideal tetrahedron Δ_i . Removing the interior of Δ_i from N_i produces a sub-ananas N_{i+1} whose boundary is pleated along β'_i and the two remaining edges of ∂N_i , proving the claim. \diamond

Observe that the two-tetrahedron triangulation of N_i gives rise to a three-tetrahedron triangulation of $N_i = \Delta_i \cup N_{i+1}$, in a geometric 2–3 move. See Figure 3. We can now apply this move to N_{i+1} , and so on, resulting in an infinite sequence of geometric triangulations of $N = N_0$. \square

Corollary 3.8. *Suppose that N is a drilled ananas, and $f : \hat{N} \rightarrow N$ is a finite cover. Then \hat{N} admits an infinite sequence of f -equivariant geometric triangulations.*




FIGURE 4 The Farey complex \mathcal{F} . Edges of the dual tree correspond to diagonal exchanges in a torus with one marked point. Every non-backtracking path in the dual tree, starting from the central triangle, can be realized via geometric 2–3 moves in a drilled ananas. Figure from Ham and Purcell [16, Figure 3.1]

The corollary follows immediately from Lemma 3.6 because every triangulation of N lifts to a triangulation of \hat{N} . We also remark that any finite cover $\hat{N} \rightarrow N$ is regular and has abelian deck group, because $\pi_1(N) \cong \mathbb{Z}^2$ is abelian.

3.1 | Connections to the Farey complex

We can now describe some additional structure in the set of triangulations of a drilled ananas N .

Definition 3.9. The *Farey complex* \mathcal{F} is a simplicial complex whose vertices are isotopy classes of arcs on a torus T^2 based at a marked point x , and whose edges correspond to arcs that are disjoint (except at x). The vertices of \mathcal{F} are commonly identified with the rational points $\mathbb{Q}\mathbb{P}^1 \subset \mathbb{R}\mathbb{P}^1$, as follows. Endow T^2 with a standard Euclidean metric, with fundamental domain a unit square. Then every loop in T^2 based at x can be pulled tight to a Euclidean geodesic of some well-defined slope $\mathbb{Q} \cup \{\infty\}$. Conversely, every rational slope defines a unique isotopy class of arc from x to x .

Triangles in \mathcal{F} correspond to (isotopy classes of) one-vertex triangulations of T^2 with the vertex at x , or equivalently to ideal triangulations of $T^2 \setminus \{x\}$. The dual 1-skeleton of \mathcal{F} is a trivalent tree, with every edge of the dual tree corresponding to a diagonal exchange. See Figure 4.

Using Definition 3.9, we can state the following stronger formulation of Lemma 3.6.

Proposition 3.10. *Let N be a drilled ananas with a geometric triangulation consisting of two acute-angled tetrahedra. Then N admits an infinite trivalent tree of geometric 2–3 moves, with vertices of the tree in natural bijection with triangles of the Farey complex \mathcal{F} .*

Proof. The proof amounts to adding some book-keeping to the proof of Lemma 3.6. Let $N_0 = N$, subdivided into ideal tetrahedra T, T' . By hypothesis, all dihedral angles of T and T' are acute. This triangulation of N defines an induced cellulation of the boundary of the horocusp $A \subset N$. In fact, this cusp cellulation is a one-vertex triangulation: the one vertex is the intersection between ∂A

and the single edge $\alpha \subset N$ into cusp A , while the two triangles are the cross-sections $\partial A \cap T$ and $\partial A \cap T'$. Let τ_0 be this triangulation of ∂A . Fix a framing of ∂A so that the three slopes occurring in τ_0 are $0/1, 1/0$, and $1/1$, as in the central triangle of Figure 4.

Observe that any one-vertex triangulation of ∂A defines an ideal triangulation of ∂N , by projecting outward from the cusp. In the opposite direction, any ideal triangulation of the boundary ∂N_i for some sub-ananas $N_i \subset N$ defines a one-vertex cusp triangulation. We will pass freely between the two viewpoints.

Now, let β_0 be a boundary edge of N_0 . In order to carry out the construction of Claim 3.7, the internal angle of β_0 needs to be less than π . But since all dihedral angles of T and T' are acute, we may choose β_0 at will. Now, the new sub-ananas $N_1 \subset N_0$, constructed as in Claim 3.7, will induce a new cusp cellulation τ_1 of ∂A , which differs from τ_0 via a diagonal exchange. Thus we may choose τ_1 to be any one of the three triangulations adjacent to τ_0 in Figure 4.

Continuing inductively, suppose that we have constructed the sub-ananas $N_i \subset N_{i-1} \subset \dots \subset N_0$, and that ∂N_i has ideal triangulation τ_i . Then, as noted in the proof of Lemma 3.6, there are two edges of ∂N_i that have interior angle less than π . The third edge necessarily has angle more than π ; it is the edge β'_{i-1} that was just created in constructing N_i . (See Figure 3(IV), with indices shifted by 1.) Thus we may choose β_i to be any edge of ∂N_i other than β'_{i-1} , which means that the new cusp triangulation τ_{i+1} is allowed to be any of the two neighbors of τ_i that are distinct from τ_{i-1} . In summary, the path τ_0, τ_1, \dots of cusp cellulations associated with N_0, N_1, \dots is allowed to be any non-backtracking path starting from τ_0 . □

In any path τ_0, τ_1, \dots of cusp cellulations constructed in the above proof, the slopes of edges approach some limiting value in \mathbb{RP}^1 , and the edges themselves approach a foliation with the limiting slope. Thus the geometric retriangulations of N can be chosen to limit to any foliation of the torus.

The strategy for proving Theorem 1.2 and Theorem 1.3 can now come into view. Given a cusped 3-manifold M , we will find a cover \hat{M} with a polyhedral decomposition \hat{P} as in Proposition 3.1, where two ideal tetrahedra fit together to form a drilled ananas. A further cover $\hat{\hat{M}}$ ensures that the other cells of \hat{P} can be subdivided into ideal tetrahedra as well. Producing these covers requires tools from subgroup separability, described in the next section.

4 | SEPARABILITY

This section begins with a review of some standard results about separable subsets and subgroups in a group G . The main content of the section is a proof of Theorem 1.4.

Definition 4.1. Let G be a group. The *profinite topology* on G is the topology whose basic open sets are cosets of finite-index normal subgroups. Since every coset of a finite-index subgroup $H \triangleleft G$ is the complement of finitely many other cosets of H , the basic open sets are also closed.

A subset $S \subset G$ is called *separable* if it is closed in the profinite topology on G .

The following characterization is standard.

Lemma 4.2. *Let G be a group and $S \subset G$ a subset. The following are equivalent.*

- (1) S is separable.

- (2) For every element $g \in G \setminus S$, there is a homomorphism $\varphi : G \rightarrow F$, where F is a finite group and $\varphi(g) \notin \varphi(S)$.

Proof. The set S is closed if and only if every $g \in G \setminus S$ is contained in a basic open set disjoint from S . But, by Definition 4.1, a basic open set containing g is precisely the preimage of an element under a homomorphism to a finite group. \square

If M is a cusped hyperbolic 3-manifold, a subgroup of $\pi_1(M)$ coming from the inclusion of a horocusp A is called *peripheral*. This peripheral subgroup of $\pi_1(M)$ is also *maximal abelian*: it is not contained in any larger abelian subgroup.

We will need to separate peripheral subgroups and their double cosets. The following separability result has been known since the 1980s, if not earlier. See, for example, Long [21, p. 484] for a proof.

Proposition 4.3. *Let $M = \mathbb{H}^3/\Gamma$ be a cusped hyperbolic 3-manifold. Then maximal abelian subgroups of Γ are separable. In particular, peripheral subgroups of Γ are separable.*

We will also need to separate peripheral double cosets, using the following theorem of Hamilton, Wilton, and Zalesskii [18, Theorem 1.4].

Theorem 4.4 (Hamilton–Wilton–Zalesskii [18]). *Let $M = \mathbb{H}^3/\Gamma$ be a finite-volume hyperbolic 3-manifold. Let H and K be abelian subgroups of Γ . Then, for every $g \in \Gamma$, the double coset $HgK = \{h g k : h \in H, k \in K\}$ is separable in Γ .*

Theorem 4.4 is the strongest separability tool needed in the proof of Theorem 1.3. The reader who is mainly interested in that result is invited to proceed directly to Section 5.

4.1 | Algebraic tools for separability

The separability results that we use in this paper, including Theorem 4.4 and Theorem 1.4, are proved using tools from algebraic number theory. To set up the proof of Theorem 1.4, we review some needed definitions and results.

A *number field* is a finite field extension of \mathbb{Q} . An extremely useful connection between number fields and 3-manifolds comes from the following result.

Theorem 4.5 (Thurston [29], Bass [3], Culler–Shalen [6, 8]). *Let $M = \mathbb{H}^3/\Gamma$ be a cusped hyperbolic 3-manifold. Then*

- (1) $\Gamma \subset \mathrm{PSL}(2, \mathbb{C})$ can be lifted to $\mathrm{SL}(2, \mathbb{C})$.
- (2) Γ can be conjugated in $\mathrm{SL}(2, \mathbb{C})$ to lie in $\mathrm{SL}(2, R) \subset \mathrm{SL}(2, k)$, where R is a finitely generated subring of a number field k .

Conclusion (1) was observed by Thurston [29, p. 98] and carefully written down by Culler and Shalen [8, Proposition 3.1.1], with a simplified proof by Culler [6, Corollary 2.2]. Conclusion (2) was observed by Thurston [29, Proposition 6.7.4] as an algebraic consequence of Mostow rigidity, and independently proved by Bass [3].

The arithmetic data associated to a ring R can be used to construct finite quotients of both rings and groups. The construction uses the following notions.

Definition 4.6. If k is a number field, let \mathcal{O}_k denote the ring of integers of k . If \mathfrak{p} is a non-zero prime ideal of \mathcal{O}_k , we let $k_{\mathfrak{p}}$ denote the \mathfrak{p} -adic completion of k and $\mathcal{O}_{k_{\mathfrak{p}}}$ the ring of integers of $k_{\mathfrak{p}}$. The ring $\mathcal{O}_{k_{\mathfrak{p}}}$ has a unique maximal ideal. The quotient of $\mathcal{O}_{k_{\mathfrak{p}}}$ by this maximal ideal is a finite field called the *residue class field* of $\mathcal{O}_{k_{\mathfrak{p}}}$. The quotient map is called the *residue class field map* with respect to \mathfrak{p} . If R is a finitely generated ring in a number field k , then $R \subset \mathcal{O}_{k_{\mathfrak{p}}}$ for all but finitely many primes \mathfrak{p} of \mathcal{O}_k . Restricting the residue class field map to R yields a finite quotient of R .

Now, suppose $\Gamma \subset \text{SL}(2, R)$, where R is a ring. Let $I \subset R$ be an ideal, such that $S = R/I$ is finite. Then the homomorphism

$$\text{SL}(2, R) \rightarrow \text{SL}(2, S)$$

where coefficients are reduced modulo I is called a *congruence quotient* of $\text{PSL}(2, R)$. We also call the composition

$$\varphi : \Gamma \hookrightarrow \text{SL}(2, R) \rightarrow \text{SL}(2, S)$$

a *congruence quotient* of Γ . (This is slightly abusive, because Γ may fail to surject $\text{SL}(2, S)$.)

We now describe three algebraic results that will be used in the proof of Theorem 1.4. The first of these is [18, Theorem 2.6].

Proposition 4.7 (Hamilton–Wilton–Zalesskii [18]). *Let R be a finitely generated ring in a number field k . By fixing a \mathbb{Q} embedding of k into \mathbb{C} , we may view $k \subset \mathbb{C}$. Let $\omega \in R$, and set $Z_{\omega} = \{m + n\omega \mid m, n \in \mathbb{Z}\}$ and $Q_{\omega} = \{m + n\omega \mid m, n \in \mathbb{Q}\}$. If $y \in R - Q_{\omega}$, then there exist a finite ring S and a ring homomorphism $\rho : R \rightarrow S$ such that $\rho(y) \notin \rho(Z_{\omega})$.*

The following technical lemma is a variant of Proposition 4.7.

Lemma 4.8. *Let R be a finitely generated ring in a number field k . By fixing a \mathbb{Q} embedding of k into \mathbb{C} , we may view $k \subset \mathbb{C}$. Let $\omega \in R - \mathbb{R}$. Then there is an infinite collection Ω of primes of \mathbb{Q} , such that for each prime $p \in \Omega$, there exist a finite field F_p of characteristic p and a ring homomorphism $\eta_p : R \rightarrow F_p$, such that $\{1, \eta_p(\omega)\}$ is linearly independent over F_p .*

Consequently, η_p has the following property. Consider an element $y_* \in Q_{\omega} = \{m + n\omega \mid m, n \in \mathbb{Q}\}$. Express y_* in lowest terms:

$$y_* = \frac{m_* + n_*\omega}{v_*},$$

where $m_*, n_* \in \mathbb{Z}$ and $v_* \in \mathbb{N}$. If $\eta_p(y_*) = \eta_p(m + n\omega)$ for some $m, n \in \mathbb{Z}$, then $v_*m \equiv m_* \pmod{p}$ and $v_*n \equiv n_* \pmod{p}$.

Proof. Since $\omega \in \mathbb{C} - \mathbb{R}$, the set $\{1, \omega\}$ is linearly independent over \mathbb{Q} . By a standard argument used in the proof of Theorem 4.4 (compare [18, p. 278]), we can preserve this property in a finite quotient. We include the details for completeness. Let L denote the normal closure of k over \mathbb{Q} and let $\bar{\tau} \in \text{Gal}(L/\mathbb{Q})$ represent complex conjugation. Since $\omega \in \mathbb{C} - \mathbb{R}$, we know ω is not fixed by $\bar{\tau}$. By the Chebotarev Density Theorem, there exist infinitely many primes p of \mathbb{Q} with unramified

extension \mathfrak{p} in L such that $\bar{\tau}$ is the Frobenius automorphism for \mathfrak{p}/p . After eliminating a finite set of primes, if necessary, we may assume that $R \subset \mathcal{O}_{L_{\mathfrak{p}}}$, where $\mathcal{O}_{L_{\mathfrak{p}}}$ denotes the ring of integers in the \mathfrak{p} -adic field $L_{\mathfrak{p}}$. Given such a prime p , let $F_{\mathfrak{p}}$ denote the residue class field of $\mathcal{O}_{L_{\mathfrak{p}}}$ and let \mathbb{F}_p denote the finite field of p elements. Let η_p be the composition of the inclusion map of R into $\mathcal{O}_{L_{\mathfrak{p}}}$ with the residue map:

$$\eta_p : R \hookrightarrow \mathcal{O}_{L_{\mathfrak{p}}} \rightarrow F_{\mathfrak{p}}.$$

Since $\bar{\tau}$ is the Frobenius automorphism of L/\mathbb{Q} with respect to \mathfrak{p}/p , $\text{Gal}(L_{\mathfrak{p}}/\mathbb{Q}_p) = \langle \bar{\tau}' \rangle$ where $\bar{\tau}' = \bar{\tau}$ on L . Since $\bar{\tau}(\omega) \neq \omega$, it follows that $\omega \notin \mathbb{Q}_p$. The Galois group of $F_{\mathfrak{p}}/\mathbb{F}_p$ is also induced by $\bar{\tau}$. This implies that $\eta_p(\omega) \notin \mathbb{F}_p$, and therefore, the set $\{1, \eta_p(\omega)\}$ is linearly independent over \mathbb{F}_p .

Next, suppose that $\eta_p(y_*) = \eta_p(m + n\omega)$ for some $m, n \in \mathbb{Z}$. Then

$$\eta_p(m_* + n_*\omega) = \eta_p(v_*y_*) = \eta_p(v_*m + v_*n\omega).$$

The above equality can be rewritten as $\eta_p(v_*m - m_*) + \eta_p(v_*n - n_*)\eta_p(\omega) = 0$. Since the set $\{1, \eta_p(\omega)\}$ is linearly independent over \mathbb{F}_p , it follows that $v_*m \equiv m_* \pmod{p}$ and $v_*n \equiv n_* \pmod{p}$. □

The third preliminary algebraic result is a combination of [17, Theorem 2.3 and Corollary 2.5].

Proposition 4.9 (Hamilton [17]). *Let R be a finitely generated ring in a number field k , let λ be a non-zero element of R that is not a root of unity, and let x_1, x_2, \dots, x_j be non-zero elements of R . Then, for every sufficiently large integer q , there exist a non-zero prime ideal \mathfrak{p} of \mathcal{O}_k , a finite field $F_{\mathfrak{p}}$, and a ring homomorphism $\sigma : R \rightarrow F_{\mathfrak{p}}$ such that $R \subset \mathcal{O}_{k_{\mathfrak{p}}}$, the multiplicative order of $\sigma(\lambda)$ is equal to q , and $\sigma(x_i) \neq 0$, for each $1 \leq i \leq j$. The field $F_{\mathfrak{p}}$ is the residue class field of $\mathcal{O}_{k_{\mathfrak{p}}}$ and the map σ is the restriction to R of the residue class field map with respect to \mathfrak{p} .*

4.2 | Conjugacy separation of peripheral cosets

We can now prove Theorem 1.4.

Proof of Theorem 1.4. Let H and K be maximal parabolic subgroups of Γ corresponding to distinct cusps of M , and let $g \in \Gamma$ be an element such that K is disjoint from every conjugate of gH . In particular, this implies $g \notin H$. Fix a non-trivial element $h_0 \in H$. Since H is a maximal abelian subgroup of Γ , and $g \notin H$, the commutator $[g, h_0] = gh_0g^{-1}h_0^{-1}$ is non-trivial. Let A be the cusp of M corresponding to K , let B be the cusp corresponding to H , and let C_1, \dots, C_ℓ be the remaining cusps. We will leave the cusp A unfilled, and will fill the remaining cusps. So long as a tuple of slopes \mathbf{s} on B, C_1, \dots, C_ℓ avoids finitely many slopes on each cusp, the Dehn filled manifold $M(\mathbf{s})$ will be hyperbolic. Thus, by Thurston’s hyperbolic Dehn surgery theorem, we can choose generators h_1 and h_2 of H that can be completed to tuples \mathbf{s}_1 and \mathbf{s}_2 where $M(\mathbf{s}_1)$ and $M(\mathbf{s}_2)$ are both hyperbolic. Then, for $j = 1, 2$, the fundamental group $\pi_1(M(\mathbf{s}_j))$ has a discrete, faithful representation to a group of isometries $\Gamma(\mathbf{s}_j) \subset \text{PSL}(2, \mathbb{C})$. By Theorem 4.5, we view $\Gamma(\mathbf{s}_j)$ as a subgroup of $\text{SL}(2, \mathbb{C})$. Let $\psi_j : \Gamma \rightarrow \Gamma(\mathbf{s}_j)$ be the quotient homomorphism induced by the inclusion $M \hookrightarrow M(\mathbf{s}_j)$. By choosing sufficiently long Dehn fillings, we may assume that $\psi_1([g, h_0])$ and $\psi_2([g, h_0])$ are non-trivial. These choices ensure the following properties for $j \in \{1, 2\}$:

- $\psi_j(K)$ is a parabolic subgroup of $\Gamma(\mathbf{s}_j)$,
- $\psi_j(H)$ is a loxodromic subgroup of $\Gamma(\mathbf{s}_j)$,
- $\psi_j(g) \notin \psi_j(H)$.

Before working with the Dehn filled manifolds, we examine the coset gH in Γ . By Theorem 4.5, we can conjugate Γ to lie in $SL(2, k)$ for some number field k . After possibly expanding k , and then conjugating Γ in $SL(2, k)$, we may assume that

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \quad h_1 = \pm \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \text{and} \quad h_2 = \pm \begin{pmatrix} 1 & \omega \\ 0 & 1 \end{pmatrix},$$

for a fixed element $\omega \in \mathbb{C} - \mathbb{R}$. Note that the traces of h_1 and h_2 are determined by the lift to $SL(2, \mathbb{C})$, and might not coincide. Thus an arbitrary element $gh_1^m h_2^n \in gH$ can be expressed as

$$gh_1^m h_2^n = \pm \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} 1 & m + n\omega \\ 0 & 1 \end{pmatrix} = \pm \begin{pmatrix} a & a(m + n\omega) + b \\ c & c(m + n\omega) + d \end{pmatrix}, \tag{4.1}$$

where the sign \pm depends on the traces of h_1, h_2 and the parity of m, n .

As above, set $Z_\omega = \{m + n\omega \mid m, n \in \mathbb{Z}\}$ and $Q_\omega = \{m + n\omega \mid m, n \in \mathbb{Q}\}$. Then

$$gH \subset \left\{ \pm \begin{pmatrix} a & ax + b \\ c & cx + d \end{pmatrix} \mid x \in Z_\omega \right\}.$$

Since $g \notin H$, and H is a maximal parabolic subgroup of Γ , we have $c \neq 0$. Thus $gh_1^m h_2^n \in gH$ is parabolic if and only if $\text{tr}(gh_1^m h_2^n) = \pm(a + d + c(m + n\omega)) \in \{\pm 2\}$. Solving for $x = m + n\omega \in Z_\omega$, let

$$y_+ = y_{+1} = \frac{2 - a - d}{c} \quad \text{and} \quad y_- = y_{-1} = \frac{-2 - a - d}{c}.$$

Then the coset gH contains a parabolic element if and only if $\{y_+, y_-\} \cap Z_\omega \neq \emptyset$. We will abuse notation slightly by thinking of the subscripts as either symbols (\pm) or numbers (± 1), as convenient.

The pair y_+ and y_- , corresponding to trace $+2$ and trace -2 , are the “problem elements” that we will need to track throughout the proof. In particular, let $R \subset k$ be the ring generated by the coefficients of the generators of Γ . By expanding R if necessary, we may assume that $c^{-1} \in R$, which implies $y_\pm \in R$.

From here, the proof proceeds as follows. For each number $i \in \{\pm 1\}$, we will construct a homomorphism $\varphi_i : \Gamma \rightarrow G_i$, where G_i is a finite group. Each φ_i will be either a congruence quotient of Γ , or the product of a congruence quotient of Γ and a congruence quotient of $\Gamma(\mathbf{s}_j)$ for some j . Then we will package these homomorphisms together to obtain

$$\varphi = \varphi_- \times \varphi_+ : \Gamma \longrightarrow G = G_- \times G_+.$$

In particular, for every $\gamma \in \Gamma$, the image $\varphi(\gamma)$ is a tuple of matrices, each with coefficients in a finite ring, and each with a well-defined trace. To complete the proof, we will see that for every $h \in H$ and every $\ell \in K$, some coordinate of $\varphi(gh)$ differs in trace from the same coordinate of $\varphi(\ell)$. This will imply that $\varphi(gh)$ cannot be conjugate to $\varphi(\ell)$.

For each $i \in \{\pm 1\}$, the definition of $\varphi_i : \Gamma \rightarrow G_i$ depends on whether y_i belongs to Q_ω . If $y_i \notin Q_\omega$, we argue as follows.

Claim 4.10. Suppose $y_i \in R - Q_\omega$. Then there is a ring homomorphism $\rho_i : R \rightarrow S_i$, where S_i is a finite ring, such that the following holds. For every pair $(m, n) \in \mathbb{Z}^2$, we have

$$\rho_i(m + n\omega) \neq \rho_i(y_i).$$

By Proposition 4.7, there exist a finite ring S_i and a ring homomorphism $\rho_i : R \rightarrow S_i$ such that $\rho_i(y_i) \notin \rho_i(Z_\omega)$. By the definition of Z_ω , this means $\rho_i(y_i) \neq \rho_i(m + n\omega)$ for any (m, n) , proving Claim 4.10. \diamond

Using ρ_i , we define a congruence quotient

$$\varphi_i : \Gamma \hookrightarrow \text{SL}(2, R) \rightarrow G_i = \text{SL}(2, S_i) \quad \text{if } y_i \notin Q_\omega, \tag{4.2}$$

completing the definition of G_i and φ_i in this case.

Alternately, if $y_i \in Q_\omega$, Lemma 4.8 provides an infinite set of primes Ω , such that for each $p \in \Omega$ there is an associated finite field F_p and ring homomorphism $\eta_p : R \rightarrow F_p$. The central claim in this case is the following.

Claim 4.11. Suppose $y_i \in Q_\omega$. Then there is a choice of Dehn filling quotient $\psi_j : \Gamma \rightarrow \Gamma(\mathbf{s}_j)$ where the coefficients of $\Gamma(\mathbf{s}_j)$ lie in a finitely generated ring T_j , a prime number $p \in \Omega$, and a ring homomorphism $\sigma_{p,i} : T_j \rightarrow E_i$, where E_i is a finite field, such that the following holds. For every pair $(m, n) \in \mathbb{Z}^2$, we have

$$\eta_p(m + n\omega) \neq \eta_p(y_i) \quad \text{or} \quad \sigma_{p,i} \circ \text{tr} \circ \psi_j(gh_1^m h_2^n) \neq \pm 2.$$

In fact, the desired homomorphism $\sigma_{p,i}$ exists for all sufficiently large $p \in \Omega$. However, we will only need $\sigma_{p,i}$ for one $p \in \Omega$.

We momentarily postpone the proof of Claim 4.11 to describe the construction of φ_i . The ring homomorphism η_p defines a congruence quotient

$$\nu_0 : \Gamma \hookrightarrow \text{SL}(2, R) \rightarrow \text{SL}(2, F_p).$$

Similarly, $\sigma_{p,i}$ defines a homomorphism ν_i factoring through a congruence quotient:

$$\nu_i : \Gamma \xrightarrow{\psi_j} \psi_j(\Gamma) = \Gamma(\mathbf{s}_j) \hookrightarrow \text{SL}(2, T_j) \rightarrow \text{SL}(2, E_i).$$

We can now define

$$\varphi_i = \nu_0 \times \nu_i : \Gamma \longrightarrow G_i = \text{SL}(2, F_p) \times \text{SL}(2, E_i) \quad \text{if } y_i \in Q_\omega, \tag{4.3}$$

completing the definition of G_i and φ_i in this case.

Proof of Claim 4.11. We begin by specifying the choice of Dehn filling quotient ψ_1 or ψ_2 . Write $y_i = (m_i + n_i\omega)/v_i$ in lowest terms, as in Lemma 4.8. If $v_i = 1$, then we set $j = 2$ and work with

the Dehn filling $\psi_j = \psi_2 : \Gamma \rightarrow \Gamma(\mathbf{s}_2)$ for concreteness (although ψ_1 would also work). Assuming $v_i \neq 1$, we have either $v_i \nmid m_i$ or $v_i \nmid n_i$. If $v_i \nmid m_i$, then we set $j = 2$ and select the Dehn filling $M(\mathbf{s}_2)$ and the quotient map $\psi_2 : \Gamma \rightarrow \Gamma(\mathbf{s}_2)$. Then $\psi_2(H)$ is an infinite cyclic loxodromic subgroup of $\Gamma(\mathbf{s}_2)$ generated by $\psi_2(h_1)$. Consequently, $\psi_2(h_1^m h_2^n) = \psi_2(h_1^m)$, where

$$v_i m \not\equiv m_i \pmod{v_i} \quad \text{because} \quad v_i \nmid m_i.$$

Similarly, if $v_i \nmid n_i$, then we set $j = 1$ and select the Dehn filling $M(\mathbf{s}_1)$ and the quotient map $\psi_1 : \Gamma \rightarrow \Gamma(\mathbf{s}_1)$. This has the effect that $\psi_1(h_1^m h_2^n) = \psi_1(h_2^n)$, where

$$v_i n \not\equiv n_i \pmod{v_i} \quad \text{because} \quad v_i \nmid n_i.$$

In either case, the above non-congruences will be used in the endgame of the proof of the claim. Because the arguments for m_i and n_i are entirely parallel, and differ only by a substitution of symbols, we assume without loss of generality that $v_i \nmid m_i$, hence $j = 2$ and we have the Dehn filling quotient $\psi_2 : \Gamma \rightarrow \Gamma(\mathbf{s}_2)$.

By Theorem 4.5, we can conjugate $\Gamma(\mathbf{s}_2)$ in $SL(2, \mathbb{C})$ such that $\Gamma(\mathbf{s}_2) \subset SL(2, T)$, where $T = T_2$ is a finitely generated ring in a number field. Moreover, we may assume that

$$\psi_2(g) = \begin{pmatrix} r & s \\ t & u \end{pmatrix} \quad \text{and} \quad \psi_2(h_1) = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix},$$

for some $r, s, t, u, \lambda \in \mathbb{C}$ with $|\lambda| \neq 1$. Then

$$\psi_2(gH) = \left\{ \begin{pmatrix} r & s \\ t & u \end{pmatrix} \begin{pmatrix} \lambda^m & 0 \\ 0 & \lambda^{-m} \end{pmatrix} \mid m \in \mathbb{Z} \right\} = \left\{ \begin{pmatrix} r\lambda^m & s\lambda^{-m} \\ t\lambda^m & u\lambda^{-m} \end{pmatrix} \mid m \in \mathbb{Z} \right\},$$

and $\text{tr} \circ \psi_2(gh_1^m h_2^n) = r\lambda^m + u\lambda^{-m}$. In particular, $\psi_2(gH)$ contains a parabolic matrix if and only if $r\lambda^m + u\lambda^{-m} = \pm 2$ for some m , which is true if and only if λ^m is a root of $rx^2 \pm 2x + u$. Let

$$\zeta = \frac{1 + \sqrt{1 - ru}}{r}, \quad \xi = \frac{1 - \sqrt{1 - ru}}{r}$$

be the roots of $rx^2 - 2x + u$. Then $\{-\zeta, -\xi\}$ are the roots of $rx^2 + 2x + u$. By expanding $T = T_2$, if necessary, we may assume that $\{\zeta, \xi\} \subset T$. Since $\psi_2(g)$ does not commute with $\psi_2(h_1)$, we have $s \neq 0$ and $t \neq 0$. Therefore, $ru \neq 1$, which implies that $\zeta \neq \xi$.

To prove the claim, we consider two cases.

Case 1: $\lambda^{2m_i} \notin \{\zeta^{2v_i}, \xi^{2v_i}\}$.

By Proposition 4.9, for all sufficiently large primes $p \in \Omega$, there exist a finite field E_i and a ring homomorphism $\sigma_{p,i} : T \rightarrow E_i$, such that

$$\sigma_{p,i}(\zeta - \xi) \neq 0, \quad \sigma_{p,i}(\lambda^{2m_i} - \zeta^{2v_i}) \neq 0, \quad \sigma_{p,i}(\lambda^{2m_i} - \xi^{2v_i}) \neq 0,$$

and the multiplicative order of $\sigma_{p,i}(\lambda)$ is equal to $2p$.

Suppose for a contradiction that there exist $m, n \in \mathbb{Z}$ such that

$$\eta_p(m + n\omega) = \eta_p(y_i) \quad \text{and} \quad \sigma_{p,i} \circ \text{tr} \circ \psi_2(gh_1^m h_2^n) = \sigma_{p,i}(r\lambda^m + u\lambda^{-m}) = \pm 2.$$

Then Lemma 4.8 implies

$$v_i m \equiv m_i \pmod{p} \Rightarrow 2v_i m \equiv 2m_i \pmod{2p}.$$

If $\sigma_{p,i}(r\lambda^m + u\lambda^{-m}) = 2$, then $\sigma_{p,i}(\lambda^m)$ is a root of $f(x) = \sigma_{p,i}(r)x^2 - 2x + \sigma_{p,i}(u)$ over E_i . Since the two distinct roots of f over E_i are equal to $\sigma_{p,i}(\zeta)$ and $\sigma_{p,i}(\xi)$, we have $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\zeta)$ or $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\xi)$. Similarly, if $\sigma_{p,i}(r\lambda^m + u\lambda^{-m}) = -2$, then $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(-\zeta)$ or $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(-\xi)$. In either case,

$$\sigma_{p,i}(\lambda^{2v_i m}) = \sigma_{p,i}(\zeta^{2v_i}) \quad \text{or} \quad \sigma_{p,i}(\lambda^{2v_i m}) = \sigma_{p,i}(\xi^{2v_i}).$$

Since $2v_i m \equiv 2m_i \pmod{2p}$ and the multiplicative order of $\sigma_{p,i}(\lambda)$ is equal to $2p$, this implies that

$$\sigma_{p,i}(\lambda^{2m_i}) = \sigma_{p,i}(\zeta^{2v_i}) \quad \text{or} \quad \sigma_{p,i}(\lambda^{2m_i}) = \sigma_{p,i}(\xi^{2v_i}),$$

a contradiction.

Case 2: $\lambda^{2m_i} \in \{\zeta^{2v_i}, \xi^{2v_i}\}$.

Without loss of generality, we may assume that $\lambda^{2m_i} = \zeta^{2v_i}$. First, suppose that $v_i = 1$. This means that $y_i = m_i + n_i\omega \in Z_\omega$, hence $gh_1^{m_i}h_2^{n_i} \in gH$ has trace ± 2 and is parabolic in Γ . Since $\lambda^{2m_i} = \zeta^{2v_i}$ and $v_i = 1$, $\lambda^{m_i} \in \{\zeta, -\zeta\}$. Therefore, $\psi_2(gh_1^{m_i}h_2^{n_i}) = \psi_2(gh_1^{m_i})$ is also a parabolic element in $\Gamma(\mathbf{s}_2)$. This is a contradiction, since the only parabolic elements of Γ that remain parabolic after the Dehn filling lie in conjugates of K , and the coset gH is disjoint from every conjugate of K . We conclude that $v_i > 1$.

Next, assume that $\lambda^{2v_i m_i} \neq \xi^{2v_i^2}$. By Proposition 4.9, for all sufficiently large primes $p \in \Omega$, there exist a finite field E_i and a ring homomorphism $\sigma_{p,i} : T \rightarrow E_i$, such that

$$\sigma_{p,i}(\zeta - \xi) \neq 0, \quad \sigma_{p,i}(\lambda^{2v_i m_i} - \xi^{2v_i^2}) \neq 0,$$

and the multiplicative order of $\sigma_{p,i}(\lambda)$ is equal to $2v_i p$. Suppose for a contradiction that there exist $m, n \in \mathbb{Z}$ such that

$$\eta_p(m + n\omega) = \eta_p(y_i) \quad \text{and} \quad \sigma_{p,i} \circ \text{tr} \circ \psi_2(gh_1^m h_2^n) = \sigma_{p,i}(r\lambda^m + u\lambda^{-m}) = \pm 2.$$

Then Lemma 4.8 implies

$$v_i m \equiv m_i \pmod{p} \Rightarrow 2v_i^2 m \equiv 2v_i m_i \pmod{2v_i p}.$$

Since $\sigma_{p,i}(r\lambda^m + u\lambda^{-m}) = \pm 2$, we have $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\zeta)$ or $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\xi)$. If $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\zeta)$, then

$$\sigma_{p,i}(\lambda^{2v_i m}) = \sigma_{p,i}(\zeta^{2v_i}) = \sigma_{p,i}(\lambda^{2m_i}).$$

Since the multiplicative order of $\sigma_{p,i}(\lambda)$ is divisible by $2v_i$, we have

$$2v_i m \equiv 2m_i \pmod{2v_i},$$

contradicting the fact that $v_i \nmid m_i$. If $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\xi)$, then, since the multiplicative order of $\sigma_{p,i}(\lambda)$ is equal to $2v_i p$ and $2v_i^2 m \equiv 2v_i m_i \pmod{2v_i p}$, we obtain

$$\sigma_{p,i}(\lambda^{2v_i m_i}) = \sigma_{p,i}(\lambda^{2v_i^2 m}) = \sigma_{p,i}(\xi^{2v_i^2}),$$

which contradicts our assumption in choosing $\sigma_{p,i}$.

Finally, assume that $\lambda^{2v_i m_i} = \xi^{2v_i^2}$. By Proposition 4.9, there exist a finite field E_i and a ring homomorphism $\sigma_{p,i} : T \rightarrow E_i$, such that $\sigma_{p,i}(\zeta - \xi) \neq 0$, and the multiplicative order of $\sigma_{p,i}(\lambda)$ is divisible by $2v_i^2$. Suppose that there exist $m \in \mathbb{Z}$ such that $\sigma_{p,i}(r\lambda^m + u\lambda^{-m}) = \pm 2$. Then $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\xi)$ or $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\xi)$. By the argument above, $\sigma_{p,i}(\lambda^m) \neq \sigma_{p,i}(\pm\xi)$. If $\sigma_{p,i}(\lambda^m) = \sigma_{p,i}(\pm\xi)$, then

$$\sigma_{p,i}(\lambda^{2v_i^2 m}) = \sigma_{p,i}(\xi^{2v_i^2}) = \sigma_{p,i}(\lambda^{2v_i m_i}).$$

Since the multiplicative order of $\sigma_{p,i}(\lambda)$ is divisible by $2v_i^2$, this implies that

$$2v_i^2 m \equiv 2v_i m_i \pmod{2v_i^2}.$$

But this contradicts the fact that $v_i \nmid m_i$, completing the proof of Claim 4.11. ◇

We can now complete the proof of the theorem. For each $i \in \{\pm 1\}$, we have defined a homomorphism $\varphi_i : \Gamma \rightarrow G_i$, using Equation (4.2) if $y_i \notin Q_\omega$ and Equation (4.3) if $y_i \in Q_\omega$. Now, define

$$\varphi = \varphi_{-1} \times \varphi_{+1} : \Gamma \longrightarrow G = G_{-1} \times G_{+1}.$$

We need to show that $\varphi(K)$ is disjoint from every conjugate of $\varphi(gH)$.

Consider an arbitrary element $gh_1^m h_2^n \in gH$, and suppose for a contradiction that $\varphi(gh_1^m h_2^n)$ is conjugate to $\varphi(\ell)$ for some $\ell \in K$. Since ℓ is parabolic, we know $\text{tr}(\ell) \in \{\pm 2\}$. Recalling the general form for $gh_1^m h_2^n$ in Equation (4.1), define a number $\epsilon = \epsilon(\ell, m, n) \in \{\pm 1\}$ so that

$$\text{tr}(gh_1^m h_2^n) = \epsilon \cdot \frac{\text{tr}(\ell)}{2} \cdot (a + d + c(m + n\omega)).$$

In other words, $\epsilon = 1$ when $\text{tr}(h_1^m h_2^n) = \text{tr}(\ell)$, and $\epsilon = -1$ otherwise. We use the coordinate φ_ϵ of φ to obtain a contradiction.

If $y_\epsilon \notin Q_\omega$, then $\varphi_\epsilon(\ell) \in \text{SL}(2, S_\epsilon)$. Since $\varphi_\epsilon(gh_1^m h_2^n)$ is conjugate to $\varphi_\epsilon(\ell)$, we have

$$\text{tr} \circ \varphi_\epsilon(gh_1^m h_2^n) = \rho_\epsilon \circ \text{tr}(gh_1^m h_2^n) = \rho_\epsilon \left(\epsilon \frac{\text{tr}(\ell)}{2} (a + d + c(m + n\omega)) \right) = \rho_\epsilon(\text{tr}(\ell)) = \text{tr} \circ \varphi_\epsilon(\ell).$$

Since ρ_ϵ is a ring homomorphism and $\epsilon \frac{\text{tr}(\ell)}{2} \in \{\pm 1\}$ is a unit, we may rearrange terms to obtain

$$\rho_\epsilon(a + d + c(m + n\omega)) = \rho_\epsilon \left(\frac{2}{\epsilon} \right) \Rightarrow \rho_\epsilon(c(m + n\omega)) = \rho_\epsilon \left(\frac{2}{\epsilon} - a - d \right).$$

But then $\rho_\epsilon(m + n\omega) = \rho_\epsilon(y_\epsilon)$, contradicting Claim 4.10.

If $y_\epsilon \in Q_\omega$, then $\varphi_\epsilon(\ell) = (\nu_0(\ell), \nu_\epsilon(\ell)) \in \text{SL}(2, F_p) \times \text{SL}(2, E_\epsilon)$, a product of two matrices. Since $\nu_0(gh_1^m h_2^n)$ is conjugate to $\nu_0(\ell)$, we obtain

$$\text{tr} \circ \nu_0(gh_1^m h_2^n) = \eta_p \circ \text{tr}(gh_1^m h_2^n) = \eta_p \left(\epsilon \frac{\text{tr}(\ell)}{2} (a + d + c(m + n\omega)) \right) = \eta_p(\text{tr}(\ell)) = \text{tr} \circ \nu_0(\ell).$$

Then the same rearrangement of terms as before gives $\eta_p(m + n\omega) = \eta_p(y_\epsilon)$. Meanwhile, since $\nu_\epsilon(gh_1^m h_2^n)$ is conjugate to $\nu_\epsilon(\ell)$, we obtain

$$\sigma_{p,\epsilon} \circ \text{tr} \circ \psi_j(gh_1^m h_2^n) = \text{tr} \circ \nu_\epsilon(gh_1^m h_2^n) = \text{tr} \circ \nu_\epsilon(\ell) = \pm 2,$$

contradicting Claim 4.11. In either case, the proof is complete. □

5 | MANIFOLDS WITH NON-RECTANGULAR CUSPS

In this section, we prove Theorem 1.3. We begin with a cusped, hyperbolic 3-manifold M containing a horocusp A . We will construct a sequence of finite covers $\hat{M} \rightarrow \check{M} \rightarrow M$, with an increasingly strong sequence of properties. The final cover \hat{M} will have infinitely many geometric ideal triangulations, implying Theorem 1.3. The construction proceeds in four steps.

- Step 1. Construct a cover $\check{M} \rightarrow M$ where A lifts to a horocusp \check{A} that has a unique shortest path to the other cusps. This is accomplished in Lemma 5.2.
- Step 2. Shrink $\check{A} \subset \check{M}$ to a small sub-horocusp \check{A}^t . By Proposition 3.1 and Corollary 3.4, the canonical cell decomposition \check{P} determined by \check{A}^t and the other cusps contains an embedded drilled ananas \check{N} consisting of one or two cells of \check{P} .
- Step 3. Construct a cover $\hat{M} \rightarrow \check{M}$ where every polyhedron in the lifted polyhedral decomposition \hat{P} has vertices at distinct cusps. This is accomplished in Lemma 5.4.
- Step 4. Now, \hat{P} can be subdivided into geometric ideal tetrahedra by Lemma 2.9, and furthermore, \hat{P} contains a cover \hat{N} of the original ananas \check{N} . If the cusp A was non-rectangular, the induced triangulation of \hat{N} is equivariant with respect to the cover of \check{N} . To conclude the proof, we use Lemma 3.6 to find infinitely many ideal triangulations of \hat{N} , hence of \hat{M} .

Step 1 builds covers using Theorem 4.4, whereas Step 3 uses Proposition 4.3. Steps 2 and 4 construct and subdivide polyhedral decompositions, but do not build any covers. The hypothesis that A is a non-rectangular cusp is used only in Step 4. See Remark 5.5 for a detailed description of how this hypothesis is used.

The following basic lemma will be used to apply the results of Section 4.

Lemma 5.1. *Let $M = \mathbb{H}^3/\Gamma$ be a cusped hyperbolic manifold. Let $\tilde{B}, \tilde{B}' \subset \mathbb{H}^3$ be horoballs that cover the same cusp in M , and let $g \in \Gamma$ be an isometry such that $g(\tilde{B}) = \tilde{B}'$. Then the set of all elements of Γ taking \tilde{B} to \tilde{B}' is of the form*

$$S = \text{Stab}_\Gamma(\tilde{B}')g \text{Stab}_\Gamma(\tilde{B}) = g \text{Stab}_\Gamma(\tilde{B}) = \text{Stab}_\Gamma(\tilde{B}')g,$$

both a left coset and a right coset of peripheral subgroups.

Proof. Let $h \in S$. Then h differs from g by pre-composition with some element $s \in \text{Stab}_\Gamma(\tilde{B})$ and post-composition with some element $s' \in \text{Stab}_\Gamma(\tilde{B}')$. In other words,

$$S = \text{Stab}_\Gamma(\tilde{B}')g \text{Stab}_\Gamma(\tilde{B}) = \{s' \cdot g \cdot s \mid s' \in \text{Stab}_\Gamma(\tilde{B}'), s \in \text{Stab}_\Gamma(\tilde{B})\},$$

proving the first equality of the lemma. Now, observe that $\text{Stab}_\Gamma(\tilde{B}') = g \text{Stab}_\Gamma(\tilde{B})g^{-1}$. Thus

$$S = g \text{Stab}_\Gamma(\tilde{B})g^{-1} \cdot g \text{Stab}_\Gamma(\tilde{B}) = g \text{Stab}_\Gamma(\tilde{B}),$$

proving the second equality of the lemma. The final equality is proved similarly. □

Now, let M be a cusped hyperbolic manifold containing a collection of disjoint, closed horocusps. Let A be one of the horocusps. Let $\gamma_1, \dots, \gamma_n$ be the set of all orthogeodesics below a certain length L that connect A to the union of the other cusps. This set is finite for any L , and non-empty when L is sufficiently large.

Step 1 of the proof is accomplished by the following lemma.

Lemma 5.2. *Let $M = \mathbb{H}^3/\Gamma$ be a cusped hyperbolic manifold and $A \subset M$ a horocusp. For $i = 1, \dots, n$, let γ_i be an orthogeodesic from ∂A to some horocusp of M . Assume that $\gamma_i \neq \gamma_j$ for $i \neq j$. Then there is a finite cover $\hat{f} : \hat{M} \rightarrow M$, where A lifts to a horocusp $\hat{A} \subset \hat{M}$ and where the path-lifts $\hat{\gamma}_1, \dots, \hat{\gamma}_n$ that start at \hat{A} lead to horocusps of \hat{M} that are distinct from one another and from \hat{A} .*

Recall from Definition 2.1 that a lift \hat{A} must cover A with degree one. Thus, for each i , the path γ_i has exactly one path-lift $\hat{\gamma}_i$ starting at \hat{A} .

Proof. Conjugate Γ in $\text{Isom}(\mathbb{H}^3)$ so that one preimage of A is a horoball \tilde{A} about ∞ . The subgroup $K = \text{Stab}_\Gamma(\tilde{A}) \cong \mathbb{Z}^2$ can be identified with $\pi_1(A)$. Choose a fundamental domain $D \subset \partial\tilde{A}$ for the action of K . For each i , let $\tilde{\gamma}_i$ be a path-lift of γ_i to \mathbb{H}^3 , whose initial point lies in D . Let \tilde{B}_i be the horoball at the forward endpoint of $\tilde{\gamma}_i$. By construction, the orthogeodesics $\tilde{\gamma}_1, \dots, \tilde{\gamma}_n$ lie in distinct K -orbits, hence the horoballs $\tilde{B}_1, \dots, \tilde{B}_n$ do also. In addition, each \tilde{B}_i is disjoint from \tilde{A} .

For every pair (i, j) with $1 \leq i < j \leq n$, let $S_{ij} \subset \Gamma$ be the set of all deck transformations that map \tilde{B}_j to \tilde{B}_i . This set may be empty (this will be the case if \tilde{B}_i and \tilde{B}_j cover distinct cusps of M). Otherwise, let $g_{ij} \in S_{ij}$ be an arbitrary element and observe that by Lemma 5.1, $S_{ij} = g_{ij} \text{Stab}_\Gamma(\tilde{B}_j)$. Let T_{ij} be the set of all deck transformations that map \tilde{B}_j to any horoball in the K -orbit of \tilde{B}_i . If $S_{ij} \neq \emptyset$, we have $T_{ij} = KS_{ij} = Kg_{ij} \text{Stab}_\Gamma(\tilde{B}_j)$, a double coset of peripheral subgroups of Γ .

In a similar fashion, let $S_{0j} \subset \Gamma$ be the set of all deck transformations that map \tilde{B}_j to \tilde{A} . If $S_{0j} \neq \emptyset$, Lemma 5.1 says that $S_{0j} = g_{0j} \text{Stab}_\Gamma(\tilde{B}_j)$ for an arbitrary element $g_{0j} \in S_{0j}$. Let $T_{0j} = KS_{0j}$. Finally, define

$$T = \bigcup_{0 \leq i < j \leq n} T_{ij} = \bigcup_{0 \leq i < j \leq n} KS_{ij} = \bigcup_{i < j, S_{ij} \neq \emptyset} Kg_{ij} \text{Stab}_\Gamma(\tilde{B}_j).$$

Theorem 4.4 says that for every non-empty S_{ij} , the double coset $T_{ij} = Kg_{ij} \text{Stab}_\Gamma(\tilde{B}_j)$ is separable in Γ . Thus each such double coset is a closed subset of Γ . Since T is a finite union of these closed sets, it follows that T itself is closed, hence separable. Observe that $1 \notin T$, because any element $g \in T_{ij}$ moves horoball \tilde{B}_j to a distinct location. (This uses the above observation that the horoballs \tilde{B}_i and \tilde{B}_j lie in distinct K -orbits.)

By Lemma 4.2, there is a homomorphism $\varphi : \Gamma \rightarrow F$, where F is a finite group, such that $1 = \varphi(1) \notin \varphi(T)$. Since $\varphi(K) = \varphi(K)^{-1}$ is a group, we have

$$\{1\} \cap \varphi(K) \cdot \varphi\left(\bigcup S_{ij}\right) = \emptyset \quad \Rightarrow \quad \varphi(K) \cap \varphi\left(\bigcup S_{ij}\right) = \emptyset.$$

Now, let $\mathring{\Gamma} = \varphi^{-1} \circ \varphi(K)$. This is a finite-index subgroup of Γ . Let $\mathring{M} = \mathbb{H}^3 / \mathring{\Gamma}$. Since $\text{Stab}_{\mathring{\Gamma}}(\tilde{A}) = K = \text{Stab}_{\mathring{\Gamma}}(\tilde{A})$, the horocusp $\mathring{A} = \tilde{A}/K \subset \mathring{M}$ is a lift of $A \subset M$.

By the above displayed equation, $\varphi(\mathring{\Gamma}) = \varphi(K)$ is disjoint from $\varphi(S_{ij})$, hence $\mathring{\Gamma}$ is disjoint from S_{ij} for every $0 \leq i < j \leq n$. Thus \tilde{B}_i and \tilde{B}_j belong to different $\mathring{\Gamma}$ -orbits and project to distinct cusps in \mathring{M} . Similarly, \tilde{A} and \tilde{B}_j belong to different $\mathring{\Gamma}$ -orbits and project to distinct cusps in \mathring{M} . Thus the geodesic arcs $\mathring{\gamma}_1, \dots, \mathring{\gamma}_n$, namely, the quotients of $\tilde{\gamma}_1, \dots, \tilde{\gamma}_n$ in \mathring{M} , lead to cusps of \mathring{M} that are distinct from one another and from \mathring{A} . □

Returning to our plan for proving Theorem 1.3, suppose that $A \subset M$ is a horocusp and that $\{\gamma_1, \dots, \gamma_n\}$ is the set of all orthogeodesics of minimal length from A to the other cusps (including itself). Following Lemma 5.2, we find a cover \mathring{M} where A lifts to \mathring{A} and where the path-lifts $\mathring{\gamma}_1, \dots, \mathring{\gamma}_n$ lead to cusps $\mathring{B}_1, \dots, \mathring{B}_n$ that are distinct from one another and from \mathring{A} . Since the \mathring{B}_i are distinct, we may adjust their sizes independently. We keep $\mathring{B} = \mathring{B}_1$ and shrink $\mathring{B}_2, \dots, \mathring{B}_n$ slightly. Now, $\mathring{\gamma}_1$ is the unique shortest path in \mathring{M} from \mathring{A} to the other horocusps.

In Step 2, we apply Proposition 3.1 to shrink \mathring{A} by a sufficiently large distance (keeping the name \mathring{A}) so that the canonical polyhedral decomposition $\mathring{\mathcal{P}}$ determined by $\mathring{A}, \mathring{B}$, and the remaining cusps meets \mathring{A} in one or two 3-cells, each with one ideal vertex at \mathring{A} . By Corollary 3.4, these cells glue together to form an embedded drilled ananas $\mathring{N} \subset \mathring{M}$.

If A is a non-rectangular cusp, then its lift \mathring{A} is also non-rectangular. Thus, by Proposition 3.1, the drilled ananas \mathring{N} consists of two isometric acute-angled tetrahedra.

Definition 5.3. Let \mathcal{P} be an ideal polyhedral decomposition of a cusped hyperbolic 3-manifold M . A *diagonal* of \mathcal{P} is a bi-infinite geodesic β that is contained in some cell of \mathcal{P} . A diagonal β is called *returning* if its endpoints are in the same horocusp of M .

Step 3 of the proof is to apply the following result due to Luo, Schleimer, and Tillmann [22, Lemmas 8 and 9].

Lemma 5.4. *Let \mathring{M} be a cusped hyperbolic 3-manifold with a polyhedral decomposition $\mathring{\mathcal{P}}$. Then there is a finite regular cover $\hat{f} : \hat{M} \rightarrow \mathring{M}$, such that the lifted polyhedral decomposition $\hat{\mathcal{P}}$ has no returning diagonals.*

Proof. This is proved in [22, Lemmas 8 and 9], using a fairly straightforward application of Proposition 4.3. □

We can now prove Theorem 1.3: a cusped hyperbolic 3-manifold M containing a non-rectangular cusp A has a cover \hat{M} with infinitely many geometric ideal triangulations, organized in a trivalent tree. We stress that each edge of this tree represents a sequence of n geometric Pachner moves as described by Corollary 3.8, where n is the degree of the local cover $\hat{N} \rightarrow \mathring{N}$ of the ananas $\mathring{N} \subset \mathring{M}$.

Proof of Theorem 1.3. Let M be a cusped hyperbolic 3-manifold, and let $A \subset M$ be a non-rectangular cusp. Consider the sequence of finite covers

$$\widehat{M} \xrightarrow{\widehat{f}} \mathring{M} \xrightarrow{\mathring{f}} M$$

constructed in Lemma 5.2 and Lemma 5.4. In particular, \mathring{M} has a polyhedral decomposition $\mathring{\mathcal{P}}$ such that two acute-angled ideal tetrahedra of $\mathring{\mathcal{P}}$ fit together to form a drilled ananas \mathring{N} . Consequently, $\partial \mathring{N}$ consists of two triangular faces of $\mathring{\mathcal{P}}$.

Let $\widehat{\mathcal{P}}$ be the polyhedral decomposition of \widehat{M} obtained by pulling back $\mathring{\mathcal{P}}$. Then $\widehat{\mathcal{P}}$ contains a submanifold \widehat{N} that covers \mathring{N} . By Lemma 5.4, $\widehat{\mathcal{P}}$ has no returning diagonals; hence, the vertices of every polyhedron $P \subset \widehat{\mathcal{P}}$ are mapped to distinct cusps of \widehat{M} .

Choose an ordering $<$ on the cusps of \widehat{M} . Then, for every polyhedron $P \subset \widehat{\mathcal{P}}$, we get a total ordering of the vertices of P . Thus, by Lemma 2.9, the iterated coning induced by $<$ subdivides $\widehat{\mathcal{P}}$ into geometric ideal tetrahedra. By construction, \widehat{N} already consists of tetrahedra, so does not need to be subdivided. Now, by Corollary 3.8, the initial geometric triangulation of \widehat{N} (which comes from lifting the two-tetrahedron triangulation of \mathring{N}) is the start of an infinite sequence of geometric ideal triangulations.

Since A is non-rectangular, the triangulation of the drilled ananas \mathring{N} consists of two acute-angled tetrahedra. Thus every path in the trivalent tree of geometric triangulations of \mathring{N} that was described in Proposition 3.10 lifts to a path of geometric triangulations of \widehat{N} , hence to a path of geometric triangulations of \widehat{M} . \square

Remark 5.5. If A is a rectangular cusp of M , the above proof of Theorem 1.3 still constructs covers $\widehat{M} \rightarrow \mathring{M} \rightarrow M$, where \widehat{M} contains a submanifold \widehat{N} that covers the drilled ananas \mathring{N} . However, this time \widehat{N} consists of rectangular pyramids that need to be subdivided into tetrahedra. The subdivision imposed by ordering the cusps of \widehat{M} may impose different choices of diagonals on the rectangles of $\partial \widehat{N}$, which would obstruct the triangulation of \widehat{N} from being equivariant with respect to the cover $\widehat{N} \rightarrow \mathring{N}$. This means that we cannot apply Corollary 3.8 to obtain an infinite sequence of triangulations. The issue of equivariance does not arise if \widehat{N} is a lift of \mathring{N} .

In the next section, we will deploy Theorem 1.4 to construct a cover \widehat{M} where \widehat{N} is indeed a lift, enabling us to handle rectangular cusps. This will have the additional benefit that each edge of the trivalent tree will represent a single geometric 2–3 move, as in Proposition 3.10.

6 | RECTANGULAR CUSPS AND DEHN FILLINGS

In this section, we prove Theorem 1.2, which extends Theorem 1.3 to manifolds with rectangular cusps, and also provides infinitely many geometric triangulations of long Dehn fillings of \widehat{M} . The proof begins in the same way as Steps 1 and 2 of the four-step outline described at the start of Section 5. In particular, we will find a cover $\mathring{M} \rightarrow M$ that contains a drilled ananas \mathring{N} . The main challenge, as mentioned in Remark 5.5, is to build further covers where \mathring{N} continues to lift but (most) returning diagonals stop being returning. Before outlining how to do this, we introduce a definition and a motivating example.

Definition 6.1. Let $M = \mathbb{H}^3/\Gamma$ be a cusped hyperbolic 3-manifold with a distinguished horocusp A and a polyhedral decomposition \mathcal{P} . Let $\tilde{A} \subset \mathbb{H}^3$ be a horoball covering A , and let $\tilde{\mathcal{P}}$ be the lifted polyhedral decomposition. An \tilde{A} -parabolic diagonal of $\tilde{\mathcal{P}}$ is a bi-infinite geodesic $\tilde{\beta}$ contained in a cell of $\tilde{\mathcal{P}}$, whose ends are in horoballs \tilde{C} and \tilde{C}' such that there is a parabolic isometry $g \in \text{Stab}_\Gamma(\tilde{A})$ with $g(\tilde{C}) = \tilde{C}'$.

An A -parabolic diagonal of \mathcal{P} is the projection $\beta \subset M$ of an \tilde{A} -parabolic diagonal of $\tilde{\mathcal{P}}$, for some horoball \tilde{A} covering A . We remark that the choice of \tilde{A} is immaterial: for any other horoball \tilde{A}' covering A , there will be some \tilde{A}' -parabolic lift of β . On the other hand, the choice of \mathcal{P} can affect the collection of A -parabolic diagonals, because it affects the set of bi-infinite geodesics that are diagonals of \mathcal{P} in the first place. We also remark that an A -parabolic diagonal is necessarily returning, according to Definition 5.3.

A horocusp $C \subset M$ is called A -problematic (relative to \mathcal{P}) if there is an A -parabolic diagonal β of \mathcal{P} whose endpoints lie in C .

Example 6.2. Suppose, as in the conclusion of Corollary 3.4, that $M = \mathbb{H}^3/\Gamma$ contains a drilled ananas N that is obtained by gluing one or two cells of \mathcal{P} . Let $A \subset M$ be a horocusp containing the cusp of N , and let $B \subset M$ be a horocusp containing the thorn of N . Now, consider a geodesic $\beta \subset \partial N$ that lies in a 2-cell of ∂N . We claim that β must be an A -parabolic diagonal of \mathcal{P} . Indeed, both endpoints of β are in the single thorn of N ; hence, every lift of β to \mathbb{H}^3 must have its endpoints in horoballs that are permuted by $\text{Stab}_\Gamma(\tilde{A})$ for an appropriate horoball \tilde{A} covering A . Thus, by Definition 6.1, B is an A -problematic cusp relative to \mathcal{P} .

Now, suppose that $\bar{M} = \mathbb{H}^3/\bar{\Gamma}$ is some finite cover of M where A lifts to \bar{A} . Since A lifts, or equivalently $\text{Stab}_\Gamma(\tilde{A}) = \text{Stab}_{\bar{\Gamma}}(\bar{\tilde{A}})$, every A -parabolic diagonal β of \mathcal{P} lifts to an \bar{A} -parabolic diagonal $\bar{\beta}$ of the lifted polyhedral decomposition $\bar{\mathcal{P}}$. In particular, some preimage of β continues to be a returning diagonal in \bar{M} .

The gist of the following outline is that Example 6.2 is a worst-case scenario that can be isolated and handled. With Theorem 1.4 and with enough care, all diagonals that are not in the boundary of a drilled ananas eventually lift to be non-returning, while the ananas continues to lift.

Now, let M be a cusped hyperbolic 3-manifold containing a horocusp A . We will take the following sequence of steps.

- Step 1. Describe a criterion on the distance from A that any A -problematic cusp must satisfy. This is accomplished in Lemma 6.3.
- Step 2. Using the criterion of Lemma 6.3, find a finite cover $\hat{M} \rightarrow M$, where A lifts to \hat{A} . This cover \hat{M} contains a polyhedral decomposition $\hat{\mathcal{P}}$ and a drilled ananas \hat{N} with its cusp in \hat{A} and its thorn in \hat{B} , such that \hat{B} is the only \hat{A} -problematic cusp of \hat{M} . In fact, all \hat{A} -parabolic diagonals of $\hat{\mathcal{P}}$ lie in $\partial \hat{N}$. See Lemma 6.4 for details.
- Step 3. Using Theorem 1.4, find a finite cover $\bar{M} \rightarrow \hat{M}$, where all of the above features hold (in particular, \hat{N} lifts to \bar{N}), and in addition every polyhedron of $\bar{\mathcal{P}}$ has some vertex in a horocusp other than \bar{B} , the thorn of \bar{N} . See Lemma 6.5 for details.
- Step 4. Using Theorem 1.4 again, find a finite cover $\check{M} \rightarrow \bar{M}$, where all of the above features hold (in particular, \bar{N} lifts to \check{N}), and in addition all returning diagonals have their endpoints in cusps that cover \bar{B} . This means that there are very few returning diagonals, and the partial order argument of Corollary 2.10 suffices to find a geometric triangulation $\check{\mathcal{T}}$ that is compatible with infinitely many geometric triangulations of \check{M} . See Lemma 6.6 for details.

Step 5. Using $H_1(\tilde{M})$, find a double cover $\hat{M} \rightarrow \tilde{M}$ where \tilde{N} has two distinct lifts, called \hat{N} and \hat{N}' . See Lemma 6.7 for details. We will replace one lift \hat{N} with a triangulated solid torus to perform a long Dehn filling $\hat{M}(s)$, while using the other lift \hat{N}' to obtain infinitely many geometric triangulations of $\hat{M}(s)$.

We now proceed to carry out these steps in detail. In the following lemma, $\mathcal{M}(T^2)$ is the moduli space of unit-area flat tori, and $\mathbb{R}_+ \mathcal{M}(T^2)$ is the moduli space of flat tori of any area.

Lemma 6.3. *There is a function $L : \mathbb{R}_+ \times \mathbb{R}_+ \mathcal{M}(T^2) \rightarrow \mathbb{R}_+$ such that the following holds for every multi-cusped hyperbolic 3-manifold M .*

Suppose that M contains a horocusp collection A, B_1, \dots, B_k and that α is an orthogeodesic from A to B_1 . Then, in the canonical polyhedral decomposition \mathcal{P} determined by A, B_1, \dots, B_k , any A -problematic horocusp B_j must satisfy $d(A, B_j) < L = L(\text{len}(\alpha), \partial A)$.

Furthermore, if A is replaced by a sub-horocusp A^t for some $t > 0$, then the distance bound L is replaced by $L + t$. In symbols,

$$L(\text{len}(\alpha) + t, \partial A^t) = L(\text{len}(\alpha), \partial A) + t.$$

One particular consequence of Lemma 6.3 is that the length bound L only depends on the horocusps A and B_1 . Although varying the sizes of B_2, \dots, B_k may have the effect of changing the polyhedral decomposition \mathcal{P} , thereby changing the collection of A -parabolic diagonals of \mathcal{P} , the conclusion of the lemma still holds for the same L .

Proof. Write $M = \mathbb{H}^3/\Gamma$, and conjugate Γ so that A is covered by a horoball \tilde{A} about ∞ . Then α has a lift $\tilde{\alpha}$ that leads from \tilde{A} to a horoball \tilde{B}_1 covering B_1 . A further conjugation, preserving the point ∞ , ensures that \tilde{B}_1 has Euclidean diameter exactly 1. Then, setting $\ell = \text{len}(\alpha) = \text{len}(\tilde{\alpha})$, it follows that $\partial\tilde{A}$ lies at Euclidean height e^ℓ . Let $K = \text{Stab}_\Gamma(\tilde{A})$. Note that the length $\ell \in \mathbb{R}_+$ and the Euclidean metric $\partial A \in \mathbb{R}_+ \mathcal{M}(T^2)$ determine the orbit of horoballs $K\tilde{B}_1$ up to Euclidean isometry (equivalently, up to a hyperbolic isometry stabilizing $\partial\tilde{A}$).

Suppose that $\tilde{\beta} \subset \mathbb{H}^3$ is an \tilde{A} -parabolic diagonal of $\tilde{\mathcal{P}}$. Let \tilde{C}, \tilde{C}' be the horoballs at the endpoints of $\tilde{\beta}$, and let \tilde{P} be a polyhedron containing $\tilde{\beta}$. Then, as in Definition 2.6, \tilde{P} contains the center of a metric ball D that is tangent to the horoballs about its vertices, including \tilde{C} and \tilde{C}' , and is disjoint from all other horoballs in the packing.

Let h denote the Euclidean diameter of \tilde{C} , which is equal to the Euclidean diameter of \tilde{C}' because $\tilde{C}' \in K\tilde{C}$. Then $d(\tilde{A}, \tilde{C}) = \ell - \log h$. We will see that when $h \ll 1$, or equivalently $d(\tilde{A}, \tilde{C}) \gg 0$, competing pressures on the diameter of D lead to a contradiction.

Let w be the shortest Euclidean translation length (along $\mathbb{C} = \partial\mathbb{H}^3 \setminus \{\infty\}$) of any element of K . Thus the Euclidean distance between the centers of \tilde{C} and \tilde{C}' is at least w . Since D must be tangent to \tilde{C} and \tilde{C}' but disjoint from \mathbb{C} , its Euclidean diameter is bounded below by a function of w and h that grows without bound as $h \rightarrow 0$. On the other hand, a ball of large Euclidean diameter whose lowest point is below Euclidean height h must intersect one of the diameter 1 horoballs in the K -orbit of \tilde{B}_1 . Compare to Figure 2. Thus every sufficiently small value $h \ll 1$ leads to a contradiction, and there is an upper bound L on $d(\tilde{A}, \tilde{C}) = \ell - \log h$. Observe that L depends only on the lattice of horoballs $K\tilde{B}_1$, hence on $\ell = \text{len}(\alpha)$ and the Euclidean metric on ∂A . Thus we may write $L = L(\text{len}(\alpha), \partial A)$.

To prove the ‘‘furthermore,’’ suppose that we replace A by a sub-horocusp A^t . This has the effect of replacing ℓ by $\ell + t$ and replacing \tilde{A} by a horoball \tilde{A}^t at Euclidean height $e^{\ell+t}$. Then the

Euclidean length w and the lattice of horoballs $K\tilde{B}_1$ both remain the same; hence, the same value of $h \ll 1$ leads to a contradiction. However, the distance $d(\tilde{A}, \tilde{C})$ has just increased by t . Thus replacing A by A^t has the effect of replacing L by $L + t$. □

We can now begin constructing covers of a cusped hyperbolic 3-manifold M containing a horocusp A . In the following lemma, corresponding to Step 2, we build a cover $\mathring{M} \rightarrow M$ that supports a polyhedral decomposition \mathring{P} that contains a drilled ananas \mathring{N} with its cusp in \mathring{A} , such that all \mathring{A} -parabolic diagonals of \mathring{P} lie in $\partial\mathring{N}$.

Lemma 6.4. *Let M be a cusped hyperbolic 3-manifold and $A \subset M$ a horocusp. Then there is a finite cover $f : \mathring{M} \rightarrow M$ such that the following hold.*

- \mathring{M} contains a horocusp collection $\mathring{A}, \mathring{B} = \mathring{B}_1, \dots, \mathring{B}_k$, where $k \geq 2$ and $\mathring{A} \subset \mathring{M}$ is a lift of A .
- There is an orthogeodesic $\mathring{\alpha}$ from \mathring{A} to \mathring{B} that is the unique shortest path from \mathring{A} to $\cup \mathring{B}_j$.
- For large $t > 0$, the polyhedral decomposition $\mathring{P} = \mathring{P}^t$ determined by $\mathring{A}^t, \mathring{B}_1, \dots, \mathring{B}_k$ contains a drilled ananas \mathring{N} , built out of one or two pyramids whose lateral edges are identified to $\mathring{\alpha}$, with its cusp in \mathring{A} and its thorn in \mathring{B} .
- For every horoball \tilde{A} covering \mathring{A} , there is a corresponding preimage \tilde{N} of \mathring{N} , such that all \tilde{A} -parabolic diagonals of \tilde{P} lie in $\partial\tilde{N}$.

The main point in Lemma 6.4 is the last bullet, as it provides a partial converse to Example 6.2.

Proof. Choose a collection of disjoint horocusps, containing A . Relative to this collection of horocusps, let α be a shortest orthogeodesic in M that starts at A . Let $L = L(\text{len}(\alpha), \partial A)$ be the bound produced by Lemma 6.3. Let $S = \{\gamma_1, \dots, \gamma_n\}$ be a set of $n \geq 2$ orthogeodesics starting at ∂A , containing all the orthogeodesics that have length at most L . We set $\gamma_1 = \alpha$. By Lemma 5.2, there is a finite cover $\mathring{M} \rightarrow M$, where A lifts to a cusp \mathring{A} , and where the $\gamma_i \in S$ have path-lifts $\mathring{\gamma}_1, \dots, \mathring{\gamma}_n$ that start on $\partial\mathring{A}$ and lead to cusps that are distinct from one another and from \mathring{A} . Let \mathring{B}_i be the horocusp at the endpoint of $\mathring{\gamma}_i$.

Let $\mathring{B} = \mathring{B}_1$ be the horocusp at the end of $\mathring{\alpha} = \mathring{\gamma}_1$. We keep \mathring{B} fixed, but shrink each of $\mathring{B}_2, \dots, \mathring{B}_n$ to ensure that $d(\mathring{A}, \mathring{B}_i) \geq L$ for $i = 2, \dots, n$. Any other horocusp of \mathring{M} , labeled B_i for $i = n + 1, \dots, k$, must already satisfy $d(\mathring{A}, B_i) \geq L$. In particular, $\mathring{\alpha}$ is the unique shortest path from \mathring{A} to any cusp of \mathring{M} .

Now, \mathring{M} and its collection of horocusps satisfies the hypotheses of Proposition 3.1. Thus, for sufficiently large t , we may replace \mathring{A} with \mathring{A}^t and build a canonical polyhedral decomposition $\mathring{P} = \mathring{P}^t$ that contains one or two ideal pyramids with a vertex in \mathring{A} and with their lateral edges glued to $\mathring{\alpha}$. By Corollary 3.4, these cells glue up to form an embedded, convex drilled ananas $\mathring{N} \subset \mathring{M}$. By construction, \mathring{N} has its cusp in \mathring{A} and its thorn in \mathring{B} .

By Lemma 6.3, any \mathring{A} -problematic cusp \mathring{B}_i in the polyhedral decomposition \mathring{P} must satisfy $d(\mathring{A}, \mathring{B}_i) < L$ or equivalently $d(\mathring{A}^t, \mathring{B}_i) < L + t$. The only horocusp satisfying these hypotheses is $\mathring{B} = \mathring{B}_1$. Thus any \mathring{A} -parabolic diagonal must have its endpoints in \mathring{B} .

Let $\mathring{\beta}$ be an \mathring{A} -parabolic diagonal of $\mathring{P} = \mathring{P}^t$. In the universal cover \mathbb{H}^3 , let \tilde{A} be a horoball covering \mathring{A} , and let \tilde{N} be the preimage of \mathring{N} containing \tilde{A}^t . The \mathring{A} -parabolic diagonal $\mathring{\beta}$ lifts to an \tilde{A} -parabolic diagonal $\tilde{\beta}$, whose endpoints must be in horoballs \tilde{B}, \tilde{B}' such that $d(\tilde{A}, \tilde{B}) < L$. By construction, $\mathring{\alpha}$ is the only orthogeodesic from \mathring{A} to \mathring{B} with length less than L . Therefore, \tilde{B} and \tilde{B}' must be full-sized horoballs connected to \tilde{A} by lifts of $\mathring{\alpha}$. Since the endpoints of $\tilde{\beta}$ are in the ideal vertices of \tilde{N} , and \tilde{N} is convex, it follows that $\tilde{\beta} \subset \tilde{N}$. In particular, $\mathring{\beta}$ must lie in one of the polyhedra comprising \mathring{N} .

To complete the proof, recall that \mathring{N} consists of either two ideal tetrahedra or one rectangular-based ideal pyramid. In either case, we think of the constituent cells as pyramids with bases along $\partial\mathring{N}$ and lateral edges glued to $\mathring{\alpha}$. Any diagonal in an ideal pyramid is either a lateral edge or contained in the base. Since the \mathring{A} -parabolic diagonal $\mathring{\beta}$ has both of its endpoints in \mathring{B} , it cannot be a lateral edge, and must be contained in the base of the ambient pyramid. Thus $\mathring{\beta} \subset \partial\mathring{N}$ and $\tilde{\beta} \subset \partial\tilde{N}$, as claimed. \square

The next lemma, corresponding to Step 3 of the outline, is our first use of Theorem 1.4. Roughly speaking, the lemma says that there is a cover $\overline{M} \rightarrow \mathring{M}$ where the drilled ananas \mathring{N} lifts, and where returning diagonals are controlled to a significant degree.

Lemma 6.5. *Let $\mathring{M} = \mathbb{H}^3/\mathring{\Gamma}$ be a cusped hyperbolic 3-manifold containing a distinguished horocusp \mathring{A} and a drilled ananas \mathring{N} whose cusp is in \mathring{A} and whose thorn is in horocusp \mathring{B} . Suppose that \mathring{P} is a polyhedral decomposition of \mathring{M} , with the following property: for every horoball \tilde{A} covering \mathring{A} , there is a corresponding preimage \tilde{N} of \mathring{N} , such that all \tilde{A} -parabolic diagonals of \tilde{P} lie in $\partial\tilde{N}$. Then there is a finite cover $f : \overline{M} \rightarrow \mathring{M}$ such that the following hold.*

- \mathring{A} lifts to a distinguished cusp \overline{A} .
- \mathring{N} lifts to a drilled ananas \overline{N} whose cusp is in \overline{A} and whose thorn is in horocusp \overline{B} .
- Every \tilde{A} -parabolic diagonal of \tilde{P} has its endpoints in \overline{B} .
- Every polyhedron $\tilde{P} \subset \tilde{M}$ has a vertex in some horocusp other than \overline{B} .

Proof. Let \tilde{A} be a horoball covering \mathring{A} , and let $K = \text{Stab}_{\mathring{\Gamma}}(\tilde{A})$. Let $\tilde{N} \subset \mathbb{H}^3$ be the preimage of \mathring{N} containing \tilde{A} . Then there is a horoball \tilde{B} covering \mathring{B} , such that all ideal vertices of \tilde{N} lie in $\tilde{A} \cup K\tilde{B}$. By hypothesis, all \tilde{A} -parabolic diagonals of \tilde{P} lie in $\partial\tilde{N}$ and have their endpoints in horoballs of $K\tilde{B}$. Let $\mathring{P}_1, \dots, \mathring{P}_n$ be the polyhedra of \mathring{P} that have all of their vertices in \mathring{B} . (If no such polyhedra exist, we may simply set $\overline{M} = \mathring{M}$ and let f be the identity map.) For each \mathring{P}_i , let $\mathring{\beta}_i$ be an edge that is not \mathring{A} -parabolic. Such an edge must exist, because all \mathring{A} -parabolic edges belong to $\partial\mathring{N}$.

For each $\mathring{\beta}_i$, choose a preimage $\tilde{\beta}_i \subset \mathbb{H}^3$. The ends of $\tilde{\beta}_i$ lie in horoballs $\tilde{B}_i, \tilde{B}'_i$, which must cover \mathring{B} because both endpoints of $\mathring{\beta}_i$ are in \mathring{B} . Thus there is an isometry $g_i \in \mathring{\Gamma} = \pi_1(\mathring{M})$ such that $g_i(\tilde{B}_i) = \tilde{B}'_i$. By Lemma 5.1, the set of all isometries in $\mathring{\Gamma}$ taking \tilde{B}_i to \tilde{B}'_i is a left coset $g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i)$. Since $\mathring{\beta}_i$ is not \mathring{A} -parabolic, the coset $g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i)$ is disjoint from all $\mathring{\Gamma}$ -conjugates of $K = \text{Stab}_{\mathring{\Gamma}}(\tilde{A})$. Equivalently, K is disjoint from all $\mathring{\Gamma}$ -conjugates of $g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i)$.

By Theorem 1.4, there is a homomorphism $\varphi_i : \mathring{\Gamma} \rightarrow G_i$, where G_i is a finite group, such that $\varphi_i(K)$ is disjoint from all G_i -conjugates of $\varphi_i(g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i))$. We can now consider the product homomorphism

$$\overline{\varphi} = (\varphi_1, \dots, \varphi_n) : \mathring{\Gamma} \longrightarrow G = G_1 \times \dots \times G_n.$$

Then, for each i , the image $\overline{\varphi}(K)$ is disjoint from all G -conjugates of $\overline{\varphi}(g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i))$.

Now, let $\overline{\Gamma} = \overline{\varphi}^{-1} \circ \overline{\varphi}(K)$, and let $\overline{M} = \mathbb{H}^3/\overline{\Gamma}$. We get a covering map $f : \overline{M} \rightarrow \mathring{M}$. Then, by construction, every $\mathring{\Gamma}$ -conjugate of $g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i)$ is disjoint from $\overline{\Gamma}$. Since $K = \text{Stab}_{\mathring{\Gamma}}(\tilde{A}) = \text{Stab}_{\overline{\Gamma}}(\tilde{A})$, the horocusp $\mathring{A} \subset \mathring{M}$ lifts to a horocusp $\overline{A} = \tilde{A}/K \subset \overline{M}$. Similarly, \mathring{N} lifts to a drilled ananas $\overline{N} = \tilde{N}/K \subset \overline{M}$. The thorn of \overline{N} is in the horocusp \overline{B} that is covered by \tilde{B} , hence \overline{B} covers \mathring{B} . This proves the first two bullets in the lemma.

For the next bullet, let $\tilde{\gamma}$ be an \tilde{A} -parabolic diagonal in \tilde{P} . Since $\mathbb{H}^3 \rightarrow \overline{M}$ is a regular cover, we may choose a lift $\tilde{\gamma}$ that is an \tilde{A} -parabolic diagonal in \tilde{P} . Recall that all \tilde{A} -parabolic diagonals in

\tilde{P} lie in $\partial\tilde{N}$ and have their endpoints in horoballs of $K\tilde{B}$. Since $K \subset \bar{\Gamma}$, it follows that $\bar{\gamma} \subset \partial\bar{N}$ has its endpoints in \bar{B} , as desired.

To prove the remaining conclusion, let \bar{P} be a polyhedron of $\bar{\mathcal{P}}$. If \bar{P} has an ideal vertex in some cusp that does not belong to $\bar{f}^{-1}(\bar{B})$, then certainly \bar{P} has a vertex that is not in \bar{B} . Otherwise, $\bar{P} = \bar{P}_i$ is a lift of some \mathring{P}_i . Let $\bar{\beta}_i \subset \bar{P}_i$ be a lift of $\mathring{\beta}_i \subset \mathring{P}_i$. We claim that the endpoints of $\bar{\beta}_i$ are in distinct cusps of \bar{M} , and in particular one endpoint is not in \bar{B} .

Let $\tilde{\gamma}_i \subset \mathbb{H}^3$ be an arbitrary preimage of $\bar{\beta}_i$, and let $\tilde{C}_i, \tilde{C}'_i$ be horoballs in the packing containing the ends of $\tilde{\gamma}_i$. Since $\tilde{\gamma}_i$ and $\tilde{\beta}_i$ both cover $\mathring{\beta}_i \subset \mathring{M}$, there is an isometry $h_i \in \mathring{\Gamma}$ such that $h_i(\tilde{\gamma}_i) = \tilde{\beta}_i$, which implies $h_i(\tilde{C}_i) = \tilde{B}_i$ and $h_i(\tilde{C}'_i) = \tilde{B}'_i$. Thus the set of all isometries in $\mathring{\Gamma}$ taking \tilde{C}_i to \tilde{C}'_i can be written as $h_i^{-1} \cdot g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i) \cdot h_i$. By construction, this conjugate of $g_i \text{Stab}_{\mathring{\Gamma}}(\tilde{B}_i)$ is disjoint from $\bar{\Gamma}$. Thus \tilde{C}_i and \tilde{C}'_i lie in distinct $\bar{\Gamma}$ -orbits, which means that the endpoints of $\bar{\beta}_i$ are in distinct cusps. This proves the claim and the lemma. \square

The next lemma, corresponding to Step 4 of the outline, builds a cover \check{M} with even stronger restrictions on the returning diagonals of the polyhedral decomposition $\check{\mathcal{P}}$.

Lemma 6.6. *Let $\bar{M} = \mathbb{H}^3/\bar{\Gamma}$ be a cusped hyperbolic 3-manifold containing a distinguished horocusp \bar{A} and a horocusp \bar{B} . Suppose that $\bar{\mathcal{P}}$ is a polyhedral decomposition of \bar{M} , such that every \bar{A} -parabolic diagonal of $\bar{\mathcal{P}}$ has its endpoints in \bar{B} . Then there is a finite cover $f: \check{M} \rightarrow \bar{M}$ where \bar{A} lifts to a distinguished cusp \check{A} , such that every returning diagonal of $\check{\mathcal{P}}$ has its endpoints in $f^{-1}(\bar{B})$.*

Proof. Let $R = \{\bar{\beta}_1, \dots, \bar{\beta}_n\}$ be the set of returning diagonals of $\bar{\mathcal{P}}$ whose endpoints are not in \bar{B} . This set is finite because $\bar{\mathcal{P}}$ has finitely many polyhedra, and each polyhedron contains finitely many diagonals. By hypothesis, every diagonal $\bar{\beta}_i \in R$ is not \bar{A} -parabolic. Assume that $R \neq \emptyset$, as otherwise we may simply take $\check{M} = \bar{M}$.

Let \bar{A} be a horoball covering \bar{A} , and let $K = \text{Stab}_{\bar{\Gamma}}(\bar{A})$. For each $\bar{\beta}_i$, choose a preimage $\tilde{\beta}_i \subset \mathbb{H}^3$. The ends of $\tilde{\beta}_i$ lie in horoballs $\tilde{B}_i, \tilde{B}'_i$, which cover the same cusp of \bar{M} because $\bar{\beta}_i$ is a returning diagonal. Thus there is an isometry $g_i \in \bar{\Gamma} = \pi_1(\bar{M})$ such that $g_i(\tilde{B}_i) = \tilde{B}'_i$. By Lemma 5.1, the set of all isometries in $\bar{\Gamma}$ taking \tilde{B}_i to \tilde{B}'_i is a left coset $g_i \text{Stab}_{\bar{\Gamma}}(\tilde{B}_i)$. Since $\bar{\beta}_i$ is not \bar{A} -parabolic, the coset $g_i \text{Stab}_{\bar{\Gamma}}(\tilde{B}_i)$ is disjoint from all $\bar{\Gamma}$ -conjugates of $K = \text{Stab}_{\bar{\Gamma}}(\bar{A})$.

As in the proof of Lemma 6.5, we use Theorem 1.4 to find a finite-index subgroup $\check{\Gamma} \subset \bar{\Gamma}$ that contains K and is disjoint from every $\bar{\Gamma}$ -conjugate of $g_i \text{Stab}_{\bar{\Gamma}}(\tilde{B}_i)$. Let $\check{M} = \mathbb{H}^3/\check{\Gamma}$. Since $K \subset \check{\Gamma}$, the horocusp $\bar{A} \subset \bar{M}$ lifts to a horocusp $\check{A} \subset \check{M}$. As in the proof of Lemma 6.5, the disjointness of $\check{\Gamma}$ and all $\bar{\Gamma}$ -conjugates of $g_i \text{Stab}_{\bar{\Gamma}}(\tilde{B}_i)$ implies that every lift $\check{\beta}_i$ of $\bar{\beta}_i$ has endpoints in distinct cusps, and is not a returning diagonal. Thus any returning diagonal in \check{M} must be the preimage of a returning diagonal in \bar{M} whose endpoints are in \bar{B} ; hence, it has endpoints in $f^{-1}(\bar{B})$. \square

Following Lemma 6.6, the manifold \check{M} has so few returning diagonals that it is possible to impose a partial order $<$ on the cusps of \check{M} such that every polyhedron $\check{P} \subset \check{\mathcal{P}}$ has a unique lowest vertex. Using Corollary 2.10, we can refine $\check{\mathcal{P}}$ to a geometric triangulation \mathcal{T} , and apply Lemma 3.6 to build infinitely many geometric triangulations of \check{M} . See Claim 6.8 below for details.

To find geometric triangulations of Dehn fillings, we need to take one more cover.

Lemma 6.7. *Let $\check{M} = \mathbb{H}^3/\check{\Gamma}$ be a cusped hyperbolic 3-manifold containing at least three cusps and a distinguished horocusp \check{A} . Then there is a double cover $\hat{f}: \hat{M} \rightarrow \check{M}$, where \check{A} has two distinct lifts.*

Proof. Since \check{M} is the interior of a compact 3-manifold with at least three boundary tori, the “half lives, half dies” lemma [19, Lemma 3.5] implies that $H_1(\check{M})$ has a \mathbb{Z}^n direct summand for $n \geq 3$. Let G be the subgroup of $H_1(\check{M})$ induced by the inclusion $\check{A} \rightarrow \check{M}$. Since G has rank at most 2, there must be a primitive, infinite-order homology class $h \in H_1(\check{M})$ such that $\langle h \rangle$ is a direct summand that is linearly independent from G . Thus we may define a projection $\pi_h : H_1(\check{M}) \rightarrow \langle h \rangle$ such that $G \subset \ker(\pi_h)$.

Now, consider the sequence of surjective homomorphisms

$$\check{\Gamma} = \pi_1(\check{M}) \xrightarrow{\text{ab}} H_1(\check{M}) \xrightarrow{\pi_h} \langle h \rangle \cong \mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z},$$

where ab is abelianization. Let $\hat{\varphi} : \check{\Gamma} \rightarrow \mathbb{Z}/2\mathbb{Z}$ be the composition, and let $\hat{\Gamma} = \ker(\hat{\varphi})$. By construction, $G \subset \ker(\pi_h)$, hence $\pi_1(\check{A}) \subset \ker(\hat{\varphi}) = \hat{\Gamma}$. By the lifting criterion, \check{A} lifts to $\hat{M} = \mathbb{H}^3/\hat{\Gamma}$. Since $\hat{f} : \hat{M} \rightarrow \check{M}$ is a regular cover of degree 2, there must be two distinct lifts. \square

We can now prove our main result, Theorem 1.2.

Proof of Theorem 1.2. Let M be a cusped hyperbolic 3-manifold containing a horocusp A . Consider the sequence of finite covers

$$\hat{M} \xrightarrow{\hat{f}} \check{M} \xrightarrow{\check{f}} \bar{M} \xrightarrow{\bar{f}} \overset{\circ}{M} \xrightarrow{f} M$$

constructed in the preceding lemmas. The cusp $A \subset M$ lifts along each covering map. Recall that by Lemma 6.4, $\overset{\circ}{M}$ has at least three cusps, hence \bar{M} does also, and we may indeed apply Lemma 6.7 to construct \hat{f} .

By Lemma 6.4, $\overset{\circ}{M}$ has a polyhedral decomposition $\overset{\circ}{P}$ such that one or two ideal 3-cells of $\overset{\circ}{P}$ fit together to form a drilled ananas $\overset{\circ}{N}$ that deformation retracts to $\overset{\circ}{A}$. The 1-skeleton of $\overset{\circ}{P}$ decomposes $\partial\overset{\circ}{N}$ into two ideal triangles or one ideal rectangle. Furthermore, since $\overset{\circ}{A}$ lifts to $\bar{A} \subset \bar{M}$ and $\check{A} \subset \check{M}$, the drilled ananas $\overset{\circ}{N}$ lifts to $\bar{N} \subset \bar{M}$ and $\check{N} \subset \check{M}$.

Claim 6.8. The polyhedral decomposition \check{P} of \check{M} can be refined to a geometric ideal triangulation $\check{\mathcal{T}}$. If $\partial\check{N}$ is a single ideal rectangle, then we may choose either diagonal of this rectangle to be an edge in $\check{\mathcal{T}}$. Finally, $\check{\mathcal{T}}$ is the start of an infinite sequence of geometric triangulations of \check{M} .

Recall, from Lemma 6.5, that the drilled ananas $\bar{N} \subset \bar{M}$ has its thorn in a horocusp \bar{B} , such that every polyhedron $\bar{P} \subset \bar{P}$ has a vertex in some horocusp apart from \bar{B} . Thus, in the lifted polyhedral decomposition \check{P} of \check{M} , every polyhedron \check{P} must have at least one vertex in a horocusp that is not in $\check{f}^{-1}(\bar{B})$. We call the cusps of $\check{f}^{-1}(\bar{B})$ blue. By Lemma 6.6, all returning diagonals of \check{P} have their endpoints in blue cusps.

Let V be the set of cusps of \check{M} . We impose a partial order $<$ on V as follows: the non-blue cusps are totally ordered in some fashion; the blue cusps are pairwise incomparable; and $\check{C} < \check{B}$ for every blue cusp \check{B} and non-blue cusp \check{C} . Since every polyhedron $P \subset \check{P}$ has at least one non-blue ideal vertex, and the non-blue vertices of P are totally ordered below the blue ones, it follows that P has a unique $<$ -minimal vertex. Thus, by Lemma 2.9, the iterated coning of \check{P} induced by $<$ produces \check{P}' , a well-defined subdivision of \check{P} into geometric ideal pyramids. By Corollary 2.10, any choice of diagonals in the non-triangular faces of \check{P}' produces a geometric ideal triangulation $\check{\mathcal{T}}$.

By construction, the thorn of \check{N} is in a blue cusp of $\check{f}^{-1}(\bar{B})$. Thus the partial order $<$ does not impose any ordering on the ideal vertices of $\partial\check{N}$. If this boundary is a single ideal rectangle, the coning induced by $<$ in Lemma 2.9 does not subdivide it, and we may choose our preferred diagonal to subdivide \check{N} into two ideal tetrahedra. By Lemma 3.6, \check{N} admits infinitely many geometric ideal triangulations, hence \check{M} does also. This proves Claim 6.8. \diamond

Next, we lift the triangulation \check{T} of \check{M} to a geometric triangulation \hat{T} of \hat{M} . Then the infinite sequence of geometric triangulations of \check{M} lifts to an infinite sequence of geometric triangulations of \hat{M} , as claimed in the statement of the theorem.

By Lemma 6.7, the horocusp $\check{A} \subset \check{M}$ has two distinct lifts to \hat{M} , which we call \hat{A} and \hat{A}' . Consequently, the drilled ananas $\check{N} \supset \check{A}$ also has two distinct lifts to \hat{M} , namely, $\hat{N} \supset \hat{A}$ and $\hat{N}' \supset \hat{A}'$. The two distinct lifts of \check{A} and \check{N} play distinct roles in the Dehn filling argument.

Claim 6.9. For all but finitely many choices of slope s on $\partial\hat{A}$, the following hold.

- The Dehn filled manifold $\hat{M}(s)$ has a hyperbolic structure where the core curve γ of the filled solid torus is isotopic to a closed geodesic γ_s .
- The ideal tetrahedra of \hat{T} remain geometric in the hyperbolic metric on $\hat{M}(s)$. Each ideal vertex of a tetrahedron of \hat{T} that used to enter cusp \hat{A} now spins about the geodesic γ_s .
- With an appropriate choice of diagonals in Claim 6.8, the union of the two tetrahedra in \hat{N} has convex boundary in $\hat{M}(s)$.

The first two bullets in the claim follow from Thurston’s hyperbolic Dehn surgery theorem [29, Chapter 4]. For a sufficiently long slope s , the hyperbolic metric on $\hat{M}(s)$ is obtained via an arbitrarily small deformation of the metric on \hat{M} . Thus, for every ideal tetrahedron $T \subset \hat{T}$, a sufficiently small deformation of the metric will keep T geometric and positively oriented. As in [29, Section 4.4], the two tips of ideal tetrahedra that enter \hat{A} will now spin about the core geodesic γ_s .

For the last bullet of the claim, suppose first that the cusp $A \subset M$ is non-rectangular (hence, so are its lifts). Then Corollary 3.4 implies that the original drilled ananas \check{N} consists of acute tetrahedra, and $\partial\check{N}$ is strictly convex at all three of its edges. The same properties are preserved in the lift \hat{N} and are still preserved in $\hat{N}(s)$ after a sufficiently small deformation of the metric.

Next, suppose that $A \subset M$ is rectangular. Then Corollary 3.4 implies that the original drilled ananas \check{N} has convex boundary, with interior angles strictly less than π at the two edges of $\partial\check{N} \cap \check{P}$, and an angle of π along the (arbitrary) diagonal of the ideal rectangle. The same properties remain true in the lift $\hat{N} \subset \hat{M}$. When we deform the metric on \hat{M} to obtain $\hat{M}(s)$, the interior angle of π may become $\pi + \epsilon$ for small ϵ , violating convexity, but then the opposite choice of diagonal on $\partial\hat{N}$ will have interior angle $\pi - \epsilon$. Thus an appropriate choice of diagonal in Claim 6.8 keeps $\partial\hat{N}$ convex in $\hat{M}(s)$, proving Claim 6.9. \diamond

To construct geometric triangulations of $\hat{M}(s)$, we need to introduce the solid-torus analog of a drilled ananas. A *filled ananas* is a 3-manifold X homeomorphic to a solid torus with one boundary point removed, and endowed with a complete hyperbolic metric with the following properties. The boundary ∂X is made up of two totally geodesic ideal triangles, with vertices at the removed point. These two ideal triangles are glued by isometry along their edges to form a standard two-triangle triangulation of a once-punctured torus, with shearing and bending allowed along the edges. Furthermore, X is subdivided into geometric ideal tetrahedra.

Claim 6.10. For all but finitely many choices of slope s on $\partial\hat{A}$, the hyperbolic manifold $\hat{M}(s)$ has a geometric triangulation $\hat{T}(s)$ with the following properties. Finitely many tetrahedra of $\hat{T}(s)$

fit together to form a filled ananas $X(s)$. Furthermore, the restriction of $\widehat{\mathcal{T}}(s)$ to the complement $\widehat{M}(s) \setminus X(s)$ is combinatorially isomorphic to the restriction of $\widehat{\mathcal{T}}$ to the complement $\widehat{M} \setminus \widehat{N}$.

This statement is due to Guéritaud and Schleimer, and closely resembles [15, Theorem 1]. Assuming that $\widehat{\mathcal{T}}$ is the canonical triangulation of \widehat{M} with respect to some choice of horocusps, they construct the triangulated filled ananas $X(s)$ and endow it with a geometric structure isometric to the completion of the two spun tetrahedra of \widehat{N} mentioned in Claim 6.9. Then they replace \widehat{N} with $X(s)$, and prove that the resulting triangulation $\widehat{\mathcal{T}}(s)$ is the canonical triangulation of $\widehat{M}(s)$. In fact, the construction of $X(s)$, which occurs in [15, Section 2] and is encapsulated in [15, Corollary 16], only uses the hypotheses that $\widehat{\mathcal{T}}$ is geometric and that \widehat{N} remains convex in $\widehat{M}(s)$. See also [16, Corollary 4.18]. The combinatorial structure of the triangulation of $X(s)$ is closely guided by the combinatorics of the Farey graph \mathcal{F} and the continued fraction expansion of the filling slope s , while the hyperbolic metric on $X(s)$ is constructed using Casson and Rivin's work on angle structures and volume optimization [11, 27]. In particular, the canonicity of $\widehat{\mathcal{T}}$ is not needed in the proof that $\widehat{\mathcal{T}}(s)$ is geometric. This proves Claim 6.10. \diamond

To complete the proof of the theorem, we have the following.

Claim 6.11. For all but finitely many choices of slope s on $\partial\widehat{A}$, the Dehn filled manifold $\widehat{M}(s)$ has an infinite sequence of geometric triangulations connected by geometric 2–3 moves.

Observe that by Claim 6.10, the geometric triangulation $\widehat{\mathcal{T}}(s)$ agrees with $\widehat{\mathcal{T}}$ on the drilled ananas $\widehat{N}' \subset \widehat{M}(s)$. By Lemma 3.6, this two-tetrahedron geometric triangulation of \widehat{N}' is the start of an infinite sequence of geometric triangulations connected by geometric 2–3 moves. \square

The following remark states a version of Theorem 1.3 for manifolds with rectangular cusps. In the statement, a 4–4 *move* is a local move on topological triangulations, which takes an octahedron that has been decomposed into 4 tetrahedra along one of its three internal diagonals and replaces it with a decomposition into 4 tetrahedra along a different internal diagonal. The move is called a *geometric 4–4 move* if both decompositions are into geometric ideal tetrahedra.

Remark 6.12. If $A \subset M$ is a rectangular cusp, the cover \widetilde{M} contains an infinite trivalent tree of geometric ideal triangulations, where one edge of the tree is a geometric 4–4 move and the remaining edges are geometric 2–3 moves. This can be seen as follows. In Claim 6.8 of the above proof, the drilled ananas \widetilde{N} consists of an ideal rectangular pyramid $\widetilde{P} \subset \widetilde{\mathcal{P}}$. After the pyramidal decomposition induced by \prec , the 3-cell \widetilde{P}' glued to \widetilde{P} is also an ideal rectangular pyramid. The two choices of diagonal for the shared face of $\widetilde{P} \cap \widetilde{P}'$ lead to ideal triangulations that differ by a geometric 4–4 move. Each of these choices can serve in Lemma 3.6 as the starting configuration in an infinite sequence of geometric ideal triangulations. In the proof of Proposition 3.10, the dual tree of the Farey complex splits in half along the edge $\frac{0}{1} - \frac{1}{0}$: half of the tree is reachable by geometric 2–3 moves if we choose the diagonal $\frac{1}{1}$, and the other half is reachable if we choose the diagonal $\frac{-1}{1}$. The union of these halves is an infinite trivalent tree of geometric 2–3 moves, with one edge replaced by a 4–4 move.

We close the paper by pointing out that the figure-8 knot complement M does not contain a drilled ananas, because M has only one cusp. Nevertheless, by the work of Dadd and Duan [9], this manifold has an infinite family of geometric triangulations. Thus the presence of a drilled ananas is a sufficient but not necessary condition. We wonder what other features will also guarantee an infinite sequence of geometric triangulations.

ACKNOWLEDGEMENTS

We thank Jessica Purcell, Saul Schleimer, and Henry Segerman for helpful discussions about triangulations. We thank Ian Agol and Matthew Stover for enlightening discussions about separating peripheral double cosets. Henry Wilton, who fielded our questions at several crucial points, deserves particular gratitude.

This project was conceived when the first author visited Oklahoma State University in November 2019, and reached a mature state when the first and third authors visited the University of Arkansas for the Redbud Topology Conference in March 2020. We thank both universities for their hospitality. From there the collaboration proceeded over Zoom, with the second author joining at this virtual stage. The Redbud conference stands out in our memories as a last hurrah of in-person discussion and collaboration, for a significant time to come.

During this project, Futer was partially supported by NSF grant DMS-1907708, while Hoffman was partially supported by Simons Foundation grant #524123.

JOURNAL INFORMATION

The *Journal of Topology* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

REFERENCES

1. H. Akiyoshi, *Finiteness of polyhedral decompositions of cusped hyperbolic manifolds obtained by the Epstein-Penner's method*, Proc. Amer. Math. Soc. **129** (2001), no. 8, 2431–2439.
2. M. D. Baker and D. Cooper, *Finite-volume hyperbolic 3-manifolds contain immersed quasi-Fuchsian surfaces*, Algebr. Geom. Topol. **15** (2015), no. 2, 1199–1228.
3. H. Bass, *Groups of integral representation type*, Pacific J. Math. **86** (1980), no. 1, 15–51.
4. S. C. Chagas and P. A. Zalesskii, *Hyperbolic 3-manifolds groups are subgroup conjugacy separable*, Ann. Sc. Norm. Super. Pisa Cl. Sci. (5) **19** (2019), no. 3, 951–963.
5. A. Champanerkar, I. Kofman, and J. S. Purcell, *Geometry of bi-periodic alternating links*, J. Lond. Math. Soc. (2) **99** (2019), no. 3, 807–830.
6. M. Culler, *Lifting representations to covering groups*, Adv. in Math. **59** (1986), no. 1, 64–70.
7. M. Culler, N. M. Dunfield, M. Goerner, and J. R. Weeks, *SnapPy, a computer program for studying the geometry and topology of 3-manifolds*, Available at <http://snappy.computop.org>
8. M. Culler and P. B. Shalen, *Varieties of group representations and splittings of 3-manifolds*, Ann. of Math. (2) **117** (1983), no. 1, 109–146.
9. B. Dadd and A. Duan, *Constructing infinitely many geometric triangulations of the figure eight knot complement*, Proc. Amer. Math. Soc. **144** (2016), no. 10, 4545–4555.
10. D. B. A. Epstein and R. C. Penner, *Euclidean decompositions of noncompact hyperbolic manifolds*, J. Differential Geom. **27** (1988), no. 1, 67–80.
11. D. Futer and F. Guéritaud, *From angled triangulations to hyperbolic structures, Interactions between hyperbolic geometry, quantum topology and number theory*, Contemporary Mathematics, vol. 541, American Mathematical Society, Providence, RI, 2011, pp. 159–182.
12. M. Goerner, *Geodesic triangulations exist for cusped Platonic manifolds*, New York J. Math. **23** (2017), 1363–1367.
13. F. Guéritaud, *On canonical triangulations of once-punctured torus bundles and two-bridge link complements*, Geom. Topol. **10** (2006), 1239–1284 (With an appendix by David Futer).
14. F. Guéritaud, *Deforming ideal solid tori*, arXiv:0911.3067, 2009.
15. F. Guéritaud and S. Schleimer, *Canonical triangulations of Dehn fillings*, Geom. Topol. **14** (2010), no. 1, 193–242.
16. S. L. Ham and J. S. Purcell, *Geometric triangulations and highly twisted links*, Algebr. Geom. Topol., to appear, arXiv:2005.11899.

17. E. Hamilton, *Finite quotients of rings and applications to subgroup separability of linear groups*, Trans. Amer. Math. Soc. **357** (2005), no. 5, 1995–2006.
18. E. Hamilton, H. Wilton, and P. A. Zalesskii, *Separability of double cosets and conjugacy classes in 3-manifold groups*, J. Lond. Math. Soc. (2) **87** (2013), no. 1, 269–288.
19. A. Hatcher, *Notes on basic 3-manifold topology*, <https://pi.math.cornell.edu/~hatcher/3M/3Mdownloads.html>, 2007.
20. N. Hoffman, K. Ichihara, M. Kashiwagi, H. Masai, S. Oishi, and A. Takayasu, *Verified computations for hyperbolic 3-manifolds*, Exp. Math. **25** (2016), no. 1, 66–78.
21. D. D. Long, *Immersions and embeddings of totally geodesic surfaces*, Bull. London Math. Soc. **19** (1987), no. 5, 481–484.
22. F. Luo, S. Schleimer, and S. Tillmann, *Geodesic ideal triangulations exist virtually*, Proc. Amer. Math. Soc. **136** (2008), no. 7, 2625–2630.
23. W. D. Neumann and D. Zagier, *Volumes of hyperbolic three-manifolds*, Topology **24** (1985), no. 3, 307–332.
24. B. E. Nimershiem, *Isometry classes of flat 2-tori appearing as cusps of hyperbolic 3-manifolds are dense in the moduli space of the torus*, Low-dimensional topology (Knoxville, TN, 1992), Conference Proceedings and Lecture Notes in Geometry and Topology, III, International Press, Cambridge, MA, 1994, pp. 133–142.
25. C. Petronio, *Ideal triangulations of hyperbolic 3-manifolds*, Boll. Unione Mat. Ital. Sez. B Artic. Ric. Mat. (8) **3** (2000), no. 3, 657–672.
26. C. Petronio and J. Porti, *Negatively oriented ideal triangulations and a proof of Thurston's hyperbolic Dehn filling theorem*, Expo. Math. **18** (2000), no. 1, 1–35.
27. I. Rivin, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. Math. (2) **139** (1994), no. 3, 553–580.
28. M. L. Sirotkina, *On the triangulation of three-dimensional hyperbolic manifolds*, Algebra i Analiz **14** (2002), no. 6, 192–204.
29. W. P. Thurston, *The geometry and topology of three-manifolds*, <http://library.msri.org/books/gt3m/>, 1980.
30. M. Wada, Y. Yamashita, and H. Yoshida, *An inequality for polyhedra and ideal triangulations of cusped hyperbolic 3-manifolds*, Proc. Amer. Math. Soc. **124** (1996), no. 12, 3905–3911.
31. H. Wilton and P. Zalesskii, *Distinguishing geometries using finite quotients*, Geom. Topol. **21** (2017), no. 1, 345–384.
32. H. Yoshida, *Ideal tetrahedral decompositions of hyperbolic 3-manifolds*, Osaka J. Math. **33** (1996), no. 1, 37–46.