Comment. Math. Helv. 87 (2012), 727–787
DOI 10.4171/CMH/267

Commentarii Mathematici Helvetici

Some virtually special hyperbolic 3-manifold groups

Eric Chesebro, Jason DeBlois* and Henry Wilton**

Abstract. Let M be a complete hyperbolic 3-manifold of finite volume that admits a decomposition into right-angled ideal polyhedra. We show that M has a deformation retraction that is a virtually special square complex, in the sense of Haglund and Wise and deduce that such manifolds are virtually fibered. We generalise a theorem of Haglund and Wise to the relatively hyperbolic setting and deduce that $\pi_1 M$ is LERF and that the geometrically finite subgroups of $\pi_1 M$ are virtual retracts. Examples of 3-manifolds admitting such a decomposition include augmented link complements. We classify the low-complexity augmented links and describe an infinite family with complements not commensurable to any 3-dimensional reflection orbifold.

Mathematics Subject Classification (2010). 20E26, 57M10.

Keywords. 3-manifold, virtual fibering, Coxeter group, virtual retract.

1. Introduction

Let $\{\mathcal{P}_i\}_{i=1}^n$ be a collection of disjoint ideal polyhedra in $\mathbb{H}^3$. A *face pairing* on $\{\mathcal{P}_i\}$ is a collection of isometries $\{\phi_f \mid f \in \mathcal{P}_i^{(2)}\ 1 \leq i \leq n\}$ of $\mathbb{H}^3$ with the following properties. If f is a face of $\mathcal{P}_i$, ϕ_f takes f onto a face f' of some $\mathcal{P}_j$, with $\phi_f(\mathcal{P}_i) \cap \mathcal{P}_j = f'$, and $\phi_{f'} = \phi_f^{-1}$. Now let M be a complete hyperbolic 3-manifold with finite volume. An *ideal polyhedral decomposition* of M is an isometry between M and a quotient $\bigsqcup_i \mathcal{P}_i / \sim$, where $\sim$ is the equivalence relation generated by a face pairing on $\{\mathcal{P}_i\}$. If the dihedral angles of every polyhedron $\mathcal{P}_i$ are all equal to $\pi/2$ then the decomposition is called an *ideal right-angled polyhedral decomposition*.

Our first result relates fundamental groups of 3-manifolds that admit ideal right-angled polyhedral decompositions to the class of right-angled Coxeter groups. A *right-angled Coxeter group* W is defined by a finite, simplicial graph Δ (called the *nerve* of W) and has an easily described presentation: the generators are the vertices; every generator is an involution; and the commutator of two generators is trivial if and only if they are joined by an edge in Δ. We will refer to the vertices of the nerve

*Second author partially supported by NSF grant DMS-1007175.
**Third author partially supported by an EPSRC Career Acceleration Fellowship.

as the *standard generating set* for W. The properties of such W discovered in [3] and [17] will particularly concern us.

Theorem 1.1. *Suppose M is a complete hyperbolic 3-manifold with finite volume that admits a decomposition into right-angled ideal polyhedra. Then $\pi_1 M$ has a subgroup of finite index isomorphic to a word-quasiconvex subgroup of a right-angled Coxeter group (equipped with the standard generating set).*

See Section 5 for the definition of word quasiconvexity. In the terminology of [19], Theorem 1.1 asserts that $\pi_1 M$ is *virtually special*. The proof relies on work of Haglund–Wise [19] defining a class of *special* cube complexes – non-positively curved cube complexes whose hyperplanes lack certain pathologies – which are locally isometric into cube complexes associated to right-angled Coxeter groups. In Section 2.1 we review the relevant definitions and in Section 2.2 describe a *standard* square complex associated with an ideal polyhedral decomposition of a hyperbolic 3-manifold.

When an ideal polyhedral decomposition is right-angled, the associated standard square complex is non-positively curved, and hyperplanes are carried by totally geodesic surfaces. We will establish these properties in Subsection 2.2 and Section 3. Separability properties of totally geodesic surfaces then imply that pathologies may be removed in finite covers. We describe these properties and prove Theorem 1.1 in Section 4.

This result has important consequences for the geometry and topology of such manifolds. The first follows directly from work of Agol [3], and confirms that the manifolds we consider satisfy Thurston's famous Virtually Fibered Conjecture.

Corollary 1.2. *Suppose M is a complete hyperbolic 3-manifold with finite volume that admits a decomposition into right-angled ideal polyhedra. Then M is virtually fibered.*

A 3-manifold that satisfies the hypotheses of Theorem 1.1 is necessarily not compact, so its fundamental group is not hyperbolic in the sense of Gromov, but rather hyperbolic *relative to* the collection of its cusp subgroups. Nonetheless, our second theorem implies that its subgroup structure shares the separability properties of its compact cousins. This generalizes [19, Theorem 1.3] to the relatively hyperbolic setting. We say a subgroup H of a group G is a *virtual retract* if H is contained in a finite-index subgroup K of G and the inclusion map $H \hookrightarrow K$ has a left inverse. (See [24] for further details of virtual retractions.)

Theorem 1.3. *Let $\mathcal{S}$ be a compact, virtually special cube complex and suppose that $\pi_1 \mathcal{S}$ is hyperbolic relative to a collection of finitely generated abelian subgroups. Then every relatively quasiconvex subgroup of $\pi_1 \mathcal{S}$ is a virtual retract.*

As in the case of [19, Theorem 1.3], the proof of Theorem 1.3 relies on a result of Haglund for separating subgroups of right-angled Coxeter groups [17, Theorem A], but it also requires new ingredients to surmount the technical obstacle that not every relatively quasiconvex subgroup is word-quasiconvex. The first is Theorem 5.3, a variation of [26, Theorem 1.7], which establishes that every relatively quasiconvex subgroup is a retract of a *fully relatively quasiconvex* subgroup (see the definition above Theorem 5.3). The second ingredient, Proposition 5.5, extends work in [21] to show that fully relatively quasiconvex subgroups satisfy the hypotheses of [17, Theorem A].

Even without any restrictions on the types of parabolic subgroups allowed, our results prove that certain subgroups of relatively hyperbolic groups are virtual retracts: see Theorem 5.8 and its corollaries for precise statements.

The consequences of Theorem 1.3 follow a long-standing theme in the study of 3-manifolds and their fundamental groups. For a group G and a subgroup H, we say H is *separable* in G if for every $g \in G - H$, there is a finite-index subgroup $K < G$ such that $H < K$ and $g \notin K$. If $G = \pi_1 M$ for some manifold M, work of G. P. Scott links separability of H with topological properties of the corresponding cover $M_H \to M$ [35]. A group is called *LERF* if every finitely generated subgroup is separable.

Corollary 1.4. *Suppose M is a complete hyperbolic 3-manifold with finite volume that admits a decomposition into right-angled ideal polyhedra. Then:*

(1) *$\pi_1 M$ is LERF.*

(2) *Every geometrically finite subgroup of $\pi_1 M$ is a virtual retract.*

We will prove Theorem 1.3 and Corollary 1.4 at the end of Section 5.

The study of LERF 3-manifold groups dates back to [35]. Although there are examples of graph manifolds with non-LERF fundamental group [9], it remains unknown whether every hyperbolic 3-manifold group is LERF. Gitik [15] constructed examples of hyperbolic 3-manifolds with totally geodesic boundary whose fundamental groups are LERF, and it is a consequence of Marden's Tameness Conjecture that her closed examples are also LERF. Agol, Long and Reid proved that the Bianchi groups are LERF [2].

It is natural to ask to what extent Theorem 1.1 describes *new* examples of 3-manifold groups that virtually embed into right-angled Coxeter groups, and more generally to what extent it describes new examples of LERF 3-manifold groups. Hitherto, there have only been a limited number of techniques for proving that finite-volume 3-manifolds are LERF. The techniques of [15] did not produce non-compact, finite-volume examples, so we shall not consider them here.

Agol, Long and Reid [2] proved that geometrically finite subgroups of right-angled, hyperbolic reflection groups are separable. They deduced a similar result

for the Bianchi groups by embedding them as totally geodesic subgroups of higher-dimensional, arithmetic right-angled reflection groups. One might naïvely suppose that the fundamental group of a 3-manifold that decomposes into right-angled polyhedra $\{\mathcal{P}_i\}$ is commensurable with the reflection group in one of the $\mathcal{P}_i$, or perhaps a union of several, and therefore that Theorem 1.1 could be deduced using the techniques of [2].

We address the above possibility in Sections 6 and 7. There we describe infinite families of hyperbolic 3-manifolds that decompose into right-angled polyhedra but are not commensurable with *any* 3-dimensional reflection orbifold. Indeed, Section 7 considers a very broad class of hyperbolic 3-manifolds, the augmented link complements (previously considered in [22] and [33], for example), that decompose into right-angled polyhedra. Our investigations there strongly support the following hypothesis: a "generic" augmented link complement is not commensurable with any 3-dimensional reflection orbifold.

If M decomposes into isometric copies of a single, highly symmetric polyhedron $\mathcal{P}$, we show in Proposition 6.1 that $\pi_1 M$ is indeed commensurable with the reflection group in the sides of $\mathcal{P}$. The lowest-complexity right-angled ideal polyhedra (measured by number of ideal vertices) are the 3- and 4-antiprisms (see Figure 2), and these are sufficiently symmetric for the hypotheses of Proposition 6.1 to apply. However, in Section 6.2, we describe hybrid examples not commensurable with reflection groups.

Theorem 1.5. *For each $n \in \mathbb{N}$, there is complete, one-cusped hyperbolic 3-manifold N_n that decomposes into right-angled ideal polyhedra, such that N_n is not commensurable with N_m for any $m \neq n$, nor to any 3-dimensional reflection orbifold.*

Recently, Haglund and Wise have proved that every Coxeter group is virtually special [18]. Since $\pi_1 N_n$ is not commensurable with any 3-dimensional reflection group, the results of [18] do not apply to it. The proof of Theorem 1.5 uses work of Goodman–Heard–Hodgson [16] to explicitly describe the commensurator of $\pi_1 N_n$.

A rich class of manifolds that satisfy the hypotheses of Theorem 1.1 consists of the *augmented links* introduced by Adams [1]. Any link L in S^3 with hyperbolic complement determines (not necessarily uniquely) an augmented link using a projection of L which is *prime* and *twist-reduced*, by adding a "clasp" component encircling each crossing region. (See Section 7 for precise definitions.) Each link with hyperbolic complement admits a prime, twist reduced diagram, and the augmented link obtained from such a diagram also has hyperbolic complement (cf. [33, Theorem 6.1]). Ian Agol and Dylan Thurston showed in an appendix to [22] that each augmented link satisfies the hypotheses of Theorem 1.1.

Example 1 (Agol–Thurston). Let M be a complete hyperbolic manifold homeomorphic to the complement in S^3 of an augmented link. Then M admits a decomposition into two isometric right-angled ideal polyhedra.

In Section 7, we describe another polyhedron, the "crushtacean", that distills the most important combinatorial features of the Agol–Thurston ideal polyhedral decomposition. We record criteria, in Lemmas 7.4 and 7.6, that describe certain situations in which one may conclude that an augmented link complement is commensurable with the reflection orbifold in the associated right-angled polyhedron. Section 7.1 describes the scissors congruence classification of the complements of augmented links with up to 5 crossing regions. Finally, in Section 7.2 we prove:

Theorem 7.10. *There is a class of augmented links $L(n)$, $n \geq 3$, such that for all but finitely many n, $M(n) \doteq S^3 - L(n)$ is not arithmetic nor commensurable with any 3-dimensional hyperbolic reflection orbifold. Moreover, at most finitely many $M(n)$ occupy any single commensurability class.*

The crushtaceans of the links of Theorem 7.10 are the famous Löbell polyhedra. We believe that the behavior recorded in the theorem is generic among augmented links, but these are particularly amenable to analysis.

While this work was in preparation, we became aware of [7] and [6], which provide other examples of virtually special hyperbolic manifolds. The former mostly concerns arithmetic lattices, while the latter deals with finite-sheeted covers of the 3-sphere that branch over the figure-eight knot.

Acknowledgements. The authors would like to thank Ian Agol, Dave Futer, Alan Reid and Matthew Stover for useful conversations. Thanks also to Jack Button for confirming some of our Alexander polynomial computations, and to Jessica Purcell for a helpful reference to [33]. Finally, we thank the referee for a careful reading and helpful comments.

2. Preliminaries

2.1. Cube complexes. In this subsection we review relevant notions about cube complexes following the treatment of Haglund–Wise [19]. Another helpful reference is [8], particularly Chapters I.7 and II.5.

Definition ([19], Definition 2.1). Let $I = [-1, 1] \subset \mathbb{R}$. A *cube complex* X is a CW-complex such that each k-cell has a homeomorphism to $I^k \subset \mathbb{R}^k$ with the property that the restriction of the attaching map to each $(k-1)$-face of ∂I^k to X^{k-1} is an isometry onto I^{k-1} followed by the inclusion of a $(k-1)$-cell. A map $f\colon X \to Y$ between cube complexes is *combinatorial* if for each k-cell $\phi\colon I^k \to X$, the map $f \circ \phi$ is a k-cell of Y following an isometry of I^k. A *square complex* is a 2-dimensional cube complex, and we will refer by *vertex*, *edge*, or *square* to the image in X of a 0-, 1- or 2-cell, respectively.

Now let X be a square complex. We will take the *link* of the vertex $(1,1) \in I^2$ to be the line segment in I^2 joining $(0,1)$ to $(1,0)$ (the midpoints of the edges abutting $(1,1)$), and the link of another vertex v to be the image of the link of $(1,1)$ under the symmetry taking it to v. The link of a vertex $v \in X$ is the 1-complex obtained by joining the links of v in the squares of X attaching to it. We say X is *simple* if for each vertex v there is a combinatorial map from the link of v to a simplicial graph. In particular, if X is simple then no two squares meet along consecutive edges.

We will say a square complex X is *nonpositively curved* if for each vertex v in X, the link of v does not contain any simple cycle with fewer than four edges. (We are taking Gromov's link condition as a definition; see eg, [8, Ch. II.5] for a discussion.) In particular, X is simple. If X is simply connected and nonpositively curved, we will say X is CAT(0). For a more general discussion, see [8], in particular Chapter II.5.

The notion of a *hyperplane* is very important in defining "special" cube complexes. Here we will specialize the definition in [19] to square complexes.

Definition ([19], Definition 2.2). The *midlines* of I^2 are the subsets $I \times \{0\}$ and $\{0\} \times I$, each parallel to two edges of X. The *center* of a square $\phi\colon I^2 \to X$ is $\phi(0,0)$, and the *midpoint* of an edge $\phi\colon I \to X$ is $\phi(0)$. A midline of I^2 meets its two *dual* edges perpendicularly at their midpoints.

Given a square complex X, we define a graph Y, the associated *midline complex*, as follows. The 0-cells of Y are the midpoints of the edges of X, and the 1-cells of Y are midlines of squares of X, attached by the restrictions of the corresponding attaching maps. A *hyperplane* of X is a component of the associated midline complex Y.

By the definition of the midline complex, each hyperplane Y has an immersion into X, taking an edge to the midline of the square that contains it. Definition 3.1 of [19] describes the following pathologies of hyperplane immersions: *self-intersection*, *one-sidedness*, *direct* or *indirect self-osculation*, or *inter-osculation*. If the hyperplanes of X do not have any such pathologies, and its one-skeleton is bipartite, we will say that X is C*-special*.

The following theorem of Haglund–Wise is our main concern.

Theorem 2.1 ([19], Lemma 4.3). *Let X be a C-special square complex. Then there exists a right-angled Coxeter group W, an injective homomorphism $\pi_1 X \hookrightarrow W$ and a $\pi_1 X$-equivariant, combinatorial, isometric embedding from the universal cover of X into the Davis–Moussong complex of W. In particular, $\pi_1 X$ is isomorphic to a word-quasiconvex subgroup of W (with respect to the standard generating set).*

The Davis–Moussong complex of a right-angled Coxeter group W is a certain CAT(0) cube complex on which W acts naturally. The reader is referred to [19] for the definition. A square complex X is called *virtually special* if X has a C-special finite-sheeted covering space. To prove Theorem 1.1, we will prove that $\pi_1 M$ is isomorphic to the fundamental group of a virtually special square complex.

We will find the notion of a regular neighborhood of a hyperplane from [19] useful.

Definition. Let $Y \to X$ be a hyperplane of a square complex X. A (*closed*) *regular neighborhood* for Y is a cellular I-bundle $p : N \to Y$ equipped with a combinatorial immersion $j : N \to X$ such that the diagram

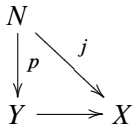

commutes. (Here the I-bundle N is given the obvious square-complex structure: the preimage of a vertex is an edge and the preimage of an edge is a square.)

Every hyperplane of a non-positively curved square complex has a regular neighborhood [19, Lemma 8.2]. The I-bundle $p\colon N \to Y$ has a section taking each $e \in Y^{(1)}$ to a midline of the square $p^{-1}(e)$. We refer to $Y \subset N$ as embedded by this section. In [19, Definition 8.7], the *hyperplane subgroup* $\pi_1 Y < \pi_1 X$ is defined as the image of j_* after an appropriate choice of basepoint.

2.2. A standard square complex. In this subsection we will take M to be a complete hyperbolic 3-manifold of finite volume, with an ideal polyhedral decomposition $\{\mathcal{P}_i\}$. For a pair of faces f and f' of polyhedra $\mathcal{P}_i$ and $\mathcal{P}_j$ such that $f' = \phi_f(f)$, we say that f and f' represent a face of the decomposition. Similarly, let $\{e_j\}_{j=1}^n$ be a sequence of edges of polyhedra $\mathcal{P}_{i_j}$ with the property that for each $j < n$, there is a face f_j of $\mathcal{P}_{i_j}$ containing e_j such that $\phi_{f_j}(e_j) = e_{j+1}$. Then we say the edges e_j represent an edge of the decomposition.

For each i, let $\bar{\mathcal{P}}_i$ be the union of $\mathcal{P}_i$ with its ideal vertices. (In the Poincaré ball model for $\mathbb{H}^3$, the ideal vertices of $\mathcal{P}_i$ are its accumulation points on ∂B^3.) Each face pairing isometry ϕ_f induces a homeomorphism from $\bar{f}$, the union of f with its ideal vertices, to $\bar{f}'$, where $f' = \phi_f(f)$.

The extended face pairings determine a cell complex $\mathcal{C}$ such that M is homeomorphic to $\mathcal{C} - \mathcal{C}^{(0)}$. The 0-cells of $\mathcal{C}$ are equivalence classes of ideal vertices under the equivalence relation generated by $v \sim \phi_f(v)$ for ideal vertices v of faces f. The 1- and 2- cells of $\mathcal{C}$ are equivalence classes of edges and faces of the $\mathcal{P}_i$ under the analogous equivalence relation, and the 3-cells are the $\bar{\mathcal{P}}_i$.

Let $\mathcal{C}'$ be the barycentric subdivision of the cell complex $\mathcal{C}$ associated to an ideal polyhedral decomposition. If v is a vertex of a cell $\bar{\mathcal{P}}$ of $\mathcal{C}'$, the open star of v in $\bar{\mathcal{P}}$ is the union of the interiors of the faces of $\bar{\mathcal{P}}$ containing v. The open star $\mathfrak{st}(v)$ of v in $\mathcal{C}'$ is the union of the open stars of v in the cells of $\mathcal{C}'$ containing it. Take $\mathfrak{st}(\mathcal{C}^{(0)})$ to be the disjoint union of the open stars in $\mathcal{C}'$ of the vertices of $\mathcal{C}$. Then $\mathcal{S}_0 \doteq \mathcal{C}' - \mathfrak{st}(\mathcal{C}^{(0)})$ is the unique subcomplex of $\mathcal{C}'$, maximal with respect to inclusion, with the property that $\mathcal{S}_0^{(0)} = (\mathcal{C}')^{(0)} - \mathcal{C}^{(0)}$.

A simplex of $\mathcal{S}_0$ is determined by its vertex set, which consists of barycenters of cells of $\mathcal{C}$. We will thus refer to each simplex of $\mathcal{S}_0$ by the tuple of cells of $\mathcal{C}$ whose barycenters are its vertices, in order of increasing dimension. For example, a simplex of maximal dimension is a triangle of the form $(\bar{e}, \bar{f}, \bar{\mathcal{P}}_i)$, where e is an edge and f a face of some ideal polyhedron $\mathcal{P}_i$ in the decomposition of M, with $e \subset f$.

Lemma 2.2. *There is a cellular deformation retraction Φ taking M to $|\mathcal{S}_0|$.*

Proof. Let v be an ideal vertex of $\bar{\mathcal{P}}_i$, and let U be the open star in $\mathcal{C}'$ of the equivalence class of v in $\mathcal{C}^{(0)}$. Let U_0 be the component of $U \cap \bar{\mathcal{P}}_i$ containing v. Then $\bar{U}_0$ is homeomorphic to the cone to v of its frontier in $\bar{\mathcal{P}}_i$, a union of triangles of $\mathcal{S}_0$. Hence there is a "straight line" deformation retraction of $\bar{U}_0 - \{v\}$ to its frontier. These may be adjusted to match up along faces of the $\mathcal{P}_i$, determining Φ. □

The standard square complex is obtained by taking a union of faces of $\mathcal{S}_0$.

Definition. Let M be a complete hyperbolic 3-manifold with a decomposition into ideal polyhedra $\{\mathcal{P}_i\}$, with associated cell complex $\mathcal{C}$ such that $M \cong \mathcal{C} - \mathcal{C}^{(0)}$, and let $\mathcal{S}_0 = \mathcal{C}' - \mathfrak{st}(\mathcal{C}^{(0)})$, where $\mathcal{C}'$ is the first barycentric subdivision of $\mathcal{C}$. Define the *standard square complex* $\mathcal{S}$ associated to $\{\mathcal{P}_i\}$, with underlying topological space $|\mathcal{S}| = |\mathcal{S}_0|$, as follows: $\mathcal{S}^{(0)} = \mathcal{S}_0^{(0)}$, $\mathcal{S}^{(1)} = \mathcal{S}_0^{(1)} - \{(\bar{e}, \bar{\mathcal{P}}_i) \mid e \subset \mathcal{P}_i\}$, and $\mathcal{S}^{(2)} = \{(\bar{e}, \bar{f}, \bar{\mathcal{P}}_i) \cup (\bar{e}, \bar{g}, \bar{\mathcal{P}}_i) \mid f, g \subset \mathcal{P}_i \text{ and } f \cap g = e\}$.

Since each 2-dimensional face $(\bar{e}, \bar{f}, \bar{\mathcal{P}}_i) \cup (\bar{e}, \bar{g}, \bar{\mathcal{P}}_i)$ of $\mathcal{S}$ is the union of two triangles of $\mathcal{S}_0$ which meet along the edge $(\bar{e}, \bar{\mathcal{P}}_i)$, it may be naturally identified with a square. Furthermore, since it is exactly the set of edges of the form $(\bar{e}, \bar{\mathcal{P}}_i)$ which are in $\mathcal{S}_0^{(1)} - \mathcal{S}^{(1)}$, $\mathcal{S}$ has the structure of a cell complex.

Lemma 2.3. *Let $\mathcal{S}$ be the standard square complex associated to an ideal polyhedral decomposition $\{\mathcal{P}_i\}$. Then $\mathcal{S}^{(1)}$ is bipartite.*

Proof. By definition, the vertices of $\mathcal{S}$ are barycenters of cells of the cell complex $\mathcal{C}$ associated to $\{\mathcal{P}_i\}$. We divide them into two classes by parity of dimension. An edge of $\mathcal{S}$ is of the form $(\bar{f}, \bar{\mathcal{P}}_i)$ for some i, where f is a face of $\mathcal{P}_i$, or $(\bar{e}, \bar{f})$, where e is an edge and f a face of some polyhedron. In either case, the endpoints belong to different classes. □

Say a cell of $\mathcal{S}$ is *external* if it is contained in $\mathcal{S} \cap \mathcal{C}^{(2)}$, and *internal* otherwise. Each square of $\mathcal{S}$ has two adjacent external edges, of the form $(\bar{e}, \bar{f})$ and $(\bar{e}, \bar{f}')$ in the notation above, and two internal edges $(\bar{f}, \bar{\mathcal{P}}_i)$ and $(\bar{f}', \bar{\mathcal{P}}_i)$. In particular, each external edge of each square is opposite an internal edge, and vice-versa.

Lemma 2.4. *As one-subcomplexes, $\mathcal{S} \cap \mathcal{C}^{(2)} = (\mathcal{C}^{(2)})' - \mathfrak{st}(\mathcal{C}^{(0)})$, where $(\mathcal{C}^{(2)})'$ is the barycentric subdivision of $\mathcal{C}^{(2)}$. In particular, Φ restricts to a deformation retraction from $\bigcup_i \partial\mathcal{P}_i$ to $|\mathcal{S} \cap \mathcal{C}^{(2)}|$.*

Proof. By definition $\mathcal{S} \cap \mathcal{C}^{(2)} = \mathcal{S}_0 \cap \mathcal{C}^{(2)}$, whence the first claim of the lemma follows. The second claim now holds because Φ is cellular. □

Lemma 2.5. *Suppose H is a hyperplane of the standard square complex associated to an ideal polyhedral decomposition $\{\mathcal{P}_i\}$ of a complete hyperbolic 3-manifold M, and let $p\colon N \to H$ be the regular neighborhood of H. N has boundary components $\partial_e N$ and $\partial_i N$, mapped by $j\colon N \to M$ to a union of external and internal edges, respectively.*

Proof. Let s be a square of $\mathcal{S}$. The vertices of s are the barycenters of $\bar{e}$, $\bar{f}$, $\bar{g}$, and $\bar{\mathcal{P}}_i$, where $\mathcal{P}_i$ is a polyhedron in the decomposition of M, e is an edge of $\mathcal{P}_i$, and f and g are the faces of $\mathcal{P}_i$ intersecting in e. One midline of s has vertices on the midpoints of the opposite edges $(\bar{e}, \bar{f})$ and $(\bar{g}, \bar{\mathcal{P}}_i)$ of s, and the other has vertices on the midpoints of $(\bar{f}, \bar{\mathcal{P}}_i)$ and $(\bar{e}, \bar{g})$. Take H to be the hyperplane containing the midline m with vertices on $(\bar{e}, \bar{f})$ and $(\bar{g}, \bar{\mathcal{P}}_i)$.

Let $s_0 = p^{-1}(m) \subset N$; then s_0 is a square which j maps homeomorphically to s. The edges of $s_0 \cap \partial N$ are mapped by j to the edges of s parallel to m. These are $(\bar{f}, \bar{\mathcal{P}}_i)$, which is internal, and $(\bar{e}, \bar{g})$, which is external. Let b_i be the edge mapped to $(\bar{f}, \bar{\mathcal{P}}_i)$ by j, let b_e be mapped to $(\bar{e}, \bar{g})$, and let $\partial_i N$ and $\partial_e N$ be the components of ∂N containing b_i and b_e, respectively. It is *a priori* possible that $\partial_i N = \partial_e N$, but we will show that $\partial_i N$ (respectively, $\partial_e N$) is characterized by the fact that its edges map to internal (resp. external) edges of $\mathcal{S}$.

Let s_1 be a square of N adjacent to s_0. Then the edge $m_1 \doteq p(s_1)$ of H is the midline of the square $s' = j(s_1)$ adjacent to s. Suppose first that s' meets s along the external edge $(\bar{e}, \bar{f})$. Then there is a polyhedron $\mathcal{P}_j$ of the decomposition with a face f' and edge $e' \subset f$ with $\phi_f(f) = f'$ and $\phi_f(e) = e'$ (ie, f and f' represent the same face of the decomposition of M, and e and e' the same edge), such that the vertices of s' are the barycenters of $\bar{e}'$, $\bar{f}'$, $\bar{g}_1$, and $\bar{\mathcal{P}}_j$. Here g_1 is the other face of $\bar{\mathcal{P}}_j$ containing e'.

Since m_1 meets m, it has an endpoint at the midpoint of $(\bar{e}', \bar{f}')$, which is identified with $(\bar{e}, \bar{f})$ in M. Then the other endpoint of m_1 is on the opposite edge $(\bar{g}_1, \bar{\mathcal{P}}_j)$ of s'. The external edge $(\bar{e}', \bar{g}_1)$ of s' which is parallel to m_1 meets the external edge $(\bar{e}, \bar{g})$ of s at the barycenter of the edge of the decomposition represented by e and e'. It follows that j maps the edge of $s_1 \cap \partial N$ adjacent to b_e to $(\bar{e}, \bar{g})$. Likewise, the edge of $s_1 \cap \partial N$ adjacent to b_i is mapped to the internal edge $(\bar{f}', \bar{\mathcal{P}}_j)$ of s'.

Now suppose s' meets s along the internal edge $(\bar{g}, \bar{\mathcal{P}}_i)$. Then there is an edge e_1 of g such that the vertices of s' are the barycenters of $\bar{e}_1$, $\bar{g}$, $\bar{f}_1$, and $\bar{\mathcal{P}}_i$. Here f_1 is the other face of $\mathcal{P}_i$ containing e_1. Then m_1 meets m at the midpoint of $(\bar{g}, \bar{\mathcal{P}}_i)$.

Since b_e is mapped by j to $(\bar{e}, \bar{g})$, the edge of $s_1 \cap \partial N$ adjacent to it is mapped to the external edge $(\bar{e}_1, \bar{g})$. It follows that the other edge of $s_1 \cap \partial N$ is mapped to the internal edge $(\bar{f}_1, \bar{\mathcal{P}}_i)$ of s' parallel to m_1.

The above establishes that the union of the set of edges of $\partial_i N$ mapped to internal edges of $\mathcal{S}$ is open and nonempty in $\partial_i N$. Since it is clearly also closed, it is all of $\partial_i N$. An analogous statement holds for $\partial_e N$, establishing the lemma. □

It is occasionally useful to think of the standard square complex associated to an ideal polyhedral decomposition as a subdivision of the "dual two-complex". If $\mathcal{C}$ is the cell complex associated to the ideal polyhedral decomposition $\{\mathcal{P}_i\}$, let $D\mathcal{C}$ be the two-complex with a vertex at the barycenter of each 3-cell of $\mathcal{C}$, for each $f \in \mathcal{C}^{(2)}$ an edge Df crossing f, and for each $e \in \mathcal{C}^{(1)}$ a face De crossed by e. The standard square complex $\mathcal{S}$ is obtained from $D\mathcal{C}$ by dividing each face along its intersections with the 2-cells of $\mathcal{C}$ which meet at the edge.

Lemma 2.6. *Suppose $\{\mathcal{P}_i\}$ is a decomposition of M into right-angled ideal polyhedra. The standard square complex $\mathcal{S}$ associated to $\{\mathcal{P}_i\}$ is non-positively curved.*

Proof. Recall that $\mathcal{S}$ is non-positively curved if and only if in the link of any vertex, each simple cycle has length at least 4. If v is a vertex of $\mathcal{S}$, a simple cycle of length k in the link of v is a sequence of squares $s_0, s_1, \ldots, s_{k-1}$ with the following properties: for each i there is an edge e_i with $v \subset e_i \subset s_i \cap s_{i+1}$ (taking $i + 1$ modulo k), and $s_i \neq s_j$ and $e_i \neq e_j$ when $i \neq j$.

Since the decomposition $\{\mathcal{P}_i\}$ is into right-angled polyhedra, the dual two-complex $D\mathcal{C}$ described above the lemma is a square complex. This follows from the fact that each edge of $\mathcal{C}$ is contained in four faces of $\mathcal{C}$. We will show that $D\mathcal{C}$ is non-positively curved; since $\mathcal{S}$ is a subdivision of $D\mathcal{C}$, it will follow that $\mathcal{S}$ is non-positively curved.

Suppose v is a vertex of $D\mathcal{C}$, and let $\{De_0, \ldots, De_{k-1}\}$ be a simple cycle in the link of v in $D\mathcal{C}$. The associated sequence of edges $\{Df_0, \ldots, Df_{k-1}\}$ determines a sequence of distinct faces $\{f_0, \ldots, f_{k-1}\}$ of the polyhedron $\mathcal{P}_i$ containing v, each meeting the next in an edge. It follows immediately from the necessary conditions of Andreev's theorem [4] and the fact that $\mathcal{P}_i$ is right-angled that every such cycle has length at least four. The conclusion of the lemma follows. □

3. Totally geodesic hyperplane groups

Fix an orientable, complete hyperbolic manifold $M = \mathbb{H}^3/\Gamma$ of finite volume, equipped with a decomposition $\{\mathcal{P}_i\}$ into right-angled ideal polyhedra. Here we have identified M with the quotient of $\mathbb{H}^3$ by a discrete group of isometries Γ, thus identifying $\pi_1 M$ with Γ. Let $\mathcal{S}$ be the standard square complex associated to the polyhedral decomposition as in Section 2.2. The goal of this section is, for each

hyperplane $H \to X$, to identify a totally geodesic surface immersed in M which "carries" H.

Since each $\mathcal{P}_i$ is right-angled and the angle in M around each edge is 2π, the equivalence class of each edge has four members. If f represents a face of the decomposition and e an edge of f, define the *flat e-neighbor of* f to be the face of the decomposition that meets f at angle π along e in M.

If $\mathcal{P}_i$ is the polyhedron containing f, let g be the other face of $\mathcal{P}_i$ containing e. Let $g' = \phi_g(g)$, a face of some polyhedron $\mathcal{P}_j$, and let $e' = \phi_g(e)$. Then e and e' represent the same edge of the decomposition, and the flat e-neighbor of f is represented by the face f_1 of $\mathcal{P}_j$ which intersects g' along e'. Let Σ_f be the collection of faces of the decomposition, minimal with respect to inclusion, satisfying the properties below.

(1) $f \in \Sigma_f$, and

(2) if $g \in \Sigma_f$ and e is an edge of g, then every flat e-neighbor of g is in Σ_f.

Note that if $g \subset \Sigma_f$ is a 2-cell then $\Sigma_f = \Sigma_g$. Furthermore, there is a sequence $\{f = f_0, f_1, \ldots, f_n = g\}$ such that for each $i > 0$ there is an edge e_i with f_i a flat e_i-neighbor of f_{i-1}. Call such a sequence a *path of flat neighbors.*

Now let $\hat{\Sigma}_f$ be the quotient of Σ_f by the following edge pairings: if g represents an element of Σ_f and e is an edge of g, glue g to its flat e-neighbor g' by the restriction of the face pairing isometry ϕ_g described above. Since each face of the decomposition has a unique flat e-neighbor along each of its edges, $\hat{\Sigma}_f$ is topologically a surface without boundary. It is connected, since any two faces in Σ_f are connected by a path of flat neighbors, and it inherits a hyperbolic structure from its faces, since the edge gluing maps are isometries.

The inclusion maps of faces $\{g \hookrightarrow \mathcal{P}_i \mid g \subset \mathcal{P}_i,\ g \in \Sigma_f\}$ determine an immersion from $\hat{\Sigma}_f$ to $\bigsqcup_i \mathcal{P}_i / \sim$. This is not necessarily an embedding because the preimage of an edge may consist of two edges of $\hat{\Sigma}_f$, each mapped homeomorphically. However, by construction it is a local isometry.

Lemma 3.1. *Let $i\colon \hat{\Sigma}_f \to M$ be the composition of the inclusion-induced map to $\bigsqcup_i \mathcal{P}_i / \sim$ with the isometry to M. Then i is a proper immersion which maps onto its image with degree one.*

Proof. If g is a face of $\mathcal{P}_i$, the inclusion $g \hookrightarrow \mathcal{P}_i$ is proper by definition. Since the collection $\{\mathcal{P}_i\}$ is finite, it follows that i is proper. By construction, the interior of each face in Σ_f is mapped homeomorphically by i, thus it has degree one onto its image. □

Since the map $i\colon \hat{\Sigma}_f \to M$ is a proper local isometry and M is complete, the hyperbolic structure on $\hat{\Sigma}_f$ is complete. Since it is contained in the union of finitely many polygons of finite area, $\hat{\Sigma}_f$ has finite area. Choosing an isometric embedding

of f in $\mathbb{H}^2$ thus determines a developing map identifying the universal cover of $\widehat{\Sigma}_f$ with $\mathbb{H}^2$, and identifying $\pi_1\widehat{\Sigma}_f$ with a subgroup Γ_f of $\mathrm{Isom}(\mathbb{H}^2)$.

Now fix a component $\tilde{f}$ of the preimage of $i(f)$ under the universal cover $\mathbb{H}^3 \to M$. This choice determines a lift $\tilde{\imath}\colon \mathbb{H}^2 \to \mathbb{H}^3$ of $i\colon \widehat{\Sigma}_f \to M$, equivariant with respect to the actions of Γ_f on $\mathbb{H}^2$ and $i_*(\pi_1\widehat{\Sigma})$ on $\mathbb{H}^3$.

Lemma 3.2. *Let $\mathcal{H}$ be the geodesic hyperplane of $\mathbb{H}^3$ containing $\tilde{f}$. Then $\tilde{\imath}$ maps $\mathbb{H}^2$ isometrically onto $\mathcal{H}$, and i_* takes $\pi_1\widehat{\Sigma}_f$ isomorphically onto* $\mathrm{Stab}_\Gamma(\mathcal{H})$.

Proof. Since i is a local isometry, $\tilde{\imath}$ maps $\mathbb{H}^2$ isometrically onto the geodesic hyperplane in $\mathbb{H}^3$ containing $\tilde{\imath}(f) = \tilde{f}$, hence $\mathcal{H}$. Since $\pi_1\widehat{\Sigma}_f$ acts faithfully on $\mathbb{H}^2$ by isometries, its action on $\mathcal{H}$, and hence all of $\mathbb{H}^3$ is also faithful. If $i_*(\pi_1\widehat{\Sigma}_f)$ were properly contained in $\mathrm{Stab}_\Gamma(\mathcal{H})$, the embedding i would factor through the covering map $\mathcal{H}/i_*(\pi_1\widehat{\Sigma}_f) \to \mathcal{H}/\mathrm{Stab}_\Gamma(\mathcal{H})$, contradicting the fact that i maps onto its image with degree one. □

Let us now take $\Gamma_f = i_*(\pi_1\widehat{\Sigma}_f)$ and $\widehat{M}_f = \mathbb{H}^3/\Gamma_f$. By Lemma 3.2, $i\colon \widehat{\Sigma}_f \to M$ lifts to an embedding $\hat{\imath}$ to $\widehat{M}_f$, such that $\widehat{M}_f$ is homeomorphic to $\hat{\imath}(\widehat{\Sigma}_f) \times \mathbb{R}$. We thus obtain the following diagram:

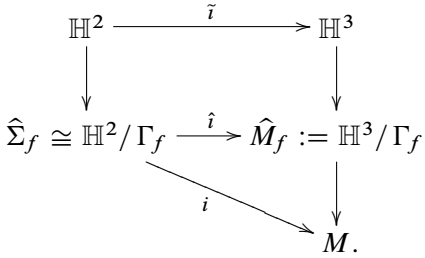

Below we will refer by $\widehat{\Sigma}_f \subset \widehat{M}_f$ to the image of $\hat{\imath}$.

Definition. Let M be a complete, orientable, hyperbolic 3-manifold of finite volume equipped with a decomposition $\{\mathcal{P}_i\}$ into right-angled ideal polyhedra, and suppose H is a hyperplane of the associated square complex, with regular neighborhood (N, p, j). Choose a midline m of H, let $s = p^{-1}(m)$, and let $\mathcal{P}_i$ contain $j(s)$. There is a unique face f of $\mathcal{P}_i$ containing the external edge of $j(s \cap \partial N)$, and we define $\widehat{\Sigma}(H) = \widehat{\Sigma}_f$, $\Gamma(H) = \Gamma_f$, and $\widehat{M}(H) = \widehat{M}_f$.

Lemma 3.3. *Using notation from the definition above, let $\hat{\mathcal{S}}$ be the standard square complex associated to the decomposition $\widehat{M}(H)$ inherits from $\{\mathcal{P}_i\}$. Then $j\colon N \to \mathcal{S}$ lifts to an immersion $\hat{\jmath}$ to $\hat{\mathcal{S}}$, taking $\partial_e N$ to a spine of $\widehat{\Sigma}(H)$, such that $\hat{\jmath}|_{\partial_e N}$ is an embedding if $\widehat{\Sigma}(H)$ is orientable, and a two-to-one cover if not.*

Corollary 3.4. *If $\Sigma(H)$ is orientable, $\pi_1 H = \Gamma(H)$; otherwise $\pi_1 H$ is the index-two orientation-preserving subgroup of $\Gamma(H)$.*

Proof of Lemma 3.3. Suppose m_0 and m_1 are two adjacent midlines of H, and let $s_0 = p^{-1}(m_0)$ and $s_1 = p^{-1}(m_1)$ in N. Take $\mathcal{P}_{i_0}$ and $\mathcal{P}_{i_1}$ to be the polyhedra containing $j(s_0)$ and $j(s_1)$, respectively, and let f_0 be the face of $\mathcal{P}_{i_0}$ and f_1 the face of $\mathcal{P}_{i_1}$ containing $j(s_0 \cap \partial_e N)$ and $j(s_1 \cap \partial_e N)$. If m_0 meets m_1 at the midpoint of an internal edge of $\mathcal{S}$, it is clear that $\mathcal{P}_{i_0} = \mathcal{P}_{i_1}$ and $f_0 = f_1$.

If m_0 meets m_1 in an external edge of $\mathcal{S}$, then $\mathcal{P}_{i_0}$ and $\mathcal{P}_{i_1}$ abut in M along a face of the decomposition. Let $g \subset \mathcal{P}_{i_0}$ represent this face of the decomposition. Then g and f_0 meet along an edge e, and $g' = \phi_g(g) \subset \mathcal{P}_{i_1}$ and f_1 meet along $e' = \phi_g(e)$. Hence if m_0 meets m_1 in an external edge of $\mathcal{S}$, there is an edge e of the decomposition of M such that f_0 and f_1 represent flat e-neighbors. It follows that a sequence of edges $m_0, m_1, \ldots, m_k$ of H, with the property that m_i is adjacent to m_{i-1} for each $i > 0$, determines a path of flat neighbors in $\Sigma(H)$. Therefore j maps $\partial_e N$ into $i(\widehat{\Sigma}(H))$.

Now let f be a face of some polyhedron $\mathcal{P}_i$ representing a face of $\Sigma(H)$. The cover $\widehat{M}(H)$ inherits a polyhedral decomposition from that of M, and since the covering map is injective on a neighborhood of $\hat{\imath}(f)$, there is a unique polyhedron $\widehat{\mathcal{P}}_i$ of this decomposition with the property that $\widehat{\mathcal{P}}_i$ projects to $\mathcal{P}_i$ and contains $\hat{\imath}(f)$. For a square s of N, we thus define $\hat{\jmath}(s)$ to be the component of the preimage of $j(s)$ contained in $\widehat{\mathcal{P}}_i$, where $\mathcal{P}_i$ is the polyhedron containing $j(s)$.

Suppose $\mathcal{P}_{i_0}$ and $\mathcal{P}_{i_1}$ contain faces f_0 and f_1, respectively, each representing a face of $\Sigma(H)$, which are flat e-neighbors for some edge e. Let $g \subset \mathcal{P}_{i_0}$ satisfy $g \cap f_0 = e$ and $\phi_g(e) = \phi_g(g) \cap f_1 \subset \mathcal{P}_{i_1}$. Since $\hat{\imath}(f_0)$ and $\hat{\imath}(f_1)$ meet in $\widehat{M}(H)$ along the preimage of e, $\widehat{\mathcal{P}}_{i_0}$ and $\widehat{\mathcal{P}}_{i_1}$ meet along the face represented by the preimage of g. For adjacent squares s_0 and s_1 in N, it follows that if $j(s_0)$ and $j(s_1)$ meet along an external edge of $\mathcal{S}$, then $\hat{\jmath}(s_0)$ and $\hat{\jmath}(s_1)$ meet along an external edge of $\widehat{\mathcal{S}}$.

If s_0 and s_1 are adjacent squares of N such that $j(s_0)$ meets $j(s_1)$ in an internal edge of $\mathcal{S}$ contained in a polyhedron $\mathcal{P}_i$, then $\hat{\jmath}(s_0)$ meets $\hat{\jmath}(s_1)$ in $\widehat{\mathcal{P}}_i$. Thus $\hat{\jmath}$ is continuous. Since j is an immersion, $\hat{\jmath}$ is an immersion as well. We claim $\hat{\jmath}$ maps $\partial_e N$ onto $\widehat{\mathcal{S}} \cap \widehat{\Sigma}(H)$.

Since $\hat{\jmath}$ is continuous, the image of $\partial_e N$ is closed in $\widehat{\mathcal{S}} \cap \widehat{\Sigma}(H)$. Now suppose e_0 and e_1 are adjacent edges of $\widehat{\mathcal{S}} \cap \widehat{\Sigma}(H)$ such that $e_0 \subset \hat{\jmath}(\partial_e N)$. Let $s_0 \subset N$ be a square such that $\hat{\jmath}(s_0)$ contains e_0, and let $m_0 = p(s_0)$ be a midline of $j(s_0)$. There is a square s of $\mathcal{S}$, containing the projection of e_1 to M, such that $s \cap j(s_0)$ is a union of edges containing the projection of $e_0 \cap e_1$. Let m_1 be the midline of s meeting m_0; then $m_1 \in H$, so by definition $s_1 = p^{-1}(m_1)$ is mapped by j to s. Now from the above it follows that $\hat{\jmath}(s_1)$ contains e_1. This implies that $\hat{\jmath}(\partial_e N)$ is open in $\widehat{\mathcal{S}} \cap \widehat{\Sigma}(H)$ and proves the claim.

Lemma 2.4 implies that $\widehat{\mathcal{S}} \cap \widehat{\Sigma}(H)$ is a spine for $\widehat{\Sigma}(H)$, hence $\hat{\jmath}$ maps $\partial_e N$ onto a spine of $\widehat{\Sigma}(H)$. Each square $s \subset N$ has the property that $s \cap \partial_e N$ is the unique

edge of s mapped by $\hat{\jmath}$ into $\hat{\Sigma}(H)$. For let $f \subset \Sigma(H)$ be the face of $\mathcal{P}_i$ containing $j(s \cap \partial_e N)$, where $\mathcal{P}_i$ contains $j(s)$, let g be the face containing the other external edge of $j(s)$, and let f_1 be the flat e-neighbor of f, where $e = f \cap g$. Then $\hat{\imath}(f)$ and $\hat{\imath}(f')$ are in $\hat{\Sigma}(H)$. If the face $\hat{g}$ of $\hat{\mathcal{P}}_i$ adjacent to $\hat{\imath}(f)$ were also in $\hat{\imath}(\hat{\Sigma}(H))$, $\hat{\imath}$ would not be an embedding.

Now suppose $\hat{\jmath}(s_0) = \hat{\jmath}(s_1)$ for squares s_0 and s_1 of N. By the property above, there is an edge e of $\hat{\mathcal{S}} \cap \hat{\Sigma}(H)$ such that $\hat{\jmath}(s_0 \cap \partial_e N) = e = \hat{\jmath}(s_1 \cap \partial_e N)$. It follows that j maps the external edge of each of s_0 and s_1 to the projection of e in M. By definition, $p(s_0)$ is the midline of $j(s_0)$ parallel to $j(s_0 \cap \partial_e N)$, and the same holds true for s_1. Thus $p(s_0) = p(s_1)$, so $s_0 = s_1$.

The paragraph above implies that $\hat{\jmath}|_{\partial_e N}$ is at worst two-to-one, since each external edge of $\hat{\mathcal{S}}$ is contained in exactly two squares. Since $\hat{M}(H)$ is orientable, if $\hat{\Sigma}(H)$ is orientable as well, then it divides any sufficiently small regular neighborhood into two components. Since N is connected and $\hat{\jmath}$ is continuous, in this case its image is on one side of $\hat{\Sigma}(H)$, so $\hat{\jmath}|_{\partial_e N}$ is an embedding.

If $\hat{\Sigma}(H)$ is nonorientable, then a regular neighborhood is connected. Thus in this case, for any edge e of $\hat{\mathcal{S}} \cap \hat{\Sigma}(H)$, both squares containing e are in the image of $\hat{\jmath}$, and the restriction to $\partial_e N$ maps two-to-one. □

The final result of this section characterizes some behaviors of hyperplanes of $\mathcal{S}$ in terms of the behavior of their associated totally geodesic surfaces. Below we say distinct hyperplanes H_1 and H_2 are *parallel* if $\Sigma(H_1) = \Sigma(H_2)$.

Proposition 3.5. *Let M be a complete, orientable hyperbolic 3-manifold equipped with a decomposition $\{\mathcal{P}_i\}$ into right-angled ideal polyhedra, with associated standard square complex $\mathcal{S}$, and let H_1 and H_2 be hyperplanes of $\mathcal{S}$. If H_1 osculates H_2 along an external edge of $\mathcal{S}$, then either*

(1) *$H_1 = H_2$ and $\Sigma(H_1)$ is nonorientable; or*

(2) *H_1 and H_2 are parallel and $\Sigma(H_1) = \Sigma(H_2)$ is orientable.*

H_1 intersects H_2 if and only if $i(\hat{\Sigma}(H_1))$ intersects $i(\hat{\Sigma}(H_2))$ at right angles.

Proof. Suppose H_1 osculates H_2 along an external edge e. Then there are squares s_1 and s_2 of $\mathcal{S}$ intersecting along e, such that the midline m_1 of s_1 parallel to e is in H_1, and the midline $m_2 \subset s_2$ parallel to e is in H_2. If f is the face of the decomposition containing e, then by definition $f \in \Sigma(H_1)$ and $f \in \Sigma(H_2)$. Since s_1 and s_2 are on opposite sides of f in M, Lemma 3.3 implies alternatives *1* and *2*.

Suppose H_1 intersects H_2 in a square s contained in some polyhedron $\mathcal{P}_i$, and for $j = 0, 1$ let m_j be the midline of s in H_j. For each j, there is a unique external edge e_j of s parallel to m_j. By definition, the faces f_1 and f_2 of $\mathcal{P}_i$ containing e_1 and e_2 are contained in $\Sigma(H_1)$ and $\Sigma(H_2)$, respectively. Since $\mathcal{P}_i$ is right-angled they meet at right angles, establishing the lemma. □

4. Embedding in Coxeter groups

Let $M = \mathbb{H}^3/\Gamma$ be a complete, orientable hyperbolic 3-manifold of finite volume, equipped with a decomposition $\{\mathcal{P}_i\}$ into right-angled ideal polyhedra. In this section we describe separability properties of hyperplane subgroups which allow pathologies to be removed in finite covers of M.

If H is a subgroup of a group G, we say H is *separable* in G if for each $g \in G - H$ there is a subgroup K, of finite index in G, such that $H < K$ and $g \notin K$. The separability result needed for the proof of Theorem 1.1 follows from [23, Lemma 1] and extends its conclusion to a slightly more general class of subgroups.

Lemma 4.1 (Cf. [23], Lemma 1). *Let $M = \mathbb{H}^3/\Gamma$ be a complete, orientable hyperbolic 3-manifold with finite volume, and let $\mathcal{H} \subset \mathbb{H}^3$ be a hyperplane such that $\mathrm{Stab}_\Gamma(\mathcal{H})$ acts on $\mathcal{H}$ with finite covolume. Then the subgroup of $\mathrm{Stab}_\Gamma(\mathcal{H})$ that acts preserving an orientation of $\mathcal{H}$ is separable in Γ.*

Proof. It follows from [23, Lemma 1] that $\mathrm{Stab}_\Gamma(\mathcal{H})$ is separable. It remains to consider the case in which $\mathrm{Stab}_\Gamma(\mathcal{H})$ is orientation-reversing on $\mathcal{H}$ and to show that the orientation-preserving subgroup is separable.

As in [23, Theorem 1], there is a finite-sheeted covering $M' \to M$ such that the immersed surface $\mathcal{H}/\mathrm{Stab}_\Gamma(\mathcal{H})$ lifts to an embedded surface Σ in M'. Because M' is orientable, the surface Σ is one-sided. Let N be a closed regular neighborhood of Σ and let M_0 be the complement of the interior of N in M'. The boundary of N is homeomorphic to $\widetilde{\Sigma}$, the orientable double cover of Σ.

The neighborhood N has the structure of a twisted interval bundle over Σ, so $\pi_1 N \cong \pi_1 \Sigma$. The double cover $\widetilde{N}$ of N obtained by pulling back the bundle structure along the covering map $\widetilde{\Sigma} \to \Sigma$ is an orientable interval bundle over $\widetilde{\Sigma}$ and hence homeomorphic to the product $\widetilde{\Sigma} \times [-1, +1]$. This homeomorphism can be chosen so that $\widetilde{\Sigma} \times \{0\}$ double covers Σ.

The inclusion map $i : \partial N \hookrightarrow N$ has precisely two lifts to $\widetilde{N}$; let $i^\pm$ be the lift that identifies ∂N with $\widetilde{\Sigma} \times \{\pm 1\}$. Construct a new manifold $\widetilde{M}$ as follows: let $M_0^\pm$ be two copies of M_0 and let $\partial^\pm N$ be the corresponding copy of ∂N in $M^\pm$; then $\widetilde{M}$ is obtained from

$$M_0^+ \sqcup \widetilde{N} \sqcup M_0^-$$

by identifying $x \in \partial^\pm N$ with $i^\pm(x)$. By construction, $\widetilde{M}$ is a double cover of M' and so a finite-sheeted cover of M. The image of $\mathcal{H}$ in $\widetilde{M}$ is precisely the orientable double cover of Σ, so $\pi_1 \widetilde{M}$ is a finite-index subgroup of Γ that contains the orientation-preserving elements of $\mathrm{Stab}_\Gamma(\mathcal{H})$ but not the orientation-reversing ones, as required. □

If H is a hyperplane of the standard square complex associated to the decomposition of M into right-angled ideal polyhedra, Lemma 3.2 and Corollary 3.4 together

describe a geodesic hyperplane $\mathcal{H}$, such that $\mathrm{Stab}_\Gamma(\mathcal{H})$ acts on it with finite covolume and $\pi_1 H$ is the subgroup which preserves an orientation of $\mathcal{H}$. Thus:

Corollary 4.2. *Suppose $M = \mathbb{H}^3/\Gamma$ is a complete, orientable hyperbolic 3-manifold of finite volume that admits a decomposition $\{\mathcal{P}_i\}$ into right-angled ideal polyhedra. If H is a hyperplane of the standard square complex associated to $\{\mathcal{P}_i\}$, then $\pi_1 H$ is separable in Γ.*

This implies, using [19, Corollary 8.9], that a hyperbolic manifold M with a right-angled ideal polyhedral decomposition has a finite cover whose associated square complex lacks most pathologies forbidden in the definition of special complexes.

Proposition 4.3. *Suppose $M = \mathbb{H}^3/\Gamma$ is a complete, orientable hyperbolic 3-manifold with finite volume that admits a decomposition into right-angled ideal polyhedra $\{\mathcal{P}_i\}$. There is a cover $M' \to M$ of finite degree such that hyperplanes of the standard square complex of M' do not self-intersect or -osculate.*

Proof. Let X be the standard square complex associated to $\{\mathcal{P}_i\}$. Lemma 2.2 implies that the inclusion $X \hookrightarrow M$ induces an isomorphism $\pi_1 X \to \Gamma$. By Corollary 4.2, each hyperplane subgroup is separable in $\pi_1 X$, so by [19, Corollary 8.9], X has a finite cover X' such that hyperplanes of X' do not self-intersect or -osculate. Let Γ' be the subgroup of $\pi_1 X$ corresponding to X', and let $M' \to M$ be the cover corresponding to Γ'. The decomposition $\{\mathcal{P}_i\}$ of M lifts to a right-angled ideal decomposition of M' with standard square complex X', proving the proposition. □

Proposition 4.3 already implies that a large class of hyperbolic 3-manifolds is virtually special. Below we will say that the decomposition $\{\mathcal{P}_i\}$ of M is *checkered* if the face pairing preserves a two-coloring – an assignment of white or black to each face f of each $\mathcal{P}_i$ such that if another face f' of $\mathcal{P}_i$ intersects f in an edge, it has the opposite color. The decompositions of augmented link complements described in the appendix to [22] are checkered, for example.

Theorem 4.4. *Suppose M is a complete hyperbolic 3-manifold with finite volume that admits a checkered decomposition into right-angled ideal polyhedra. Then $\pi_1 M$ has a subgroup of finite index that is isomorphic to a word-quasiconvex subgroup of a right-angled Coxeter group.*

Proof. Let $M = \mathbb{H}^3/\Gamma$ be a complete hyperbolic 3-manifold of finite volume with a decomposition $\{\mathcal{P}_i\}$ into right-angled polyhedra. If the decomposition is checkered, and f represents a face of the decomposition, it is easy to see that for each edge $e \subset f$, the flat e-neighbor of f has the same color as f. It follows that each face of the surface Σ_f described in Section 3 has the same color as f. If H is a hyperplane

of the square complex X associated to $\{\mathcal{P}_i\}$, we will say H is white if all faces of $\Sigma(H)$ are white, and black if they are black.

By Proposition 3.5, a hyperplane intersects only hyperplanes of the opposite color and osculates only hyperplanes of the same color along an external edge. If hyperplanes H_0 and H_1 osculate along an internal edge, let s_0 and s_1 be squares of $\mathcal{S}$, meeting along an internal edge e, with parallel midlines $m_0 \in H_0$ and $m_1 \in H_1$. Then e is of the form $(\bar{g}, \overline{\mathcal{P}_i})$, where $\mathcal{P}_i$ is the polyhedron containing s_0 and s_1 and g is a face of $\mathcal{P}_i$. The edges of s_0 and s_1 opposite e are contained in faces f_0 and f_1 of $\mathcal{P}_i$ in $\hat{\Sigma}(H_0)$ and $\hat{\Sigma}(H_1)$, respectively. Then each of f_0 and f_1 intersects g, so the color of f_0 and f_1 is opposite that of g. It follows that hyperplanes of $\mathcal{S}$ do not inter-osculate.

By Proposition 4.3, M has a finite cover M' such that hyperplanes of the square complex X' associated to the lifted ideal polyhedral decomposition of M' do not self-intersect or -osculate. The lifted ideal polyhedral decomposition of M' inherits the checkered property from that of M, so by the above, hyperplanes of X' do not inter-osculate. In addition, Lemma 2.6 implies that X' is nonpositively curved, Lemma 2.5 implies that each hyperplane is two-sided, and Lemma 2.3 implies that $X'^{(1)}$ is bipartite. Thus X' is C-special, and by Theorem 2.1, the subgroup $\Gamma' < \Gamma$ corresponding to M' embeds as a word-quasiconvex subgroup of a right-angled Coxeter group. □

In fact, we will show below that every right-angled decomposition determines a twofold cover whose associated decomposition is checkered. This uses the lemma below, which is a well-known consequence of Andreev's theorem.

Lemma 4.5. *Let $\mathcal{P} \subset \mathbb{H}^3$ be a right-angled ideal polyhedron of finite volume. There are exactly two checkerings of the faces of $\mathcal{P}$.*

Theorem 1.1 follows quickly from this lemma and Theorem 4.4.

Proof of Theorem 1.1. Suppose $\{\mathcal{P}_i\}_{i=1}^n$ is a right-angled ideal decomposition of M. Let $\{\mathcal{P}_i^{(0)}, \mathcal{P}_i^{(1)}\}_{i=1}^n$ be a collection of disjoint right-angled polyhedra such that for each i, $\mathcal{P}_i^{(0)}$ and $\mathcal{P}_i^{(1)}$ are each isometric to $\mathcal{P}_i$, and the faces of $\mathcal{P}_i^{(0)}$ have the opposite checkering of the faces of $\mathcal{P}_i^{(1)}$. Here we take for granted that we have fixed marking isometries $\mathcal{P}_i^{(j)} \to \mathcal{P}_i$ for each $j \in \{0, 1\}$, so that each face f of $\mathcal{P}_i$ has fixed correspondents $f^{(0)} \subset \mathcal{P}_i^{(0)}$ and $f^{(1)} \subset \mathcal{P}_i^{(1)}$.

For each i and each face f of $\mathcal{P}_i$, we determine face pairing isometries $\phi_{f^{(0)}}$ and $\phi_{f^{(1)}}$ for $\{\mathcal{P}_i^{(0)}, \mathcal{P}_i^{(1)}\}$ using the following requirements: each $\phi_{f^{(j)}}$, $j \in \{0, 1\}$ must commute with ϕ_f under the marking isometries, and each must preserve color. Thus if $f' = \phi_f(f)$ and $f'^{(0)}$ has the same color as $f^{(0)}$, we take $\phi_{f^{(j)}}(f^{(j)}) = f'^{(j)}$ for each j; otherwise we take $\phi_{f^{(j)}}(f^{(j)}) = f'^{(1-j)}$

Let $\widetilde{M}$ be the quotient of $\{\mathcal{P}_i^{(0)}, \mathcal{P}_i^{(1)}\}_{i=1}^n$ by the face pairing isometries described above. By construction, $\widetilde{M}$ is a double cover of M, and it is easy to see that $\widetilde{M}$ is disconnected if and only if the original decomposition $\{\mathcal{P}_i\}$ admits a checkering. If it did, Theorem 4.4 would apply directly to M, so we may assume that it does not. Then, by Theorem 4.4, the conclusion of Theorem 1.1 applies to $\widetilde{M}$; hence it applies as well to M. □

5. Virtual retractions and quasiconvexity

This section contains the proof of Theorem 1.3. We will need to work with various different definitions of quasiconvexity for subgroups. These definitions all coincide in the case of a Gromov-hyperbolic group because Gromov-hyperbolic metric spaces enjoy a property sometimes known as the Morse Property, which asserts that quasi-geodesics are uniformly close to geodesics. In our case, M has cusps and therefore $\Gamma = \pi_1 M$ is not Gromov hyperbolic but rather relatively hyperbolic. One of the results we use to circumvent this difficulty, Proposition 5.5, makes use of [13, Theorem 1.12], which the authors call the 'Morse Property for Relatively Hyperbolic Groups'.

Definition. Let X be a geodesic metric space. A subspace Y is *quasiconvex* if there exists a constant κ such that any geodesic in X between two points of Y is contained in the κ-neighborhood of Y.

We will apply this notion in two contexts. If U is a CAT(0) cube complex with base vertex v and a group G acts properly discontinuously by combinatorial isometries on U then we consider the one-skeleton $X = U^{(1)}$ with the induced length metric (where each edge has length one). We say that a subgroup H is *combinatorially quasiconvex* if Hv is a quasiconvex subspace of X. In fact, combinatorial quasiconvexity is independent of the choice of basepoint if the action of G on U is special [19, Corollary 7.8].

On the other hand, given a group G with a generating set S we can consider the Cayley graph $\mathrm{Cay}_S(G)$. A subgroup H is *word quasiconvex* if H is a quasiconvex subspace of $\mathrm{Cay}_S(G)$.

Let W be a right-angled Coxeter group with standard generating set S and let U be the universal cover of the Davis–Moussong complex for W. The one-skeleton of U is very closely related to $\mathrm{Cay}_S(W)$: the edges of the Cayley graph come in pairs; identifying these pairs gives $U^{(1)}$. Furthermore, the image of the universal cover of a special cube complex under the isometry defined by Haglund and Wise to the Davis–Moussong complex of W is a convex subcomplex [19, Lemma 7.7]. We therefore have the following relationship between combinatorial quasiconvexity and word quasiconvexity in special cube complexes.

Remark. Suppose that G is the fundamental group of a C-special cube complex $\mathcal{S}$, so that G is isomorphic to a word-quasiconvex subgroup of a right-angled Coxeter group W [19]. If H is a subgroup of G, then H is combinatorially quasiconvex in G (with respect to the action of G on the universal cover of $\mathcal{S}$) if and only if H is word quasiconvex in W (with respect to the standard generating set).

The idea is to prove Theorem 1.3 by applying the following theorem of Haglund.

Theorem 5.1 ([17], Theorem A). *Let W be a right-angled Coxeter group with the standard generating set and let H be a word-quasiconvex subgroup. Then H is a virtual retract of W.*

Theorem A of [17] is not stated in this form. Nevertheless, as observed in the paragraph following Theorem A, this is what is proved.

Corollary 5.2 (Cf. [19], Corollary 7.9). *If G is the fundamental group of a compact, virtually special cube complex and H is a combinatorially quasiconvex subgroup of G then H is a virtual retract of G.*

Proof. Let G' be a special subgroup of finite index in G. It is clear that $H' = H \cap G'$ is combinatorially quasiconvex in G'. By the above remark, H' is word-quasiconvex in the right-angled Coxeter group W, so H' is a virtual retract of W and hence of G' by Theorem 5.1. By [19, Theorem 4.4], G is linear. We can now apply the argument of [24, Theorem 2.10] to deduce that H is a virtual retract of G. □

The reader is referred to [26] and [21] for definitions of *relatively hyperbolic* groups and *relatively quasiconvex* subgroups, which are the subject of Theorem 1.3. (See Proposition 5.4 below for a characterization of relative quasiconvexity.) In order to deduce Theorem 1.3 from Corollary 5.2, it would be enough to show that every relatively quasiconvex subgroup of the relatively hyperbolic fundamental group of a C-special cube complex is combinatorially quasiconvex. Unfortunately, this may be false. For instance, the diagonal subgroup of $\mathbb{Z}^2$ with the standard generating set is not quasiconvex. The next theorem, a minor modification of a result of [26], gets round this difficulty.

Definition. Suppose a group G is hyperbolic relative to a finite set of subgroups $\mathcal{P}$. Then a relatively quasiconvex subgroup is called *fully relatively quasiconvex* if for every $P \in \mathcal{P}$ and every $g \in G$, either $H \cap gPg^{-1}$ is trivial or $H \cap gPg^{-1}$ has finite index in gPg^{-1}.

Theorem 5.3 (Cf. [26] Theorem 1.7). *Suppose that G is hyperbolic relative to $\mathcal{P}$ and that every $P \in \mathcal{P}$ is finitely generated and abelian. If Q is a relatively quasiconvex subgroup of G then G has a fully relatively quasiconvex subgroup H such that Q is a retract of H.*

Proof. In the proof of [26, Theorem 1.7], the authors construct a sequence of relatively quasiconvex subgroups

$$Q = Q_0 \subseteq Q_1 \subseteq \cdots \subseteq Q_n = H$$

with H fully relatively quasiconvex. We recall a few details of the construction of Q_k from Q_{k-1}. We will modify this construction slightly so that Q_{k-1} is a retract of Q_k for each k. For some maximal infinite parabolic subgroup K_k of Q_{k-1}, there is $P_k \in \mathcal{P}$ and $f_k \in G$ such that $K_k \subseteq f_k P_k f_k^{-1}$. Manning and Martinez-Pedroza find a finite-index subgroup R_k of $f_k P_k f_k^{-1}$ that contains K_k and excludes a certain finite set F. We shall impose an extra condition on R_k that is easily met when P_k is abelian, namely that K_k should be a direct factor of R_k. Just as in [26], the next subgroup in the sequence is now defined as $Q_k = \langle Q_{k-1}, R_k \rangle$, and just as in that setting it follows that Q_k is relatively quasiconvex.

It remains only to show that Q_{k-1} is a retract of Q_k. By assertion (1) of [26, Theorem 3.6], the natural map

$$Q_{k-1} *_{K_k} R_k \to Q_k$$

is an isomorphism. But K_k is a direct factor of R_k and so there is a retraction $R_k \to K_k$, which extends to a retraction $Q_k \to Q_{k-1}$ as required. □

In light of Theorem 5.3, to prove Theorem 1.3 it will suffice to show that when G is the relatively hyperbolic fundamental group of a non-positively curved cube complex, its fully relatively quasiconvex subgroups are combinatorially convex. This is the content of Proposition 5.5 below.

Hruska has extensively investigated various equivalent definitions of relative hyperbolicity and relative quasiconvexity [21]. Corollary 8.16 of [21] provides a characterization of relative quasiconvexity in terms of geodesics in the Cayley graph. Unfortunately, to prove Theorem 1.3 we need to work in the one-skeleton of the universal cover of a cube complex. This is not actually a Cayley graph unless the cube complex in question has a unique vertex. It is, however, quasi-isometric to the Cayley graph. Therefore, we will need a quasigeodesic version of Hruska's Corollary 8.16. Fortunately, we shall see that Hruska's proof goes through.

In what follows, S is any choice of finite generating set for G and d is the usual length metric on $\mathrm{Cay}_S(G)$. For any $g \in G$ write $l(g)$ for $d(1, g)$, the word length of g with respect to S. For $x \in \mathrm{Cay}_S(G)$ we denote by $B(x, R)$ the open ball of radius R about x. We define

$$N_R(Y) = \bigcup_{y \in Y} B(y, R)$$

for any subspace $Y \subseteq \mathrm{Cay}_S(G)$ and any $R > 0$. To keep notation to a minimum we will work with τ-quasigeodesics, which are more usually defined as (τ, τ)-quasi-

geodesics. That is, a path c is a τ-quasigeodesic if

$$\frac{1}{\tau}|s-t|-\tau \leq d(c(s),c(t)) \leq \tau|s-t|+\tau$$

for all suitable s and t. We will always assume that our quasigeodesics are continuous, which we can do by [8, Lemma III.H.1.11]. The following definition is adapted from [21].

Definition (Cf. [21], Definition 8.9). Let H be a subgroup of G. Let c be (the image of) a quasigeodesic in $\mathrm{Cay}_S(G)$. If $x \in c$ is not within distance R of the endpoints of c and

$$B(x,R) \cap c \subseteq N_\epsilon(gP)$$

for some $g \in G$ and $P \in \mathcal{P}$ then x is called (ϵ, R)-*deep* in gP. If $x \in c$ is not (ϵ, R)-deep in any such coset gP then x is called an (ϵ, R)-*transition point* of c.

The next proposition characterizes relatively quasiconvex subgroups in terms of quasigeodesics in the Cayley graph. Roughly, it asserts that every point on a quasigeodesic between elements of H is either close to H or is close to some peripheral coset gP.

Proposition 5.4 (Cf. [21], Corollary 8.16). *Suppose G is hyperbolic relative to $\mathcal{P}$ and H is a subgroup of G. Then H is relatively quasiconvex in G if and only if for every τ there are constants ϵ, R, κ such that the following two properties hold.*

(1) *For any continuous τ-quasigeodesic c in $\mathrm{Cay}_S(G)$, any connected component $\bar{c}$ of the set of all (ϵ, R)-deep points of c is (ϵ, R)-deep in a unique peripheral left coset gP; that is, there exists a unique $P \in \mathcal{P}$ and $gP \in G/P$ such that every $x \in \bar{c}$ is (ϵ, R)-deep in gP and no $x \in \bar{c}$ is (ϵ, R)-deep in any other peripheral left coset.*

(2) *If the quasigeodesic c joins two points of H then the set of (ϵ, R)-transition points of c is contained in $N_\kappa(H)$.*

The statement of [21, Corollary 8.16] only deals with the case when c is a geodesic. However, the necessary results of Section 8 of [21] also hold in the quasigeodesic case.

The following proposition completes the proof of Theorem 1.3.

Proposition 5.5. *Let G be finitely generated and relatively hyperbolic. Suppose that G acts properly discontinuously and cocompactly by isometries on a geodesic metric space X. Fix a basepoint $v \in X$. For any fully relatively quasiconvex subgroup $H \subseteq G$ there exists a constant ν such that any geodesic between two points of the orbit Hv lies in the ν-neighborhood of Hv. In particular, if G is the fundamental group of a non-positively curved cube complex then, taking X to be the one-skeleton of the universal cover, it follows that H is combinatorially quasiconvex.*

Proposition 5.4 implies that, to prove Proposition 5.5, it is enough to prove that deep points of quasigeodesics between points of H lie in a bounded neighborhood of H. The key technical tool is the following lemma, which is nothing more than the Pigeonhole Principle.

Lemma 5.6. *Let G be a finitely generated group. Fix a choice of finite generating set and the corresponding word metric on G. If H, K are subgroups and $H \cap K = 1$, then*

$$\#(H \cap N_r(K)) < \infty$$

for any $r > 0$.

Proof. For a contradiction, suppose $h_i \in H \cap N_r(K)$ are distinct for all $i \in \mathbb{N}$. For each i, there is $k_i \in K$ with $d(h_i, k_i) < r$. Let $g_i = h_i^{-1} k_i$, so $l(g_i) < r$. The ball of radius r in G is finite, so $g_i = g_j$ for some $i \neq j$ by the Pigeonhole Principle. But now

$$h_i h_j^{-1} = h_i g_i g_j^{-1} h_j^{-1} = k_i k_j^{-1}$$

is a non-trivial element of $H \cap K$, a contradiction. □

It follows that only short elements of H can be close to parabolic left cosets for which H intersects the stabilizer trivially.

Lemma 5.7. *Suppose G is hyperbolic relative to $\mathcal{P}$ and H is any subgroup of G. Let $g \in G$ and $P \in \mathcal{P}$ be such that $H \cap gPg^{-1} = 1$. For any $r > 0$ there exists finite $\lambda = \lambda(r, gP)$ such that if $h \in N_r(gP) \cap H$ then $l(h) \leq \lambda$.*

Proof. Choose g of minimal word length in gP and set $k = l(g)$. For any $p \in P$, $d(gp, gpg^{-1}) = k$ and it follows that

$$N_r(gP) \subseteq N_{k+r}(gPg^{-1})$$

by the triangle inequality. Therefore, by Lemma 5.6 with $K = gPg^{-1}$, $N_r(gP) \cap H$ is finite and so

$$\lambda = \max\{l(h) \mid h \in N_r(gP) \cap H\}$$

is as required. □

We are now ready to prove Proposition 5.5.

Proof of Proposition 5.5. Consider a geodesic b in X joining two points of Hv. We need to show that b is contained in a uniformly bounded neighborhood of Hv.

By the Švarc–Milnor Lemma, G has a finite generating set S and X is quasi-isometric to the Cayley graph $\mathrm{Cay}_S(G)$. The geodesic b maps to some τ-quasi-geodesic in $\mathrm{Cay}_S(G)$, which we denote c. Furthermore, we can assume that c is

continuous by [8, Lemma III.H.1.11]. It is therefore enough to show that c is contained in a uniformly bounded neighborhood of H in the word metric d on $\mathrm{Cay}_S(G)$.

Let ϵ, R and κ be as in Proposition 5.4. By assertion 2 of Proposition 5.4, the (ϵ, R)-transition points of c are contained in the κ-neighborhood of H. Therefore, it remains to show that the (ϵ, R)-deep points of c are contained in a uniformly bounded neighborhood of H.

Let $\bar{c}$ be a connected component of the set of all (ϵ, R)-deep points of c. By definition, every $x \in \bar{c}$ is in the ϵ-neighborhood of some peripheral left coset gP. By assertion 1 of Proposition 5.4, the component $\bar{c}$ is contained between two (ϵ, R)-transition points of c, which we shall denote y_1 and y_2. We can take these points to be arbitrarily close to $\bar{c}$, and hence we can assume that $d(y_i, gP) \leq \epsilon$ for $i = 1, 2$. On the other hand, by assertion 2 of Proposition 5.4, there exist $h_1, h_2 \in H$ such that $d(h_i, y_i) < \kappa$ for $i = 1, 2$. Therefore, $h_i \in N_{\epsilon+\kappa}(gP)$ for $i = 1, 2$.

Let $h_0 = h_1^{-1}h_2$ and let $g_0 = h_1^{-1}g$, so $h_0 \in N_{\epsilon+\kappa}(g_0P)$ and, without loss of generality, $l(g_0) \leq \epsilon + \kappa$. There are two cases to consider, depending on whether h_0 is long or short. Let

$$\lambda_{\max} = \max\{\lambda(\epsilon + \kappa, gP) \mid P \in \mathcal{P}, l(g) \leq \epsilon + \kappa, H \cap gPg^{-1} = 1\}$$

where $\lambda(\epsilon + \kappa, gP)$ is provided by Lemma 5.7. In the first case, $l(h_0) \leq \lambda_{\max}$ so $d(h_1, h_2) \leq \lambda_{\max}$ and therefore $d(y_1, y_2) < \lambda_{\max} + 2\kappa$. Because c is a τ-quasigeodesic it follows that for every $x \in \bar{c}$, for some $i = 1, 2$, we have that

$$d(x, y_i) < \lambda' = \frac{\tau^2}{2}(\lambda_{\max} + 2\kappa + \tau) + \tau$$

and so $d(x, h_i) < \lambda' + \kappa$.

In the second case, $l(h_0) > \lambda_{\max}$ and so $H \cap g_0Pg_0^{-1} \neq 1$ by Lemma 5.7. Therefore $H \cap g_0Pg_0^{-1}$ has finite index in $g_0Pg_0^{-1}$ because H is fully relatively quasiconvex. For each $g \in G$ and $P \in \mathcal{P}$ for which $H \cap gPg^{-1}$ has finite index in P, let $\mu = \mu(gP)$ be a number such that $gPg^{-1} \subseteq N_\mu(H \cap gPg^{-1})$. Set

$$\mu_{\max} = \max\{\mu(gP) \mid P \in \mathcal{P}, l(g) \leq \epsilon + \kappa, |gPg^{-1} : H \cap gPg^{-1}| < \infty\}.$$

Therefore

$$g_0Pg_0^{-1} \subseteq N_{\mu_{\max}}(H)$$

and so

$$g_0P \subseteq N_{\mu_{\max}+\epsilon+\kappa}(H)$$

because $l(g_0) \leq \epsilon + \kappa$. For each $x \in \bar{c}$ we have $h_1^{-1}x \in N_\epsilon(g_0P)$ and so $h_1^{-1}x \in N_{\mu_{\max}+2\epsilon+\kappa}(H)$. Therefore $x \in N_{\mu_{\max}+2\epsilon+\kappa}(H)$.

In summary, we have shown the following: the (ϵ, R)-transition points of the geodesic c are contained in the κ-neighborhood of H; the short (ϵ, R)-deep components of c are contained in the $(\lambda' + \kappa)$-neighborhood of H; and the long (ϵ, R)-deep

components of c are contained in the $(\mu_{\max}+2\epsilon+\kappa)$-neighborhood of H. Therefore, c is completely contained in the ν-neighborhood of H, where

$$\nu = \max\{\kappa, \lambda' + \kappa, \mu_{\max} + 2\epsilon + \kappa\}.$$

This completes the proof. □

We have assembled all the tools necessary to prove Theorem 1.3.

Proof of Theorem 1.3. Let Q be a relatively quasiconvex subgroup of $G = \pi_1 \mathcal{S}$. By Theorem 5.3, there exists a fully relatively quasiconvex subgroup H of G such that Q is a retract of H. Let X be the one-skeleton of the universal cover of $\mathcal{S}$, equipped with the induced length metric. By Proposition 5.5, for any basepoint v the orbit Hv is quasiconvex in X; that is, H is a combinatorially quasiconvex subgroup of G. Therefore, by Corollary 5.2, H is a virtual retract of G and so Q is also a virtual retract of G, as required. □

Corollary 1.4 now follows easily.

Proof of Corollary 1.4. Let $\Gamma = \pi_1 M$. As pointed out in [10], to prove that Γ is LERF it is enough to prove that Γ is GFERF – that is, that the geometrically finite subgroups are separable. Furthermore, by [17, Proposition 3.28], it is enough to prove that the geometrically finite subgroups of G are virtual retracts.

First, suppose that M is orientable. Let Q be a geometrically finite subgroup of Γ. By [26, Theorem 1.3], for instance, Γ is hyperbolic relative to its maximal parabolic subgroups and Q is a relatively quasiconvex subgroup of Γ. The maximal parabolic subgroups of Γ are isomorphic to $\mathbb{Z}^2$. By Theorem 1.1, Γ is the fundamental group of a virtually special cube complex $\mathcal{S}$, so Q is a virtual retract of Γ by Theorem 1.3.

If M is nonorientable then we can pass to a degree-two orientable cover M' with fundamental group Γ'. As above, we see that for every geometrically finite subgroup Q of Γ, the intersection $Q' = Q \cap \Gamma'$ is a virtual retract of Γ'. Now, by the proof of [24, Theorem 2.10], it follows that H is a virtual retract of Γ. □

We take this opportunity to note that the combination of Proposition 5.5 and Corollary 5.2 shows that many subgroups of virtually special relatively hyperbolic groups are virtual retracts, even without any hypotheses on the parabolic subgroups. Indeed, we have the following.

Theorem 5.8. *Let $\mathcal{S}$ be a compact, virtually special cube complex and suppose that $\pi_1 \mathcal{S}$ is relatively hyperbolic. Then every fully relatively quasiconvex subgroup of $\pi_1 \mathcal{S}$ is a virtual retract.*

Recall that an element γ of a relatively hyperbolic group is called *hyperbolic* if it is not conjugate into a parabolic subgroup. Denis Osin has shown that cyclic

subgroups generated by hyperbolic elements are strongly relatively quasiconvex [31, Theorem 4.19]. In the torsion-free case this implies *a fortiori* that such subgroups are fully relatively quasiconvex.

Corollary 5.9. *Let $\mathcal{S}$ be a compact, virtually special cube complex and suppose that $\pi_1\mathcal{S}$ is relatively hyperbolic. For any hyperbolic element $\gamma \in \pi_1\mathcal{S}$, the cyclic subgroup $\langle\gamma\rangle$ is a virtual retract of $\pi_1\mathcal{S}$.*

Combining Theorem 5.8 with [26, Theorem 1.7], we obtain a slightly weaker version of Theorem 1.3 that holds when the peripheral subgroups are only assumed to be LERF and slender. (A group is *slender* if each subgroup is finitely generated.)

Corollary 5.10. *Let $\mathcal{S}$ be a compact, virtually special cube complex and suppose that $\pi_1\mathcal{S}$ is hyperbolic relative to a collection of slender, LERF subgroups. Then every relatively quasiconvex subgroup of $\pi_1\mathcal{S}$ is separable and every fully relatively quasiconvex subgroup of $\pi_1\mathcal{S}$ is a virtual retract.*

This result would apply if $\pi_1\mathcal{S}$ were the fundamental group of a finite-volume negatively curved manifold of dimension greater than three, in which case the parabolic subgroups would be non-abelian but nilpotent. Note that the full conclusion of Theorem 1.3 does not hold in this case: nilpotent groups that are not virtually abelian contain cyclic subgroups that are not virtual retracts.

6. Examples

In this section we describe many hyperbolic 3-manifolds that decompose into right-angled ideal polyhedra. Our aim is to display the large variety of situations in which Theorem 1.1 applies, and to explore the question of when a manifold that decomposes into right-angled ideal polyhedra is commensurable with a right-angled reflection orbifold. When this is the case, the results of this paper follow from previous work, notably that of Agol–Long–Reid [2]. The theme of this section is that this occurs among examples of lowest complexity, but that one should not expect it to in general.

Lemma 6.1 describes when one should expect a manifold M that decomposes into right-angled ideal polyhedra to be commensurable with a right-angled reflection orbifold. This is the case when all of the polyhedra decomposing M are isometric to a single right-angled ideal polyhedron $\mathcal{P}$, which furthermore is highly symmetric. A prominent example which satisfies this is the Whitehead link complement, which is commensurable with the reflection orbifold in the regular ideal octahedron.

The octahedron (also known as the 3-antiprism, see Figure 1) is the simplest right-angled ideal polyhedron, as measured by the number of ideal vertices. Propositions 6.3 and 6.4 imply that any manifold that decomposes into isometric copies of

the right-angled ideal octahedron or, respectively, the 4-antiprism, is commensurable with the corresponding reflection orbifold. On the other hand, in Section 6.2 we will describe an infinite family of "hybrid" hyperbolic 3-manifolds N_n, each built from both the 3- and 4-antiprisms, that are not commensurable with any 3-dimensional hyperbolic reflection orbifold. We use work of Goodman-Heard-Hodgson [16] here to explicitly identify the commensurator quotients for the N_n.

6.1. The simplest examples.. It may initially seem that a manifold that decomposes into right-angled polyhedra should be commensurable with the right-angled reflection orbifold in one or a collection of the polyhedra. This is not the case in general; however, the technical lemma below implies that it holds if all of the polyhedra are isometric and sufficiently symmetric.

Lemma 6.1. *Let M be a complete hyperbolic 3-manifold with a decomposition $\{\mathcal{P}_i\}$ into right-angled ideal polyhedra. For a face $f \in \mathcal{P}_i$, let γ_f be reflection in the hyperplane containing f. If for each such face, $\phi_f \circ \gamma_f$ is an isometry to the polyhedron $\mathcal{P}_j$ containing $\phi_f(f)$, then $\pi_1 M$ is contained in $\Gamma \rtimes \mathrm{Sym}(\mathcal{P}_1)$, where Γ is the reflection group in $\mathcal{P}_1$ and $\mathrm{Sym}(\mathcal{P}_1)$ is its symmetry group.*

Proof. Let M be a hyperbolic 3-manifold satisfying the hypotheses of the lemma. There is a "dual graph" to the polyhedral decomposition $\{\mathcal{P}_i\}$ with a vertex for each i, such that the vertex corresponding to $\mathcal{P}_i$ is connected by an edge to that corresponding to $\mathcal{P}_j$ for every face f of $\mathcal{P}_i$ such that $\phi_f(f)$ is a face of $\mathcal{P}_j$. Let $\mathcal{T}$ be the tiling of $\mathbb{H}^3$ by Γ-translates of $\mathcal{P}_1$. A maximal tree T in the dual graph determines isometries taking the $\mathcal{P}_i$ into $\mathcal{T}$ as follows.

Suppose f is a face of $\mathcal{P}_1$ that corresponds to an edge of T. Then by hypothesis $\phi_f^{-1}(\mathcal{P}_i) = \gamma_f(\mathcal{P}_1)$, where $\mathcal{P}_i$ contains $\phi_f(f)$. For arbitrary i, let α be an embedded edge path in T from the vertex corresponding to $\mathcal{P}_1$ to that of $\mathcal{P}_i$, and suppose $\mathcal{P}_{i_0}$ corresponds to the vertex with distance one on α from that of $\mathcal{P}_i$. We inductively assume that there exists an isometry ϕ_{i_0} such that $\phi_{i_0}(\mathcal{P}_{i_0})$ is a Γ-translate of $\mathcal{P}_1$. Let f be the face of $\mathcal{P}_{i_0}$ corresponding to the edge of T between $\mathcal{P}_{i_0}$ and $\mathcal{P}_i$. Then $\phi_{i_0}\gamma_f\phi_{i_0}^{-1} = \gamma_{\phi_{i_0}(f)} \in \Gamma$, so by hypothesis,

$$\phi_{i_0} \circ \phi_f^{-1}(\mathcal{P}_i) = \phi_{i_0}\gamma_f(\mathcal{P}_{i_0}) = (\phi_{i_0}\gamma_f\phi_{i_0}^{-1})(\phi_{i_0}(\mathcal{P}_{i_0}))$$

is a Γ-translate of $\mathcal{P}_1$.

Now for each i, after replacing $\mathcal{P}_i$ by $\phi_i(\mathcal{P}_i)$ we may assume that there is some $\gamma_i \in \Gamma$ such that $\mathcal{P}_i = \gamma_i(\mathcal{P}_1)$. For a face f of $\mathcal{P}_i$, let $\mathcal{P}_j$ be the polyhedron containing $\phi_f(f)$. Then by hypothesis $\gamma_j^{-1}\phi_f\gamma_f\gamma_i \in \mathrm{Sym}(\mathcal{P}_1)$. Therefore $\phi_f \in \Gamma \rtimes \mathrm{Sym}(\mathcal{P}_1)$; thus the lemma follows from the Poincaré polyhedron theorem. □

A natural measure of the complexity of a right-angled ideal polyhedron is its number of ideal vertices. By this measure, the two simplest right-angled ideal poly-

hedra are the 3- and 4-antiprisms, pictured in Figure 1. (The general definition of a k-antiprism, $k \geq 5$ should be evident from the figure.)

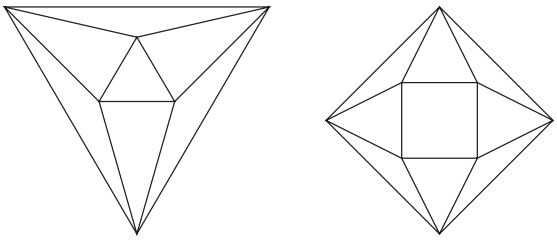

Figure 1. The 3- and 4-antiprisms.

Lemma 6.2. *The only right-angled ideal polyhedra with fewer than ten vertices are the* 3- *and* 4-*antiprisms.*

Proof. By a *polyhedron* we mean a 3-complex with a single 3-cell whose underlying topological space is the 3-dimensional ball, such that no two faces that share an edge e have vertices in common other than the endpoints of e. By Andreev's theorem, there is a right-angled ideal polyhedron in $\mathbb{H}^3$ with the combinatorial type of a given polyhedron if and only if each vertex has valence 4, there are no prismatic 3- or 4-circuits, and the following criterion holds: given faces f_0, f_1, and f_2 such that f_0 and f_2 each share an edge with f_1, f_0 and f_2 have no vertices in common with each other but not f_1. (A prismatic k-circuit is a sequence of k faces $f_0, f_1, \ldots, f_{k-1}$ such that no three faces have a common vertex but for each i, f_i shares an edge with f_{i-1} and f_{i+1}, taking indices modulo k.)

If f is a k-gon face of a right-angled ideal polyhedron $\mathcal{P}$, the final criterion above implies that $\mathcal{P}$ has at least $2k$ ideal vertices, since each face that abuts f contributes at least one unique vertex to $\mathcal{P}$. Thus any right-angled ideal polyhedron with fewer than 10 ideal vertices has only triangular and quadrilateral faces. Let v, e, and f be the number of vertices, edges and faces of $\mathcal{P}$, respectively. Since each vertex has valence 4, we have $4v = 2e$. If $\mathcal{P}$ has only triangular faces, then $2e = 3f$, and an Euler characteristic calculation yields

$$v - e + f = \frac{3f}{4} - \frac{3f}{2} + f = 2.$$

Therefore in this case $f = 8$, and it is easy to see that $\mathcal{P}$ must be the 3-antiprism.

If $\mathcal{P}$ has a quadrilateral face f and only 8 vertices, then by the final criterion of the first paragraph all faces adjacent to it are triangles. The union of f with the triangular faces adjacent to it is thus a subcomplex that is homeomorphic to a disk and contains all vertices of $\mathcal{P}$. It follows that $\mathcal{P}$ is the 4-prism. Since each vertex of

a right-angled ideal polyhedron is 4-valent, the number of vertices is even, and the lemma follows. □

It is well known that the 3-antiprism $\mathcal{P}$, better known as the octahedron, is *regular*: there is a symmetry exchanging any two ordered triples (v, e, f) where $v \subset e \subset f$ are faces of dimension 0, 1, and 2, respectively. Now suppose M is a manifold with a decomposition into polyhedra $\{\mathcal{P}_i\}$ such that for each i, there is an isometry $\gamma_i \colon \mathcal{P} \to \mathcal{P}_i$. If $\mathcal{P}_i$ and $\mathcal{P}_j$ are polyhedra in this decomposition, containing faces f and f', respectively, such that $\phi_f(f) = f'$, then $\gamma_j^{-1}\phi_f\gamma_i$ takes one face of $\mathcal{P}$ isometrically to another; hence it is realized by a symmetry σ of $\mathcal{P}$. It follows that $\gamma_{f'} \circ \gamma_j \sigma \gamma_i^{-1} = \phi_f$. Thus Lemma 6.1 implies:

Proposition 6.3. *Let Γ_1 be the group generated by reflections in the sides of the octahedron $\mathcal{P}$, and let Σ_1 be its symmetry group. If M is a complete hyperbolic manifold that decomposes into copies of $\mathcal{P}$, then $\pi_1 M < \Gamma_1 \rtimes \Sigma_1$. In particular, $\pi_1 M$ is commensurable to Γ_1.*

The 4-antiprism does not have quite enough symmetry to directly apply Lemma 6.1, but its double across a square face is the cuboctahedron, the semi-regular polyhedron pictured on the right-hand side of Figure 2. The cuboctahedron has a symmetry exchanging any two square or triangular faces, and each symmetry of each face extends over the cuboctahedron.

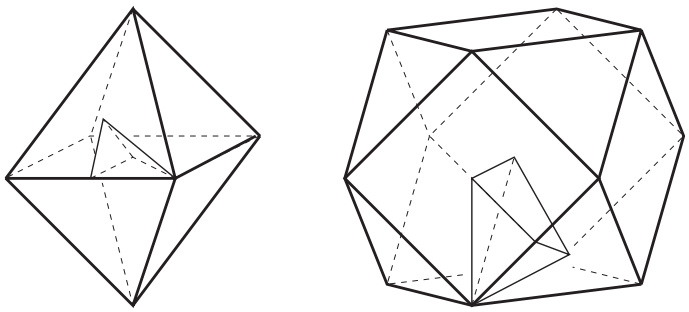

Figure 2. The ideal octahedron $\mathcal{P}_1$ and cuboctahedron $\mathcal{P}_2$.

Proposition 6.4. *Let Γ_2 be the group generated by reflections in the sides of the cuboctahedron, and let Σ_2 be its group of symmetries. If M is a complete hyperbolic 3-manifold that decomposes into copies of the cuboctahedron, then $\pi_1(M) < \Gamma_2 \rtimes \Sigma_2$. If M decomposes into 4-antiprisms, then $\pi_1(M)$ has an index-2 subgroup contained in $\Gamma_2 \rtimes \Sigma_2$.*

Proof. Since face pairing isometries must in particular preserve combinatorial type, it follows from Lemma 6.1 as argued above Proposition 6.3 that if M decomposes into copies of the cuboctahedron, then $\pi_1(M) < \Gamma_2 \rtimes \Sigma_2$.

Opposite square faces of the 4-antiprism inherit opposite colors from any checkering. Thus if a hyperbolic 3-manifold M has a checkered decomposition into right-angled ideal 4-antiprisms, they may be identified in pairs along, say, dark square faces, yielding a decomposition into right-angled ideal cuboctahedra. The proof of Theorem 1.1 shows that if the decomposition of M is not checkered, there is a twofold cover $\widetilde{M} \to M$ that inherits a checkered decomposition. Hence if M decomposes into 4-antiprisms, $\widetilde{M}$ decomposes into copies of the cuboctahedron. The final claim of the proposition follows. □

The results of [20] imply that for $j = 1, 2$, $\Gamma_j \rtimes \Sigma_j$ is isomorphic to the arithmetic group $\mathrm{PGL}_2(\mathcal{O}_j)$, where $\mathcal{O}_j$ is the ring of integers of $\mathbb{Q}(\sqrt{-j})$.

The fundamental domain for $\mathrm{Sym}(\mathcal{P}_1)$ pictured in Figure 2 intersects $\partial\mathcal{P}_1$ in a $(2, 3, \infty)$ triangle. We refer by Λ to the group generated by reflections in the sides of this triangle. The fundamental domain for $\mathrm{Sym}(\mathcal{P}_2)$ intersects a triangular face in a $(2, 3, \infty)$ triangle as well; thus Λ embeds in $\Gamma_j \rtimes \mathrm{Sym}(\mathcal{P}_j)$ for $j = 1$ and 2. The lemma below records an observation we will find useful in the following sections.

Lemma 6.5. *For $j = 1, 2$, let $\mathcal{T}_j$ be the tiling of $\mathbb{H}^3$ by Γ_j-conjugates of $\mathcal{P}_j$. The action of $\Gamma_j \rtimes \mathrm{Sym}(\mathcal{P}_j)$ is transitive on the set of all geodesic planes that contain a triangular face of a tile of $\mathcal{T}_j$.*

This lemma follows from the fact, evident by inspection of the fundamental domains in Figure 2, that $\mathrm{Sym}(\mathcal{P}_j)$ acts transitively on triangular faces of $\mathcal{P}_j$.

6.2. A family of one-cusped manifolds. In this section, we exhibit an infinite family $\{N_n\}$ of pairwise incommensurable manifolds that are not commensurable to *any* 3-dimensional reflection group. Each of these manifolds has a single cusp, and they are constructed using an explicit right-angled ideal polyhedral decomposition.

Definition. For $n \geq 2$, let $\{\mathcal{P}_i\}_{i=1}^{n+2}$ be a collection of right-angled ideal polyhedra embedded in $\mathbb{H}^3$ with the following properties.

(1) $\mathcal{P}_i$ is an octahedron if $i \in \{1, n+2\}$, and a cuboctahedron otherwise.

(2) There is an ideal vertex $\hat{v}$ shared by all the polyhedra.

(3) $\mathcal{P}_i \cap \mathcal{P}_j$ if and only if $i = j \pm 1$.

(4) If $\mathcal{P}_i$ and $\mathcal{P}_j$ meet, then they share a triangular face.

Define $\mathcal{D}_n = \bigcup_{i=1}^{n+2} \mathcal{P}_i$.

An isometric copy in $\mathbb{H}^3$ of such a collection is determined by an embedding of $\mathcal{P}_1$, a choice of $\hat{v}$, and a choice of triangular face $\mathcal{P}_1 \cap \mathcal{P}_2$. If we use the upper half space model for $\mathbb{H}^3$ then $\mathrm{Isom}^+(\mathbb{H}^3)$ is identified with $\mathrm{PSL}_2(\mathbb{C})$, by isometrically extending the action by Möbius transformations on $\partial\mathbb{H}^3 = \mathbb{C} \cup \{\infty\}$. Using this model, we apply an isometry so that $\hat{v} = \infty$ we can project the faces of the $\mathcal{P}_i$'s to $\partial\mathbb{H}^3$ to get a cell decomposition of $\mathbb{C}$. This decomposition is pictured for $n = 2$ in Figure 3.

Each 2-cell in the figure corresponds to a face of some $\mathcal{P}_i$ which is not shared by any other $\mathcal{P}_j$. Shade half of the faces of $\mathcal{P}_1$ and $\mathcal{P}_{n+2}$ gray and label them A, B, C, D, E, F, G, and H as indicated in the figure. Label the square face of $\mathcal{P}_2$ which shares an edge with B (respectively A, D) as X_1 (respectively Y_1, Z_1). Label the square face opposite X_1 as X_1' and so on. Now use the parabolic translation c that takes $\mathcal{P}_2$ to $\mathcal{P}_3$ to translate the labeling to the other cuboctahedra, adding one to the subscript every time we apply c.

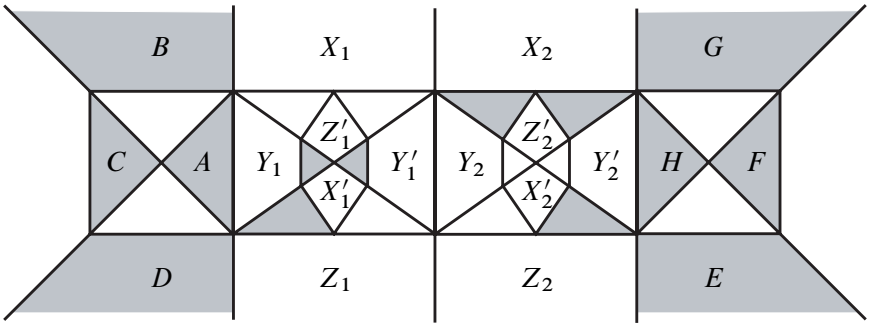

Figure 3. $\mathcal{D}_2$.

Define the isometries $\mathsf{a}, \mathsf{b}, \mathsf{f}, \mathsf{g}, \mathsf{x}, \mathsf{y}, \mathsf{z} \in \mathrm{Isom}^+(\mathbb{H}^3)$ as follows. The isometry taking A to B so that their shared vertex is taken to the vertex shared by B and C is a. The isometry taking C to D so that their shared vertex is taken to the vertex shared by B and D is b. The isometry taking E to F so that their shared vertex is taken to the vertex shared by F and G is f. The isometry taking G to H so that their shared vertex is taken to the vertex shared by H and F is g. The isometry taking Y_1' to X_1 so that their shared vertex is taken to the vertex shared by X_1 and Z_1' is x. The isometry taking Z_1' to Z_1 so that the vertex shared by Z_1' and Y_1' is taken to the vertex shared by X_1' and Z_1 is y. The isometry taking X_1' to Y_1 so that their shared vertex is taken to the vertex shared by Y_1 and Z_1 is z.

The set S_n defined below is a collection of face pairings for $\{\mathcal{P}_i\}_1^{n+2}$. Here we take $\mathsf{x}^{\mathsf{c}} = \mathsf{c}\mathsf{x}\mathsf{c}^{-1}$.

$$S_n = \left\{\mathsf{a}, \mathsf{b}, \mathsf{f}, \mathsf{g}, \mathsf{x}, \mathsf{y}, \mathsf{z}, \mathsf{x}^{\mathsf{c}}, \mathsf{y}^{\mathsf{c}}, \mathsf{z}^{\mathsf{c}}, \ldots, \mathsf{x}^{\mathsf{c}^{n-1}}, \mathsf{y}^{\mathsf{c}^{n-1}}, \mathsf{z}^{\mathsf{c}^{n-1}}\right\}.$$

By examining the combinatorics of these face pairings, one deduces that the quotient by these side pairings is a complete hyperbolic manifold N_n with finite volume and a single cusp. (See, for instance, [34, Theorem 11.1.6].) By Poincaré's polyhedron theorem [34, Theorem 11.2.2], $\Delta_n = \langle S_n \rangle$ is discrete and $\mathcal{D}_n$ is a fundamental domain for Δ_n. Furthermore, in the manner of [11], one can write down explicit matrices in $\mathrm{PSL}_2(\mathbb{C})$ which represent these isometries and see that the trace field for Δ_n is $\mathbb{Q}(i, \sqrt{2})$. Hence, $N_n \cong \mathbb{H}^3/\Delta_n$ is non-arithmetic.

Definition. The *commensurator* of $\Gamma < \mathrm{Isom}(\mathbb{H}^3)$ is defined as

$$\mathrm{Comm}(\Gamma) \doteq \{\mathsf{g} \in \mathrm{Isom}(\mathbb{H}^3) \mid [\Gamma : \mathsf{g}\Gamma\mathsf{g}^{-1} \cap \Gamma] < \infty\}.$$

It is easy to see that every group commensurable with Γ is contained in $\mathrm{Comm}(\Gamma)$. A well-known theorem of Margulis asserts that if Γ is discrete and acts with finite covolume, then $\mathrm{Comm}(\Gamma)$ is itself discrete if and only if Γ is not arithmetic (see [27, (1) Theorem]).

Let $G_n = \mathrm{Comm}(\Delta_n)$ and $O_n = \mathbb{H}^3/G_n$. Since Δ_n is a non-arithmetic Kleinian group, G_n is discrete and O_n is an orbifold. We will use the techniques of Goodman–Hodgson–Heard [16] to prove the following proposition.

Proposition 6.6. *Every element of G_n is orientation preserving. Hence, Δ_n is not commensurable to any 3-dimensional reflection group.*

Theorem 1.5 will follow immediately from the proposition above upon observing that the N_n are pairwise incommensurable. This follows most easily from a *Bloch invariant* computation. The Bloch invariant of a hyperbolic 3-manifold M is a sum of parameters, each an element of $\mathbb{C}^*$, of a tetrahedral decomposition of M, considered as an element of $\mathcal{P}(\mathbb{C})$. For a field k, the *Pre-Bloch group* $\mathcal{P}(k)$ is the quotient of the free $\mathbb{Z}$-module on $k - \{0, 1\}$ by a "five-term relation" that can be geometrically interpreted as relating different decompositions of the union of two tetrahedra. The Bloch group $\mathcal{B}(k)$ is a subgroup of $\mathcal{P}(k)$; see e.g. [28].

We will use the decomposition of N_n into a collection of 2 right-angled ideal octahedra and n cuboctahedra. These may each be divided into tetrahedra yielding a decomposition of N_n. The parameters of the tetrahedra contained in the octahedron sum to an element $\beta_1 \in \mathcal{B}(\mathbb{Q}(i))$, and those of the cuboctahedron sum to an element $\beta_2 \in \mathcal{B}(\mathbb{Q}(i\sqrt{2}))$. It can be showed that β_1 and β_2 are linearly independent in $\mathcal{B}(\mathbb{Q}(i, \sqrt{2}))$, and this in turn implies that the invariants $2 \cdot \beta_1 + n \cdot \beta_2$ of the N_n are pairwise linearly independent. Hence the N_n are pairwise incommensurable; see [11, Proposition 4.5] for an analogous proof.

In proving Proposition 6.6, we give a partial description of the commensurator G_n. We use the algorithm of [16] to perform such computations here and in Section 7.2, so we briefly introduce the set-up below. The *Lorentz inner product* on $\mathbb{R}^4$ is the

degenerate bilinear pairing

$$\langle \boldsymbol{v}, \boldsymbol{w} \rangle = v_1 w_1 + v_2 w_2 + v_3 w_3 - v_4 w_4.$$

The *hyperboloid model* of $\mathbb{H}^3$ is the set $\{\boldsymbol{v} \mid \langle \boldsymbol{v}, \boldsymbol{v} \rangle = -1,\ v_4 > 0\}$ equipped with the Riemannian metric on tangent spaces determined by the Lorentz inner product. The *positive light cone* is the set $L^+ = \{\boldsymbol{v} \mid \langle \boldsymbol{v}, \boldsymbol{v} \rangle = 0,\ v_4 \geq 0\}$. The *ideal boundary* $\partial\mathbb{H}^3$ is identified with the set PL^+ of equivalence classes of $\boldsymbol{v} \in L^+$, where $\boldsymbol{v} \sim \boldsymbol{w}$ if $\boldsymbol{w} = \lambda \boldsymbol{v}$ for $\lambda \in \mathbb{R}^+$.

Given a vector $\boldsymbol{v} \in L^+$, we say the set $H_{\boldsymbol{v}} = \{\boldsymbol{w} \in \mathbb{H}^3 \mid \langle \boldsymbol{v}, w \rangle = -1\}$ is a *horosphere centered at* $v = [\boldsymbol{v}]$. If $\alpha \in \mathbb{R}^+$ the horosphere $H_{\alpha \boldsymbol{v}}$ is a horosphere centered at the same ideal point as $H_{\boldsymbol{v}}$ and if $\alpha \leq 1$ then $H_{\boldsymbol{v}}$ is contained in the horoball determined by $\alpha \boldsymbol{v}$. This correspondence between vectors in L^+ and horospheres in $\mathbb{H}^3$ is a bijection. Hence, we call the vectors in L^+ *horospherical vectors*.

The group $\mathrm{Isom}(\mathbb{H}^3)$ is the subgroup $\mathrm{O}_0(3,1) \subset \mathrm{GL}_4(\mathbb{R})$ (acting by matrix multiplication) which preserves the Lorentz inner product and the sign of the last coordinate of each vector in $\mathbb{R}^4$.

Suppose $M = \mathbb{H}/\Lambda$ is a complete finite volume hyperbolic orbifold with k cusps. For each cusp c_i of M, choose a horospherical vector $\boldsymbol{v}_i$ for which $H_{\boldsymbol{v}_i}$ projects to a cross section of c_i under the covering map $\mathbb{H}^3 \to M$. Then $V = \Lambda \cdot \{\boldsymbol{v}_i\}_1^k$ is Λ-invariant and determines a Λ-invariant set of horospheres. The convex hull C of V in $\mathbb{R}^4$ is called the *Epstein–Penner convex hull*. Epstein and Penner show that ∂C consists of a countable set of 3-dimensional faces F_i, where each F_i is a finite sided Euclidean polyhedron in $\mathbb{R}^4$. Furthermore, this decomposition of ∂C projects to a Λ-invariant tiling $\mathcal{T}$ of $\mathbb{H}^3$ [14, Proposition 3.5 and Theorem 3.6]. If M is a manifold then the quotient of this tiling by Λ gives a cell decomposition of M. We refer to the tiling as a *canonical tiling* for M and to the cell decomposition as a *canonical cell decomposition* of M. If we make a different choice for $\{\boldsymbol{v}_i\}_1^k$ by multiplying each vector by a common positive scalar then the resulting Epstein–Penner convex hull differs from C by multiplication by this scalar. The combinatorics of the boundary of this scaled convex hull is identical to that of C and projects exactly to the tiling $\mathcal{T}$. Hence, we obtain all possibilities for canonical tilings using initial sets of the form $\{\boldsymbol{v}_1, \alpha_2 \boldsymbol{v}_2, \ldots, \alpha_k \boldsymbol{v}_k\}$.

Consider the group of symmetries $\mathrm{Sym}(\mathcal{T}) \subset \mathrm{Isom}(\mathbb{H}^3)$. Since $\mathcal{T}$ is Λ-invariant $\Lambda \subset \mathrm{Sym}(\mathcal{T})$. On the other hand, $\mathrm{Sym}(\mathcal{T})$ acts on the set V of horospherical vectors. It follows that $\mathrm{Sym}(\mathcal{T})$ is discrete [16, Lemma 2.1] and therefore $\mathbb{H}^3/\Lambda \to \mathbb{H}^3/\mathrm{Sym}(\mathcal{T})$ is a finite cover between orbifolds.

Suppose that Λ is non-arithmetic. Since $\mathrm{Comm}(\Lambda)$ is the unique maximal discrete group that contains Λ, then $\mathrm{Sym}(\mathcal{T}) \subset \mathrm{Comm}(\Lambda)$ for every canonical tiling $\mathcal{T}$. Furthermore, every canonical tiling for $\mathrm{Comm}(\Lambda)$ is also a canonical tiling for Λ, hence $\mathrm{Comm}(\Lambda) = \mathrm{Sym}(\mathcal{T})$ for some canonical tiling $\mathcal{T}$ for Λ.

We say that a set $\{\mathcal{P}_i\}$ of ideal polyhedra Λ-*generate* the tiling $\mathcal{T}$ if every tile of $\mathcal{T}$ is of the form $\gamma \mathcal{P}_i$ for some $\gamma \in \Lambda$ and some i. The canonical tilings can be

determined using elementary linear algebra. According to [16, Lemma 3.1], a set $\{\mathcal{P}_i\}$ of ideal polyhedra Λ-generates the canonical tiling associated to the set V if

(1) $\Lambda \cdot \{\mathcal{P}_i\}$ is a tiling of $\mathbb{H}^3$,

(2) given any vertex of any $\mathcal{P}_i$ there is a horospherical vector $\boldsymbol{v} \in V$ so that the vertex lies at the center of the horosphere $H_{\boldsymbol{v}}$,

(3) the set of horospherical vectors corresponding to the vertices of any given $\mathcal{P}_i$ lie on a single plane in $\mathbb{R}^4$,

(4) if $\mathcal{P}_i$ and $\gamma\mathcal{P}_j$ are two tiles that meet in a common face then the Euclidean planes in $\mathbb{R}^4$ determined by the two tiles meet convexly.

The last two conditions can be re-phrased using linear algebra. If $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_s\}$ are the horospherical vectors for $\mathcal{P}_i$ and $\boldsymbol{w}$ is a horospherical vector for a neighboring tile which is not shared by $\mathcal{P}_i$ then there exists a *normal* vector for $\mathcal{P}_i$, $\boldsymbol{n} \in \mathbb{R}^4$ such that

(3) (coplanar) $\boldsymbol{n} \cdot \boldsymbol{v}_i = 1$ for every $i = 1, \ldots s$, and

(4) (positive tilt) $\boldsymbol{n} \cdot \boldsymbol{w} > 1$,

where $\cdot$ denotes the standard Euclidean inner product. Note that these conditions are invariant under $\mathrm{Isom}(\mathbb{H}^3)$, for if $\boldsymbol{n} \cdot \boldsymbol{v} = \alpha$ and $A \in \mathrm{Isom}(\mathbb{H}^3)$ then $(\boldsymbol{n}A^{-1}) \cdot A\boldsymbol{v} = \alpha$.

Proposition 6.7. *Let $\Delta_n < \mathsf{O}_0(3,1)$ be determined by the following embedding of the $\mathcal{P}_i$ in $\mathbb{H}^3$: the isometry group of $\mathcal{P}_2$ fixes $(0,0,0,1)^\top$, the ideal vertex $\hat{v}$ shared by the $\mathcal{P}_i$ is $[\hat{\boldsymbol{v}}]$, where $\hat{\boldsymbol{v}} = (2,0,0,2)^\top$, and $\mathcal{P}_1 \cap \mathcal{P}_2$ has ideal vertices $[\hat{\boldsymbol{v}}], [\boldsymbol{v}_9], [\boldsymbol{v}_4]$, where $\boldsymbol{v}_4 = (1,1,-\sqrt{2},2)^\top$ and $\boldsymbol{v}_9 = (1,-1,-\sqrt{2},2)^\top$. Let $\mathcal{T}_n$ be the tiling of $\mathbb{H}^3$ determined by $V_n = \Delta_n \cdot \{\hat{\boldsymbol{v}}\}$. The tiles of $\mathcal{T}_n$ are the Δ_n-orbits of the $\mathcal{P}_i$.*

Proof. If X is a $4 \times n$ matrix we denote the i^{th} column of X by x_i. When the columns of X lie in L^+ and the convex hull of the corresponding ideal points is an ideal polyhedron we call the polyhedron $\mathcal{P}_X$. Consider the matrices

$$M = \begin{pmatrix} 2 & 1 & 0 & 1 & 0 & -1 & -2 & -1 & 1 & -1 & -1 & 1 \\ 0 & 1 & 2 & 1 & -2 & -1 & 0 & -1 & -1 & 1 & 1 & -1 \\ 0 & \sqrt{2} & 0 & -\sqrt{2} & 0 & \sqrt{2} & 0 & -\sqrt{2} & -\sqrt{2} & -\sqrt{2} & \sqrt{2} & \sqrt{2} \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \end{pmatrix},$$

and

$$N = \begin{pmatrix} \sqrt{2} & 0 & 0 & -\sqrt{2} & 0 & 0 \\ 0 & \sqrt{2} & 0 & 0 & -\sqrt{2} & 0 \\ 0 & 0 & \sqrt{2} & 0 & 0 & -\sqrt{2} \\ \sqrt{2} & \sqrt{2} & \sqrt{2} & \sqrt{2} & \sqrt{2} & \sqrt{2} \end{pmatrix}.$$

The columns of M and N are horospherical vectors and represent horospheres centered about the ideal vertices of a regular ideal cuboctahedron and octahedron respectively. These matrices are chosen so that, for $X = M, N$, the isometries in $\mathrm{Isom}(\mathcal{P}_X)$ all fix $(0,0,0,1)^\top \in \mathbb{H}^3$ and the columns of X are $\mathrm{Isom}(\mathcal{P}_X)$-invariant. Furthermore, if h is the orientation preserving hyperbolic isometry that takes the triangular face (n_1, n_2, n_3) of $\mathcal{P}_N$ to the triangular face (m_1, m_9, m_4) of $\mathcal{P}_M$ so that $\mathsf{h}(\mathcal{P}_N) \cap \mathcal{P}_M$ is exactly this face, then our choice of horospheres agree on this intersection. That is, $\mathsf{h}(n_1, n_2, n_3) = (m_1, m_9, m_4)$.

Let $\mathcal{P}_1 = \mathsf{h}(\mathcal{P}_N)$ and $\mathcal{P}_2 = \mathcal{P}_M$. Embed the remaining polyhedra in $\{\mathcal{P}_i\}_1^{n+2}$, as described above, so that the common ideal vertex is the center of the m_1 horosphere. Choose horospherical vectors for the $\mathcal{P}_i$'s so that they are $\mathrm{Isom}(\mathcal{P}_i)$-invariant and to coincide with the horospherical vectors of $\mathcal{P}_{i\pm1}$ wherever ideal vertices are shared.

Notice that the face pairings of $\mathcal{P}_i$ in S_n are all compositions of elements of $\mathrm{Isom}(\mathcal{P}_i)$ with parabolics that fix an ideal vertex of $\mathcal{P}_i$. Since we have chosen our horospherical vectors to be $\mathrm{Isom}(\mathcal{P}_i)$-invariant, it follows that our choice of horospheres is compatible with the face pairings in S_n. Hence, the choice of horospheres descends to a choice of horospherical torus in N_n and therefore determines a canonical cell decomposition of N_n and a canonical tiling of $\mathbb{H}^3$ whose symmetry group is G_n. To prove the proposition, we need to show that this tiling is $\mathcal{T}_n$.

Take $\boldsymbol{n} = (0,0,0,1/2)^\top$. Then $\boldsymbol{n} \cdot m_i = 1$ for $i = 1, \ldots, 12$ and $\sqrt{2}\boldsymbol{n} \cdot n_i = 1$ for $i = 1, \ldots, 6$. Therefore by Goodman–Hodgson–Heard's criterion (3), the horospherical vertices of $\mathsf{k}(\mathcal{P}_i)$ are coplanar for every $\mathsf{k} \in \Delta_n$. It remains only to show that condition (4) holds for adjacent pair of cuboctahedra that meet along a triangular face, an adjacent pair of cuboctahedra that meet along a square face, and an octahedron adjacent to a cuboctahedron.

If Q is a cuboctahedron adjacent to $\mathcal{P}_M$ sharing the triangular face (m_1, m_9, m_4) with $\mathrm{Isom}(Q)$-invariant horospherical vectors which agree with (m_1, m_9, m_4) then $w = (7, 1, -5\sqrt{2}, 10)^\top$ is a horospherical vector for Q which is not shared by $\mathcal{P}_M$. We have $\boldsymbol{n} \cdot w = 5 > 1$. If Q is a cuboctahedron adjacent to $\mathcal{P}_M$ sharing the square face (m_1, m_2, m_3, m_4) with $\mathrm{Isom}(Q)$-invariant horospherical vectors which agree with (m_1, m_2, m_3, m_4) then $w = (3, 5, -\sqrt{2}, 6)^\top$ is a horospherical vector for Q which is not shared by $\mathcal{P}_M$. We have $\boldsymbol{n} \cdot w = 3 > 1$. The octahedron $\mathsf{h}(\mathcal{P}_N)$ is adjacent to $\mathcal{P}_M$ sharing the face (m_1, m_9, m_4). Its vectors are invariant under the isometry group of $\mathsf{h}(\mathcal{P}_N)$ and they agree with those of $\mathcal{P}_M$ along the shared face. The vector $w = (2 + 2\sqrt{2}, 0, -2 - 2\sqrt{2}, 4 + 4\sqrt{2})^\top$ is a horospherical vector for $\mathsf{h}(\mathcal{P}_N)$ which is not shared by $\mathcal{P}_M$. We have $\boldsymbol{n} \cdot w = 2 + \sqrt{2} > 1$. □

For $i = 2, n+1$, shade each face of $\mathcal{P}_i$ gray if it is identified with a face of an octahedron in the quotient. For the other cuboctahedra $\mathcal{P}_i$, color each triangular face red if it is identified with a face of $\mathcal{P}_{i-1}$. (In Figure 3, every white triangular face of $\mathcal{P}_3$ should be colored red.)

The tiles of $\mathcal{T}_n$ inherit a coloring from the coloring of the $\mathcal{P}_i$'s. We can further classify the triangular faces in cuboctahedral tiles of $\mathcal{T}$ into *type I* and *type II* triangles. A face of a cuboctahedral tile T is type I if it has exactly one ideal vertex that is shared by a triangular face of T of the opposite color. Triangular faces of cuboctahedra that are not type I are type II.

Proof of Proposition 6.6. Suppose $\mathsf{h} \in G_n - \Delta_n$. By [16], h is a symmetry for the tiling $\mathcal{T}_n$. The polyhedron $\mathcal{D}_n$ is a fundamental domain for Δ_n, so by composing h with some element of Δ_n, we may assume that $\mathsf{h}(\mathcal{P}_2) \in \{\mathcal{P}_i\}_1^{n+2}$. It is clear that h must preserve the set of gray faces in the tiling, hence $\mathsf{h}(\mathcal{P}_2)$ is either $\mathcal{P}_2$ or $\mathcal{P}_{n+1}$.

The isometry h must also preserve the types of the triangular faces of cuboctahedra. By examining the combinatorics of the face pairings in S_n, we see that every cuboctahedron in the tiling has exactly two vertices that are shared by a pair of type I triangles. There is one such vertex for each of the two triangular colors on the tile. Let v be the vertex of $\mathcal{P}_2$ which is shared by the two gray type I triangles of $\mathcal{P}_2$ and w the vertex shared by the two white type I triangles. If $\mathsf{h}(\mathcal{P}_2) = \mathcal{P}_2$ then, by considering the coloring of $\mathcal{P}_2$ we see that h must be the order-2 elliptic fixing v and w. If, on the other hand, we have $\mathsf{h}(\mathcal{P}_2) = \mathcal{P}_{n+1}$ then $\mathsf{h}(v)$ must be the vertex shared by the two gray type I triangles of $\mathcal{P}_{n+1}$ and $\mathsf{h}(w)$ must be the vertex shared by the two red type I triangles of $\mathcal{P}_{n+1}$. The gray pattern on $\mathcal{P}_{n+1}$ forces h to be orientable. □

7. Augmented links

A rich class of examples that satisfy the hypotheses of Theorem 1.1 is that of the *augmented links*. These were introduced by Adams [1] and further studied in e.g. [22], [32], and [33]. In this section we will describe their construction and, in Section 7.1, classify up to scissors congruence the complements of augmented links with at most 5 twist regions. We will discuss when an augmented link complement is commensurable with a right-angled reflection orbifold, and in Section 7.2 describe an infinite family of augmented link complements that do not have this property.

A link L in S^3 with hyperbolic complement determines (not necessarily uniquely) an augmented link using a projection of L which is *prime* and *twist-reduced*. We will regard a projection of L as a 4-valent graph in the plane, together with crossing information at each vertex, and use the term *twist region* to denote either a maximal collection of bigon regions of the complement arranged end-to-end or an isolated crossing that is not adjacent to any bigon.

A projection is prime if there is no simple closed curve γ in the projection plane intersecting it in exactly two points, with the property that each component of the complement of γ contains a crossing. A projection is twist-reduced if for every simple closed curve γ in the projection plane which intersects it in four points, such that two points of intersection are adjacent to one crossing and the other two are adjacent to

another, there is a single twist region containing all crossings in one component of the complement of γ.

An augmented link is obtained from a prime, twist reduced projection by encircling each twist region with a single unknotted component, which we call a *clasp*. This process is illustrated in Figure 4 for the figure-8 knot, pictured on the left-hand side with its twist regions in boxes. The augmented link that it determines is pictured in the middle of the figure. Each link with hyperbolic complement admits a prime, twist reduced diagram, and the augmented link obtained from such a diagram also has hyperbolic complement (a direct proof of this fact is given in Theorem 6.1 of [33]). Thus every hyperbolic link complement in S^3 is obtained by Dehn surgery on some cusps of the complement of an augmented link.

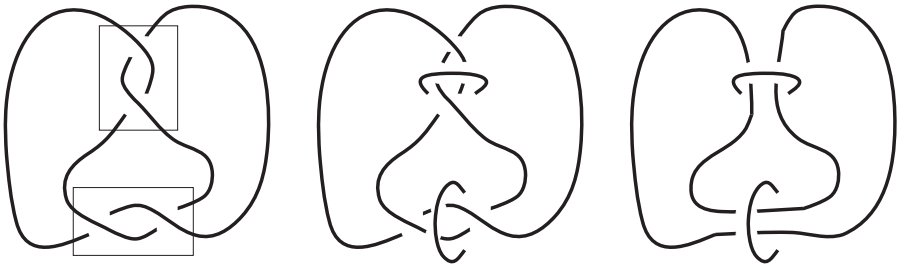

Figure 4. Augmenting the figure-8 knot.

Each clasp of an augmented link L bounds a disk that has two points of transverse intersection with L. Given such a disk D, a family of homeomorphisms of $S^3 - L$ is determined by cutting along the twice-punctured open disk $D - L$ and re-gluing by a rotation of angle $n \cdot 2\pi$, where $n \in \mathbb{Z}$. This adds or subtracts $2n$ crossings to the twist region of L encircled by the clasp bounding D. It follows that the link on the right-hand side of Figure 4 has a complement homeomorphic to that of the link in the middle. The complements of two augmented links that differ by only a single crossing in a twist region are not necessarily homeomorphic; however, we will see below that they are scissors congruent. We also have:

Lemma 7.1. *Let L be an augmented link. Reflection through the projection plane determines an automorphism of $S^3 - L$.*

This is because while such a reflection changes the sign of each crossing, it does not change the parity of the number of crossings per twist region.

Given an augmented link projection, the appendix to [22] describes a decomposition of its complement into two isometric ideal polyhedra. These polyhedra may be checkered so that each white face lies in the projection plane and each dark face is

an ideal triangle in a "vertical" twice-punctured disk. This is illustrated in Figure 5 for an augmented link with two twist regions.

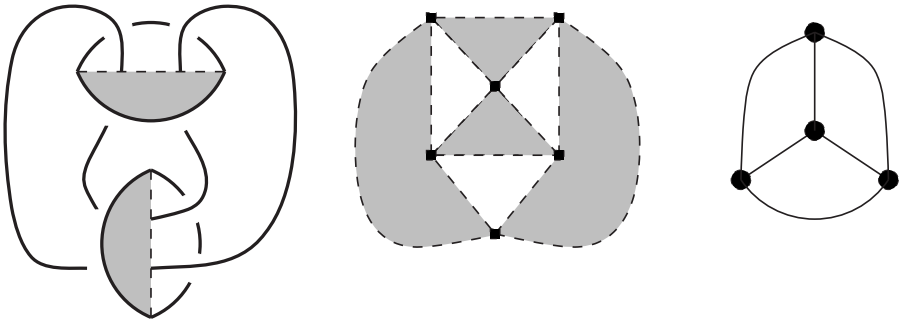

Figure 5. An augmented link, the associated polyhedron, and its crushtacean.

On the left-hand side of the figure, the dotted lines divide each twice-punctured clasp disk into the union of two ideal triangles. We arrange for these disks to meet the projection plane transversely in the dotted lines, so the darkened ideal triangles lie above the projection plane and the others below it. Cutting the link complement along the clasp disks and the projection plane yields two ideal polyhedra, one above and one below the projection plane, with edges coming from the dotted arcs. These are isomorphic by reflection through the projection plane. Flattening the two-skeleton of the polyhedron above it onto the plane yields the polyhedron in the middle of the figure, an ideal octahedron, where each of the darkened half-disks on the left-hand side gives rise to two ideal triangles and the link itself has been shrunken to darkened rectangles at the vertices. (See also [32, Figure 3].)

If L is an augmented link, after removing all crossings in each twist region, we call the polyhedron produced by cutting along the projection plane and clasp disks the *ideal polyhedron associated to* L. This polyhedron may be checkered by coloring black the triangular faces that lie in clasp disks and white the faces that lie in the projection plane. Note also that each black triangular face has a unique ideal vertex corresponding to a clasp. The following lemma summarizes the construction of the appendix to [22], in our language.

Lemma 7.2. *If L is an augmented link with hyperbolic complement, there is a right-angled checkered ideal polyhedron $\mathcal{P}$ in $\mathbb{H}^3$ combinatorially isomorphic to the ideal polyhedron associated to L. For a face f of $\mathcal{P}$, let ρ_f denote reflection in the plane containing f. Fix a white face f_0 of $\mathcal{P}$, and let $\bar{\mathcal{P}} = \rho_{f_0}(\mathcal{P})$, $\bar{f} = \rho_{f_0}(f)$ for each face f of $\mathcal{P}$, and $\bar{v} = \rho_{f_0}(v)$ for each ideal vertex. Then the quotient of $\mathcal{P} \cup \bar{\mathcal{P}}$ by the following face pairing gives a right-angled ideal decomposition of $S^3 - L$.*

(1) *If $f \neq f_0$ is a white face of $\mathcal{P}$, let $\phi_f = \rho_{f_0} \circ \rho_f$, taking f to $\bar{f} \subset \bar{\mathcal{P}}$.*

(2) *If f is a black triangular face of $\mathcal{P}$, let f' be the black face of $\mathcal{P}$ that shares the ideal vertex v of f corresponding to a clasp.*

   (a) *If the corresponding twist region has an even number of crossings, let ϕ_f be the unique orientation-preserving isometry with $\phi_f(f) = f'$, $\phi_f(v) = v$, and $\phi_f(\mathcal{P}) \cap \mathcal{P} = f'$.*

   (b) *If the corresponding twist region has an odd number of crossings, let ϕ_f be the unique orientation-preserving isometry with $\phi_f(f) = \bar{f}'$, $\phi_f(v) = \bar{v}$, and $\phi_f(\mathcal{P}) \cap \bar{\mathcal{P}} = \bar{f}'$.*

Furthermore, ρ_{f_0} induces the isometry of $S^3 - L$ supplied by Lemma 7.1. In particular, $\phi_{\bar{f}} = \rho_{f_0} \circ \phi_f \circ \rho_{f_0}$ for each face f of $\mathcal{P}$.

For another discussion of the content of Lemma 7.2, see [32, §2.3]. In particular, Figure 4 there clarifies the different gluings producing twist regions with even vs. odd numbers of crossings. The last sentence of the lemma is not covered in [22]; however it follows easily from the discussion above.

On the right-hand side of Figure 5 is the compact polyhedron obtained from the checkered ideal octahedron by the following rule: it has a vertex corresponding to every dark face and an edge joining each pair of vertices that correspond to dark faces which share ideal vertices. We will call this the *crushtacean* of L, since it may be regarded as obtained by crushing the darkened faces of the associated right-angled polyhedron to points. We note that each vertex of the crushtacean has valence 3, since each dark face is an ideal triangle. The right-angled ideal polyhedron associated to L is recovered by truncation from its crushtacean.

For an alternative perspective on obtaining the crushtacean and a connection with Andreev's theorem, we refer the reader to Section 6 of [33], in particular page 487. The one-skeleton of the crushtacean is the graph Γ dual to the nerve γ of the circle packing defined there. We thank Jessica Purcell for pointing this out.

Figure 6 illustrates two augmented links with the same underlying polyhedron, each depicted draped over the one-skeleton of its crushtacean, the 6-prism. (More generally, for $k \geq 3$ we will call the *k-prism* the polyhedron combinatorially isomorphic to the cartesian product of a k-gon with an interval.) Since the associated right-angled ideal polyhedron is obtained by truncating vertices of the crushtacean, its ideal vertices occur at midpoints of edges. Each triangular face resulting from truncation is paired with one of its neighbors across an ideal vertex producing a clasp; thus for each vertex of the crushtacean, exactly one edge which abuts it is encircled by a clasp. Each other edge carries a single strand of the "horizontal" component of the augmented link.

Since the ideal polyhedron $\mathcal{P}$ associated to an augmented link is canonically obtained from its crushtacean, each symmetry of the crushtacean determines a combinatorial symmetry of $\mathcal{P}$. Together with Mostow rigidity, this implies:

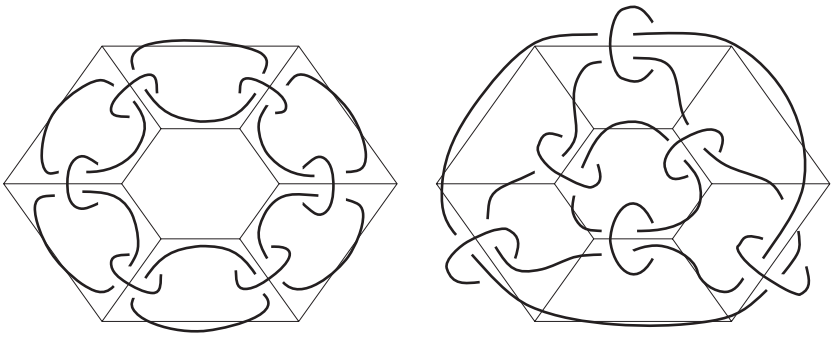

Figure 6. Two augmented links with crushtacean the 6-prism.

Lemma 7.3. *Let L be an augmented link, $\mathcal{P}$ the associated right-angled ideal polyhedron in $\mathbb{H}^3$, and $\mathcal{C}$ its crushtacean. There is a canonical injection* $\mathrm{Sym}(\mathcal{C}) \to \mathrm{Sym}(\mathcal{P})$.

Lemma 7.3 implies that the complement of an augmented link with a highly symmetric crushtacean may be commensurable with the reflection group in the associated right-angled polyhedron.

Lemma 7.4. *Let L be an augmented link, $\mathcal{P}$ the associated right-angled polyhedron, and $\mathcal{C}$ its crushtacean, and suppose $\mathcal{C}$ has the property that for each clasp component K of L, corresponding to an edge e of $\mathcal{C}$ with vertices v and v',*

(1) *if K encloses a twist region with an even number of crossings, there is a reflective involution of $\mathcal{C}$ preserving e and exchanging v with v';*

(2) *if K encloses a twist region with an odd number of crossings, there is a rotational involution of $\mathcal{C}$ preserving e and exchanging v with v'.*

Then $\pi_1(S^3 - L) < \Gamma_{\mathcal{P}} \rtimes \mathrm{Sym}(\mathcal{P})$, where $\Gamma_{\mathcal{P}}$ is the group generated by reflections in $\mathcal{P}$ and $\mathrm{Sym}(\mathcal{P})$ *is the group of symmetries of $\mathcal{P}$.*

Proof. Lemma 7.3 implies that for each edge e of $\mathcal{C}$ corresponding to a clasp K of L, there is an involution ι_e of $\mathcal{P}$ that exchanges the triangular faces f and f' corresponding to v and v', and fixes the ideal vertex that they share. This involution is a reflection or 180-degree rotation in case (1) or (2) above, respectively.

We now use the notation of Lemma 7.2, and record that case (2a) there is the same as case (1) above. In this case, $\iota_e \circ \rho_f$ realizes the orientation-preserving isometry ϕ_f there. In case (2) above, corresponding to case (2b) of Lemma 7.2, the required isometry ϕ_f is realized by $\rho \circ \iota_e \circ \rho_f$. □

Lemma 7.4 implies for instance that the link on the left-hand side of Figure 6 is commensurable with the reflection group in the corresponding right-angled polyhedron, but it does not apply to the link on the right-hand side on account of the twist region with a single crossing. On the other hand, the commensurability classes of some links are entirely determined by their crushtaceans.

Corollary 7.5. *Suppose L is an augmented link such that the crushtacean of L is a regular polyhedron. Then $\pi(S^3 - L)$ is commensurable with the reflection group in the sides of the corresponding right-angled polyhedron.*

In some cases the crushtacean of an augmented link may not have much symmetry, but it may be built from highly symmetric polyhedra. In such cases the link may have hidden symmetries. We will say a crushtacean is *decomposable* if it contains a prismatic 3-cycle – that is, a sequence of three faces so that any two intersect along an edge but all three do not share a common vertex – and *indecomposable* otherwise.

If $\mathcal{C}$ is a decomposable crushtacean, we decompose along a prismatic 3-cycle by selecting a simple closed curve γ which lies in the union of the faces of the cycle and intersects each of the edges of the cycle once, and using the following procedure: cut along γ, separate the components that result, and complete each by replacing γ with a single vertex containing the endpoints of all the three edges intersecting it. This is illustrated for the triangular prism in Figure 7, with the dotted curve on the left-hand side representing γ. Decomposing results in a disjoint union of two tetrahedra.

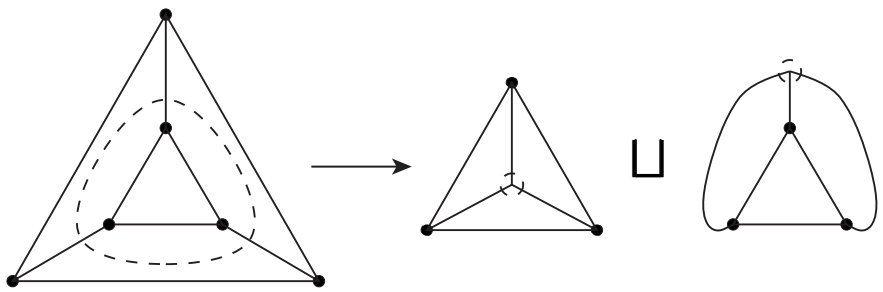

Figure 7. Decomposing the 3-prism into two tetrahedra.

Suppose L is a link with a decomposable crushtacean $\mathcal{C}$, and let f_0, f_1, and f_2 determine a prismatic 3-cycle of $\mathcal{C}$. Then the corresponding faces in the associated right-angled ideal polyhedron $\mathcal{P}$, obtained by truncating vertices of $\mathcal{C}$, do not pairwise intersect but each two share an ideal vertex. It is an elementary fact of hyperbolic geometry that there is a single hyperplane $\mathcal{H}$ which perpendicularly intersects the hyperplanes containing each of f_0, f_1 and f_2. Cutting $\mathcal{P}$ along $\mathcal{H}$ decomposes it into two new right-angled ideal polyhedra, each with an ideal triangular face contained

in $\mathcal{H}$. Their crushtaceans are obtained by decomposing $\mathcal{C}$ along the prismatic cycle determined by f_0, f_1, and f_2.

Lemma 7.6. *Suppose L is an augmented link such that the crushtacean of L decomposes into a disjoint union of copies of $\mathcal{C}$, where $\mathcal{C}$ is a regular polyhedron. Then $\pi_1(S^3 - L)$ is contained in $\Gamma_{\mathcal{P}} \rtimes \mathrm{Sym}(\mathcal{P})$, where $\mathcal{P}$ is the right-angled ideal polyhedron obtained from $\mathcal{C}$ by truncating vertices.*

Proof. There is a tiling $\mathcal{T}$ of $\mathbb{H}^3$ consisting of $\Gamma_{\mathcal{P}}$-translates of $\mathcal{P}$. If $\mathcal{C}_0$ is the crushtacean of L, then the hypothesis and the description above the lemma establish that the associated right-angled polyhedron $\mathcal{P}_0$ is a union of tiles of $\mathcal{T}$. Checkering $\mathcal{P}_0$ so that dark faces are triangles obtained by truncating vertices of $\mathcal{C}_0$, we claim that for each pair of dark faces f and f' which share an ideal vertex v, there exist in $\Gamma_{\mathcal{P}} \rtimes \mathrm{Sym}(\mathcal{P})$ both a reflective and a rotational involution of $\mathbb{H}^3$ exchanging f and f' and fixing v. We will prove the claim by induction on the number of tiles comprising $\mathcal{P}_0$. The case of one tile, $\mathcal{C}_0 = \mathcal{C}$, follows as in the proof of Lemma 7.4 from the fact that $\mathcal{C}$ is regular.

Suppose that $\mathcal{P}_0$ is the union of more than one tile, and let $\gamma(\mathcal{P})$ be a $\Gamma_{\mathcal{P}}$-translate of $\mathcal{P}$ such that $\mathcal{P}_0$ is the union of $\gamma(\mathcal{P})$ and a polyhedron $\mathcal{P}_1 \subset \mathcal{T}$ across a face f which is an ideal triangle. The checkering of $\mathcal{P}_0$ determines checkerings of each of $\mathcal{P}_1$ and $\gamma(\mathcal{P})$ by declaring f to be dark. The claim holds for $\mathcal{P}_1$ by induction and for $\gamma(\mathcal{P})$ by the base case. Thus it only remains to verify the claim for dark faces of $\mathcal{P}_0$ sharing an ideal vertex, one of which lies in $\mathcal{P}_1$ and one in $\gamma(\mathcal{P})$.

Suppose f_0 and f_1 are dark faces, of $\gamma(\mathcal{P})$ and $\mathcal{P}_1$ respectively, which share an ideal vertex v in $\mathcal{P}_0$. Then each of f_0 and f_1 shares v with f. Let ρ_0 (respectively, ρ_1) be a reflective involution in $\Gamma_{\mathcal{P}} \rtimes \mathrm{Sym}(\mathcal{P})$ fixing v and exchanging f_0 (resp. f_1) with f, and let ι_0 and ι_1 be rotational involutions satisfying the same description. Then $\rho_1 \circ \rho_0$ and $\iota_1 \circ \rho_0$ are isometries of infinite order taking f_0 to f_1. This can be discerned by considering their actions on a horosphere centered at v, intersected by $\gamma(\mathcal{P})$ and $\mathcal{P}_1$ in adjacent rectangles. The first acts on this cross section as a translation and the second as a glide reflection. If ρ_{f_1} is reflection in the hyperplane containing f_1, it follows that $\rho_{f_1} \circ \rho_1 \circ \rho_0$ and $\rho_{f_1} \circ \iota_1 \circ \rho_0$ satisfy the conclusion of the claim.

The conclusion of the lemma now follows from Lemma 7.2. □

7.1. Examples with low complexity. The most natural measure of complexity of an augmented link is the number of twist regions, which is equal to half the number of dark faces of the associated right-angled polyhedron, or half the number of vertices of its crushtacean. Here we will classify the augmented link complements with up to five twist regions up to *scissors congruence*. We will say that finite-volume hyperbolic 3-manifolds are scissors congruent if they can be cut into identical collections of ideal polyhedra. It is natural for us to use this invariant because many different augmented

links may be produced by different choices of face pairing on the same underlying right-angled polyhedron.

Lemma 7.7. *The indecomposable crushtaceans with at most ten vertices are the tetrahedron, the cube* (*or* 4*-prism*)*, and the* 5*-prism.*

Proof. The only indecomposable crushtacean with a triangular face is the tetrahedron, since the family of faces adjacent to a triangular face determines a prismatic 3-cycle unless they share a common vertex. On the other hand, if a crushtacean $\mathcal{C}$ with at most ten vertices has a face which is a k-gon for $k \geq 6$, then two edges which emanate from distinct vertices of this face must share a common endpoint. That $\mathcal{C}$ is decomposable follows from the claim below.

Claim. *Suppose f is a face of a crushtacean $\mathcal{C}$, and e_0 and e_1 are distinct edges of $\mathcal{C}$, each with one endpoint on f, which share a vertex v. Then e_0 and e_1 bound a triangle face of $\mathcal{C}$ together with an edge of f.*

Proof of claim. The set $f \cup e_0 \cup e_1$ cuts $\partial\mathcal{C}$ into two disks. Let D be the closure of the disk that does not intersect the edge $e_2 \neq e_0, e_1$ with an endpoint at v. There is a face $f' \subset D$ of $\mathcal{C}$ which has v as a vertex and e_0 and e_1 as edges. Then f' intersects f along an edge e_0' with an endpoint at $e_0 \cap f$ and also along an edge e_1' with an endpoint at $e_1 \cap f$. But since f and f' cannot meet along more than one edge, we must have $e_0' = e_1'$. Thus since $e_0 \cup e_1 \cup e_0'$ forms a simple closed curve in the boundary of f', $f' = D$ is a triangle. □

Thus if $\mathcal{C}$ is indecomposable and not a tetrahedron, with at most ten vertices, then every face of $\mathcal{C}$ is a quadrilateral or pentagon. Let j be the number of quadrilateral faces and k the number of pentagon faces, and let v and e be the number of vertices and edges, respectively. Since each vertex is 3-valent, we have $3v = 2e$, and since each edge bounds two faces we have $2e = 4j + 5k$. Computing the Euler characteristic thus yields:

$$v - e + (j + k) = \frac{4j + 5k}{3} - \frac{4j + 5k}{2} + (j + k) = \frac{j}{3} + \frac{k}{6} = 2.$$

Using the equation above we find that $j + k/2 = 6$. Since we require that $\mathcal{C}$ have at most ten vertices, the vertex and edge equations yield $4j + 5k \leq 30$. Thus using the fact that j and k are non-negative integers, we find that either $j = 6$ and $k = 0$ (and hence $v = 8$) or $j = 5$ and $k = 2$ (and $v = 10$). The cube and the 5-prism respectively realize these possibilities. It remains to show that these are the unique crushtaceans with the prescribed numbers of quadrilateral and pentagon faces.

In general, if a crushtacean $\mathcal{C}$ has a k-gon face which is adjacent to only quadrilaterals, then $\mathcal{C}$ is the k-prism. This immediately implies that the only crushtacean

with six quadrilateral faces and no pentagons is a cube. Similarly, if $\mathcal{C}$ is an indecomposable crushtacean with two pentagonal faces and five quadrilaterals, then $\mathcal{C}$ is a 5-prism unless the pentagonal faces are adjacent. In the latter case, we note that the union of the pentagonal faces has eight vertices, and by the claim above and indecomposability, the three "free" edges emanating from one of them have distinct vertices. Hence $\mathcal{C}$ has at least eleven vertices, a contradiction. Therefore the 5-prism is the only indecomposable crushtacean with five quadrilateral faces and two pentagons. □

Lemma 7.8. *If $\mathcal{C}$ is a decomposable crushtacean with at most ten vertices, a maximal sequence of decompositions yields a disjoint union of up to four tetrahedra or of a single tetrahedron and a single cube.*

Proof. Suppose $\mathcal{C}$ is a decomposable crushtacean, and let $\mathcal{C}_0$ and $\mathcal{C}_1$ be obtained by decomposing $\mathcal{C}$ along a prismatic 3-cycle. If v, v_0, and v_1 are the numbers of vertices of $\mathcal{C}$, $\mathcal{C}_0$, and $\mathcal{C}_1$, respectively, then from the description of decomposition one finds that

$$v + 2 = v_0 + v_1.$$

It is easy to see that each crushtacean has at least four vertices, and that the tetrahedron is the unique such with exactly four. Thus by the equation above, any crushtacean with six vertices decomposes into two tetrahedra. (By the classification of indecomposable crushtaceans, every crushtacean with six vertices is decomposable.) If $\mathcal{C}$ is a decomposable crushtacean with eight vertices, we thus find that a sequence of two decompositions yields a disjoint union of three tetrahedra.

Finally, suppose that $\mathcal{C}$ is a decomposable crushtacean with ten vertices, and decompose it along a prismatic 3-cycle into crushtaceans $\mathcal{C}_0$ and $\mathcal{C}_1$ with $v_0 \leq v_1$ vertices, respectively. Then either $v_0 = v_1 = 6$ or $v_0 = 4$ and $v_1 = 8$. In the former case, the above implies that neither $\mathcal{C}_0$ nor $\mathcal{C}_1$ is indecomposable; hence each decomposes into a disjoint union of two tetrahedra. In the case $v_0 = 4$ and $v_1 = 8$, $\mathcal{C}_0$ is a tetrahedron. If $\mathcal{C}_1$ is indecomposable, it is a cube; otherwise, a sequence of two decompositions cuts it into a disjoint union of three tetrahedra. □

The scissors congruence classification of augmented links with up to five twist regions is now readily obtained. Below let L be an augmented link.

- If the crushtacean of L decomposes into a disjoint union of tetrahedra, then $S^3 - L$ is a union of right-angled ideal octahedra. It thus follows from Lemma 7.6 and the results of [20] that $\pi_1(S^3 - L) < \mathrm{PGL}_2(\mathcal{O}_1)$. This holds in particular for all augmented links with at most three twist regions, or for any with four twist regions and a decomposable crushtacean.
- If L has four twist regions and an indecomposable crushtacean, then $S^3 - L$ is a union of two right-angled ideal cuboctahedra, and by Corollary 7.5 and the results of [20], $\pi_1(S^3 - L) < \mathrm{PGL}_2(\mathcal{O}_2)$.

In particular, the commensurability class of an augmented link with at most four twist regions is determined by its crushtacean, and each such link falls into one of two commensurability classes. The augmented links with five twist regions display more variability.

- If L has five twist regions and an indecomposable crushtacean $\mathcal{C}$, then $\mathcal{C}$ is the 5-prism. In most cases, we have $\pi_1(S^3 - L) < \Gamma_{\mathcal{P}} \rtimes \mathrm{Sym}(\mathcal{P})$, where $\mathcal{P}$ is the associated right-angled polyhedron, the double of the 5-antiprism across one of its pentagon faces. This holds by Lemma 7.4, unless L has a twist region with an odd number of crossings that corresponds to an edge of a pentagon face of $\mathcal{C}$.
- If L has five twist regions and a decomposable crushtacean that does not decompose into tetrahedra, then $S^3 - L$ is a union of two right-angled octahedra and two cuboctahedra. Two such links are pictured in Figure 8. Using the techniques of [11, §4.3], one can show that the horizontal component that runs across all vertices of the crushtacean on the right-hand side has cusp parameter that is $\mathrm{PGL}_2(\mathbb{Q})$-inequivalent to the parameters of all cusps of the left-hand link. Hence their complements are incommensurable.

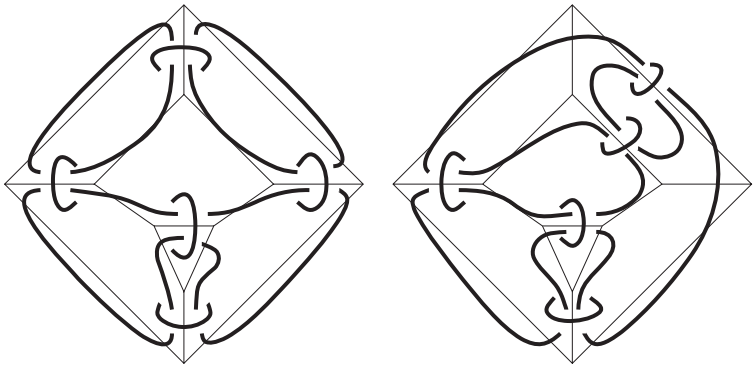

Figure 8. Augmented links with 5 twist regions and a decomposable crushtacean.

From the classification above, we find that an augmented link with at most five twist regions is almost determined up to commensurability by its crushtacean. This is primarily because the indecomposable crushtaceans with at most ten vertices have so much symmetry. Already among those with twelve vertices, we find an example with less symmetry. This is pictured on the left-hand side of Figure 9. On the right-hand side is an augmented link that has this polyhedron as a crushtacean.

Lemma 7.9. *The indecomposable crushtaceans with twelve vertices are the* 6*-prism and the polyhedron on the left-hand side of Figure* 9*.*

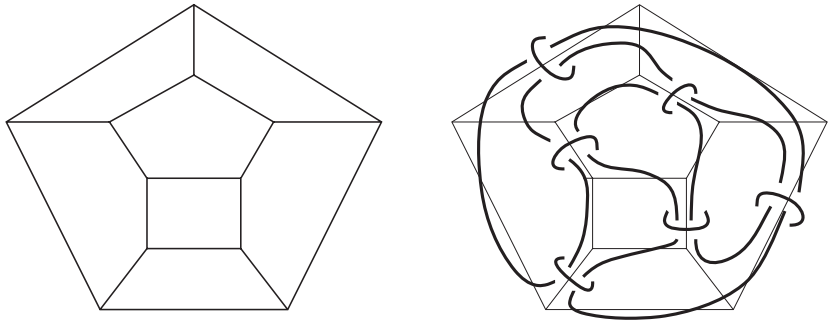

Figure 9. An indecomposable crushtacean with 12 vertices, and an augmented link built on it.

Proof. Reasoning as in the proof of Lemma 7.7, we find that a crushtacean with twelve vertices and a face which is a k-gon for $k > 6$ is decomposable, and that such a crushtacean with a hexagonal face is the 6-prism. Thus as in the proof of that lemma, we are left to consider crushtaceans with all quadrilateral and pentagon faces. If j is the number of quadrilateral and k the number of pentagonal faces, an Euler characteristic calculation again yields $j + k/2 = 6$. Counting vertices in this case yields $4j + 5k = 36$, and solving these two equations yields $j = 4$ and $k = 4$.

Let $\mathcal{C}$ be an indecomposable crushtacean with twelve vertices and 4 each of quadrilateral and pentagon faces. Then every pentagon face of $\mathcal{C}$ is adjacent to at least one other pentagon face.

Claim. *No vertex of $\mathcal{C}$ is shared by three pentagon faces.*

Proof of claim. Suppose v is a vertex with this property, and let v_0, v_1, and v_2 be the vertices adjacent to v in the one-skeleton of $\mathcal{C}$. Then for $i \in \{0, 1, 2\}$, let f_i be the face of $\mathcal{C}$ which contains v_i but not v. We may assume without loss of generality that f_0 and f_1 are quadrilaterals (at least two must be).

Consider the subcomplex of $\partial\mathcal{C}$ which is the union of f_0, f_1, and the pentagon faces containing v. If any edges on the boundary of this subcomplex were identified in $\partial\mathcal{C}$, then it would have a prismatic k-cycle for $k \leq 3$; hence this subcomplex is a disk embedded in $\partial\mathcal{C}$. It contains all twelve vertices, and sixteen out of the eighteen edges of $\mathcal{C}$. But it is easy to see that any way of joining the four "free" vertices by two edges in the complement yields a triangular face, contradicting indecomposability. □

One may also rule out the possibility of a quadrilateral face which meets only pentagonal faces – the union of these faces would be an embedded disk containing all twelve vertices but only fourteen edges – and to establish that each pentagonal face meets at least two other pentagonal faces. Thus the pentagonal faces form a prismatic

4-cycle of $\mathcal{C}$, neither of whose complementary regions can be occupied by a single quadrilateral. It follows that $\mathcal{C}$ is as pictured in Figure 9. □

7.2. Löbell links. For $n \geq 3$, we will denote by $\mathcal{L}(n)$ the nth Löbell polyhedron. This is the unique polyhedron with vertices of valence 3 and faces consisting of n-gons F and F', and $2n$ pentagons, such that F has distance 3 from F' in the dual graph. The Löbell polyhedron $\mathcal{L}(4)$ is pictured on the left-hand side of Figure 10, under a link that has it as a crushtacean. We denote this link $L(4)$. There is an evident rotational symmetry of $(S^3, L(4))$, with order 4 and quotient the link on the right-hand side of Figure 10. An additional component, the fixed axis of this rotation,

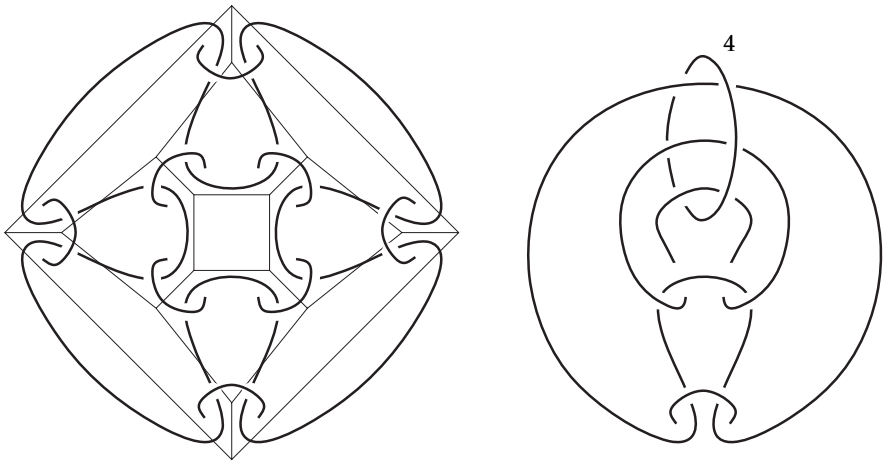

Figure 10. The Löbell link $L(4)$ and its 4-fold cyclic quotient.

has been added to the diagram and labeled with 4. For arbitrary $n \geq 3$, we define $L(n)$ to be the link with crushtacean $\mathcal{L}(n)$ that n-fold branched covers the diagram on the right-hand side. The main result of this section is:

Theorem 7.10. *For all but finitely many $n \geq 4$, $M(n) \doteq S^3 - L(n)$ is not arithmetic nor commensurable with any 3-dimensional hyperbolic reflection orbifold. Moreover, at most finitely many $M(n)$ occupy any commensurability class.*

Remark. Since $\mathcal{L}(5)$ is the dodecahedron, $L(5)$ falls under the purview of Corollary 7.5 and so is commensurable with a right-angled reflection orbifold. Therefore the stipulation "all but finitely many" above is necessary. We do not know of any $M(n)$ that is arithmetic, however. We note also that $\mathcal{L}(3)$ decomposes into two tetrahedra and a cube, whereas $\mathcal{L}(n)$ is indecomposable for $n > 3$.

Proving the theorem requires identifying the commensurator quotient of $M(n)$. We begin by identifying the symmetry group of $\mathcal{L}(n)$.

Fact. *For $n \neq 5$, the symmetry group of $\mathcal{L}(n)$ has presentation*

$$\Sigma(n) = \langle\, \mathsf{a}, \mathsf{b}_n, \mathsf{s} \mid (\mathsf{b}_n)^n = \mathsf{s}^2 = \mathsf{a}^2 = 1, \mathsf{s}\mathsf{b}_n\mathsf{s} = (\mathsf{b}_n)^{-1}, \mathsf{a}\mathsf{b}_n\mathsf{a} = (\mathsf{b}_n)^{-1}, \mathsf{a}\mathsf{s}\mathsf{a} = \mathsf{b}_n\mathsf{s} \,\rangle.$$

The subgroup $\langle \mathsf{a}, \mathsf{b}_n \rangle$ preserves orientation, and s reverses it. The subgroup $\langle \mathsf{b}_n, \mathsf{s} \rangle$ preserves each n-gon face, and a exchanges them.

Proof. Since $n \neq 5$, $\mathcal{L}(n)$ has exactly two n-gon faces F and F'. Let $e_0, e_1, \ldots, e_{n-1}$ be a cyclic ordering of the edges of F; ie, for each i, e_i shares a vertex with e_{i+1}, where $i+1$ is taken modulo n. The union of F with the pentagonal faces of $\mathcal{L}(n)$ that abut it is a disk D embedded in $\partial\mathcal{L}(n)$, with boundary consisting of $2n$ edges that can be cyclically ordered $f_1, f_2, \ldots, f_{2n}$ as follows: for $0 \leq i < n$, let F_i be the pentagonal face of $\mathcal{L}(n)$ containing e_i and let $f_{2i+1} \subset F_i \cap \partial D$ and $f_{2(i+1)} \subset F_{i+1} \cap \partial D$ be the unique pair of edges that share a vertex (with $i+1$ taken modulo n).

We now let b_n be the rotational symmetry of F taking e_i to e_{i+1} for each i, and take s to be the reflection of F preserving e_0 and exchanging e_i with e_{n-i} for $0 < i < n$. It is easy to see that these extend to a rotation and reflection of $\mathcal{L}(n)$, respectively, yielding the subgroup $\langle \mathsf{b}_n, \mathsf{s} \rangle$ described above (we refer to the extensions by the same name).

There is a symmetry a of the embedded circle $f_1 \cup f_2 \cup \cdots \cup f_{2n}$ that preserves f_1 and f_{n+1}, exchanging endpoints of each, and exchanges f_i with f_{2n+2-i} for $1 < i \leq n$. This extends to a rotational symmetry of $\mathcal{L}(n)$ taking F to F'. In particular, for $0 \leq i < n$, we can take F_i' to be the pentagonal face adjacent to F' that contains f_{2i+1} and $f_{2(i+1)}$. Then a takes F_i to F'_{n-i}.

The relations on b_n, s, and a follow by considering their actions on F. Since every automorphism of $\mathcal{L}$ either exchanges F and F' or preserves each, there is a map to $\mathbb{Z}/2\mathbb{Z} = \{\pm 1\}$ taking such an element to -1 or 1, respectively. The subgroup $\langle \mathsf{b}_n, \mathsf{s} \rangle$ is contained in the kernel of this map; since it is the entire symmetry group of F, it is the entire kernel. Hence the entire symmetry group of $\mathcal{L}(n)$ is generated by $\langle \mathsf{b}_n, \mathsf{s} \rangle$ and a, which maps to -1. □

A fundamental domain for the action on $\mathcal{L}(n)$ of the cyclic group $\langle \mathsf{b}_n \rangle$ is depicted on the left-hand side of Figure 11, cut out by the dotted line segments. These should be interpreted as meeting at the point at infinity, in addition to the center of F. The segment that runs through the edge joining endpoints of e_0 and f_{2n} is fixed by the reflection $\mathsf{s}\mathsf{b}_n$, and the other is fixed by $\mathsf{b}_n\mathsf{s}$.

Recall that by Lemma 7.3, each symmetry of $\mathcal{L}(n)$ determines a symmetry of the right-angled ideal polyhedron $\mathcal{P}(n)$ obtained by truncating vertices of $\mathcal{L}(n)$. In particular, $\mathsf{s}\mathsf{b}_n$ and $\mathsf{b}_n\mathsf{s}$ determine reflective symmetries of $\mathcal{P}(n)$. Cutting along the

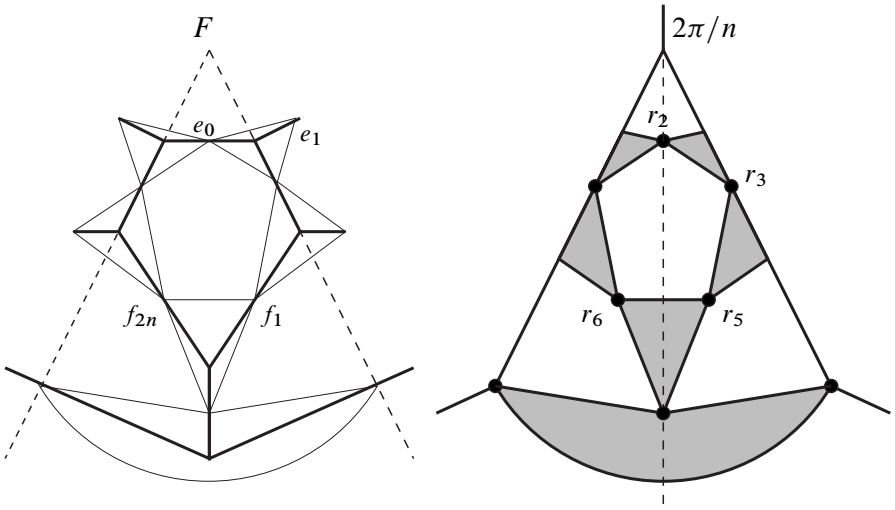

Figure 11. A fundamental domain for the action of $\langle \mathsf{b}_n \rangle$ on $\mathcal{L}(n)$, and the corresponding sub-polyhedron $\mathcal{O}(n)$ of $\mathcal{P}(n)$.

mirrors of these reflections yields the polyhedron $\mathcal{O}(n)$ pictured on the right-hand side of the figure. The three edges with "free" ends should again be interpreted as meeting at the point at infinity. The darkened vertices of $\mathcal{O}(n)$ are ideal; the remaining vertices, each the midpoint of an edge of $\mathcal{P}(n)$, are not.

The intersection of the mirror of s with $\partial\mathcal{O}(n)$ is the dotted axis on the right-hand side of Figure 11. Clearly, s restricts to an isometry of $\mathcal{O}(n)$. Although a does not preserve $\mathcal{O}(n)$, it does preserve the sub-polyhedron, obtained by cutting along the mirror of s, that contains the ideal vertex labeled r_5. Indeed, it acts on this polyhedron as a 180-degree rotation fixing r_5 and the midpoint of the edge labeled $2\pi/n$, exchanging each of r_3 and r_2 with an unlabeled ideal vertex.

Since $\mathcal{P}(n)$ is right-angled, each edge of $\mathcal{O}(n)$ that is contained in one of $\mathcal{P}(n)$ has dihedral angle $\pi/2$. Since the mirrors of sb_n and $\mathsf{b}_n\mathsf{s}$ meet each edge of $\mathcal{P}(n)$ transversely, each edge of $\mathcal{O}(n)$ that is the intersection of $\partial\mathcal{P}(n)$ with a mirror of one of these reflections has dihedral angle $\pi/2$ as well. Thus the only edge of $\mathcal{O}(n)$ with a dihedral angle different than $\pi/2$ is the intersection of the mirrors of sb_n and $\mathsf{b}_n\mathsf{s}$, labeled $2\pi/n$ at the top of the figure. That this is the dihedral angle follows from the fact that the product of these reflections is the rotation $(\mathsf{b}_n)^2$, through an angle of $2 \cdot 2\pi/n$.

Each symmetry of $\mathcal{L}(n)$, $n \neq 5$, exchanges edges enclosed by clasps of $L(n)$; hence the corresponding isometry of $\mathcal{P}(n)$ induces one of $M(n) = S^3 - L(n)$. Since $\mathcal{O}(n)$ is a fundamental domain for the action of the rotation group $\langle \mathsf{b}_n \rangle$ on $\mathcal{L}(n)$, Lemma 7.2 implies $\mathcal{O}(n) \cup \overline{\mathcal{O}}(n)$ is a fundamental domain for the action on $\mathbb{H}^3$ of

the orbifold fundamental group of $O(n) = M(n)/\langle \mathsf{b}_n \rangle$. Here $\bar{\mathcal{O}}(n) \doteq \mathsf{d}_1(\mathcal{O}(n))$, where d_1 is the reflection through the white face of $\mathcal{O}(n)$ whose sole ideal vertex is r_2. Using the further symmetries a and s of $\mathcal{P}(n)$, we thus obtain the lemma below.

Lemma 7.11. *Let d_2 be the reflection through the white face of $\mathcal{O}(n)$ with ideal vertices r_2, r_3, $\mathsf{s}(r_3)$, r_5, r_6, and let c be the parabolic isometry fixing r_3 and taking r_2 to r_5. Then $O(n)$ is isometric to $\mathbb{H}^3/\Gamma(n)$, where*

$$\Gamma(n) = \langle\, \mathsf{d}_1\mathsf{d}_2, \mathsf{d}_1\mathsf{d}_2^{\mathsf{a}}, \mathsf{d}_1\mathsf{d}_2^{\mathsf{sa}}, \mathsf{d}_1\mathsf{d}_1^{\mathsf{a}}, \mathsf{b}_n, \mathsf{c}, \mathsf{c}^{\mathsf{a}}, \mathsf{c}^{\mathsf{s}}, \mathsf{b}_n^{\mathsf{d}_1}, \mathsf{c}^{\mathsf{d}_1}, \mathsf{c}^{\mathsf{d}_1\mathsf{a}}, \mathsf{c}^{\mathsf{d}_1\mathsf{s}} \,\rangle.$$

Furthermore, the isometry of $O(n)$ visible on the right-hand side of Figure 10 as reflection through the projection plane is induced by d_1.

Let L be the link in S^3 that is the union of the fixed locus of $O(n)$ with the other components pictured on the right-hand side of Figure 10. Then $O(n)$ is obtained from $S^3 - L$ by $(n, 0)$-Dehn filling on the added component, where the meridian here is chosen to lie in the projection plane and the longitude bounds a 3-punctured disk. Because the singular locus of $O(n)$ is the image of the edge e of $\mathcal{O}(n)$ with dihedral angle $2\pi/n$, $S^3 - L$ is obtained from $\mathcal{O}(n) - e$ by the restriction of the face pairings described in Lemma 7.11. Thus Poincaré's polyhedron theorem implies:

Lemma 7.12. *Let $\mathcal{O}$ be the all-right polyhedron in $\mathbb{H}^3$ homeomorphic to $O(n)-e$, and let a, b, c, d_1 and d_2 have the same combinatorial descriptions as the correspondingly-named isometries determined by $\mathcal{O}(n)$. Let*

$$\Gamma_L = \langle \mathsf{d}_1\mathsf{d}_2, \mathsf{d}_1\mathsf{d}_2^{\mathsf{a}}, \mathsf{d}_1\mathsf{d}_2^{\mathsf{sa}}, \mathsf{d}_1\mathsf{d}_1^{\mathsf{a}}, \mathsf{b}, \mathsf{c}, \mathsf{c}^{\mathsf{a}}, \mathsf{c}^{\mathsf{s}}, \mathsf{b}^{\mathsf{d}_1}, \mathsf{c}^{\mathsf{d}_1}, \mathsf{c}^{\mathsf{d}_1\mathsf{a}}, \mathsf{c}^{\mathsf{d}_1\mathsf{s}} \rangle.$$

Then $S^3 - L$ is homeomorphic to $\mathbb{H}^3/\Gamma_L$.

The only aspect of this lemma that requires comment is that Andreev's theorem implies that there is a right-angled polyhedron $\mathcal{O}$ with the requisite combinatorial description. An ideal vertex of $\mathcal{O}$ replaces the edge of $\mathcal{O}(n)$ with dihedral angle $2\pi/n$. Thus b is parabolic, rather than elliptic like b_n.

Denote by r_7 the ideal vertex of $\mathcal{O}$ fixed by b; that is, r_7 replaces the edge of $\mathcal{O}(n)$ with dihedral angle $2\pi/n$. The polyhedron obtained by cutting along the mirror of s, that has r_5 as an ideal vertex, has 180-degree rotational symmetry a fixing r_5 and r_7. Therefore a single geodesic plane contains the ideal vertices r_2, r_5, r_7, and $\mathsf{a}(r_2)$. Let $\mathcal{Q}_0$ be the polyhedron with r_3 as an ideal vertex that is obtained by cutting along this plane.

An ideal polyhedron $\mathcal{Q}$ may be obtained from $\mathcal{Q}_0$ as follows. The geodesic plane $[r_5, r_3, \mathsf{a}(r_2)]$ containing r_5, r_3, and $\mathsf{a}(r_2)$ cuts off a tetrahedron $\mathcal{T}$, with a finite vertex opposite this plane, from the remainder of $\mathcal{Q}_0$. Let $\mathcal{Q} = (\overline{\mathcal{Q}_0 - \mathcal{T}}) \cup \mathsf{c}^{-1}(\mathcal{T})$. Since all edges abutting each finite vertex of $\mathcal{Q}_0$ have dihedral angle $\pi/2$, the finite vertices

of $\mathcal{Q}_0$, which are identified in $\mathcal{Q}$, lie in the interior of an edge of $\mathcal{Q}$. We have depicted $\mathcal{Q}_0$ and $\mathcal{Q}$ on the left- and right-hand sides of Figure 12, respectively, coloring black the face of $\mathcal{Q}$ in $[r_5, r_3, \mathsf{a}(r_2)]$ and its image under c^{-1}.

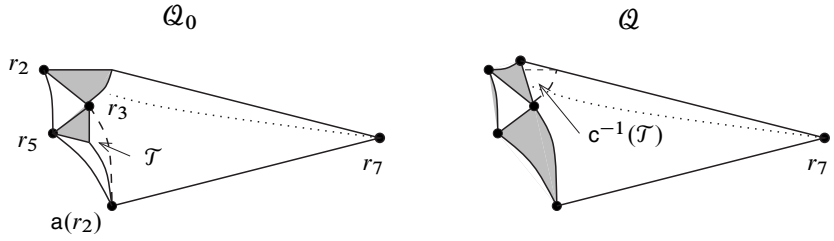

Figure 12. $\mathcal{Q}_0$ and $\mathcal{Q}$.

The lemma below follows from Poincaré's polyhedron theorem and the descriptions from Lemma 7.12 of face pairing isometries on $\mathcal{O} \cup \mathsf{d}_1(\mathcal{O})$ yielding $S^3 - L$.

Lemma 7.13. *Let* $\Gamma = \langle \mathsf{a}, \mathsf{c}, \mathsf{d}_1, \mathsf{d}_2, \mathsf{d}_3 \doteq \mathsf{asa} \rangle$ *be generated by face pairings for* $\mathcal{Q}$. *Then* $\mathbb{H}^3/\Gamma$ *is a three-cusped hyperbolic* 3*-orbifold, and* $\Gamma_L \lhd \Gamma$ *with index* 8.

The isometry d_3 defined in Lemma 7.13 acts as reflection in the face of $\mathcal{Q}$ containing r_7, $\mathsf{a}(r_2)$, r_3, and $\mathsf{c}^{-1}\mathsf{a}(r_2)$, since a takes this face into the mirror of s. That the other generators act as face pairings follows from previous observations. The index computation uses the fact that $\mathcal{O}$ is the union of 4 isometric copies of $\mathcal{Q}$; namely, $\mathcal{O} = \mathcal{Q} \cup \mathsf{a}(\mathcal{Q}) \cup \mathsf{s}(\mathcal{Q} \cup \mathsf{a}(\mathcal{Q}))$. In verifying that each generator for Γ_L lies in Γ, it is helpful to note that $\mathsf{b} = \mathsf{d}_1\mathsf{s} \in \Gamma$.

The key result in the proof of Theorem 7.10 is the proposition below.

Proposition 7.14. Γ *is its own commensurator.*

We defer the proof of Proposition 7.14 for now, and first apply it.

Proof of Theorem 7.10. Since the orbifold fundamental group $\Gamma(n)$ of $O(n)$ contains the elliptic element b_n, with order n, its invariant trace field $k\Gamma(n)$ contains the trace of $(\mathsf{b}_n)^2$ and thus $\mathbb{Q}(\cos(2\pi i \frac{2}{n}))$ (cf. [25, §3.3] for the definition and properties of the invariant trace field). This is a degree-two subfield of the cyclotomic field $\mathbb{Q}(\zeta_k)$, where $k = n$ if n is odd and $k = n/2$ otherwise. Thus $\liminf_{n\to\infty}[k\Gamma(n) : \mathbb{Q}]$ is infinite. It follows that at most finitely many $O(n)$ belong to any one commensurability class. Furthermore, at most finitely many are arithmetic, since non-compact arithmetic hyperbolic 3-manifolds have quadratic invariant trace fields.

Throwing away the arithmetic $\Gamma(n)$, Margulis' theorem implies that $\mathrm{Comm}(\Gamma(n))$ is a finite extension of $\Gamma(n)$ for the remaining n. We remarked above Lemma 7.11 that each symmetry of $\mathcal{L}(n)$ determines an isometry of $M(n) = S^3 - L(n)$. In particular,

there are isometries determined by a and s, and since $\langle \mathsf{b}_n \rangle \lhd \langle \mathsf{a}, \mathsf{b}_n, \mathsf{s} \rangle$, these generate a group of isometries of $O(n) = M(n)/\langle \mathsf{b}_n \rangle$ with order 4. By Lemma 7.11, d_1 determines an additional isometry of $O(n)$, that can easily be seen to commute with $\langle \mathsf{a}, \mathsf{s} \rangle$. Thus $\mathrm{Comm}(\Gamma(n))$ contains the degree-8 extension $\langle \Gamma(n), \mathsf{a}, \mathsf{s}, \mathsf{d}_1 \rangle$ of $\Gamma(n)$.

As the right-hand side of Figure 10 makes clear, $O(n)$ is obtained from $S^3 - L$ by $(n, 0)$-Dehn filling on a fixed component. Therefore the hyperbolic Dehn surgery theorem implies that the $O(n)$ converge geometrically to the hyperbolic structure on $S^3 - L$, and in particular, their volumes approach its from below. (See e.g. [5, §E.5] for background on the hyperbolic Dehn surgery theorem.) Furthermore, the explicit descriptions above imply that the $\Gamma(n)$ converge algebraically to Γ_L, and the $\langle \Gamma(n), \mathsf{a}, \mathsf{s}, \mathsf{d}_1 \rangle$ to Γ.

If on an infinite subsequence, $\langle \Gamma(n), \mathsf{a}, \mathsf{s}, \mathsf{d}_1 \rangle$ were contained in $\mathrm{Comm}(\Gamma(n))$ properly, then a further subsequence of the $\mathrm{Comm}(\Gamma(n))$ would converge to a discrete group Γ_0 with covolume a proper fraction of that of Γ. This follows from the fact that the Chabauty topology on discrete subgroups of $\mathrm{PSL}_2(\mathbb{C})$ with bounded covolume is compact, see e.g. [5, Corollary E.1.7]. In this case, since $\langle \Gamma(n), \mathsf{a}, \mathsf{s}, \mathsf{d}_1 \rangle \to \Gamma$ and limits are unique in this topology (see e.g. [5, Lemma E.1.1]), we would have $\Gamma < \Gamma_0$ properly, contradicting Proposition 7.14. Thus for all but finitely many n, $\mathrm{Comm}(\Gamma(n)) = \langle \Gamma(n), \mathsf{a}, \mathsf{s}, \mathsf{d}_1 \rangle$.

Fixing a horosphere $\mathcal{H}$ centered at the ideal vertex r_3 of $\mathcal{O}(n)$, a fundamental domain for the action on $\mathcal{H}$ of its stabilizer in $\mathrm{Comm}(\Gamma(n))$ is thus the rectangle $\mathcal{O}(n) \cap \mathcal{H}$. Two parallel sides of this rectangle are given by the intersection of $\mathcal{H}$ with the white sides of $\mathcal{O}(n)$ containing r_3. One of these, contained in the side with ideal vertices $r_2, r_3, \mathsf{s}(r_3), r_5$, and r_6, is stabilized by the reflection $\mathsf{d}_2 \in \Gamma(n)$ defined in Lemma 7.11. The other is stabilized by the reflection $\mathsf{d}_3 \doteq \mathsf{asa} \in \mathrm{Comm}(\Gamma(n))$, defined in analogy with the identically-named element of Γ from Lemma 7.13. For the other pair of parallel sides of this rectangle, the parabolic c fixing r_3 acts by translation taking one to the other.

The stabilizer of $\mathcal{H}$ in $\mathrm{Comm}(\Gamma(n))$ is thus $\langle \mathsf{c}, \mathsf{d}_2, \mathsf{d}_3 \rangle$. If Γ' were a reflection group commensurable with $\Gamma(n)$, then $\mathrm{Stab}_{\Gamma'}(r_3)$ would be a reflection group contained in $\langle \mathsf{c}, \mathsf{d}_2, \mathsf{d}_3 \rangle$, acting on $\mathcal{H}$ with finite coarea. But since c translates parallel to the lines fixed by d_2 and d_3, every reflection in $\langle \mathsf{c}, \mathsf{d}_2, \mathsf{d}_3 \rangle$ fixes a line parallel to the lines fixed by d_2 and d_3. Hence no reflection subgroup of $\langle \mathsf{c}, \mathsf{d}_2, \mathsf{d}_3 \rangle$ acts on $\mathcal{H}$ with finite coarea. Therefore $\mathrm{Comm}(\Gamma(n))$ is not commensurable with a reflection group. □

Proving Proposition 7.14 requires an explicit description of Γ. This will follow from the lemma below, which describes an embedding of $\mathcal{Q}$ in the upper half-space model for $\mathbb{H}^3$.

Lemma 7.15. *There is an isometric embedding of $\mathcal{Q}$ in $\mathbb{H}^3$ determined by the following ideal vertices: $r_2 = -1 + i$, $r_3 = 0$, $r_5 = (\sqrt{3} + i)/2$, $r_7 = \infty$.*

Proof. Our description of $\mathcal{Q}$ includes the following facts: its edge joining the ideal vertex r_7 to $\mathsf{c}^{-1}\mathsf{a}(r_2)$ has a dihedral angle of $\pi/2$, and there are two quadrilateral faces with ideal vertices $r_7, \mathsf{a}(r_2), r_5, r_2$ and $r_7, \mathsf{a}(r_2), r_3, \mathsf{c}^{-1}\mathsf{a}(r_2)$, respectively. We will choose an embedding of $\mathcal{Q}$ that sends the latter face into the geodesic plane of $\mathbb{H}^3$ with ideal boundary $\mathbb{R} \cup \{\infty\}$, taking r_3 to 0 and r_7 to ∞ in particular.

We have pictured such an embedding in Figure 13. The ideal vertices $\mathsf{c}^{-1}\mathsf{a}(r_2)$ and $\mathsf{a}(r_2)$ go to points x and z, respectively, in $\mathbb{R}$ on either side of $r_3 = 0$. We take $x < 0$ and $z > 0$. Since the edge joining r_7 to $\mathsf{c}^{-1}\mathsf{a}(r_2)$ has dihedral angle $\pi/2$, the image of r_2 is of the form $x + iy$ for some $y \in \mathbb{R}$. We may assume $y > 0$, by reflecting through $\mathbb{R}$ if necessary. The final ideal vertex r_5 lies somewhere on the line segment joining r_2 with $\mathsf{a}(r_2)$, since it is in the ideal boundary of a plane containing r_2, r_7, and $\mathsf{a}(r_2)$. Its coordinates are determined by the fact that a preserves this plane, fixing r_5 and r_7.

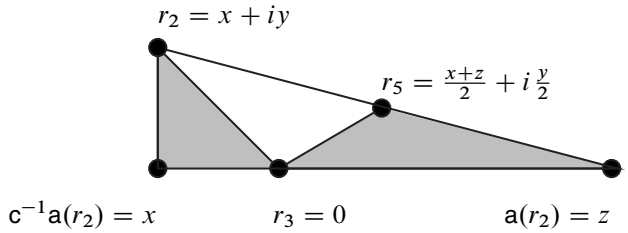

Figure 13. An embedding of $\mathcal{Q}$ in $\mathbb{H}^3$.

We have darkened the triangles in $\mathbb{C}$ that lie under the dark faces of $\mathcal{Q}$ after the embedding described above. The parabolic isometry c takes one to the other, fixing r_3, thus it is of the form $\left(\begin{smallmatrix} 1 & 0 \\ w & 1 \end{smallmatrix}\right)$ for some $w \in \mathbb{C}$. Using the fact that c takes $\mathsf{c}^{-1}\mathsf{a}(r_2) = x$ to $\mathsf{a}(r_2) = z$, a computation implies $w = (x - z)/xz$. Another computation, using the fact that $\mathsf{c}(r_2) = r_5$, determines $z = -x(\sqrt{3} + 1)$.

We are free to choose $x < 0$, since one choice may be changed to another by applying a hyperbolic isometry fixing 0 and ∞ to $\mathcal{Q}$. Choosing $x = -1$ yields

$$\mathsf{c}^{-1}\mathsf{a}(r_2) = x = -1, \qquad r_2 = x + iy = -1 + i,$$
$$\mathsf{a}(r_2) = z = \sqrt{3} + 1, \quad r_5 = \frac{\sqrt{3} + i}{2}.$$

This is the embedding described in the statement. □

A few additional parabolic fixed points that will be useful below we name as follows: let $r_1 = \mathsf{d}_3(r_2)$, $r_4 = \mathsf{d}_3(r_5)$, and $r_6 = \mathsf{d}_3^{\mathsf{a}}(r_5)$. Note that r_2, r_4, r_5, and r_6 are each Γ-equivalent to r_1.

In proving Proposition 7.14, it will be convenient to use a different embedding of $\mathcal{Q}$ than that described in Lemma 7.15 above. Let us apply a Möbius transformation

taking r_1, r_2, and r_3 to 0, 1, and ∞, respectively. Such a map is given by $z \mapsto \frac{1+i}{2} + i/z$. This takes the other ideal vertices to

$$r_4 = i\frac{1+\sqrt{3}}{2}, \quad r_5 = 1 + i\frac{1+\sqrt{3}}{2}, \quad r_6 = 1 + i\frac{3+\sqrt{3}}{6}, \quad r_7 = \frac{1+i}{2}.$$

The representation of Γ determined by the embedding described above is related to that determined by the embedding of Lemma 7.15 by conjugation by

$$\begin{pmatrix} -i\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2}(1-i) \\ -\frac{\sqrt{2}}{2}(1+i) & 0 \end{pmatrix}.$$

Since $\mathcal{Q}$ is a fundamental domain for Γ, each cusp of $\mathbb{H}^3/\Gamma$ corresponds to a point on $\partial\mathbb{H}^3$ that is Γ-equivalent to an ideal vertex of $\mathcal{Q}$. Inspection of the face pairings of Lemma 7.13 thus reveals that $\mathbb{H}^3/\Gamma$ has exactly three cusps. We let c_1 correspond to the points of $\Gamma \cdot r_1$, c_2 to $\Gamma \cdot r_7$, and c_3 to $\Gamma \cdot r_3$.

Our explicit description of $\mathcal{Q}$ allows computation of the invariant trace field and cusp parameters. This implies:

Lemma 7.16. *Γ is non-arithmetic. The cusps c_1 and c_2 are commensurable to each other and are not commensurable to c_3.*

Proof. An explicit description of generators for Γ, as may be obtained from Lemma 7.15, enables direct computation of the invariant trace field (see [25, §3.5]). Performing this calculation, we find that Γ has trace field $\mathbb{Q}(i, \sqrt{3})$. Alternatively, the link L may be entered into the computer program Snappea, and the resulting triangulation data into Snap, yielding the same description (see [12]). Since every non-compact arithmetic hyperbolic 3-manifold has an imaginary quadratic invariant trace field, Γ is not arithmetic.

Using the embedding described in Lemma 7.15, we find that an index-8 subgroup of $\mathrm{Stab}_\Gamma(\infty)$ is generated by $z \mapsto z + 2(2+\sqrt{3})$ and $z \mapsto z + 2i$; thus the parameter of the associated cusp c_2 is $\mathrm{PGL}_2(\mathbb{Q})$-equivalent to $i(2+\sqrt{3})$ (cf. [11, §4.3]). After re-embedding as above, the stabilizer of ∞ corresponds to the cusp c_3. An index-2 subgroup of this lattice is generated by $\mathsf{c}\colon z \mapsto z + i\frac{1+\sqrt{3}}{2}$ and the product of reflections $\mathsf{d_2d_3}\colon z \mapsto z + 1$. Thus the parameter of c_3 is $\mathrm{PGL}_2(\mathbb{Q})$-equivalent to $i(1+\sqrt{3})$. A similar computation reveals that c_1 has the same parameter as c_2. Since the complex modulus is a complete commensurability invariant for lattices in $\mathbb{C}^2$, and $i(1+\sqrt{3})$ is not $\mathrm{PGL}_2(\mathbb{Q})$-equivalent to $i(2+\sqrt{3})$, the lemma follows. □

From Margulis' theorem, we immediately obtain:

Corollary 7.17. *$\mathrm{Comm}(\Gamma)$ is a finite extension of Γ, and the minimal orbifold $O \doteq \mathbb{H}^3/\mathrm{Comm}(\Gamma)$ has either two or three cusps.*

In particular, if O has two cusps then c_1 and c_2 are identified by the covering map $\mathbb{H}^3/\Gamma \to O$. We have used the algorithm of Goodman–Heard–Hodgson [16] to compute $\mathrm{Comm}(\Gamma)$. Recall that we introduced the setting for this algorithm in Section 6.2 between the statements of Propositions 6.6 and 6.7.

Let

$$\boldsymbol{v}_1 = (-2,\, 2,\, -1,\, 3)^\top, \quad \boldsymbol{v}_3 = (0,\, 0,\, -3,\, 3)^\top,$$
$$\boldsymbol{v}_7 = \left(0,\, 0,\, 9-4\sqrt{3},\, 9-4\sqrt{3}\right)^\top.$$

These vectors are chosen so that there is an isometry Φ from the upper half space model to the hyperboloid model which takes the parabolic fixed point r_i to the center of the horosphere $H_{\boldsymbol{v}_i}$ when $i = 1, 3, 7$. Under Φ, the isometries a, b, c, d_1, d_2, d_3 correspond to the matrices $\mathsf{A}, \mathsf{B}, \mathsf{C}, \mathsf{D}_1, \mathsf{D}_2, \mathsf{D}_3 \in \mathrm{O}_0(3,1)$ listed below:

$$\mathsf{A} = \begin{pmatrix} -1 & 0 & -1/2 & 1/2 \\ 0 & -1 & \sqrt{3}/2 & -\sqrt{3}/2 \\ -1/2 & \sqrt{3}/2 & 1/2 & 1/2 \\ -1/2 & \sqrt{3}/2 & -1/2 & 3/2 \end{pmatrix}, \qquad \mathsf{D}_1 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & -1 & 1 \\ 0 & -1 & \frac{1}{2} & \frac{1}{2} \\ 0 & -1 & -\frac{1}{2} & \frac{3}{2} \end{pmatrix},$$

$$\mathsf{B} = \begin{pmatrix} 1 & 0 & -1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1/2 & 1/2 \\ 1 & 0 & -1/2 & 3/2 \end{pmatrix}, \qquad \mathsf{D}_2 = \begin{pmatrix} -1 & 0 & 2 & 2 \\ 0 & 1 & 0 & 0 \\ 2 & 0 & -1 & -2 \\ -2 & 0 & 2 & 3 \end{pmatrix},$$

$$\mathsf{C} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & -1-\sqrt{3} & -1-\sqrt{3} \\ 0 & 1+\sqrt{3} & -1-\sqrt{3} & -2-\sqrt{3} \\ 0 & -1-\sqrt{3} & 2+\sqrt{3} & 3+\sqrt{3} \end{pmatrix}, \qquad \mathsf{D}_3 = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Thus, Φ allows us to also think of Γ as a subgroup of $\mathrm{O}_0(3,1)$. Each $\boldsymbol{v}_i$ is a horospherical vector for the cusp c_i of $\mathbb{H}^3/\Gamma$ so $\{\boldsymbol{v}_1, \boldsymbol{v}_3, \boldsymbol{v}_7\}$ determines a Γ-invariant set V as above. We have $\{\boldsymbol{v}_i\}_1^7$ given by $\boldsymbol{v}_i = \Phi(r_i)$ and these vectors may be calculated explicitly by applying appropriate isometries from Γ. We have that $\boldsymbol{v}_i$ is the i^{th} column of the matrix

$$\begin{pmatrix} -2 & 2 & 0 & -2 & 2 & 6 & 0 \\ 2 & 2 & 0 & -2\sqrt{3} & -2\sqrt{3} & -2\sqrt{3} & 0 \\ -1 & -1 & -3 & -3 & -3 & -1 & 9-4\sqrt{3} \\ 3 & 3 & 3 & 5 & 5 & 7 & 9-4\sqrt{3} \end{pmatrix}.$$

As discussed above, we obtain all possibilities for canonical tilings associated to Γ by using initial sets of the form $\{\boldsymbol{v}_1, \beta\boldsymbol{v}_7, \gamma\boldsymbol{v}_3\}$ where $\beta, \gamma \in \mathbb{R}^+$. We write $\mathcal{H}(\beta, \gamma)$

to denote the set $\Gamma \cdot \{\boldsymbol{v}_1, \beta \boldsymbol{v}_7, \gamma \boldsymbol{v}_3\}$ and $\mathcal{T}(\beta, \gamma)$ to denote the associated canonical tiling.

Recall that $O = \mathbb{H}^3/\mathrm{Comm}(\Gamma)$ has either 2 or 3 cusps. If O has 3 cusps then, for any pair (β, γ), $\mathcal{H}(\beta, \gamma)$ descends to cusp cross sections of O and so $\mathrm{Comm}(\Gamma) = \mathrm{Sym}(\mathcal{T}(\beta, \gamma))$. If O has 2 cusps then there is some $\mathsf{g} \in \mathrm{Comm}(\Gamma)$ and β_0 with $\mathsf{g}(\boldsymbol{v}_1) = \beta_0 \boldsymbol{v}_7$. We have $\mathrm{Comm}(\Gamma) = \mathrm{Sym}(\mathcal{T}(\beta_0, \gamma))$ for any $\gamma \in \mathbb{R}^+$. Therefore, it suffices to compute the triangulations $\mathcal{T}(\beta, 1)$ for $\beta \in \mathbb{R}^+$. Either there exists a unique β_0 so that $\mathrm{Sym}(\mathcal{T}(\beta_0, 1))$ contains an isometry taking $\boldsymbol{v}_1$ to $\beta_0 \boldsymbol{v}_7$ or there is no such β. In the first case, O has 2 cusps and $\mathrm{Comm}(\Gamma) = \mathrm{Sym}(\mathcal{T}(\beta_0, 1))$. In the second case, O has 3 cusps and $\mathrm{Comm}(\Gamma) = \mathrm{Sym}(\mathcal{T}(\beta, 1))$ for every β.

Lemma 7.18. *O has 3 cusps and $\mathrm{Comm}(\Gamma) = \mathrm{Sym}(\mathcal{T}(\beta, 1))$ for every β.*

Proof. The proof follows by showing that there does not exist a unique β so that $\mathrm{Sym}(\mathcal{T}(\beta_0, 1))$ contains an isometry taking $\boldsymbol{v}_1$ to $\beta_0 \boldsymbol{v}_7$. We first describe the canonical triangulations as β decreases from ∞ to 0. The interval $(0, \infty)$ has a finite cell decomposition so that if two values for β are chosen from the same cell then they determine the same canonical triangulation. As β moves to the boundary of a 1-cell there is a pair of neighboring tiles T_1 and T_2 so that the tilt at their common face changes from positive to zero. At the boundary value, these two tiles merge to form a tile in the new canonical triangulation. The decomposition of $(0, \infty)$ and the associated tilings of $\mathbb{H}^3$ are described in Tables 1–3. The triangulations $\mathcal{T}(\beta, 1)$ can be checked by repeatedly verifying the coplanar and positive tilt conditions on sets of Γ-generating tiles. In the tables, we let $[p_1, \ldots, p_k]$ denote the convex hull in $\mathbb{H}^3$ of a collection $\{p_1, \ldots, p_k\} \subset \partial\mathbb{H}^3$.

Table 1. The data that determine the first two canonical tilings.

	β	Γ-generating tiles
$\mathcal{T}_1$	$\beta > \frac{3}{11}(4 + 3\sqrt{3})$	$\mathcal{P}_1 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$
		$\mathcal{P}_2 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$
		$\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{CA}(\boldsymbol{v}_2), \mathsf{CA}(\boldsymbol{v}_3)]$
		$\mathcal{P}_4 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{CA}(v_2)]$
		$\mathcal{P}_5 = [\boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{CA}(v_2), \mathsf{C}(\boldsymbol{v}_7)]$
$\mathcal{T}_2$	$\beta = \frac{3}{11}(4 + 3\sqrt{3}) \sim 2.51$	$\mathcal{P}_1 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$
		$\mathcal{P}_2 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$
		$\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{CA}(\boldsymbol{v}_2), \mathsf{CA}(\boldsymbol{v}_3)]$
		$\mathcal{P}_4 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{CA}(v_2), \mathsf{C}(\boldsymbol{v}_7)]$

Table 2. More canonical tilings.

	β	Γ-generating tiles
$\mathcal{T}_3$	$\frac{1}{22}(21+13\sqrt{3}) < \beta < \frac{3}{11}(4+3\sqrt{3})$	$\mathcal{P}_1 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{C}(\boldsymbol{v}_7)]$ $\mathcal{P}_2 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$ $\mathcal{P}_4 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{C}(\boldsymbol{v}7), \mathsf{CA}(\boldsymbol{v}_2)]$ $\mathcal{P}_5 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{CA}(\boldsymbol{v}_2), \mathsf{CA}(\boldsymbol{v}_3)]$
$\mathcal{T}_4$	$\beta = \frac{1}{22}(21+13\sqrt{3}) \sim 1.978$	$\mathcal{P}_1 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{C}(\boldsymbol{v}_7)]$ $\mathcal{P}_2 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$ $\mathcal{P}_4 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{C}(\boldsymbol{v}7), \mathsf{CA}(\boldsymbol{v}_2), \mathsf{CA}(\boldsymbol{v}_3)]$
$\mathcal{T}_5$	$\frac{1}{11}(9+4\sqrt{3}) < \beta < \frac{1}{22}(21+13\sqrt{3})$	$\mathcal{P}_1 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{C}(\boldsymbol{v}_7)]$ $\mathcal{P}_2 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$ $\mathcal{P}_4 = [\boldsymbol{v}_3, \boldsymbol{v}_7, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$
$\mathcal{T}_6$	$\beta = \frac{1}{11}(9+4\sqrt{3}) \sim 1.45$	$\mathcal{P}_1 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{C}(\boldsymbol{v}_7)]$ $\mathcal{P}_2 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \boldsymbol{v}_7, \mathsf{A}(\boldsymbol{v}_2), \mathsf{A}(\boldsymbol{v}_3)]$
$\mathcal{T}_7$	$\frac{1}{121}(72+43\sqrt{3}) < \beta < \frac{1}{11}(9+4\sqrt{3})$	$\mathcal{P}_1 = [\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_7]$ $\mathcal{P}_2 = [\boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_5, \boldsymbol{v}_7]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_5, \boldsymbol{v}_7, \mathsf{A}(\boldsymbol{v}_2)]$ $\mathcal{P}_4 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \mathsf{A}(\boldsymbol{v}_2)]$
$\mathcal{T}_8$	$\beta = \frac{1}{121}(72+43\sqrt{3}) \sim 1.21$	$\mathcal{P}_1 = [\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_7]$ $\mathcal{P}_2 = [\boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_5, \boldsymbol{v}_7]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \boldsymbol{v}_7, \mathsf{A}(\boldsymbol{v}_2)]$
$\mathcal{T}_9$	$(-21+13\sqrt{3})^{-1} < \beta < \frac{1}{121}(72+43\sqrt{3})$	$\mathcal{P}_1 = [\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_7]$ $\mathcal{P}_2 = [\boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_5, \boldsymbol{v}_7]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \boldsymbol{v}_7]$ $\mathcal{P}_4 = [\boldsymbol{v}_2, \boldsymbol{v}_5, \boldsymbol{v}_6, \boldsymbol{v}_7]$
$\mathcal{T}_{10}$	$\beta = (-21+13\sqrt{3})^{-1} \sim 0.659$	$\mathcal{P}_1 = [\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_7]$ $\mathcal{P}_2 = [\boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_5, \boldsymbol{v}_7]$ $\mathcal{P}_3 = [\boldsymbol{v}_3, \boldsymbol{v}_4, \boldsymbol{v}_5, \boldsymbol{v}_7]$ $\mathcal{P}_4 = [\boldsymbol{v}_2, \boldsymbol{v}_5, \boldsymbol{v}_6, \boldsymbol{v}_7, \mathsf{D}_2(\boldsymbol{v}_7)]$

Table 3. The remaining tilings.

	β	Γ-generating tiles
$\mathcal{T}_{11}$	$\frac{1}{143}(48+25\sqrt{3}) < \beta < (-21+13\sqrt{3})^{-1}$	$\mathcal{P}_1 = [v_3, v_4, v_5, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_2, v_3, v_5, v_7]$ $\mathcal{P}_3 = [v_3, v_4, v_5, v_7]$ $\mathcal{P}_4 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_5 = [v_2, v_5, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{12}$	$\beta = \frac{1}{143}(48+25\sqrt{3}) \sim 0.638$	$\mathcal{P}_1 = [v_3, v_4, v_5, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_2, v_3, v_5, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_3 = [v_3, v_4, v_5, v_7]$ $\mathcal{P}_4 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{13}$	$\frac{1}{11}(6-\sqrt{3}) < \beta < \frac{1}{143}(48+25\sqrt{3})$	$\mathcal{P}_1 = [v_3, v_4, v_5, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_3, v_4, v_5, v_7]$ $\mathcal{P}_3 = [v_3, v_5, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_4 = [v_2, v_3, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_5 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{14}$	$\beta = \frac{1}{11}(6-\sqrt{3}) \sim 0.39$	$\mathcal{P}_1 = [v_3, v_4, v_5, v_7, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_3, v_5, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_3 = [v_2, v_3, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_4 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{15}$	$\frac{1}{33}(3+5\sqrt{3}) < \beta < \frac{1}{11}(6-\sqrt{3})$	$\mathcal{P}_1 = [v_3, v_5, v_7, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_4, v_5, v_7, \mathsf{C}(v_7)]$ $\mathcal{P}_3 = [v_3, v_5, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_4 = [v_2, v_3, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_5 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{16}$	$\beta = \frac{1}{33}(3+5\sqrt{3}) \sim 0.353$	$\mathcal{P}_1 = [v_3, v_5, v_7, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_4, v_5, v_7, \mathsf{C}(v_7), \mathsf{D}_1^{\mathsf{C}}(v_7)]$ $\mathcal{P}_3 = [v_3, v_5, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_4 = [v_2, v_3, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_5 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{17}$	$\frac{1}{143}(24+7\sqrt{3}) < \beta < \frac{1}{33}(3+5\sqrt{3})$	$\mathcal{P}_1 = [v_3, v_5, v_7, \mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_5, v_7, \mathsf{C}(v_7), \mathsf{D}_1^{\mathsf{C}}(v_7)]$ $\mathcal{P}_3 = [v_3, v_5, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_4 = [v_2, v_3, v_7, \mathsf{D}_2(v_7)]$ $\mathcal{P}_5 = [v_5, v_6, v_7, \mathsf{D}_2(v_7)]$
$\mathcal{T}_{18}$	$\beta = \frac{1}{143}(24+7\sqrt{3}) \sim 0.252$	$\mathcal{P}_1 = [v_3, v_5, v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_5, v_7, \mathsf{AC}(v_7), \mathsf{D}_2(v_7)]$
$\mathcal{T}_{19}$	$\frac{1}{33}(6-\sqrt{3}) < \beta < \frac{1}{143}(24+7\sqrt{3})$	$\mathcal{P}_1 = [v_3, v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_5, v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7)]$ $\mathcal{P}_3 = [v_5, v_7, \mathsf{AC}(v_7), \mathsf{D}_2(v_7)]$
$\mathcal{T}_{20}$	$\beta = \frac{1}{33}(6-\sqrt{3}) \sim 0.129$	$\mathcal{P}_1 = [v_3, v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_5, v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7), \mathsf{AC}(v_7)]$
$\mathcal{T}_{21}$	$\beta < \frac{1}{33}(6-\sqrt{3})$	$\mathcal{P}_1 = [v_3, v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7)]$ $\mathcal{P}_2 = [v_7, \mathsf{C}(v_7), \mathsf{D}_2(v_7), \mathsf{D}_2\mathsf{C}(v_7), \mathsf{AC}(v_7)]$ $\mathcal{P}_3 = [v_5, \mathsf{C}(v_7), \mathsf{D}_2\mathsf{C}(v_7), \mathsf{AC}(v_7)]$

From our earlier observations, it remains only to check that there are no symmetries of the even numbered tilings that interchange vertices of $\Gamma.\boldsymbol{v}_1$ with those of $\Gamma.\boldsymbol{v}_7$. The arguments for each of the cases are very similar, we start with $\mathcal{T}_2$ as a model case. Recall that $\boldsymbol{v}_2$, $\boldsymbol{v}_4$, $\boldsymbol{v}_5$, and $\boldsymbol{v}_6$ are each Γ-equivalent to $\boldsymbol{v}_1$.

Suppose there exists $\gamma \in \mathrm{Sym}(\mathcal{T}_2)$ exchanging $\Gamma.\boldsymbol{v}_1$ with $\Gamma.\boldsymbol{v}_7$. Then $\gamma(\mathcal{P}_4)$ is a tile of $\mathcal{T}_2$ with exactly five vertices. $\mathcal{P}_4$ is the unique generating tile with five vertices so there exists $\gamma' \in \Gamma$ with $\gamma'\gamma(\mathcal{P}_4) = \mathcal{P}_4$. Since $\gamma' \in \Gamma$ it preserves the cusp classes of the vertices of tiles in $\mathcal{T}_2$. On the other hand, since γ exists, the minimal orbifold must have exactly two cusps, hence $\gamma'\gamma$ must exchange the vertices of $\mathcal{P}_4$ in $\Gamma.\boldsymbol{v}_1$ with those in $\Gamma.\boldsymbol{v}_7$. But our explicit description implies that there are three of the former and only one of the latter, a contradiction.

The same sort of argument also works for the remaining even numbered triangulations with the exception of $\mathcal{T}_{12}$ and $\mathcal{T}_{14}$. Consider the case of $\mathcal{T}_{12}$. Suppose there is $\gamma \in \mathrm{Sym}(\mathcal{T}_{12})$ exchanging $\Gamma.\boldsymbol{v}_1$ with $\Gamma.\boldsymbol{v}_7$. Arguing as before, we have an element $\delta \in \mathrm{Comm}(\Gamma)$ with $\delta(\mathcal{P}_2) = \mathcal{P}_2$ and which interchanges its vertices in $\Gamma.\boldsymbol{v}_1$ with those in $\Gamma.\boldsymbol{v}_7$. Since $\mathcal{P}_2$ has two vertices in $\Gamma.\boldsymbol{v}_1$ and two in $\Gamma.\boldsymbol{v}_7$, we have not yet arrived at a contradiction. But such a δ still cannot exist since it can be seen that the two vertices in $\Gamma.\boldsymbol{v}_7$ are connected by an edge but those in $\Gamma.\boldsymbol{v}_1$ are not.

The argument for $\mathcal{T}_{14}$ follows the outline of the argument for $\mathcal{T}_{12}$. Here $\mathcal{P}_1$ is the unique generating tile with five vertices, its two vertices in $\Gamma.\boldsymbol{v}_1$ are connected by an edge, and those in $\Gamma.\boldsymbol{v}_7$ are not. □

Proof of Proposition 7.14. By Lemma 7.18, we have $\mathrm{Comm}(\Gamma) = \mathrm{Sym}(\mathcal{T}_{18})$ so to prove the theorem we need to show $\mathrm{Sym}(\mathcal{T}_{18}) = \Gamma$. We already know that $\Gamma \subset \mathrm{Sym}(\mathcal{T}_{18})$.

Suppose that $\gamma \in \mathrm{Sym}(\mathcal{T}_{18}) - \Gamma$ is non-trivial. Since $\mathcal{T}_{18}$ is Γ-generated by $\mathcal{P}_1$ and $\mathcal{P}_2$ and these two polyhedra have different numbers of vertices, we may assume that $\gamma(\mathcal{P}_1) = \mathcal{P}_1$. By composing γ with d_2, if necessary, we may assume that γ is orientation preserving. $\mathcal{P}_1$ has one vertex in $\Gamma.\boldsymbol{v}_1$, one in $\Gamma.\boldsymbol{v}_3$, and four in $\Gamma.\boldsymbol{v}_7$; thus by Lemma 7.18, γ fixes $\boldsymbol{v}_5$ (which is in $\Gamma.\boldsymbol{v}_1$) and $\boldsymbol{v}_3$.

Using our embedding in the upper half space model, the vertices of $\mathcal{P}_1$ in $\Gamma.\boldsymbol{v}_7$ are taken to

$$r_7 = \frac{1+i}{2}, \qquad \mathsf{c}(r_7) = \frac{1+i(2+\sqrt{3})}{2},$$
$$\mathsf{d}_2\mathsf{c}(r_7) = \frac{3+i(2+\sqrt{3})}{2}, \qquad \mathsf{d}_2(r_7) = \frac{3+i}{2}.$$

Since γ is an elliptic isometry preserving $\mathcal{P}_1$ and fixing $r_3 = \infty$ and $r_5 = 1 + \frac{i}{2}(1+\sqrt{3})$, it must act as a cyclic permutation on the set $\{r_7, \mathsf{c}(r_7), \mathsf{d}_2\mathsf{c}(r_7), \mathsf{d}_2(r_7)\}$. But it is easy to see that the axis of γ is not perpendicular to the plane that these points span, so this is impossible. □

References

[1] C. C. Adams, Augmented alternating link complements are hyperbolic. In *Low-dimensional topology and Kleinian groups* (Coventry/Durham, 1984), London Math. Soc. Lecture Note Ser. 112, Cambridge University Press, Cambridge 1986, 115–130. Zbl 0632.57008 MR 0903861

[2] I. Agol, D. D. Long, and A. W. Reid, The Bianchi groups are separable on geometrically finite subgroups. *Ann. of Math.* (2) **153** (2001), no. 3, 599–621. Zbl 1067.20067 MR 1836283

[3] I. Agol, Criteria for virtual fibering. *J. Topol.* **1** (2008), no. 2, 269–284. Zbl 1148.57023 MR 2399130

[4] E. M. Andreev, Convex polyhedra of finite volume in Lobačevskiĭ space. *Mat. Sb.* (*N.S.*) **83** (1970), no. 125, 256–260. Zbl MR 0273510

[5] R. Benedetti and C. Petronio, *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin 1992. Zbl 0768.51018 MR 1219310

[6] N. Bergeron, Virtual fibering of certain cover of $\mathbb{S}^3$, branched over the figure eight knot. *C. .R Math. Acad. Sci. Paris* **346** (2008), no. 19–20, 1073–1078. Zbl MR

[7] N. Bergeron, F. Haglund, and D. Wise, Hyperplane sections in arithmetic hyperbolic manifolds. *J. London Math. Soc.* (2) **83** (2011), no. 2, 431–448 Zbl 1236.57021 MR 2776645

[8] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer-Verlag, Berlin 1999. Zbl 0988.53001 MR 1744486

[9] R. G. Burns, A. Karrass, and D. Solitar, A note on groups with separable finitely generated subgroups. *Bull. Austral. Math. Soc.* **36** (1987), no. 1, 153–160. Zbl 0613.20018 MR 0897431

[10] R. D. Canary, Marden's tameness tonjecture: history and applications. In *Geometry, analysis and topology of discrete groups*, ed. by L. Ji, K. Liu, L. Yang and S.T. Yau, Higher Education Press, Beijing 2008, 137–162. Zbl 1149.57028 MR 2464394

[11] E. Chesebro and J. DeBlois, Algebraic invariants, mutation, and commensurability of link complements. arXiv:1202.0765v2 [math.GT].

[12] D. Coulson, O. Goodman, C. Hodgson, and W. Neumann, Computing arithmetic invariants of 3-manifolds. *Experiment. Math.* **9** (2000), no. 1, 127–152. Zbl 1002.57044 MR 1101.20025

[13] C. Drutu and M. Sapir, Tree-graded spaces and asymptotic cones of groups. With an appendix by Denis Osin and Sapir. *Topology* **44** (2005), no. 5, 959–1058. Zbl 1101.20025 MR 2153979

[14] D. B. A. Epstein and R. C. Penner, Euclidean decompositions of noncompact hyperbolic manifolds. *J. Differential Geom.* **27** (1988), no. 1, 67–80. Zbl 0611.53036 MR 0918457

[15] R. Gitik, Doubles of groups and hyperbolic LERF 3-manifolds. *Ann. of Math.* (2) **150** (1999), no. 3, 775–806. Zbl 0944.20023 MR 1740992

[16] O. Goodman, D. Heard, and C. Hodgson, Commensurators of cusped hyperbolic manifolds. *Experiment. Math.* **17** (2008), no. 3, 283–306. Zbl 05500548 MR 2455701

[17] F. Haglund, Finite index subgroups of graph products. *Geom. Dedicata* **135** (2008), 167–209. Zbl 1195.20047 MR 2413337

[18] F. Haglund and D. T. Wise, Coxeter groups are virtually special. *Adv. Math.* **224** (2010), no. 5, 1890–1903. Zbl 1195.53055 MR 2646113

[19] F. Haglund and D. T. Wise, Special cube complexes. *Geom. Funct. Anal.* **17** (2008), no. 5, 1551–1620. Zbl 1155.53025 MR 2377497

[20] A. Hatcher, Hyperbolic structures of arithmetic type on some link complements. *J. London Math. Soc.* (2) **27** (1983), no. 2, 345–355. Zbl 0516.57001 MR 0692540

[21] G. C. Hruska, Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.* **10** (2010), no. 3, 1807–1856. Zbl 1202.20046 MR 2684983

[22] M. Lackenby, The volume of hyperbolic alternating link complements. With an appendix by Ian Agol and Dylan Thurston. *Proc. London Math. Soc.* (3) **88** (2004), no. 1, 204–224. Zbl 1041.57002 MR 2018964

[23] D. D. Long, Immersions and embeddings of totally geodesic surfaces. *Bull. London Math. Soc.* **19** (1987), no. 5, 481–484. Zbl 0596.57011 MR 0898729

[24] D. D. Long and A. W. Reid, Subgroup separability and virtual retractions of groups. *Topology* **47** (2008), no. 3, 137—159. Zbl 1169.57003 MR 2414358

[25] C. Maclachlan and A. W. Reid, *The arithmetic of hyperbolic 3-manifolds*. Grad. Texts in Math. 219, Springer-Verlag, New York 2003. Zbl 1025.57001 MR 1937957

[26] J. Fox Manning and E. Martínez-Pedroza, Separation of relatively quasiconvex subgroups. *Pacific J. Math.* **244** (2010), no. 2, 309–334. Zbl 1201.20024 MR 2587434

[27] G. A. Margulis, *Discrete subgroups of semisimple Lie groups*. Ergeb. Math. Grenzgeb. (3) 17, Springer-Verlag, Berlin 1991. Zbl 0732.22008 MR 1090825

[28] W. Neumann, Hilbert's 3rd problem and invariants of 3-manifolds. In *The Epstein birthday schrift*, Geom. Topol. Monogr. 1, Geometry & Topology Publications, Coventry 1998, 383–411. Zbl 0902.57013 MR 1668316

[29] W. Neumann and A. W. Reid, Arithmetic of hyperbolic manifolds. In *Topology '90*, Ohio State Univ. Math. Res. Inst. Publ. 1, de Gruyter, Berlin 1992, 273–310. Zbl 0777.57007 MR 1184416

[30] W. Neumann and D. Zagier, Volumes of hyperbolic three-manifolds. *Topology* **24** (1985), no. 3, 307–332. Zbl 0589.57015 MR 0815482

[31] D. V. Osin, Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.* **179** (2006), no. 843. Zbl 1093.20025 MR 1093.20025

[32] J. S. Purcell, Volumes of highly twisted knots and links. *Algebr. Geom. Topol.* **7** (2007), 93–108. Zbl 1135.57005 MR 2289805

[33] J. S. Purcell, Cusp shapes under cone deformation. *J. Differential Geom.* **80** (2008), no. 3, 453–500. Zbl 1182.57015 MR 2472480

[34] J. G. Ratcliffe, *Foundations of hyperbolic manifolds*. Grad. Texts in Math. 149, Springer-Verlag, New York 1994. Zbl 0809.51001 MR 1299730

[35] P. Scott, Subgroups of surface groups are almost geometric. *J. London Math. Soc.* (2) **17** (1978), no. 3, 555–565. Zbl 0412.57006 MR 0811778

[36] W. P. Thurston, *The geometry and topology of* 3*-manifolds*. Mimeographed lecture notes, 1979. Zbl 0483.57007 MR 0662424

[37] D. T. Wise, Subgroup separability of the figure 8 knot group. *Topology* **45** (2006), no. 3, 421–463. Zbl 1097.20030 MR 2218750

Received September 16, 2009

Eric Chesebro, Department of Mathematical Sciences, University of Montana, Missoula, MT 59812-0864, U.S.A.

E-mail: Eric.Chesebro@mso.umt.edu

Jason DeBlois, Department of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260, U.S.A.

E-mail: jdeblois@pitt.edu

Henry Wilton, Department of Mathematics, University College London, Gower Street, London, WC1E 6BT, England

E-mail: hwilton@math.ucl.ac.uk