Canonical decompositions of hyperbolic 3-orbifolds

by

Mark Fincher

Bachelor of Science, University of North Texas, 2016

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Mark Fincher

It was defended on

August 2nd 2022

and approved by

Jason DeBlois, University of Pittsburgh

Nathan Dunfield, University of Illinois at Urbana-Champaign

Armin Schikorra, University of Pittsburgh

Carl Wang-Erickson, University of Pittsburgh

Canonical decompositions of hyperbolic 3-orbifolds

Mark Fincher, PhD

University of Pittsburgh, 2022

This thesis describes the theory behind *Sym*, software created by the author for computations with finite-volume cusped hyperbolic 3-orbifolds. The main purpose of Sym, in its current form, is to compute canonical (Epstein-Penner) decompositions of these orbifolds. This was originally motivated by a joint project between the author, his advisor, and his advisor's other graduate students to create a census of orbifolds commensurable to the figure-eight knot complement.

Underlying Sym is a non-standard notion of an orbifold triangulation, in which tetrahedra may be labeled with groups of symmetries acting on them. This allows us to consider fully ideal hyperbolic triangulations of orbifolds, which we attempt to treat in the same way that SnapPy treats ideal triangulations of manifolds. *SnapPy* is powerful existing software for hyperbolic 3-manifolds and some orbifolds, originally developed by Weeks and now maintained by Culler, Dunfield, and Goerner.

The way SnapPy finds canonical decompositions of hyperbolic manifolds is complicated both theoretically and computationally, and relies on influential work by Epstein, Penner, Weeks, and others. The main goal of this thesis is to extend that work to orbifolds. A key idea we develop is an orbifold version of Pachner moves, which are moves which change an orbifold triangulation locally.

Table of Contents

List of Figures

Preface

I dedicate this work to my grandfather, Dr. Bobby Fincher, known as “Buddy” by his grandchildren. Thank you for showing me how to be a mathematician.

1.0 Introduction

The *canonical decomposition* of a finite-volume cusped hyperbolic 3-manifold M is a decomposition of M into hyperbolic ideal polyhedra which is uniquely determined by the hyperbolic metric. By the Mostow-Prasad rigidity theorem, this decomposition in fact only depends on the topology of M. The canonical decomposition was defined (in greater generality) by Epstein and Penner in the 1980s [15]. It carries deep information about the manifold. In particular, the self-isometry group can be read off from it, and two finite-volume cusped hyperbolic 3-manifolds are isometric if and only if they have identical canonical decompositions.

A theme in the intersection of hyperbolic geometry and 3-manifold topology is that many powerful ideas are also computationally tractable. Canonical decompositions give a nice example of this. Weeks described how to compute them in the early 1990s in [43], and implemented his algorithm in his computer program SnapPea, now maintained as SnapPy with extended functionality [13]. For an application of this, as well as other breakthroughs in low-dimensional topology, consider the following paragraph, which would have astounded any mathematician if they were to read it 60 years ago.

Suppose we have two knots K and K' in S^3, and we would like to know if they are actually the same. In other words, can we move around K within S^3, without breaking it or passing any part of it through itself, to turn it into K'? By [21], the knots are the same if and only if their knot complements $S^3 - K$ and $S^3 - K'$ are homeomorphic. Thurston showed most knot complements admit finite-volume cusped hyperbolic metrics [41, Corollay 2.5]. We can draw K and K' in a computer, load them to SnapPy, and if they have hyperbolic structures then SnapPy should be able to find them. Assuming that is the case, it can then compute the canonical decompositions, compare them with each other, and confidently tell us whether $S^3 - K$ and $S^3 - K'$ are homeomorphic or not, and hence whether or not the knots are the same.

An *orbifold* is a topological space locally modelled on quotients of Euclidean space by finite groups of diffeomorphisms, so it is a generalization of a manifold. For hyperbolic orbifolds, we can state this generalization in a more straightforward way. A hyperbolic 3-manifold is the quotient of $\mathbb{H}^3$ by the action of a discrete, torsion-free group of isometries of $\mathbb{H}^3$. A hyperbolic 3-orbifold is the same thing, except the group is allowed to have torsion. We would like to know if the story of canonical decompositions for hyperbolic manifolds extends to hyperbolic orbifolds, in particular we would like to answer the following.

1. Does the definition of "canonical decomposition" extend from finite-volume cusped hyperbolic 3-manifolds to finite-volume cusped hyperbolic 3-orbifolds, along with its useful properties?
2. If so, can we generalize SnapPy's algorithm for finding canonical decompositions of manifolds to orbifolds?

The approach we take to answering the first question is to modify what we mean by "polyhedral decomposition" to be suitable for orbifolds. Loosely speaking, we define a *hyperbolic orbifold polyhedral decomposition* to be a collection of hyperbolic ideal polyhedra with face gluing data (where a face can be glued to itself), edges labelled by integers representing rotation groups, and with polyhedra labelled with symmetry groups. Often all polyhedra are tetrahedra, in which case we are really working with a kind of non-standard triangulation.

Theorem 7. *Let Q be a finite-volume cusped hyperbolic 3-orbifold. The Epstein-Penner construction still applies to Q, and the canonical decomposition it provides is a hyperbolic orbifold polyhedral decomposition. The self-isometry group of Q is equal to the group of combinatorial symmetries of this decomposition. Two orbifolds are isometric if and only if they have combinatorially isomorphic canonical decompositions.*

The fact that the Epstein-Penner construction makes sense for orbifolds will not be a surprise to anyone familiar with it. The question then is, what exactly is the object which the construction gives? This is why we introduce orbifold polyhedral decompositions. There has been at least one other interpretation of what a canonical decomposition of an orbifold should be defined to be, by Damian Heard in his computer program *Orb*. We briefly describe Heard's interpretation, to the best of our knowledge, at the end of Section 1.1.

SnapPy's algorithm for finding the canonical decomposition of a manifold is called *canonize*. As an answer to the second question, we have:

Theorem 12. *There is a "canonize" algorithm for orbifolds, which generalizes SnapPy's corresponding algorithm for manifolds. Implemented in Sym, it has successfully computed the certified canonical decompositions of all* 336 *orbifolds which cover* $\mathbb{H}^3/\operatorname{PGL}(2, O_3)$ *up to covering degree* 48.

Unfortunately, SnapPy's canonize is not guaranteed to successfully find the canonical decomposition, but in practice it always does. Sym is also not guaranteed to succeed, but it also has succeeded for all examples it has been tried on.

The goal of this thesis is to explain Theorems 7 and 12. In Section 1.1, we give an outline of our "canonize for orbifolds" algorithm, written for experts familiar with manifold canonical decompositions. In Chapter 2, we record some of the relevant background material about orbifolds, manifolds, and hyperbolic geometry. In Chapter 3, we give our definition of an orbifold triangulation, in which some of the tetrahedra may be labelled with symmetries, which is the basic data structure Sym uses to represent orbifolds. In Chapter 4, we define the canonical decomposition of an orbifold, proving Theorem 7. Our canonize algorithm is split into two parts, which we describe in detail in Chapters 5 and 6. Finally, in Chapter 7 we describe the role of Sym in helping to create a census of orbifolds commensurable to the figure-eight knot complement, which is a joint project with the author's PhD advisor and this advisor's other students. In particular, we discuss the 336 orbifolds on which we have tested our canonize algorithm.

1.1 Canonize: a brief outline

The canonical decomposition of a finite-volume cusped hyperbolic 3-manifold $\mathbb{H}^3/\Gamma$ is defined using a convex hull construction in Minkowski space. Really, the construction associates to Γ a unique Γ-invariant tiling of $\mathbb{H}^3$ by ideal hyperbolic polyhedra, and because Γ has no torsion, this tiling projects to a decomposition of $\mathbb{H}^3/\Gamma$. If Γ does have torsion, so

$\mathbb{H}^3/\Gamma$ is an orbifold, then the construction still gives a unique Γ-invariant ideal tiling. But it will not project to a decomposition of $\mathbb{H}^3/\Gamma$ if any of the polyhedra are fixed by torsion elements of Γ. Essentially our approach is to accept that and just work in the universal cover instead.

From that perspective, we assume we have some Γ-invariant tiling by ideal hyperbolic tetrahedra which is perhaps not the canonical tiling, and our goal is to change it into the canonical tiling by some finite sequence of moves. We require that the result of each move is still a Γ-invariant tiling by ideal hyperbolic tetrahedra. The standard way of changing a simplicial complex, without changing the underlying topological space, is with *Pachner moves*, first defined in [33]. For our purposes we create new kinds of Pachner moves which are sensitive to the orbifold structure, allowing us to modify the tiling in a Γ-invariant way. The choice of where to apply these moves is guided by *tilt* computations, as described by Weeks [43].

We hope that after a sequence of orbifold Pachner moves we arrive at the canonical tiling. However, because of their computational convenience we are only using tilings by tetrahedra, but the canonical tiling could have polyhedra which are not tetrahedra. In that case, we attempt to find the *canonical re-triangulation* of the canonical decomposition. This is what SnapPy does for manifolds, but extending the algorithm to our orbifold setting takes work.

While we can always imagine we are in the universal cover, in practice we work with objects which we call *orbifold triangulations* or, more generally, *orbifold polyhedral decompositions*. Our definition of an orbifold triangulation is not standard, because it is not a triangulation of the underlying space. With this new definition of an orbifold triangulation, we are able to triangulate orbifolds by tetrahedra which are fully ideal, which is often not possible with a standard triangulation.

Apart from its application here to hyperbolic geometry, we feel that our definition of an orbifold triangulation more harmoniously blends the combinatorics of tetrahedra with the local group structures of an orbifold. It seems to the author that this is useful in order to have a theory of Pachner moves for orbifolds.

Damian Heard took a different approach to orbifold canonical decompositions in his thesis [23], which he implemented in his computer program Orb (which appears to no longer be maintained). To a given orientable finite-volume cusped hyperbolic 3-orbifold Q there is an associated *pared hyperbolic manifold*, an infinite-volume hyperbolic manifold obtained from Q by making the orders of its isotropy groups go to infinity. In [24], Kojima generalized the Epstein-Penner construction to work for a broader class of hyperbolic 3-manifolds, including these pared manifolds, and in [17] Frigerio-Petronio gave an algorithm to compute this Kojima decomposition. To Orb, the canonical decomposition of Q is this, the Kojima decomposition of the pared manifold associated to Q.

2.0 Background

2.1 Orbifold definition

Orbifolds were first defined by Satake in [38],[39]. He called them $V-$manifolds. It was Thurston (with input from his students) who later chose the name “orbifold”. An orbifold is one natural generalization of a manifold and, like manifolds, they appear in many different areas of math. See the introduction of [4] for a broad perspective on orbifolds and their importance in topology, algebraic geometry, and physics. For low-dimensional topologists, standard references for the definitions below are [7], [12], and Chapter 13 of Thurston’s original notes [40].

Definition 1. An *n-dimensional orbifold* Q is a pair $(X_Q, \mathcal{U})$, where X_Q is a paracompact, Hausdorff, topological space, called the *underlying space* of Q, and $\mathcal{U}$ is an *orbifold atlas.* An orbifold atlas is a collection of *charts* $\{(U_i, \tilde{U}_i, \phi_i, \Gamma_i)\}$, where each U_i is an open subset of X_Q, $\tilde{U}_i$ is an open subset of $\mathbb{R}^n$, ϕ_i is a continuous map from $\tilde{U}_i$ to U_i, and Γ_i is a finite group of diffeomorphisms of $\tilde{U}_i$, and these all satisfy:

1. The U_i’s cover X_Q.
2. If $U_i \cap U_j$ is non-empty, then it is part of a chart.
3. Each ϕ_i factors through a homeomorphism from $\tilde{U}_i/\Gamma_i$ to U_i.
4. If $U_i \subset U_j$ then there is a smooth embedding $\psi : \tilde{U}_i \to \tilde{U}_j$ such that $\phi_j \circ \psi = \phi_i$.

We assume all group actions are effective, meaning the only group element acting as the identity map is the identity element. Additionally, it is convenient (for instance, for defining covering maps) to assume that all orbifold atlases are maximal, so we assume that too.

The definition says that each U_i is identified with a *local model* $\tilde{U}_i/\Gamma_i$ and condition 4 ensures these identifications do not conflict with each other, although it takes some work to see how it accomplishes this. Namely, by the following remark, proven in the appendix of [31], condition 4 implies that if $U_i \subset U_j$ then we can just assume the local models satisfy $\tilde{U}_i \subset \tilde{U}_j$ and $\Gamma_i \subset \Gamma_j$.

Remark 1. In the situation of condition 4, ψ gives rise to an injective homomorphism $\lambda : \Gamma_i \to \Gamma_j$ such that ψ is equivariant with respect to λ. Further, if $\gamma \in \Gamma_j$ is such that $\gamma(\psi(\tilde{U}_i)) \cap \psi(\tilde{U}_i) \neq \emptyset$ then $\gamma \in \lambda(\Gamma_i)$.

For $x \in X_Q$ contained in a chart U_i, with $\tilde{x} \in \tilde{U}_i$ such that $\phi_i(\tilde{x}) = x$, define the *local group* of x to be the stabilizer of $\tilde{x}$ in $\tilde{U}_i$ with respect to the action of Γ_i. Using Remark 1, it can be seen that the local group is well-defined up to isomorphism. Say that x is *singular* if its local group is non-trivial and *regular* if its local group is trivial. The *singular locus* of Q is the set of all singular points.

An orbifold is *locally orientable* if it has an atlas for which each Γ_i is a group of orientation preserving diffeomorphisms. It is *orientable* if, in addition, all inclusion maps $U_i \subset U_j$ induce orientation preserving maps $\tilde{U}_i \to \tilde{U}_j$. As with manifolds, an orientable orbifold can be given an orientation. In this work, all orbifolds are assumed to be orientable and oriented.

A standard kind of orbifold is a global quotient space M/Γ, where M is a differentiable manifold and Γ is a group of diffeomorphisms of M acting *properly discontinuously*. Γ acting properly discontinuously means that any points $x, y \in M$ have neighborhoods U_x, U_y such that

$$\{\gamma \in \Gamma : \gamma(U_x) \cap U_y \neq \emptyset\}$$

is finite. This property implies that each point $x \in M$ has a neighborhood U_x which is preserved by the stabilizer of x in Γ and disjoint from its other translates under Γ. The restriction of the quotient map $p : M \to M/\Gamma$ to U_x is then used as a chart map for an orbifold atlas on M/Γ.

We would now like to define a map of orbifolds. For motivation, suppose M, M' are manifolds and Γ, Γ' are groups of diffeomorphisms acting properly discontinuously on M and M' respectively. If there is a homomorphism $\lambda : \Gamma \to \Gamma'$ and a $\lambda-$equivariant smooth map $g : M \to M'$, then g induces a continuous map $M/\Gamma \to M'/\Gamma'$. We would like this induced map to be an example of an orbifold map. So, for orbifolds Q, Q' we define an *orbifold map* from $Q \to Q'$ to be a continuous map $f : X_Q \to X_{Q'}$ such that for every $x \in X_Q$, there is a chart U_i containing x, another chart U'_j containing $f(U_i)$, and a smooth map $\tilde{f} : \tilde{U}_i \to \tilde{U}'_j$ which is equivariant with respect to some homomorphism $\Gamma_i \to \Gamma'_j$ and descends to $f\big|_{U_i}$. We say that f is an *orbifold isomorphism* if each $\tilde{f} : \tilde{U}_i \to \tilde{U}'_j$ and $\Gamma_i \to \Gamma'_j$

is a diffeomorphism and group isomorphism respectively. An *orbifold automorphism* is an orbifold isomorphism from an orbifold to itself.

For the most part, in this work we are interested in three-dimensional orientable orbifolds. We conclude this section with a simple characterization of them, which is more down-to-earth than Definition 1.

Theorem 1 (Thm 2.5, [12]). *Let Q be an orientable 3-orbifold. Then the underlying space X_Q is an orientable manifold and the singular locus consists of edges of order $k \geq 2$ and vertices where 3 edges meet. An edge labelled k corresponds to a rotation group of order k. At a vertex, the three edges have orders $(2, 2, k)$ where $k \geq 2$, $(2, 3, 3), (2, 3, 4)$, or $(2, 3, 5)$. Conversely, every such labelled graph in an orientable 3-manifold describes an orientable 3-orbifold.*

2.2 Fundamental groups and covering theory

A *covering* $\tilde{Q}$ of Q is a continuous map $p : X_{\tilde{Q}} \to X_Q$, called a *covering map*, satisfying that every $x \in X_Q$ is contained in a chart U such that each component V_i of $p^{-1}(U)$ is a chart of $\tilde{Q}$, the local model of U is $\tilde{U}/\Gamma$, that of V_i is $\tilde{U}/\Gamma_i$ where $\Gamma_i \subset \Gamma$, and the following diagram commutes.

$$\begin{array}{ccc} \tilde{U}/\Gamma_i & \xrightarrow{\sim} & V_i \\ \downarrow & & \downarrow p \\ \tilde{U}/\Gamma & \xrightarrow{\sim} & U \end{array}$$

The map $\tilde{U}/\Gamma_i \to \tilde{U}/\Gamma$ is the natural quotient map. Recall that we assume all orbifold atlases are maximal, which is why it makes sense that some local chart for V_i might have such a particular form. Note that a covering map is an example of an orbifold map.

The *deck transformation group* of a covering $p : \tilde{Q} \to Q$ is the group of all orbifold automorphisms $f : \tilde{Q} \to \tilde{Q}$ such that $p = p \circ f$. A *universal cover* of Q is a cover $p : \tilde{Q} \to Q$ such that for any other cover $q : \tilde{Q}' \to Q$, there is a cover $r : \tilde{Q} \to \tilde{Q}'$ such that $p = q \circ r$. Two covers $p : \tilde{Q} \to Q$ and $q : \tilde{Q}' \to Q$ are *equivalent* if there is an orbifold isomorphism $f : \tilde{Q} \to \tilde{Q}'$ such that $p = q \circ f$.

Theorem 2 (Thurston)**.** *Every connected orbifold Q has a unique, up to equivalence, universal cover, which we denote $\tilde{Q}$.*

We define the *fundamental group* of Q, denoted $\pi_1(Q)$, as the deck transformation group of the universal cover of Q.

The deck transformation group of a cover $p : Q' \to Q$ acts properly discontinuously on Q', so the quotient of Q' under the action of the deck group has a natural orbifold structure. We say that $p : Q' \to Q$ is *regular* if that quotient is isomorphic to Q. Continuing the analogy with covering space theory for manifolds, we have the following theorem.

Theorem 3 ([7])**.** *Let Q be an orbifold. There is a one-to-one correspondence between conjugacy classes of subgroups of $\pi_1(Q)$ and equivalence classes of (connected) covers of Q. A cover corresponds to a normal subgroup if and only if it is regular.*

2.3 Geometric structures

Let X be a real analytic manifold and G a group of real analytic diffeomorphisms of X. Q is a (G, X)-orbifold if it is locally modelled on quotients of open subsets of X by finite subgroups of G. In other words, for each chart $(U_i, \tilde{U}_i, \phi_i, \Gamma_i)$ in the orbifold atlas, we should have $\tilde{U}_i \subset X$ and Γ_i a finite subgroup of G preserving $\tilde{U}_i$. We also require that each ϕ from condition 4 of the definition of an orbifold atlas be the restriction of a function in G.

If Q is a (G, X)-orbifold, we say it has a "geometric structure". Namely, a (G, X) structure. Often, X is $\mathbb{R}^n, \mathbb{H}^n$, or S^n, and G is the group of Euclidean, hyperbolic, or spherical isometries, respectively. The reason we care about Q having a (G, X) structure is that it potentially allows us to view Q as the quotient of X by some discrete subgroup of G. For more details on the following theorem, see Theorem 2.26 from [12].

Theorem 4. *If Q is a (G, X)-orbifold where (G, X) is Euclidean, spherical, or hyperbolic geometry, and if the induced metric on Q is complete, then Q is isomorphic to X/Γ for some discrete group of isometries $\Gamma \subset G$.*

The story of geometric structures on compact 2-orbifolds is particularly clear. A *teardrop* is a 2-orbifold whose underlying space is S^2 and whose singular locus is a single point, with local group an order n rotation where $n > 1$. A *spindle* is a 2-orbifold whose underlying space is S^2 and whose singular locus consists of two points, with local groups the rotations of orders n and m where $n \neq m, n, m > 1$.

Theorem 5 (Theorem 2.22, [12])**.** *Let Q be a compact 2-orbifold which is not a teardrop nor a spindle, nor the quotient of these by an involution. Then exactly one of the following is true.*

1. *Q admits a Euclidean structure.*
2. *Q admits a hyperbolic structure.*
3. *Q admits a spherical structure.*

Since Q is closed in the theorem, the induced metric must be complete. So the theorem implies that every closed 2-orbifold, with the exception of the orbifolds listed, can be obtained by quotienting $\mathbb{R}^2, \mathbb{H}^2$, or S^2 by some discrete group of isometries.

2.4 Hyperbolic geometry

Three dimensional hyperbolic space is defined to be the unique 3-dimensional Riemmanian manifold which is complete, simply-connected, and with constant sectional curvature -1. It is denoted $\mathbb{H}^3$. There are several models of $\mathbb{H}^3$ which are common to use. In this section, we will work in the upper half space model. In Chapter 4, we will work in the hyperboloid model. There are many references for a detailed introduction to hyperbolic geometry, for instance Thurston's book [42], Benedetti and Petronio's book [6], or Martelli's book [28].

2.4.1 Upper half space

The *upper half space model* of $\mathbb{H}^3$ is the manifold

$$\{(x, y, t) \in \mathbb{R}^3 : t > 0\},$$

equipped with the Riemannian metric

$$\langle v, w\rangle_{(x,y,t)} = \frac{v_1w_1 + v_2w_2 + v_3w_3}{t^2}$$

This is a conformal model of $\mathbb{H}^3$, meaning hyperbolic angles are the same as Euclidean angles.

We can regard the xy-plane, $\{(x, y, 0) : x, y \in \mathbb{R}\}$, as $\mathbb{C}$. Then the boundary $\partial\mathbb{H}^3$ of $\mathbb{H}^3$ in the upper half space model is $\partial\mathbb{H}^3 = \mathbb{C} \cup \{\infty\} = S^2$. Recall that the conformal automorphisms of S^2 are the Mobius transformations. These are maps of the form

$$z \mapsto \frac{az + b}{cz + d},$$

where $a, b, c, d \in \mathbb{C}$ and $ad - bc = 1$. It turns out that they extend to isometries of the upper half space model and that every orientation-preserving isometry of the upper half space model comes in this way. Hence the group of orientation-preserving isometries of $\mathbb{H}^3$, $\mathrm{Isom}^+(\mathbb{H}^3)$, is identified with $\mathrm{PSL}(2, \mathbb{C})$. Viewed as the group of Mobius transformations, it is easy to see that $\mathrm{Isom}^+(\mathbb{H}^3)$ acts properly and transitively on the set of distinct triples of points of $\partial\mathbb{H}^3$.

There are three kinds of orientation-preserving isometries: *elliptic*, *loxodromic*, and *parabolic*.

An elliptic isometry fixes a geodesic point-wise and rotates the rest of hyperbolic space around it. It is conjugate to a Mobius transformation of the form $z \mapsto e^{i\theta}z$ for $\theta \in [0, 2\pi)$, which extends to upper half space as $(z, t) \mapsto (e^{i\theta}z, t)$.

A loxodromic isometry translates along a unique geodesic axis and acts on the rest of $\mathbb{H}^3$ with a combination of twisting and translating, like a corkscrew. Such an isometry is conjugate to a Mobius transformation of the form $z \mapsto \lambda e^{i\theta}z$ where $\lambda > 0$ and $\theta \in [0, 2\pi)$, which extends to upper half space as $(z, t) \mapsto (\lambda e^{i\theta}z, \lambda t)$.

Parabolic isometries are characterized by the fact that they fix a unique point of $\partial\mathbb{H}^3$. They are conjugate to a Mobius transformation fixing ∞ of the form $z \mapsto z + a$ for $a \in \mathbb{C}$, which extends to upper half space as $(z, t) \mapsto (z + a, t)$.

The geodesics in the upper half space model are vertical Euclidean lines and Euclidean half circles with center in the plane $t = 0$. Geodesic planes are vertical Euclidean planes and Euclidean half-spheres with center in the plane $t = 0$. A third kind of object is a *horosphere*.

In this model, a *horosphere centered at* $z \in \mathbb{C}$ is a Euclidean sphere tangent to z. A *horosphere centered at* ∞ is a horizontal plane. Horospheres can be defined intrinsically in terms of the hyperbolic metric, so they are preserved by isometries. Each horosphere divides $\mathbb{H}^3$ into two regions. The region which has one point at infinity is called a *horoball*.

2.4.2 Hyperbolic structures on orbifolds

In Section 2.3, we defined geometric structures on orbifolds. We now specify to hyperbolic structures.

Every complete hyperbolic 3-orbifold is of the form $\mathbb{H}^3/\Gamma$ for some discrete group $\Gamma \subset \mathrm{Isom}(\mathbb{H}^3)$. An isometry $\mathbb{H}^3/\Gamma \to \mathbb{H}^3/\Gamma'$ is a map which lifts to an isometry $\mathbb{H}^3 \to \mathbb{H}^3$, conjugating Γ to Γ'. The self-isometry group of $\mathbb{H}^3/\Gamma$ is isomorphic to $\mathrm{N}(\Gamma)/\Gamma$, where $\mathrm{N}(\Gamma)$ is the normalizer in $\mathrm{Isom}(\mathbb{H}^3)$ of Γ.

We only consider orientable orbifolds, so we assume $\Gamma \subset \mathrm{Isom}^+(\mathbb{H}^3)$. Such an orbifold is a manifold if and only if Γ acts *freely* on $\mathbb{H}^3$, meaning no point of $\mathbb{H}^3$ is fixed by any element of Γ, which is equivalent to Γ not containing any elliptics. If Γ does contain elliptics, then discreteness implies they must have finite order, i.e. be torsion elements.

All 3-orbifolds in this work are orientable, finite-volume, non-compact, complete hyperbolic 3-orbifolds, which for convenience we just call "finite-volume cusped" or "cusped". The word *cusped* refers to the *thin* parts of such an orbifold. Every hyperbolic orbifold has a *thick-thin* decomposition which divides the orbifold into those points with injectivity radius less than a fixed constant, called the Margulis constant, and those with larger injectivity radius [14]. The points with smaller injectivity radius make up the thin part of the orbifold, which is a disjoint union of *cusps*, i.e. orbifolds of the form $Q^2 \times [0, \infty)$, where Q^2 is a closed Euclidean 2-orbifold. An orbifold which is finite-volume but non-compact must have a positive, finite number of cusps.

The following hugely important theorem implies that if a 3-orbifold has a finite-volume complete hyperbolic structure, then that structure is unique.

Theorem 6 (Mostow-Prasad rigidity [32] [34])**.** *Suppose $n > 2$. Let Γ and Γ' be discrete finite co-volume subgroups of* $\mathrm{Isom}(\mathbb{H}^n)$ *and $\phi : \Gamma \to \Gamma'$ an abstract group isomorphism.*

Then ϕ must be conjugation, i.e. there must be some $g \in \mathrm{Isom}(\mathbb{H}^n)$ such that $\phi(\gamma) = g\gamma g^{-1}$ for all $\gamma \in \Gamma$.

2.4.3 Polyhedra, tetrahedra, and tilings

A *hyperbolic polyhedron* is the intersection of finitely many closed half-spaces in $\mathbb{H}^3$. We only work with finite-volume polyhedra. It is a special feature of hyperbolic geometry that there exist finite-volume non-compact hyperbolic polyhedra.

The simplest of these is the *hyperbolic ideal tetrahedron*, which is the hyperbolic convex hull of any four distinct points in $\partial\mathbb{H}^3$. We say that its vertices are *ideal*. Any polyhedron whose vertices lie in $\partial\mathbb{H}^3$ we also call ideal.

Any face of a hyperbolic ideal tetrahedron can be mapped to the ideal triangle which is the hyperbolic convex hull of $0, 1, \infty$. The fourth ideal vertex of the tetrahedron we can assume maps to some $z \in \mathbb{C}$ with non-negative imaginary part. The *edge parameter* of the edge from 0 to ∞ is z. The edge parameter of any edge is defined in this way, by mapping it to the geodesic from 0 to ∞, and the other points to 1 and some $w \in \mathbb{C}$ with non-negative imaginary part. Any one of the edge parameters uniquely determines the tetrahedron up to orientation-preserving isometry. An ideal tetrahedron is *flat*, i.e. contained in a plane, if and only if all of its edge parameters are real. The relationships between edge parameters is given in figure 1.

The group of symmetries of any ideal tetrahedron has a subgroup isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. For a non-flat tetrahedron, this subgroup's non-trivial elements can be seen as follows. For any pair of opposite edges of a non-flat ideal tetrahedron, there must be a unique geodesic of minimal length connecting them, which must intersect both edges at right angles. Hence π rotation around this geodesic is a symmetry of the tetrahedron. This implies that opposite edges must have equal dihedral angles (this could already be deduced from figure 1). The corresponding symmetries of any flat ideal tetrahedron can be seen in a similar way.

A symmetry which an ideal tetrahedron might or might not have is the order 3 rotation fixing a vertex. If it does have this symmetry, then in fact it must be a *regular ideal tetrahedron*, meaning all edges have equal dihedral angles. This is because the order 3

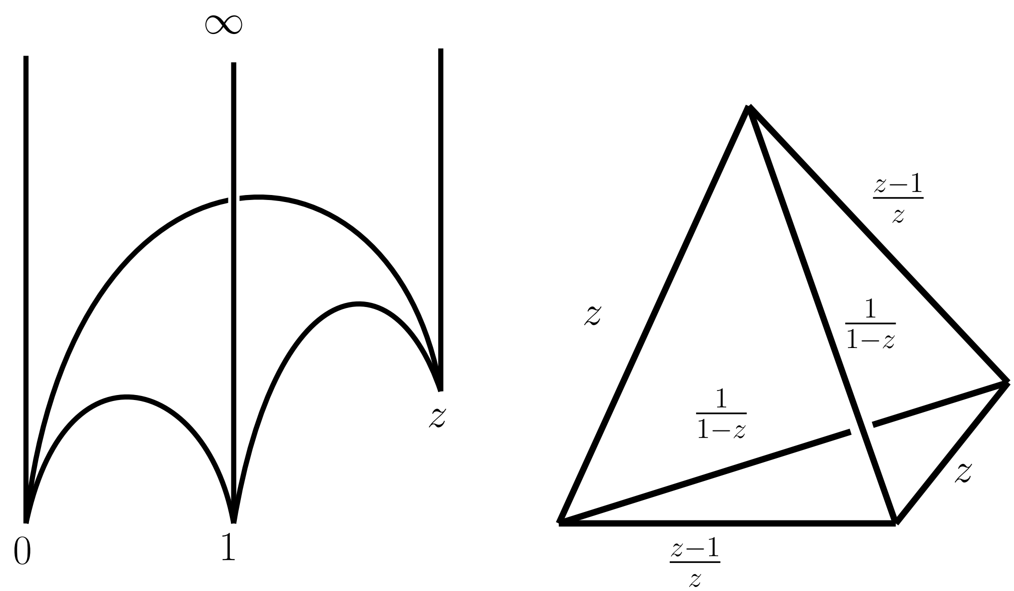

Figure 1: The edge parameters.

rotations combined with the order 2 rotations described in the previous paragraph generate the full group of orientation preserving symmetries of a tetrahedron, which acts transitively on edges. Up to isometry there is only one regular ideal tetrahedron, which we can take to be the convex hull of $0, 1, \infty$, and $\frac{1}{2} + \frac{\sqrt{3}i}{2}$.

Definition 2. A *tiling* of $\mathbb{H}^3$ is a decomposition of $\mathbb{H}^3$ into a locally finite set of hyperbolic polyhedra which intersect only in common faces. If all polyhedra are ideal, then we say the tiling is *ideal*.

Definition 3. A *tiling-group pair* $(\mathcal{T}, \Gamma)$ consists of a tiling of $\mathbb{H}^3$, $\mathcal{T}$, and a discrete group, $\Gamma < \mathrm{Isom}^+(\mathbb{H}^3)$, preserving $\mathcal{T}$.

A *geometric ideal triangulation* of a finite-volume cusped hyperbolic 3-manifold $M = \mathbb{H}^3/\Gamma$ is a triangulation of M by non-flat hyperbolic ideal tetrahedra. Because ideal tetrahedra are missing their vertices, this is not a standard kind of triangulation as seen in other areas of topology. The vertices could be added back in to get a non-manifold compactification of M. We can lift the triangulation to $\mathbb{H}^3$ to get an ideal tiling-group pair consisting of tetrahedra. It is typical to switch back and forth between the two perspectives—triangulation of the quotient manifold and Γ-invariant triangulation of $\mathbb{H}^3$.

Thurston came up with a method for finding complete hyperbolic structures on certain non-compact topological 3-manifolds by finding geometric ideal triangulations of them. First, try to get a triangulation of the manifold by topological ideal tetrahedra. If successful, then assign a complex variable to each tetrahedron, representing a possible edge parameter. If we choose the edge parameters in such a way that the total dihedral angle around each quotient edge is 2π, then these hyperbolic ideal tetrahedra induce a geometric structure on the manifold. Finding these edge parameters ends up amounting to solving a system of complex-variable algebraic equations, called the *consistency equations*. For the structure to be complete, it turns out the edge parameters just need to satisfy another set of equations, called the *completeness equations*. SnapPy implements this method for hyperbolic link complements.

2.5 Pachner moves

Pachner moves are used to locally change a triangulation of a manifold, without changing the manifold itself. In this section we describe a few of them.

There are two natural triangulations of a triangular bipyramid. See figure 2. One triangulation has 2 tetrahedra and the other has 3. The $2 \to 3$ Pachner move changes from the triangulation with 2 tetrahedra to the one with 3, while the $3 \to 2$ move does the reverse. The triangulation of the boundary of the bipyramid is not changed.

Of course, there are usually many such triangulated bipyramids within a triangulation. In particular, anywhere we have two distinct tetrahedra glued together along a common face, we can do a $2 \to 3$ move. Conversely, anywhere three distinct tetrahedra are glued in a cycle around a common edge, we can do a $3 \to 2$ move. There is a slight annoyance when working with a geometric triangulation. In the geometric setting, we can talk about the bipyramid being convex or not. If it is not convex, then a $2 \to 3$ move does not make geometric sense; the newly introduced edge connecting opposite vertices is not contained within the bipyramid. For this reason, we check for convexity before doing a geometric $2 \to 3$ move.

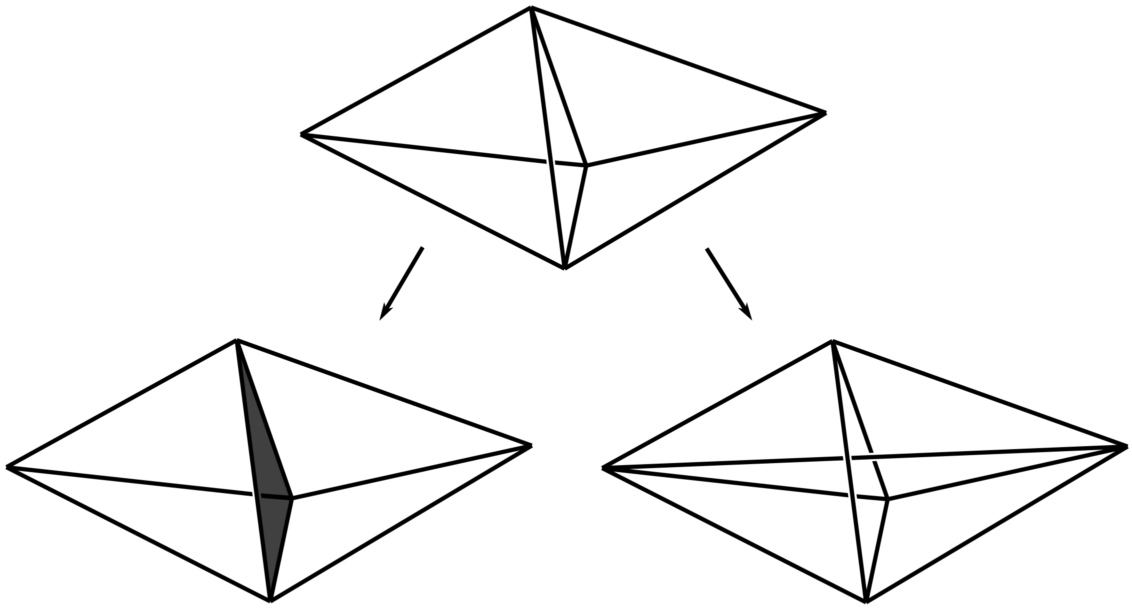

Figure 2: Two triangulations of a triangular bipyramid.

Two other Pachner moves are the $2 \to 0$ move and its inverse, the $0 \to 2$ move. The $2 \to 0$ move takes two tetrahedra which are glued together around a valence 2 edge, and collapses them onto each other. See figure 3. The union of the two tetrahedra is a "pillow". In this picture, the valence 2 edge is the vertical edge. Faces A and B belong to one tetrahedron, while A' and B' belong to the other. A is collapsed onto A', B onto B'. The $2 \to 0$ move is also called a *cancellation move*.

To do the $0 \to 2$ move, find a a pair of faces in the triangulation which have a common edge and insert a pillow there.

Note that in a geometric triangulation, the two tetrahedra of a pillow are necessarily flat. Often flat tetrahedra are undesirable, which makes the $2 \to 0$ move useful since it gets rid of them.

Pachner moves, also called *bistellar flips*, were introduced by Pachner in [33] to show that any two triangulations of a closed piece-wise linear manifold are related by some finite sequence of them. See [25] for another account of that work, or Matveev's book [29, Theorem 1.2.5]. It is difficult to obtain good bounds for the number of Pachner moves needed to get

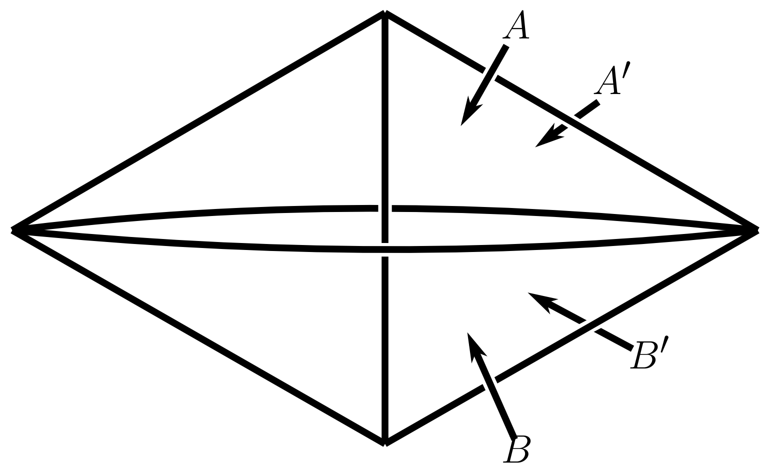

Figure 3: The two-to-zero move.

from one triangulation of a manifold to another. For instance, Mijatovic showed the number of Pachner moves needed to change a triangulation of S^3 with n tetrahedra into a standard triangulation is bounded above by $6 \cdot 10^6 n^2 2^{5 \cdot 10^4 n^2}$ [30]. Yet, experimentally it seems much fewer moves can be used [10].

3.0 Orbifold triangulations

A triangulation of an orientable 3-orbifold is usually defined to be a decomposition of the underlying space into tetrahedra such that the singular locus lies in the 1-skeleton. See [12], for instance. We introduce a more broad definition of an orbifold triangulation, which we formally define in this chapter. The idea is to allow the singular locus into the rest of the triangulation, by labelling each tetrahedron with some symmetry group and allowing faces to be glued to themselves. The underlying space of the orbifold is the quotient of the union of the tetrahedra under the equivalence relation generated by the symmetries and the face gluing maps. The local group structure is determined by the symmetries of the tetrahedra, any faces which are glued to themselves, and positive integer labels on the edges of the tetrahedra which represent orders of rotation groups.

This definition is combinatorial. We will build on it by considering an additional hyperbolic structure on the triangulation. This will just mean viewing the tetrahedra as ideal hyperbolic tetrahedra in such a way that the induced hyperbolic structure on the orbifold is complete. With this approach we can have fully ideal hyperbolic triangulations of orbifolds, i.e. triangulations by hyperbolic tetrahedra, all of whose vertices are ideal. With the more traditional definition of an orbifold triangulation, this is not always possible.

3.1 Combinatorial definitions

We work in dimension 3 and with orientable orbifolds in mind, although the definitions can be made in greater generality. In what follows, any map from one k-simplex to another (possibly the same) k-simplex is implicitly a simplicial isomorphism, i.e. really just a bijection from the set of vertices of the first simplex to the set of vertices of the second.

Definition 4. An *orientable 3-orbifold triangulation* K, which we will simply call an *orbifold triangulation*, is a tuple $K = (\mathcal{T}, \mathcal{G}, \mathcal{S}, \mathcal{E})$, where:

1. $\mathcal{T}$ is a set of finitely many oriented tetrahedra.
2. (Gluing maps) $\mathcal{G}$ is a function which assigns to each face f of each $T \in \mathcal{T}$ some $\mathcal{G}(f, T)$, which is either an orientation-reversing map from f to some other (possibly the same) face of a tetrahedron in $\mathcal{T}$, or 0 if f is not attached to anything.
3. (Symmetries) $\mathcal{S}$ is a function which assigns to each $T \in \mathcal{T}$ a group of orientation preserving symmetries $\mathcal{S}(T)$ of T. Note that this is a subgroup of the group of all even permutations of a set with 4 elements.
4. (Edge labels) $\mathcal{E}$ is a function which assigns to each edge e of each $T \in \mathcal{T}$ a positive integer $\mathcal{E}(e, T)$, which we think of as the order of a finite rotation group with axis e.

This data is required to satisfy the conditions:

(a) If $\mathcal{G}(f, T)$ maps f to a face f' of a tetrahedron T', then $\mathcal{G}(f', T')$ maps f' to f and $\mathcal{G}(f', T') = \mathcal{G}(f, T)^{-1}$.

(b) If an edge e of a tetrahedron T is mapped to e' of T' by a symmetry or face gluing map, then $\mathcal{E}(e, T) = \mathcal{E}(e', T')$.

(c) For any particular face f of a tetrahedron T, the orbit of f under $\mathcal{S}(T)$ either has exactly one face f' such that $\mathcal{G}(f', T) \neq 0$, or no such face.

(d) Suppose that for a tetrahedron T, $\mathcal{S}(T)$ contains the order 3 group of rotations fixing a face f. Suppose further that f is glued to a face f' of a tetrahedron T'. Then $\mathcal{S}(T')$ must contain the order 3 group of rotations fixing f'.

Define $\hat{X}(K)$ to be the topological space which is the quotient of the disjoint union of the tetrahedra under the equivalence relation generated by the gluing maps and symmetries.

We will sometimes describe an orbifold triangulation in terms of a *numbering*, i.e an assignment to each tetrahedron a unique number from 0 to $n-1$ (assuming there are n tetrahedra) and an assignment to each vertex of each tetrahedron a unique number from 0 to 3. Tetrahedron i we will denote as T_i, vertex i as v_i, the edge connecting v_i and v_j as $e_{i,j}$, and the face opposite v_i as f_i.

An example of an orbifold triangulation is in figure 4. It has a single tetrahedron. The tetrahedron is labelled with the full group of symmetries, f_2 is glued to itself by rotation around the blue axis while all other faces are glued to nothing, and all edge labels are 6. On

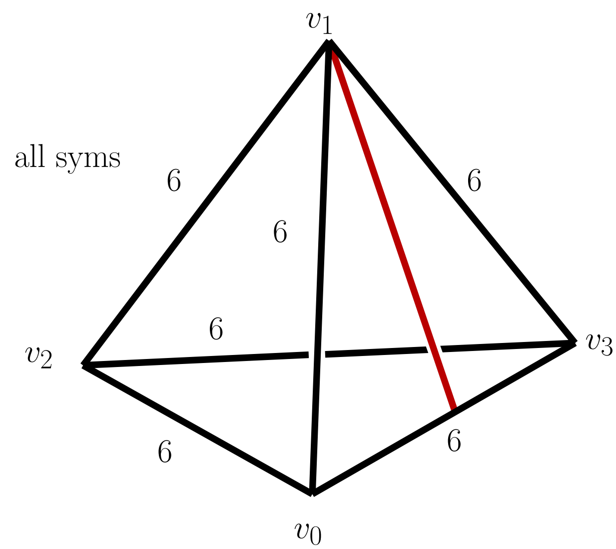

Figure 4: An example of an orbifold triangulation.

the other side of f_2 we visualize a copy of this tetrahedron, obtained by rotating the original by π around the blue axis. Although the other faces are technically not glued to anything, we think of them as being glued to other copies of the tetrahedron, obtained by applying symmetries to the tetrahedron which correspondingly move the copy attached to f_2. As in this figure, we often use the convention that a red line represents a rotation axis, either for a face gluing map or a symmetry.

Let us expand somewhat on this example. We care about the equivalence relation generated by the face gluing maps and the symmetries. It might be possible to change the face gluing data while still generating the same equivalence relation. We could make many changes to the face gluing data in figure 4 without affecting the equivalence relation. For instance, we could change how f_2 is glued to itself, say by rotation around the axis connecting v_0 to $e_{1,3}$. Or we could glue f_2 to f_1 in any orientation reversing way. Or we could glue all faces to themselves in any orientation reversing ways. Taken together with the symmetries, all these choices generate the same equivalence relation. However, in our definition of an orbifold triangulation, the latter two choices are disallowed by condition (c). This shows the point of that condition; we record as little face gluing data as possible, as long as it still gives the correct equivalence relation. This especially makes sense when working with computers. However, it is helpful sometimes to have the opposite perspective, which is to imagine all possible face gluings implied by the symmetries. In the case of figure 4, that literally is all possible face gluings.

We say K is *without boundary* if for any face f of any tetrahedron T there exists $\gamma \in \mathcal{S}(T)$ such that $G(\gamma(f), T) \neq 0$. Otherwise, it is *with boundary.* Unless stated otherwise, we assume all of our orbifold triangulations are without boundary.

Definition 5. A *combinatorial map* (or just "map") of orbifold triangulations, $F : K = (\mathcal{T}, \mathcal{G}, \mathcal{S}, \mathcal{E}) \to K' = (\mathcal{T}', \mathcal{G}', \mathcal{S}', \mathcal{E}')$, is a map from the disjoint union of the tetrahedra in $\mathcal{T}$ to the disjoint union of the tetrahedra in $\mathcal{T}'$ satisfying:

1. Suppose e is an edge of $T \in \mathcal{T}$ and $F(e) = e'$, an edge in $F(T) = T'$. Then $\mathcal{E}(e, T)$ divides $\mathcal{E}(e', T')$.
2. Suppose $F(T) = T' \in \mathcal{T}'$. Then $F\big|_T \circ \mathcal{S}(T) \circ F\big|_T^{-1} < \mathcal{S}(T')$.
3. Suppose f_1 is a face of $T_1 \in \mathcal{T}$ and $\mathcal{G}(f_1, T)$ maps f_1 to f_2 in T_2. Suppose further that $F(T_1) = T_1' \in \mathcal{T}'$ and $F(T_2) = T_2' \in \mathcal{T}'$. There must exist $\gamma_1' \in \mathcal{S}(T_1')$ and $\gamma_2' \in \mathcal{S}(T_2')$ such that for $f_1' = \gamma_1'(F(f_1))$, we have
$$\mathcal{G}(f_1', T_1') \circ \gamma_1' \circ F\big|_{f_1} = \gamma_2' \circ F\big|_{f_2} \circ \mathcal{G}(f_1, T).$$

Two combinatorial maps $F_1, F_2 : K \to K'$ are *equivalent* if for each $T \in \mathcal{T}$, $F_1(T) = F_2(T)$ and there exists $\gamma \in \mathcal{S}(F_1(T))$ such that $F_2\big|_T = \gamma \circ F_1\big|_T$.

Remark 2. With the notion of equivalence given above, the set of combinatorial maps from an orbifold triangulation to itself, i.e. the orbifold triangulation's "combinatorial automorphisms", forms a group.

$\mathcal{G}$, $\mathcal{S}$, and $\mathcal{E}$ suggest a natural way of trying to put an orbifold structure on $\hat{X}(K)$. We might actually have to remove some vertices from $\hat{X}(K)$ to get an orbifold. To explain this, we now define the "link" of a vertex of an orbifold triangulation.

Let v be a quotient vertex of $\hat{X}(K)$. Represent the set of vertices of K mapping to v under the quotient map as $\{(v_i, T_j)\}$, where the pair (v_i, T_j) corresponds to vertex v_i of tetrahedron T_j. Each (v_i, T_j) has a corresponding triangular cross section $t_{i,j}$, and the set of all triangular cross sections is naturally decorated with combinatorial data coming from Q. In particular, each vertex of $t_{i,j}$ is labelled with the corresponding edge label of T_j, if $\mathcal{S}(T_j)$ has the order 3 symmetry fixing v_i then $t_{i,j}$ is labelled with the order 3 symmetry group, and the face gluings of the tetrahedra induce edge gluings of the triangles. The quotient of $\{t_{i,j}\}$ under all symmetries and edge gluings admits a natural orbifold structure. We call this orbifold the *link* of v.

Let $X(K) = \hat{X}(K) - V$, where V is the set of quotient vertices of $\hat{X}(K)$ whose links do not admit spherical geometric structures.

Proposition 1. *For K an orbifold triangulation without boundary, $X(K)$ admits a natural orientable orbifold structure inherited from K. A combinatorial map $f : K \to K'$ induces an orbifold covering map $X(F) : X(K) \to X(K')$.*

Proof. Let x be a point in a tetrahedron T of K. We can define a chart around the image of x in $X(K)$, depending on whether x lies in the interior of T, the interior of a face, the interior of an edge, or is a vertex with a spherical link. In all cases, the local group actions are orthogonal.

If x is in the interior of T, we may choose a small open neighborhood $U \subset T$ of x such that for all $\gamma \in \mathcal{S}(T)$, if $\gamma(x) = x$ then $\gamma(U) = U$, and otherwise $\gamma(U) \cap U = \emptyset$. The projection of this U to $X(K)$ gives an orbifold chart, whose local group is the stabilizer of x in $\mathcal{S}(T)$.

Now suppose x lies in the interior of a face f of T. Because K is without boundary, f must either be glued to itself or some other face, perhaps after applying a symmetry of T. The kind of chart we take for x depends on whether f is glued to itself or not, and where x is located in f. For example, the case of largest local group is when f is glued to itself, $\mathcal{S}(T)$ has the order 3 symmetries rotating f, and x is the fixed point of these rotations. Then a chart corresponding to x is an open ball with the order 6 dihedral group acting on it, with fixed point x. The other cases are examined similarly.

If x is in the interior of an edge e of T, there are again several possibilities for a chart. Of course if the edge label $\mathcal{E}(e, T) = k > 1$, then the local group of x should at least contain order k rotations. It could also be the case that a face of T is glued to itself by a face gluing map which restricts to an involution of e, having a unique fixed point p, or $\mathcal{S}(T)$ could have the order two symmetry which preserves e, restricting to e as that same involution. Or it could be that some other tetrahedron which has an edge glued to e has such a face gluing map or symmmetry. In this case, if x is p we can make a chart corresponding to x whose local group is a dihedral group of order $2k$. If x is not p, its local group will just be the order k rotations, regardless of these involutions.

Suppose x is a vertex of T with a spherical link. Then, as is standard, we can choose a chart for x whose local group is the fundamental group of the link. In particular, by Theorem 4, the link is the quotient of $\mathbb{S}^2$ by a finite group of isometries, which naturally extend over the unit ball to give a local model.

It is immediate from the definition that a combinatorial map $f : K \to K'$ induces a continuous map $X(f) : X(K) \to X(K')$. By the construction given above of the natural orbifold structures for $X(K)$ and $X(K')$, $X(f)$ must be a covering map. $\square$

From now on, we view $X(K)$ as an orbifold, rather than just a topological space.

3.2 Polyhedral decompositions and philosophy of local modifications

We can naturally extend the definition of an orbifold triangulation to an *orbifold polyhedral decomposition*, as a set of polyhedra with face gluing data, symmetries, and edge labels. Just as with triangulations, there is a notion of a combinatorial map and a quotient orbifold. As the definition of an orbifold polyhedral decomposition is similar to the definition of an orbifold triangulation, given in in the previous section, we will not write it down in more detail.

We might try to turn a polyhedral decomposition into a triangulation by subdividing each polyhedron into tetrahedra. For example, in figure 5 we have a triangular bipyramid labelled with the order 3 symmetry group fixing the opposite vertices, and we show two natural subdivisions. Recall from Section 2.5 that they correspond to the $2 \to 3$ and $3 \to 2$ Pachner moves. For the decomposition into two tetrahedra, we must give each of the two new tetrahedra the corresponding order 3 symmetry group. For the decomposition into three tetrahedra, all three of them belong to the same orbit under the action of the symmetry

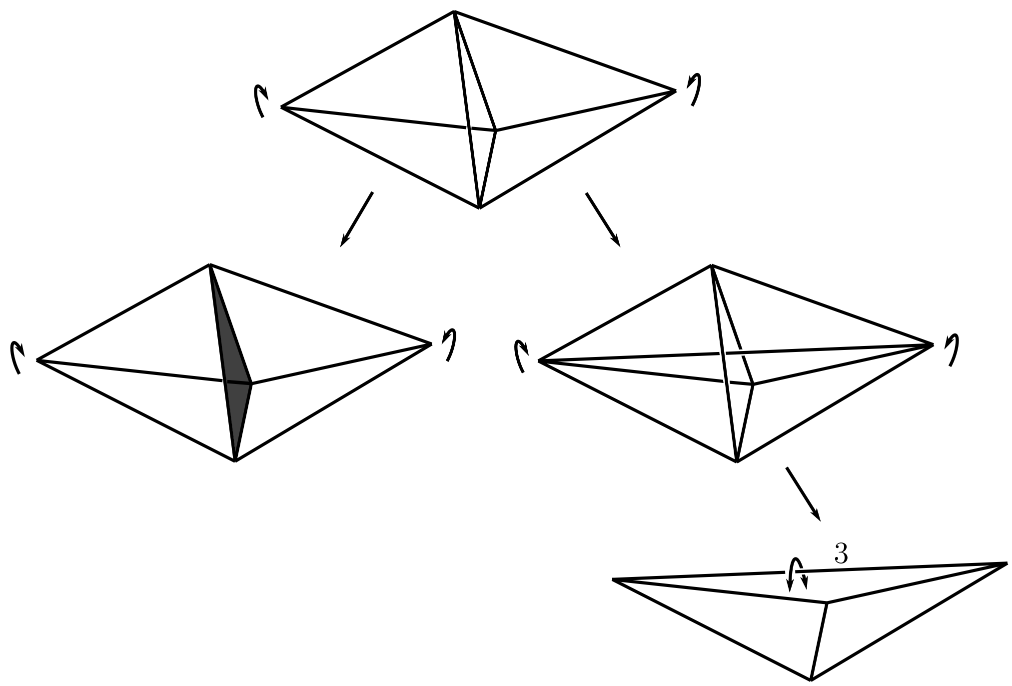

Figure 5: Triangulating a polyhedron which has symmetries.

group. Hence, we should just take one of these tetrahedra, label the edge in the interior of the bipyramid with the edge label 3, and glue the two faces adjacent to that edge as indicated in the figure.

In order for this to make sense, it is important that the subdivision be invariant with respect to the symmetries of the polyhedron. In figure 6, if the square base has no symmetries (or even an order 2 symmetry), then the indicated subdivision is possible. But if it is labelled with the order 4 symmetry, then no triangulation of it is invariant, if the triangulation is required to have the same vertex set. If we are okay with introducing new vertices, then the barycentric sub-division of a polyhedron is always invariant with respect to any symmetries. The newly introduced vertices are finite, not ideal, which might not be desirable.

In both examples, the result is a different representation of the same orbifold. We will often locally modify an orbifold triangulation in this way:

1. Recognize some part of the triangulation as an invariant triangulation of a polyhedron

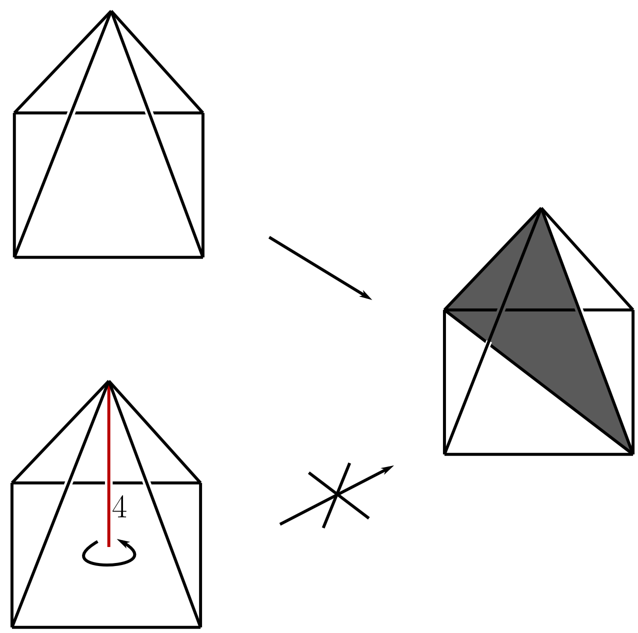

Figure 6: It may not be possible to invariantly triangulate a polyhedron.

with symmetries.

2. Triangulate that polyhedron in a different way. We may or may not choose to change the vertex set.
3. The result is a new orbifold triangulation of the same orbifold.

This is a general description of an orbifold Pachner move, which we will discuss in more detail in Chapters 5 and 6.

3.3 Hyperbolic orbifold triangulations

We can view an orbifold triangulation K as "geometric" if we assign to each tetrahedron the structure of some hyperbolic ideal tetrahedron. However, we are not guaranteed that the hyperbolic tetrahedra consistently glue up to induce a well-defined and complete hyperbolic structure on $X(K)$. We could try to define some kind of consistency and completeness equations in analogy with the manifold case, which would hopefully guarantee a complete hyperbolic structure on $X(K)$. However, for us it is more convenient and on theme to work with "geometric" orbifold triangulations arising from ideal tiling-group pairs (defined in Section 2.4.3).

Let $(\mathcal{T}, \Gamma)$ be an ideal tiling-group pair with Γ finite co-volume and orientation-preserving.

Lemma 1. *The action of Γ on the set of polyhedra of $\mathcal{T}$ has finitely many orbits.*

Proof. Let $Q = \mathbb{H}^3/\Gamma$, let $p : \mathbb{H}^3 \to Q$ be the universal covering map, and let $Q_{\geq\epsilon}$ be the ϵ-thick part of Q, where ϵ is any positive number smaller than the Margulis constant. Because $Q_{\geq\epsilon}$ is compact, there is a compact ball $B \subset \mathbb{H}^3$ such that $Q_{\geq\epsilon} \subset p(B)$. For every polyhedron P of $\mathcal{T}$, we must have that $p(P) \cap Q_{\geq\epsilon} \neq \emptyset$. This means there must exist $\gamma \in \Gamma$ such that $\gamma(P) \cap B \neq \emptyset$. Because $\mathcal{T}$ is locally finite, the set of all polyhedra intersecting B is finite, so we conclude that the action of Γ on $\mathcal{T}$ must have finitely many orbits. $\square$

Choose one polyhedron from each of these finitely many orbits, $P_1, \ldots, P_n$. The group Γ gives instructions to define an orbifold polyhedral decomposition K such that $X(K)$ is $\mathbb{H}^3/\Gamma$. Each face of each P_i is adjacent to some polyhedron which is mapped to a P_j by an element of Γ, which gives the face gluing data. The stabilizer in Γ of each P_i is a finite group which becomes the symmetry group of P_i. The orders of the stabilizers in Γ of each edge of each P_i become the edge labels. K is uniquely determined by $(\mathcal{T}, \Gamma)$ up to combinatorial isomorphism. Any orbifold polyhedral decomposition K which arises from an ideal tiling-group pair in this way we call a *hyperbolic orbifold polyhedral decomposition*. If the ideal tiling is actually a triangulation, hence K is an orbifold trangulation, then we call K a *hyperbolic orbifold triangulation*.

Note that if all edges are labeled 1, all symmetry groups are trivial, and no face is glued to itself, then a hyperbolic orbifold triangulation is a geometric ideal triangulation of a manifold. For that reason, any geometric ideal triangulation of a finite-volume cusped hyperbolic 3-manifold is an example of a hyperbolic orbifold triangulation. For another example, take the tetrahedron in figure 4 to be regular. This turns out to be the orbifold $\mathbb{H}^3/\operatorname{PGL}_2(O_3)$, which we will have more to say about in Chapter 7.

We should emphasize that, in a hyperbolic orbifold triangulation, the symmetries a tetrahedron is labelled with will necessarily be isometries. So, for example, a tetrahedron which is not regular cannot be labelled with all possible combinatorial symmetries. On the other hand, the symmetry group attached to the tetrahedron need not be the full group of isometries of the tetrahedron, e.g. a regular tetrahedron might only be labelled with the trivial group.

There is a correspondence between isometries preserving tiling-group pairs and combinatorial maps.

Proposition 2. *Let $(\mathcal{T}, \Gamma)$ and $(\mathcal{T}', \Gamma')$ be finite co-volume orientable ideal tiling-group pairs with corresponding hyperbolic orbifold polyhedral decompositions K and K'. Let $g \in \mathrm{Isom}(\mathbb{H}^3)$ be an isometry which maps $\mathcal{T}$ to $\mathcal{T}'$ and conjugates Γ to a subgroup of Γ'. Then g induces a combinatorial map $f : K \to K'$ such that g is a lift of $X(f)$. Conversely, every combinatorial map $f : K \to K'$ is induced by such an isometry g. In the specific case that $\mathcal{T} = \mathcal{T}'$ and $\Gamma = \Gamma'$, the group of combinatorial automorphisms of K is isomorphic to G/Γ, where G is defined by*

$$G = \{g \in \mathrm{Isom}(\mathbb{H}^3) : g \textit{ preserves } \mathcal{T} \textit{ and normalizes } \Gamma\}.$$

Proof. Let $(\mathcal{T}, \Gamma), (\mathcal{T}', \Gamma'), K, K'$ and g be as in the statement of the proposition. Because g maps $\mathcal{T}$ to $\mathcal{T}'$ and conjugates Γ to Γ', each Γ-orbit of $\mathcal{T}$ maps to a Γ'-orbit of $\mathcal{T}'$. Hence g induces a map from the tetrahedra of K to the tetrahedra of K'. Again using that g conjugates Γ to Γ', we easily see that this satisfies the criteria for being a combinatorial map, which we call f. By construction, g is a lift of $X(f)$.

Now suppose $f : K \to K'$ is a combinatorial map. Then $X(f) : X(K) \to X(K')$ is a covering map. So we can pull back the hyperbolic structure on $X(K') = \mathbb{H}^3/\Gamma'$ to a hyperbolic structure on $X(K) = \mathbb{H}^3/\Gamma$. By Mostow-Prasad rigidity, this hyperbolic structure must be isometric to the original hyperbolic structure on $\mathbb{H}^3/\Gamma$, so there must be an isometry $g \in \mathrm{Isom}(\mathbb{H}^3)$ such that g conjugates Γ to a subgroup of Γ' and g is a lift of $X(f)$. This implies that g maps $\mathcal{T}$ to $\mathcal{T}'$.

To prove the last part of the proposition, let A be the group of combinatorial automorphisms of K. By the first part of this proof there is a surjective map $G \to A$. Clearly this map is a group homomorphism and its kernel is Γ, so we are done. $\square$

4.0 Defining the canonical decomposition

Let M be a finite-volume cusped hyperbolic 3-manifold with p cusps. Let $C_1, \ldots, C_p$ be a choice of cusp neighborhoods for those cusps. In [15], Epstein and Penner gave a construction which associates to $C_1, \ldots C_p$ a particular decomposition of M into hyperbolic ideal polyhedra, which we call the *Epstein-Penner decomposition corresponding to* $C_1, \ldots, C_p$. Different choices of cusp neighborhoods can in general give different decompositions, although we make the following remarks.

1. Uniformly shrinking or expanding the cusp neighborhoods gives the same decomposition. For that reason, we think of there being a $(p-1)$-parameter, rather than p-parameter, family of Epstein-Penner decompositions of M. In particular, if M has one cusp then it has exactly one Epstein-Penner decomposition.
2. For $p > 1$, there are infinitely many choices of the $p-1$ parameters. However, they do not correspond to all different decompositions. There are only finitely many possible Epstein-Penner decompositions which can arise, see [5].

The Epstein-Penner decomposition resulting from choosing all cusp neighborhoods to have the same volume is called the *canonical decomposition*. It is uniquely determined by the hyperbolic metric, since any dependence on cusp neighborhoods has been thrown out. The self-isometry group of a finite-volume cusped hyperbolic 3-manifold can be read off

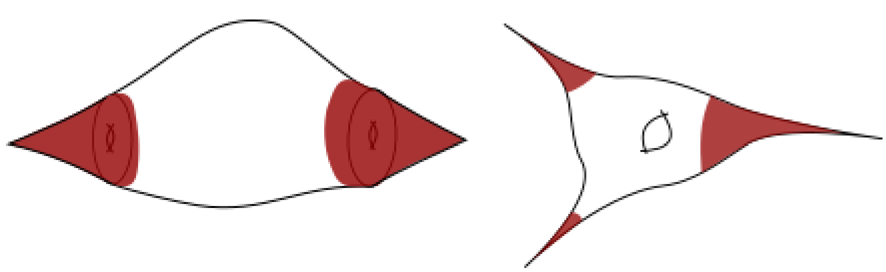

Figure 7: Cusp neighborhoods of cusped hyperbolic manifolds.

from its canonical decomposition, as it is just the decomposition's group of combinatorial automorphisms. Furthermore, two such manifolds are isometric if and only if they have the same canonical decomposition.

The goal of Section 4.1 is to describe Epstein-Penner's construction, and in the process to show that the construction still makes sense for orbifolds by proving the following theorem.

Theorem 7. *Let Q be a finite-volume cusped hyperbolic 3-orbifold. The Epstein-Penner construction still applies to Q, and the canonical decomposition it provides is a hyperbolic orbifold polyhedral decomposition. The self-isometry group of Q is equal to the group of combinatorial automorphisms of this decomposition. Two orbifolds are isometric if and only if they have combinatorially isomorphic canonical decompositions.*

In Section 4.2, we define the *tilt* of a face of a tetrahedron, following Weeks. We can use tilts to determine if a given hyperbolic orbifold triangulation is the canonical decomposition or not. To account for the fact that the canonical decomposition might not be a triangulation, we define *proto-canonical triangulations* in Section 4.3. In that case, we would like to work with a naturally defined sub-triangulation of the canonical decomposition, which we define in Section 4.4.

4.1 The convex hull construction

The construction is valid in any dimension, but our outline will be for dimension 3. We work in the hyperboloid model of Minkowski space. A good introduction to this model is in Ratcliffe's book [35]. Four-dimensional Minkowski space, M^4, is $\mathbb{R}^4$ equipped with the bilinear form

$$\langle x, y\rangle = x_0y_0 + x_1y_1 + x_2y_2 - x_3y_3.$$

The hyperboloid

$$\{x \in M^4 : \langle x, x\rangle = -1\}$$

has two sheets, one for $x_3 \geq 1$, and the other for $x_3 \leq -1$. Denote the upper sheet as $\mathbb{H}^3$. It is a model for 3-dimensional hyperbolic space if we give it the Riemannian metric induced by $\langle \cdot, \cdot \rangle$ (the bilinear form induces a positive definite inner product on each tangent space of the upper sheet).

The *light cone* L is defined to be

$$L = \{x \in M^4 : \langle x, x \rangle = 0.\}$$

The *positive light cone* L^+ is the set of points of L with $x_3 > 0$. To every point $v \in L^+$ there corresponds a horoball $\{x \in \mathbb{H}^3 : 0 > \langle x, v \rangle \geq -1\}$, and this correspondence is a bijection.

The Lie group of linear isomorphisms of M^4 preserving $\langle \cdot, \cdot \rangle$ is called $\mathrm{O}(3,1)$. Of course every such map either takes $\mathbb{H}^3$ to itself or swaps it with the lower sheet. Let $\mathrm{O}^+(3,1)$ be the subgroup which preserves $\mathbb{H}^3$. Then $\mathrm{O}^+(3,1)$ is actually equal to the isometry group of $\mathbb{H}^3$. Let $\mathrm{SO}^+(3,1) < \mathrm{O}^+(3,1)$ be the subgroup which preserves the orientation of M^4. Then it will also preserve the orientation of $\mathbb{H}^3$ and so is equal to $\mathrm{Isom}^+(\mathbb{H}^3)$.

Let $\Gamma < \mathrm{O}^+(3,1)$ be discrete such that $\mathbb{H}^3/\Gamma$ has finite volume. Then Γ acts in a nice, discrete, i.e. properly discontinuous, way on $\mathbb{H}^3$, but acts ergodically on $\partial\mathbb{H}^3$, which can be thought of in this model as the projectivization of L^+. In general, Γ also acts in a non-discrete, "heavily mixing" way on L^+, however we do have the following surprising fact.

Theorem 8 ([15], Thm 2.4)**.** *Suppose $v \in L^+$ has non-trivial stabilizer in Γ. Then the orbit of v under Γ is discrete in M^4 and does not accumulate at* 0.

Such a v must be fixed by some parabolic isometry in Γ. So in that case $\mathbb{H}^3/\Gamma$ has at least one cusp, and in fact the horoball $\{x \in \mathbb{H}^3 : 0 > \langle x, v \rangle \geq -1\}$ projects to a cusp neighborhood of that cusp (perhaps not embedded). For any $t > 0$, tv will also be fixed by that parabolic isometry, and its horoball will project to another neighborhood of that cusp, different from the one for v if $t \neq 1$. In particular, the area of the boundary of the cusp neighborhood is scaled by $1/t^2$. In summary, a cusp neighborhood in $\mathbb{H}^3/\Gamma$ corresponds to the orbit under Γ of some $v \in L^+$, and that orbit is as described by Theorem 8.

Assume from now on that $\mathbb{H}^3/\Gamma$ has finite volume and at least one cusp. Label the cusps $C_1, \ldots, C_p$. Each C_i corresponds to a Γ-orbit in the light cone, $\Gamma.v_i$. By Theorem 8, the set V defined by $V = \cup_i(\Gamma.v_i)$ is discrete and does not accumulate at 0. Let C be the closed convex hull of V.

Theorem 9 (Proposition 3.5, [15]). *The boundary of C in M^4 is the union of $\{tv : v \in V, t \geq 1\}$ and a countable set of 3-dimensional faces $F_1, F_2, \ldots$. Each F_i is the convex hull of a finite number of points in V. The set of faces $\{F_i\}$ is locally finite in the interior of the light cone. The affine hull A_i of F_i is Euclidean, and the intersection of A_i with L^+ is spherical with respect to the Euclidean structure on A_i.*

Now the idea is to "project" ∂C to $\mathbb{H}^3$, mapping each F_i to an ideal polyhedron and resulting in a Γ-invariant tiling of $\mathbb{H}^3$, i.e. an ideal tiling-group pair. We record this more precisely as follows.

Fact 1. *For a fixed i, let $v_0, v_1, \ldots, v_n \subset V$ be the points whose convex hull is F_i. The set $F_i - V$ projects injectively along rays through the origin to a convex ideal polyhedron P_i in $\mathbb{H}^3$ whose ideal vertices are the projective classes of the v_j. It follows from the fact that the affine hull A_i of F_i is Euclidean that P_i cannot be flat. The projection of each $F_i - V$ for all i results in a Γ-invariant ideal tiling of $\mathbb{H}^3$ by non-flat ideal polyhedra.*

We call this the *canonical Γ-invariant tiling with respect to $C_1, \ldots, C_p$*. Since shrinking or expanding all these neighborhoods by the same amount corresponds to multiplying every element of V by the same scalar, which does not change the tiling, we can assume that one cusp neighborhood is fixed. Thus, if $\mathbb{H}^3/\Gamma$ has p cusps, then there is a $p-1$ parameter family of Γ-invariant tilings. The tiling resulting from choosing all cusp neighborhoods to have the same volume we simply call *the canonical Γ-invariant tiling*.

If Γ acts freely on $\mathbb{H}^3$, i.e. $\mathbb{H}^3/\Gamma$ is a manifold, then the universal covering map $\mathbb{H}^3 \to \mathbb{H}^3/\Gamma$ restricts to an embedding on the interior of each polyhedron of a canonical tiling. Hence the tiling descends to a decomposition of the manifold. This is the Epstein-Penner decomposition of $\mathbb{H}^3/\Gamma$ corresponding to $C_1, \ldots, C_p$..

If Γ does not act freely, then $\mathbb{H}^3/\Gamma$ is not a manifold, but an orbifold with non-trivial singular locus, and the canonical tiling does not necessarily map down to a decomposition of the underlying space of $\mathbb{H}^3/\Gamma$. For instance, Γ could contain an isometry which maps one of the polyhedra in the tiling to itself non-trivially, hence the restriction of the universal covering map to that polyhedron is not an embedding. In the case that Γ does not act freely, we still have the $p-1$ parameter family of Γ-invariant canonical tilings, but in what sense do they give decompositions of the quotient orbifold?

As discussed in Section 3.3, an ideal tiling-group pair corresponds to a unique orbifold polyhedral decomposition. So, let us call the orbifold polyhedral decomposition corresponding to the canonical Γ-invariant tiling the *canonical decomposition* of $\mathbb{H}^3/\Gamma$. This proves the first part of Theorem 7. To prove the rest, we use the following lemma.

Lemma 2. *Let $Q = \mathbb{H}^3/\Gamma, Q' = \mathbb{H}^3/\Gamma'$ be finite-volume cusped hyperbolic 3-orbifolds. Let $\mathcal{T}$ and $\mathcal{T}'$ be the canonical invariant tilings of Γ and Γ' respectively. For any isometry $g \in \mathrm{Isom}(\mathbb{H}^3)$ lifting an isometry $Q \to Q'$, g must map $\mathcal{T}$ to $\mathcal{T}'$.*

Proof. Fix cusp neighborhoods of Q and Q' all of the same volume. Lift the cusp neighborhoods of Q to a horoball packing H of $\mathbb{H}^3$, and lift the cusp neighborhoods of Q' to a horoball packing H' of $\mathbb{H}^3$. Because g is a lift of an isometry $Q \to Q'$, and because the original cusp neighborhoods all had the same volume, g must map H to H'. Hence it maps the points on the light cone corresponding to H to the points on the light cone corresponding to H', so it preserves their convex hulls and therefore maps $\mathcal{T}$ to $\mathcal{T}'$. $\square$

Let K and K' be the canonical decompositions of finite-volume cusped hyperbolic 3-orbifolds Q and Q' respectively. Suppose there is a combinatorial isomorphism from K to K'. By Proposition 2, this must be an isometry. Conversely, suppose Q is isometric to Q'. The isometry must lift to $\mathbb{H}^3$, and by the previous lemma, it must preserve the canonical tilings. Hence, by Proposition 2, the original isometry of orbifolds must have actually been a combinatorial map from K to K', hence K and K' are combinatorially isomorphic.

Now we want to show the self-isometry group of $Q = \mathbb{H}^3/\Gamma$ is equal to the group of combinatorial automorphisms of K. By Proposition 2 again, the group of combinatorial automorphisms of K is isomorphic to G/Γ, where $G \subset \mathrm{Isom}(\mathbb{H}^3)$ is the group of isometries which normalize Γ and preserve the tiling corresponding to K. As K is the canonical decomposition, by the lemma every isometry of $\mathbb{H}^3$ lifting a self-isometry of Q must preserve the tiling corresponding to K. Hence G is the full normalizer of Γ. Therefore the group of combinatorial automorphisms of K is isomorphic to the normalizer of Γ mod Γ, but that is exactly the self-isometry group of Q.

4.2 Tilt

For motivation for the following definitions, see Section 3 of the paper by Weeks [43].

Let v_0, v_1, v_2, v_3 be points on L^+ which are linearly independent (equivalently, we suppose the ideal tetrahedron in $\mathbb{H}^3$ whose vertices v_0, v_1, v_2, v_3 correspond to is not flat). Denote their convex hull $\mathrm{conv}\{v_0, v_1, v_2, v_3\}$ as T. Let A be the affine hull of T, i.e. $A = v_0 + \mathrm{span}\{v_1 - v_0, v_2 - v_0, v_3 - v_0\}$. Define $p_T \in M^4$ to be the unique vector satisfying $\langle p_T, v\rangle = -1$ for all $v \in A$. For each i, let f_i be the face opposite v_i. Define $n_{T,f_i} \in M^4$ to be the unit vector orthogonal to $\mathrm{span}(\{v_0, v_1, v_2, v_3\} \setminus \{v_i\})$. and satisfying $\langle n_{T,f_i}, v_i\rangle < 0$.

Definition 6. The *tilt of T relative to f_i* is $\langle p_T, n_{T,f_i}\rangle$.

Suppose $\Gamma < \mathrm{Isom}^+(\mathbb{H}^3)$ is discrete and finite co-volume, $C_1, \ldots, C_p$ are fixed cusp neighborhoods of $\mathbb{H}^3/\Gamma$, and $\mathcal{T}$ is a Γ-invariant tiling of $\mathbb{H}^3$ by hyperbolic ideal tetrahedra. Let $V \subset L^+$ be the points on the light cone corresponding to $C_1, \ldots, C_p$, as defined in the previous section. There is a bijection between the ideal vertices of $\mathcal{T}$ in $\partial\mathbb{H}^3$ and the set V. For each tetrahedron T of $\mathcal{T}$, define the *V-lift* of T to be $\mathrm{conv}\{v_0, v_1, v_2, v_3\}$, where v_0, v_1, v_2, v_3 are the elements of V corresponding to the vertices of T. Define the *tilt* of a tetrahedron T of $\mathcal{T}$ relative to one of its faces f to be the tilt of the V-lift of T relative to the V-lift of f. Denote this number as $\mathrm{tilt}(T, f)$.

Theorem 10 (Weeks)**.** *$\mathcal{T}$ is the canonical Γ-invariant tiling with respect to $C_1, \ldots, C_p$ if for any tetrahedra T, T' of $\mathcal{T}$ having a shared face, f, we have*

$$\mathrm{tilt}(T, f) + \mathrm{tilt}(T', f) < 0.$$

If all tilt sums are non-positive and at least one is 0, then the canonical tiling has some polyhedra which are not tetrahedra. The canonical tiling is obtained from $\mathcal{T}$ by removing all faces with tilt sum equal to 0.

Each cusp neighborhood C_i of $\mathbb{H}^3/\Gamma$ lifts to horoballs in $\mathbb{H}^3$. The boundary of such a horoball, a horosphere, intersects each tetrahedron of $\mathcal{T}$ in a *vertex cross section.* This is just a triangle, and the metric induced on it by the hyperbolic metric is Euclidean. Weeks

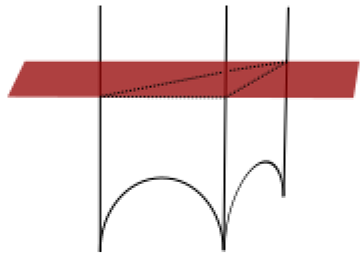

Figure 8: The horosphere intersects the ideal tetrahedron in a vertex cross section.

showed that we can compute the tilts of a tetrahedron in terms of its geometry and the geometry of its vertex cross sections.

Theorem 11 ([43], Thm 5.1)**.** *Let T be a tetrahedron of $\mathcal{T}$. The tilt of T relative to each of its faces may be computed as*

$$\begin{pmatrix} t_0 \\ t_1 \\ t_2 \\ t_3 \end{pmatrix} = \begin{pmatrix} 1 & -\cos\theta_{01} & -\cos\theta_{02} & -\cos\theta_{02} \\ -\cos\theta_{10} & 1 & -\cos\theta_{12} & -\cos\theta_{13} \\ -\cos\theta_{20} & -\cos\theta_{21} & 1 & -\cos\theta_{23} \\ -\cos\theta_{30} & -\cos\theta_{31} & -\cos\theta_{32} & 1 \end{pmatrix} \begin{pmatrix} R_0 \\ R_1 \\ R_2 \\ R_3 \end{pmatrix}$$

where t_i is the tilt relative to the face opposite vertex i, R_i is the circumradius of vertex cross section i, and θ_{ij} is the dihedral angle of the edge from vertex i to vertex j.

Warning: recall from the beginning of this section that tilt is not defined for flat tetrahedra. Hence, we cannot compute the tilts of flat tetrahedra. This is even reflected in the formula above, as a flat tetrahedron has flat triangles as vertex cross sections, which have “infinite circumradius”. However, that is okay, because a canonical tiling cannot have flat tetrahedra, as we noted after Theorem 9. This means that if we want to check if a tiling-group pair is canonical, and it has some flat tetrahedra, then checking tilts is not necessary—we already know it is not canonical. Nevertheless, flat tetrahedra do play an important role when hunting for canonical decompositions, which we will see very soon.

4.3 Proto-canonical triangulations

In this section, we fix an ideal tiling-group pair $(\mathcal{T}, \Gamma)$ with $\mathcal{T}$ a triangulation, Γ orientation-preserving and finite co-volume, and K the associated hyperbolic orbifold triangulation. Up to now, we have thought of tilt as a quantity associated to a face of a tetrahedron of $\mathcal{T}$. We can also work in K. We would like to know, just from looking at K, if $\mathcal{T}$ is the canonical Γ-invariant tiling. First, construct equal-area cusp cross-sections for K.

1. Pick an arbitrary vertex v of an arbitrary tetrahedron T of K. Place a cross section at that vertex, i.e. choose any triangle from the Euclidean triangle similarity class of vertex cross sections at v. This triangle will determine the other cusp cross sections for this cusp, as follows.
2. Any vertex of T which is mapped to v by a symmetry of T is assigned its cusp cross section, by pulling back the cross section of v using the symmetry. Suppose a face f of T, containing v, is glued to a face f' of a tetrahedron T'. This determines a vertex cross section for the vertex v' of T' which is mapped to v by this face gluing.
3. Determine the other vertex cross sections for this cusp in a similar way, using symmetries and face gluing maps. After building all cross sections for this cusp, do the same for all other cusps. The result is a choice of cusp neighborhoods, in terms of the data of its boundary as a collection of Euclidean triangles for each vertex of K.
4. Because of our arbitrary choice of an initial cross section for each cusp, the cusp neighborhoods created do not necessarily have the same volume. Equivalently, the areas of the boundaries of the cusps are not necessarily the same. Compute the areas of each cusp boundary and scale so that all cusp boundaries have area 1 (or some other number, if convenient). For a manifold, the area of the cusp boundary is just the sum of the areas of all the vertex cross sections. For an orbifold, we have to be mindful of the symmetries. If a vertex lies in a tetrahedron having the order 3 symmetry fixing that vertex, we must divide the area of its vertex cross section by 3. Similarly, only one vertex from each orbit of symmetries contributes area to the cusp boundary.

Once we have computed vertex cross sections bounding equal-volume cusp neighborhoods with the method just described, we can compute tilts using Theorem 11. Let f be a quotient face of K. Let f_i be a face of a tetrahedron T_j such that f_i maps to f in the quotient map. Glued to f_i is some face $f_{i'}$ of a tetrahedron $T_{j'}$. Define the *tilt sum* of f to be $\text{tilt}(T_j, f_i) + \text{tilt}(T_{j'}, f_{i'})$.

Definition 7. Let f be a quotient face of K. Say that f is *concave/convex/transparent* if its tilt sum is positive/negative/zero. Let e be a quotient edge of K. Say that e is *concave/convex/transparent* if every face containing e is concave/convex/transparent.

By Section 4.2, we have:

Remark 3. If every quotient face of K is convex, then K lifts to the canonical Γ-invariant tiling, so K is the canonical decomposition of $\mathbb{H}^3/\Gamma$.

If the canonical decomposition is not a triangulation, then of course any given hyperbolic orbifold triangulation could not be the canonical decomposition. However, it could be a sub-division of the canonical decomposition. In the manifold case, the authors of [16] call such a triangulation a *geometric proto-canonical triangulation*. Before we can define "proto-canonical" for orbifolds, we need to highlight the kind of flat tetrahedra which we allow to appear.

Definition 8. Let T_0 be a flat tetrahedron in the hyperbolic orbifold triangulation K. We say T_0 is *admissible* if either of the following two conditions are satisfied.

1. T_0 has an order 2 symmetry group and $\text{tilt}(T_1, f) + \text{tilt}(T_2, f') < 0$, where T_1, T_2, f, f' are described as follows. Up to re-numbering the vertices of T_0 and adjusting its face gluing data using the symmetry, T_0 is glued to the non-flat tetrahedron T_1 along vertices v_0, v_2, and v_3, and T_0 is glued to the non-flat tetrahedron T_2 along vertices v_0, v_1, v_2, as in figure 9. Let f be the face of T_1 glued to f_1 of T_0 and f' the face of T_2 glued to f_0 of T_0.
2. T_0 has an order 4 symmetry group and $\text{tilt}(T_1, f) < 0$, where T_1 and f are defined as in the previous condition.

We must accept the possibility that, for a given canonical Γ-invariant tiling, any other Γ-invariant tiling by ideal tetrahedra which is a subdivision of the canonical tiling has flat tetrahedra. For example, suppose a polyhedron in the canonical tiling is an ideal cube, and Γ contains an order 2 rotation γ which maps a face of the cube to itself by rotation around

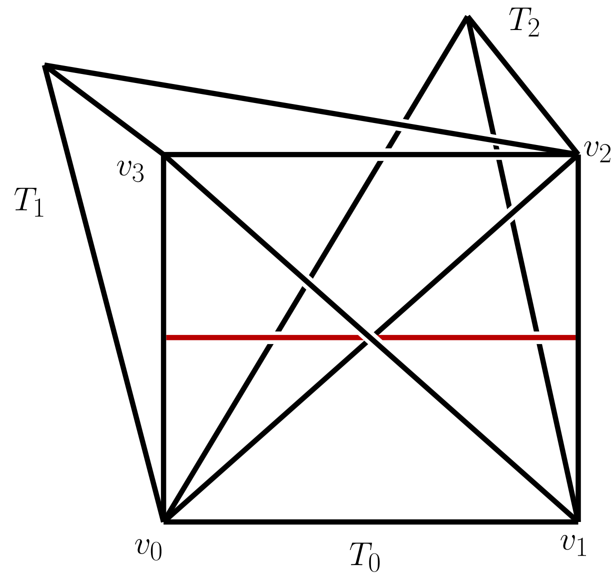

Figure 9: The setup for an admissible flat tetrahedron.

an axis connecting a pair of opposite edges of the face. An ideal triangulation of the cube gives a triangulation of the face, with the same four vertices, and this triangulation cannot possibly be invariant with respect to γ. However, we could insert a flat tetrahedron for the face of the cube, with an order 2 symmetry corresponding to γ, similar to figure 9. This is an admissible flat tetrahedron, and it fixes the problem of being invariant with respect to γ.

Definition 9. We say a hyperbolic orbifold triangulation K is *proto-canonical* if:

1. Every quotient face of $X(K)$ which does not belong to a flat tetrahedron is convex or transparent, and
2. Every flat tetrahedron is admissible.

We now make a remark about how this definition connects to the definition of proto-canonical in [16]. By our definition, if $X(K)$ is a manifold, then K cannot have any flat tetrahedra at all if it is proto-canonical. By contrast, the authors of [16] allow proto-canonical triangulations to have some flat tetrahedra. Hence our definition of “proto-canonical” is actually a generalization of their definition of “geometric proto-canonical”.

4.4 The canonical re-triangulation

Suppose the canonical decomposition is not a triangulation. We would still like to view it as a triangulation in some way, because this is more convenient when it comes to programming. For that reason, we define the *canonical re-triangulation* of the canonical decomposition. This is the generalization to orbifolds of the definition for manifolds given in [16]. The point is that a proto-canonical triangulation is some arbitrary sub-division of this canonical decomposition, and we would like to have a canonical sub-divison instead.

Let P be a convex hyperbolic ideal polyhedron with an assigned finite group of orientation preserving isometries Γ of it. We define a *barycenter* of P to be any point in the interior of P which is fixed by every element of Γ. Such a point must exist, see Part 2, Chapter 2, Proposition 2.7 of [9]. If Γ is trivial, then we allow the barycenter to be any point in the interior of P.

Definition 10. Denote the polyhedra of the canonical decomposition as $P_1, \ldots, P_n$. For each P_i, introduce a finite vertex at a barycenter, then cone the boundary of P_i to this vertex. We now have a sub-division of P_i which is invariant with respect to the symmetries of P_i. Some of the cells of this sub-division may be mapped to each other by symmetries. Take one cell from each orbit, call these choices $C_{i,1}, \ldots,$ and note that each $C_{i,j}$ has exactly one face which is an external face of P_i. This external face may be glued to itself, or to another face of P_i, or to a face of some other cell. So there is some $C_{k,\ell}$ such that it and $C_{i,j}$ are glued to each other at this external face (possible after applying a symmetry of P_k). Remove this external face to get a cell which is just the union of $C_{k,\ell}$ and $C_{i,j}$, which we call a *diamond*. See figure 10. The diamond inherits symmetries from $C_{k,\ell}$ and $C_{i,j}$, and if $k = i$ and $\ell = j$ then it inherits the symmetry corresponding to the face gluing map. Over all such pairs of $C_{k,\ell}$ and $C_{i,j}$, create all diamonds. We call the result the *diamond sub-division* of the canonical decomposition. Triangulate each diamond in the natural way, by connecting the finite vertices with an edge, and having one tetrahedron for each edge which belonged to the external face. Then, as usual, if the diamond has symmetries, take just one tetrahedron from each orbit. The resulting orbifold triangulation is defined to be the *canonical re-triangulation*.

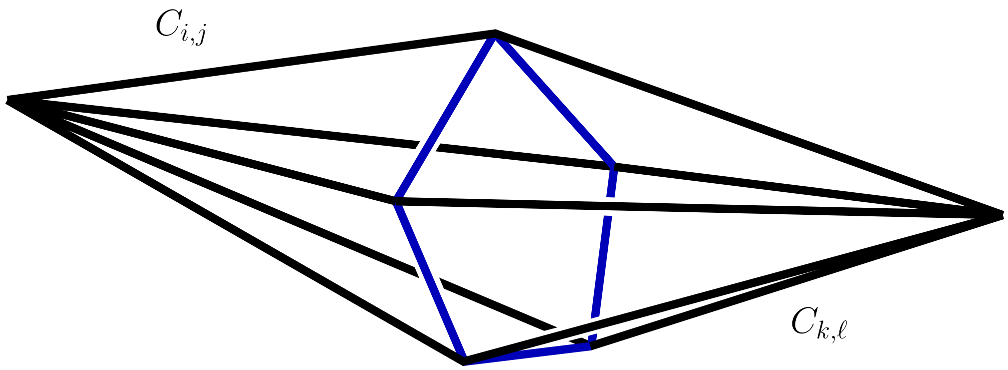

Figure 10: A diamond.

We now give an example. Consider the triangulated orbifold of figure 11. The tetrahedron on the left has all symmetries, and the one on the right has just the symmetries which preserve their shared face. The face of the tetrahedron on the right is glued to itself by rotation around the red axis. If we make both tetrahedra regular, then this is a hyperbolic orbifold triangulation (it is O^5 in the census of chapter 7). It is proto-canonical, but not canonical, because the quotient face shared by the two tetrahedra is transparent and the one other quotient face is convex. If we add a copy of the right tetrahedron to each other face of the left tetrahedron, we get a triangulated cube. See figure 12. Since the inner faces are transparent, the canonical decomposition must be the cube, equipped with all the symmetries which preserve the inner tetrahedron, with each face glued to itself by π rotation around a diagonal, and with edge labels corresponding to the labels of the tetrahedra.

We now compute the canonical re-triangulation of this canonical decomposition. There is a single cell, the cube P. When we introduce a finite vertex and cone to it, we get six new cells, one for each face of P. They all belong to the same orbit, so we just take one of them, call it C. Its data is described in figure 13.

The external face is glued to itself. To make the diamond, we glue a copy of C to itself along that face and make sure to add the symmetry corresponding to the face gluing map. After triangulating the diamond, all four tetrahedra belong to the same orbit, so we just take one of them. Its data is in figure 14. This is the canonical re-triangulation.

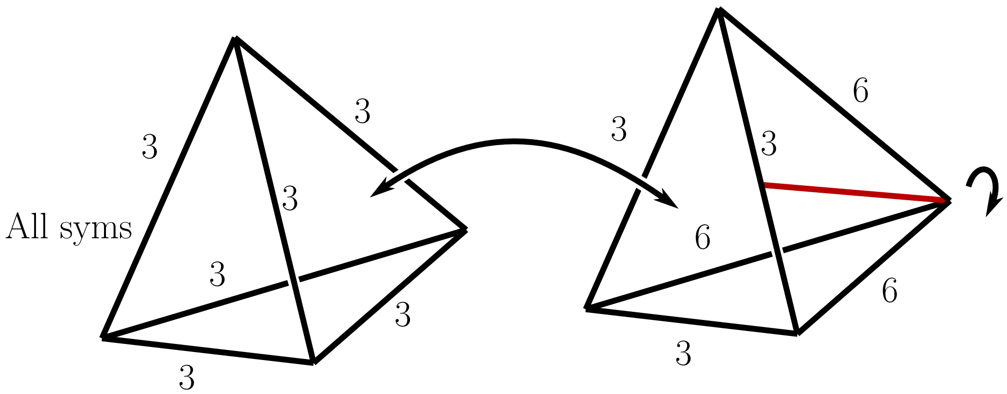

Figure 11: The proto-canonical triangulation of O5.

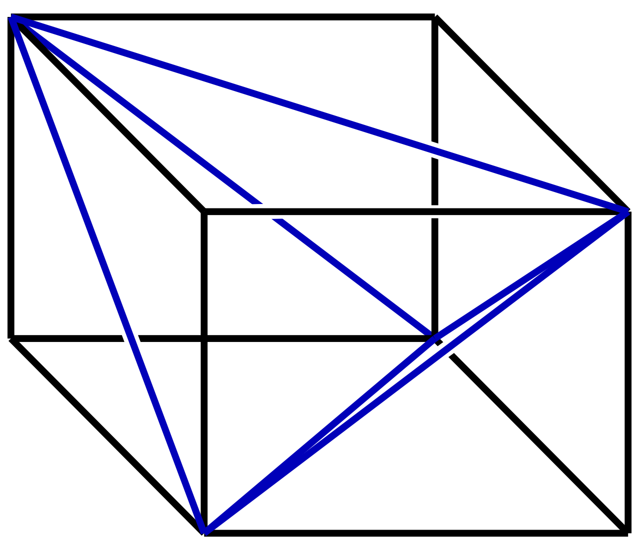

Figure 12: A triangulation of a cube is determined by the choice of inner tetrahedron.

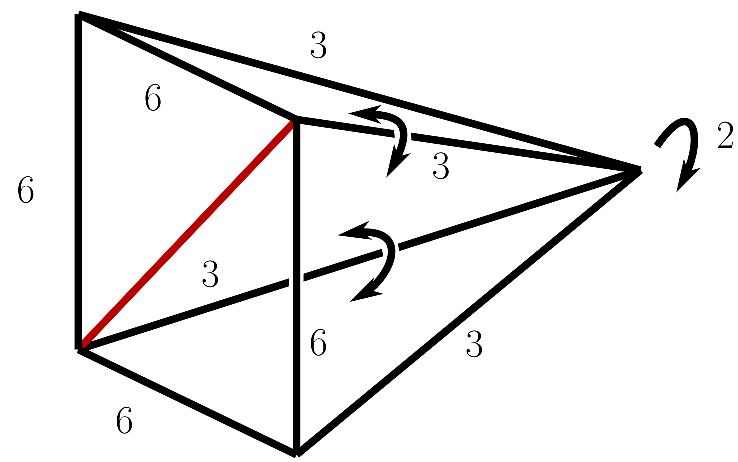

Figure 13: The orbifold polyhedral decomposition which results from coning the cube's boundary.

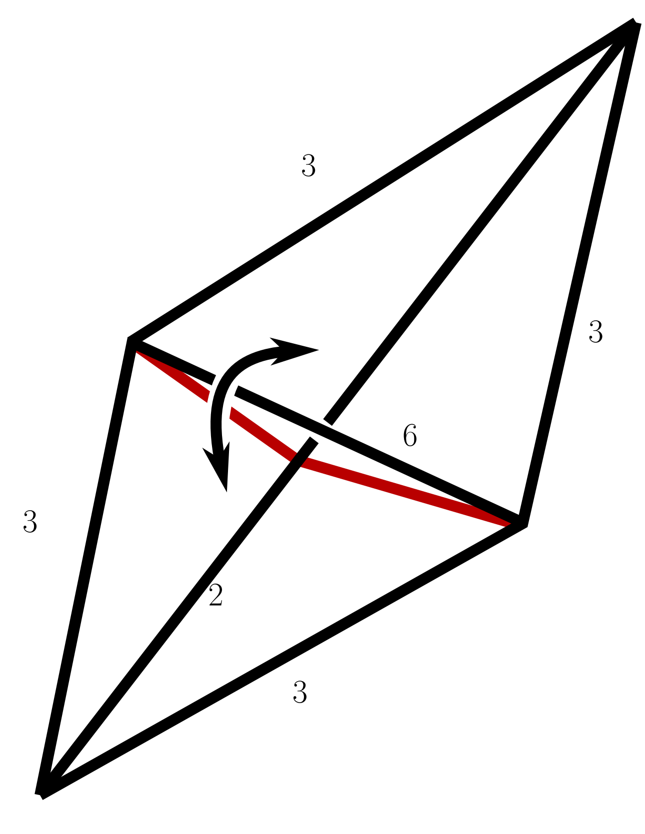

Figure 14: The canonical re-triangulation.

5.0 Canonize part one: finding a proto-canonical triangulation

Now that we have defined the canonical decomposition of an orbifold, we turn to the task of computing it. Our goal is to generalize SnapPy's algorithm for computing canonical decompositions of manifolds, which is called *canonize*. Over the course of this chapter and the next, we describe how we have done this, justifying the following theorem.

Theorem 12. *There is a "canonize" algorithm for orbifolds, which generalizes SnapPy's corresponding algorithm for manifolds. Implemented in Sym, it has successfully computed the certified canonical decompositions of all* 336 *orbifolds which cover* $\mathbb{H}^3/\operatorname{PGL}(2, O_3)$ *up to covering degree* 48.

For an explanation of what we mean by *certified*, see the end of the next section.

Our "canonize" follows very closely the structure of SnapPy's "canonize". The input of the algorithm is a hyperbolic orbifold triangulation of some finite-volume cusped hyperbolic 3-orbifold Q. If the canonical decomposition of Q is a triangulation, then the output of the algorithm should be that triangulation. Otherwise, the output should be the canonical re-triangulation of the canonical decomposition, as defined in Section 4.4. Like SnapPy, our algorithm is not guaranteed to work, although it has for every example we have tried.

For the most part, it really is a generalization. If we plug in a hyperbolic orbifold triangulation which happens to be a triangulation of a manifold, then it should try to compute the canonical decomposition in the same way as SnapPy. The only difference is that SnapPy will try to randomly change the triangulation if it gets stuck, whereas ours will halt and raise an error.

Just like SnapPy's canonize, we split our canonize algorithm into two parts. In *canonize part one*, we attempt to turn a given hyperbolic orbifold triangulation into a proto-canonical triangulation using a set of *orbifold Pachner moves*. If this proto-canonical triagulation is not the canonical decomposition, i.e. the canonical decomposition is not a triangulation, then we attempt to turn this proto-canonical triangulation into the canonical re-triangulation in *canonize part two*. We explain canonize part one in this chapter, and canonize part two in the next.

5.1 The algorithm

Let K be a hyperbolic orbifold triangulation. First, we set equal-volume cusp neighborhoods, then we compute tilts as described in Section 4.3. From the tilts, we know if K is proto-canonical or not. If it is proto-canonical, then we are done with canonize part one.

Otherwise, there is either a quotient face which is concave or a flat tetrahedron which is not admissible. In either case, our plan is to make some modification to K, turning it into a new triangulation K' of the same orbifold which we hope is closer to being proto-canonical. The tools we have to make these moves are orbifold Pachner moves, described in the next sections.

The structure of the canonize part one algorithm is then very simple. Search for anywhere in $K = K_0$ where some orbifold Pachner move would give us progress, meaning doing the move would remove a concave face/edge, or would remove a non-admissible flat tetrahedron. If there is such a move, do it, re-compute tilts for this new triangulation K_1, and again try to do any move which would make progress. Continue in this way, possibly creating many new triangulations $K_1, K_2, \ldots$. If at some point we cannot make progress, hopefully it is because our current triangulation K_n is already proto-canonical. Check this, and if so then we are done. If not, and there are no moves we can do to make progress, then we give up and raise an error.

As is often the case in computational mathematics, we need to worry about rounding errors. Edge parameters are complex numbers, and if we naively represent them as pairs of rounded floating point numbers, then there could be errors when we compute tilts. This is already an issue which SnapPy has to deal with for manifolds, and there are no new complications coming from working with orbifolds. For SnapPy's approach to dealing with this problem, we refer the reader to its documentation [13]. In this work, we are most interested in finding canonical decompositions of orbifolds commensurable to the figure-eight knot complement, as described in Chapter 7. For such orbifolds, it turns out that the edge parameters of the tetrahedra of any hyperbolic orbifold triangulation must lie in the finite field extension $\mathbb{Q}(\sqrt{-3})$. It is easy to represent the elements of this field in an exact form. Because of the simplicity of the tilt formula Theorem 11, we can then represent the tilts of such tetrahedra in an exact form as well. The result is that, in this commensurability class, we can be sure if we have a proto-canonical triangulation or not. These methods are described in Fominykh-Garoufalidis-Goerner-Tarkaev-Vesnin [16], where they certify canonical decompositions of tetrahedral manifolds. Sym uses Goerner's implementation of this exact arithmetic. As a result, when Sym finds a proto-canonical triangulation of an orbifold in this commensurability class, it is *certified*, meaning we know with certainty that it is proto-canonical, and there were no rounding errors. For orbifolds not in this commensurabilty class, other methods should be used, and will be incorporated into Sym in the near future.

We now describe the orbifold Pachner moves used by canonize part one. We note that in order to try to remove non-admissible flat tetrahedra we use either a cancellation move or a special $4 \to 4$ move. All other moves in this chapter are used to remove a concave face/edge which does not belong to a flat tetrahedron.

5.2 The two-to-three and three-to-two moves

The $2 \to 3$ and $3 \to 2$ moves re-triangulate a triangular bipyramid. We already showed the two possible triangulations in Section 2.5. In the orbifold setting, the bipyramid might have some symmetries. We will describe the cases of the $2 \to 3$ move, where each case corresponds to a group of symmetries of the bipyramid. The $3 \to 2$ move reverses the $2 \to 3$ move, so there is a self-explanatory $3 \to 2$ move for each case of the $2 \to 3$ move.

1. Case 0. In this case, the bipyramid has no non-trivial symmetries. Two distinct tetrahedra are glued along a pair of faces, and we re-triangulate as in figure 15.
2. Case 1. Two distinct tetrahedra are glued along a pair of faces, each tetrahedron has the order 3 symmetry group fixing the vertex opposite the face they share. Hence their union is a triangular bipyramid with an order 3 symmetry group. The three tetrahedra resulting from the re-triangulation are all in the same orbit of that symmetry group, so

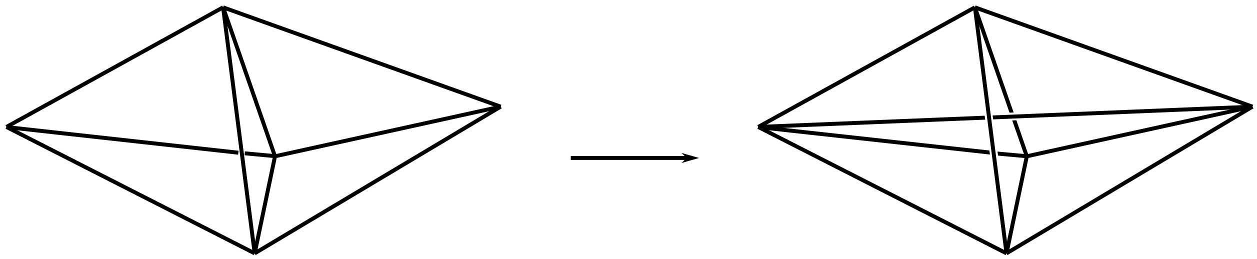

Figure 15: The two-to-three move without symmetries.

we take only one of them. See figure 16. Note that the edge which is internal to the bipyramid gets labelled 3.

3. Case 2. Suppose a face of a tetrahedron is glued to itself. Attach another copy of the tetrahedron to this face, giving a bipyramid with an order 2 symmetry group coming from the face gluing map. Of the three tetrahedra created after re-triangulation, two of them are mapped to to each other by this symmetry. So we only take one of those two. See figure 17. The bottom tetrahedron has an order two symmetry. This rotation axis extends through the face of the other tetrahedron, meaning that face is glued to itself by rotation around that axis.
4. Case 3. This is when both of the two previous cases occur. We have a triangular bipyramid with its full symmetry group. The $2 \to 3$ move results in a single tetrahedron as in figure 18. It has an order 2 symmetry group, an edge labelled 3, and a face which is glued to

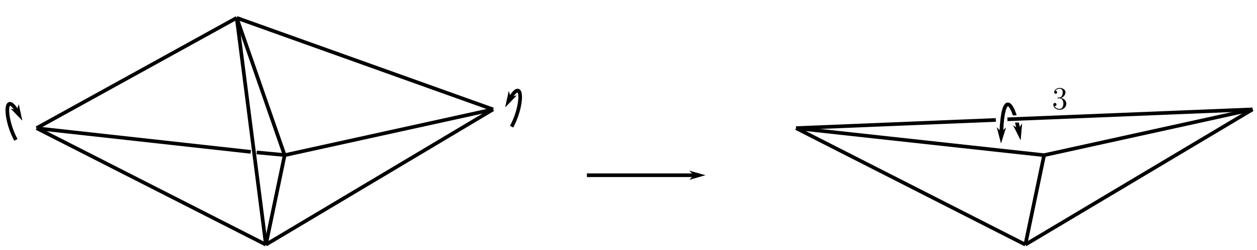

Figure 16: Case one of the orbifold two-to-three move.

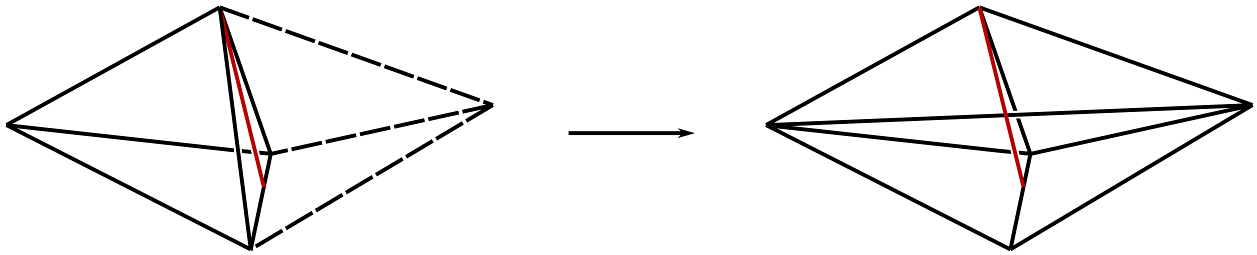

Figure 17: Case two of the orbifold two-to-three move.

itself.

We will now discuss the geometry of the $2 \to 3$ move. We do not consider a $2 \to 3$ move if either of the tetrahedra are flat. We can arrange them in the upper half space model as (x_1, x_2, x_3, y_1) and (x_1, x_2, x_3, y_2) as in figure 19. We assume that x_1 is at $0 \in \mathbb{C}$, x_2 is at $1 \in \mathbb{C}$, x_3 is at ∞, y_1 is at $w_1 \in \mathbb{C}$ where w_1 is the edge parameter of (x_1, x_3) in (x_1, x_2, x_3, y_1), and y_2 is at $1 - w_2$, where w_2 is the edge parameter of (x_2, x_3) in (x_1, x_2, x_3, y_2). If we are in cases 2 or 3, in which the face is glued to itself, we assume (x_1, x_2, x_3, y_1) is the original tetrahedron and its face (x_1, x_2, x_3) is glued to itself by the Mobius transformation $z \mapsto 1 - z$. In that case, $w_1 = w_2$.

As briefly mentioned in Section 2.5, the $2 \to 3$ move is not possible if the union of the two tetrahedra is not convex. Convexity is equivalent to the sums of the two dihedral

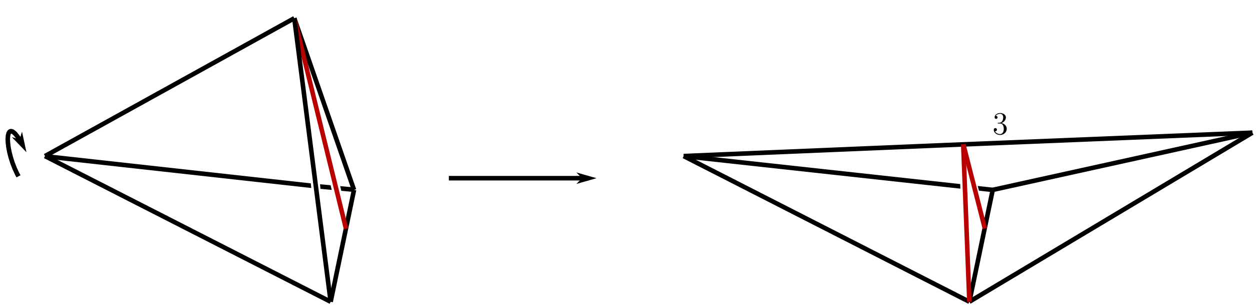

Figure 18: Case three of the orbifold two-to-three move.

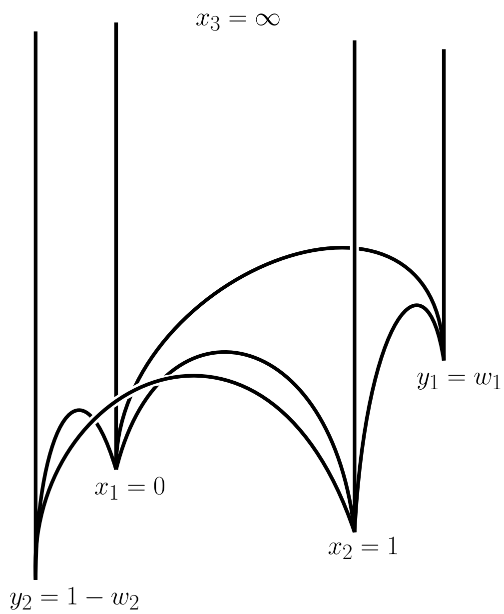

Figure 19: The two-to-three setup in upper half space.

angles contributing to each of $(x_1, x_2), (x_2, x_3), (x_1, x_3)$ being less than or equal to π. This is easy to check with edge parameters. For example, if a is the edge parameter of (x_1, x_2) in (x_1, x_2, x_3, y_1) and b is the edge parameter of (x_1, x_2) in (x_1, x_2, x_3, y_2), then the angle at (x_1, x_2) is less than or equal to π if and only if the imaginary part of ab is non-negative.

We want to determine the geometries of the tetrahedra T_1, T_2, T_3, where

$$T_1 = (x_2, x_3, y_1, y_2)$$
$$T_2 = (x_1, x_3, y_1, y_2)$$
$$T_3 = (x_1, x_2, y_1, y_2)$$

To do this, we will find the edge parameters of (y_1, y_2) in each of T_1, T_2, T_3. Call these edge parameters z_1, z_2, z_3 respectively.

To get z_1, let f be the Mobius transformation satisfying $1 - w_2 \mapsto 0, \infty \mapsto 1, w_1 \mapsto \infty$. Then $f(z) = \frac{z-1+w_2}{z-w_1}$, so $z_1 = f(1) = w_2/(1-w_1)$.

For z_2, let f map $1 - w_2 \mapsto 0, 0 \mapsto 1, w_1 \mapsto \infty$. Then $f(z) = \frac{z-1+w_2}{z-w_1}\frac{-w_1}{-1+w_2}$, and so $z_2 = f(\infty) = -w_1/(-1+w_2)$.

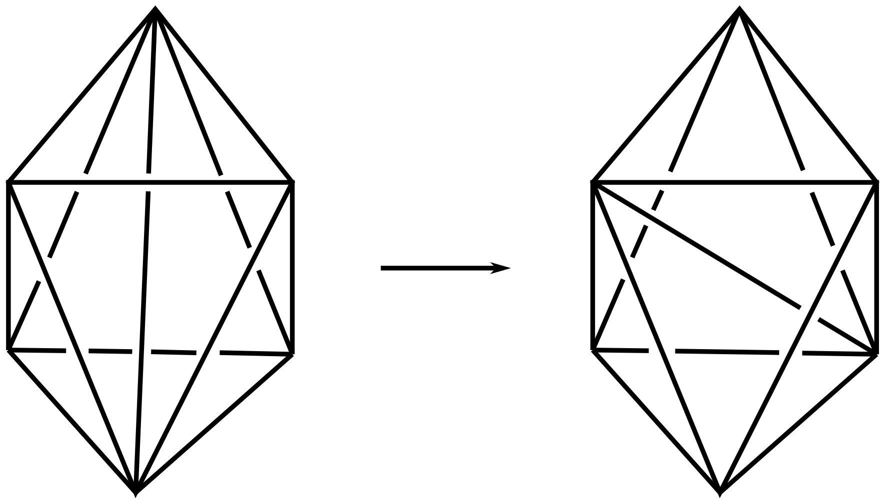

Figure 20: The four-to-four move with no symmetries.

For z_3, let f map $1 - w_2 \mapsto 0, 1 \mapsto 1, w_1 \mapsto \infty$. Then $f(z) = \frac{z-1+w_2}{z-w_1}\frac{1-w_1}{w_2}$, so $z_3 = f(0) = \frac{(1-w_2)^2}{w_1 w_2}$.

Depending on the case of the $2 \to 3$ move, we might not create all of T_1, T_2, T_3, so we might not need all three of z_1, z_2, z_3. It should now be clear how to determine the geometry of the $3 \to 2$ move, which we do not show here.

5.3 The four-to-four move

The $4 \to 4$ move re-triangulates an octahedron, usually as in figure 20. Each triangulation in the figure consists of four tetrahedra arranged in a cycle around an inner edge, and the re-triangulation just changes the inner edge.

We can imagine an orbifold version of the $4 \to 4$ move for every group of symmetries of an octahedron. We will focus on one particular case in this section, which is pictured in figure 21. The symmetry group has a single non-trivial element, the order two rotation around the red axis.

We assume we start with the triangulation of the octahedron on the left of figure 21, so

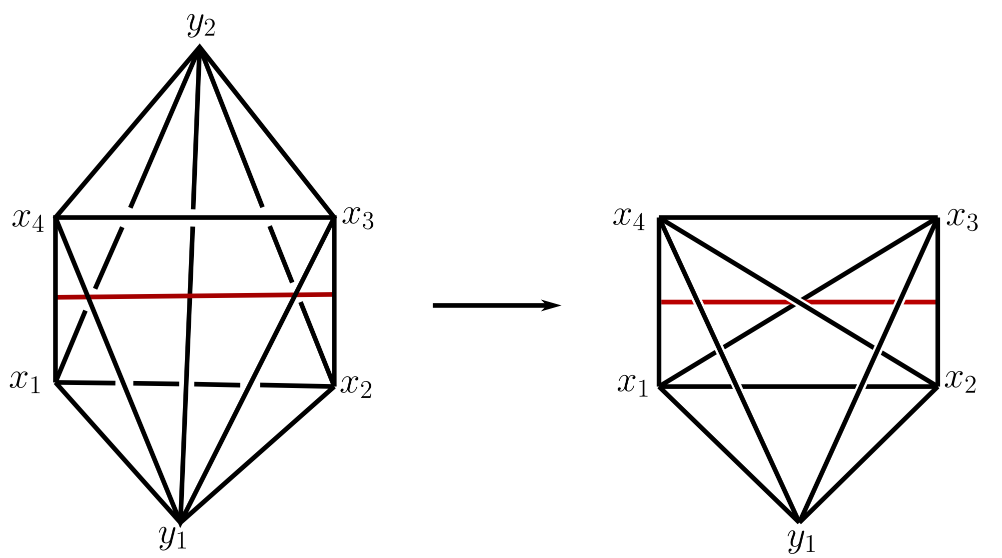

Figure 21: The four-to-four move with an order two symmetry.

we have tetrahedra

$$(y_1, y_2, x_2, x_3), (y_1, y_2, x_3, x_4), \text{ and } (y_1, y_2, x_1, x_4).$$

Note that (y_1, y_2, x_1, x_2) is identified with (y_1, y_2, x_3, x_4) by the symmetry, which is why we can assume it is not there. The usual $4 \to 4$ move, as in figure 20, results in a triangulation of the octahedron which is not invariant with respect to this symmetry. What we should instead do is take just two of the four new tetrahedra, say (y_1, x_1, x_2, x_4) and (y_1, x_2, x_3, x_4), and introduce a new tetrahedron, (x_1, x_2, x_3, x_4), having the symmetry which swaps x_1 with x_4 and x_2 with x_3. It could be that (x_1, x_2, x_3, x_4) is flat, as suggested by the picture, but this does not have to be the case.

We should actually be somewhat careful about the picture of the setup. In figure 21 we assume that, from the perspective of y_1, (x_2, x_4) is "in front of" (x_1, x_3). This might not be geometrically true. If not, then we can rotate the picture 90 degrees and re-label to match the previous picture, i.e. so that x_1 is in the bottom left, x_2 bottom right, etc. Then (x_2, x_4) is in front of (x_1, x_3) again, but now the rotation axis is vertical, and swaps x_1 with x_2 and

x_3 with x_4. In summary, there are two cases: when the symmetry axis is horizontal, and when it is vertical. If (x_1, x_2, x_3, x_4) is flat, we can choose either one.

In the remainder of this section, we determine the geometries of the tetrahedra created by this $4 \to 4$ move, where the octahedron has the order 2 symmetry just described. Their geometries will be the same regardless of whether we are in the "horizontal" or "vertical" case, with the only difference between the two cases being what symmetry group (x_1, x_2, x_3, x_4) gets labelled with.

Let

$$\begin{aligned}
z_0 &= \text{parameter of } (y_1, x_4) \text{ in } (y_1, y_2, x_1, x_4) \\
z_1 &= \text{parameter of } (y_1, y_2) \text{ in } (y_1, y_2, x_2, x_3) \\
z_2 &= \text{parameter of } (y_1, x_4) \text{ in } (y_1, y_2, x_3, x_4) \\
z_3 &= \text{parameter of } (y_1, y_2) \text{ in } (y_1, y_2, x_1, x_2).
\end{aligned}$$

Place the octahedron in the upper half space model as in figure 22. We put y_1 at 0, x_3 at 1, x_4 at ∞, x_1 at w_0, x_2 at w_1, and y_2 at w_2. Note that, in this figure, only the edges of the octahedron are shown—no internal edges of a triangulation of the octahedron are shown. We want to determine the values of w_0, w_1, and w_2 in terms of z_0, z_1, and z_2. We do not actually need z_3, because it is determined by the other edge parameters because of the symmetry. But we will consider it in the next section, when we reverse this move. By definition of z_2, we have $w_2 = z_2$. To get w_0 and w_1 we must use isometries.

Let f be the Mobius transformation satisfying $0 \mapsto 0, z_2 \mapsto 1, \infty \mapsto \infty$. Then $f(z) = z/z_2$. By definition of z_0, we have $z_0 = f(w_0)$. This implies $w_0 = z_0 z_2$.

Now let f satisfy $0 \mapsto 0, 1 \mapsto 1, z_2 \mapsto \infty$. Then $f(z) = (1 - z_2)z/(z - z_2)$. And we have $f(w_1) = z_1 \implies w_1 = z_1 z_2/(z_1 + z_2 - 1)$.

Now we can use w_0, w_1, and w_2 to get the shapes of the tetrahedra in the new triangulation of the octahedron. As explained before, in this case the $4 \to 4$ move results in the tetrahedra

$$(y_1, x_1, x_2, x_4), (y_1, x_2, x_3, x_4), (x_1, x_2, x_3, x_4).$$

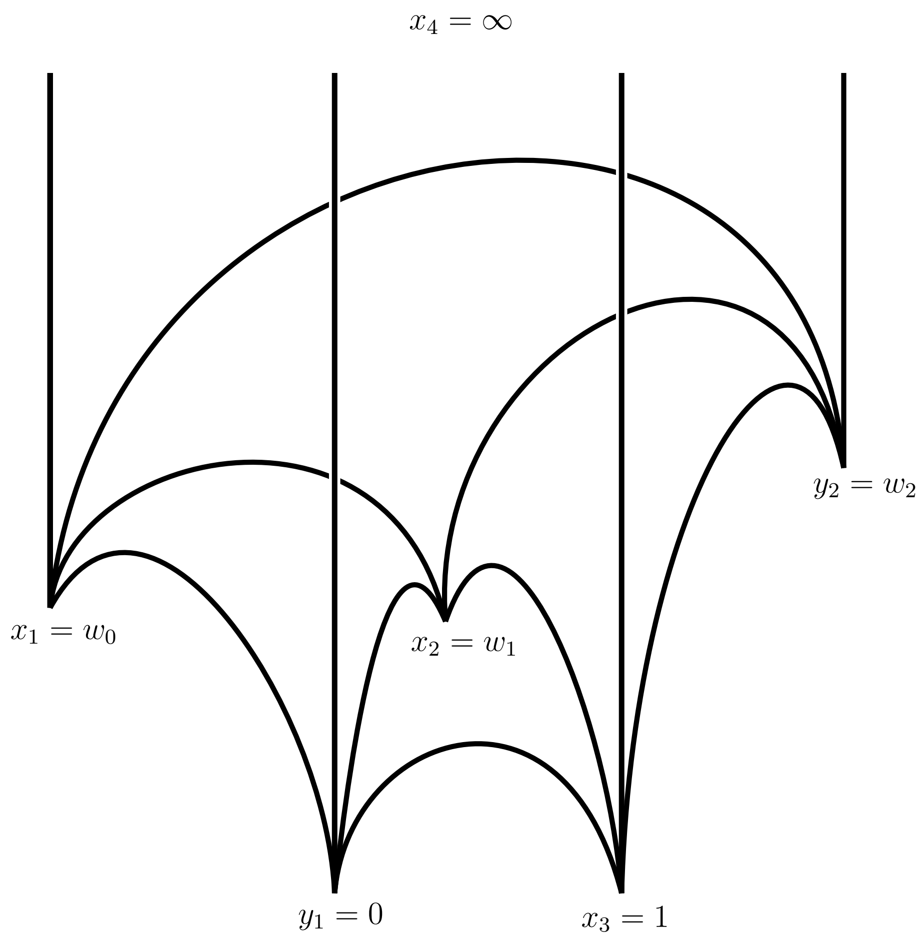

Figure 22: The coordinates of a hyperbolic ideal octahedron.

Let

$$u_0 = \text{shape of } (y_1, x_4) \text{ in } (y_1, x_1, x_2, x_4)$$

$$u_1 = \text{shape of } (y_1, x_4) \text{ in } (y_1, x_2, x_3, x_4)$$

$$u_2 = \text{shape of } (x_2, x_4) \text{ in } (x_1, x_2, x_3, x_4)$$

Let f satisfy $0 \mapsto 0$, $w_1 \mapsto 1$, and $\infty \mapsto \infty$. Then $f(z) = z/w_1$, and $u_0 = f(w_0) \implies u_0 = z_0(z_1 + z_2 - 1)/z_1$.

By definition, $u_1 = w_1 = z_1 z_2/(z_1 + z_2 - 1)$.

Now let f satisfy $w_1 \mapsto 0, 1 \mapsto 1, \infty \mapsto \infty$. Then $f(z) = (z - w_1)/(1 - w_1)$, and $u_2 = f(w_0)$, so

$$u_2 = \frac{(z_1 + z_2 - 1)z_0 z_2 - z_1 z_2}{z_1 + z_2 - 1 - z_1 z_2}.$$

5.4 The special four-to-four move

We may need to do the reverse of the order two symmetry $4 \to 4$ move described in the previous section. We call this the *special* $4 \to 4$ *move*. We assume we have three tetrahedra as in the right of figure 21, and we want to change it to the triangulation on the left. Similar to the discussion of the previous section, there are really two cases of the special $4 \to 4$ move: when the symmetry axis is horizontal or when it is vertical. See figure 23.

Our goal in this section is to explain the geometry of this move. In other words, to explain how we determine z_0, z_1, z_2, z_3 in terms of u_0, u_1, u_2, u_3, where we use the labelling of edge parameters introduced in the previous section. As before, we first get w_0, w_1, w_2, defined as in figure 22.

From the previous section, we have

$$w_1 = u_1, \quad w_0 = u_0 w_1 = u_0 u_1.$$

We determine w_2 using the symmetry. If the axis is horizontal, then by checking that it swaps the correct ideal vertices, we see that it is the Mobius transformation

$$z \mapsto \frac{w_0 z - w_0 + w_1 - w_0 w_1}{z - w_0}.$$

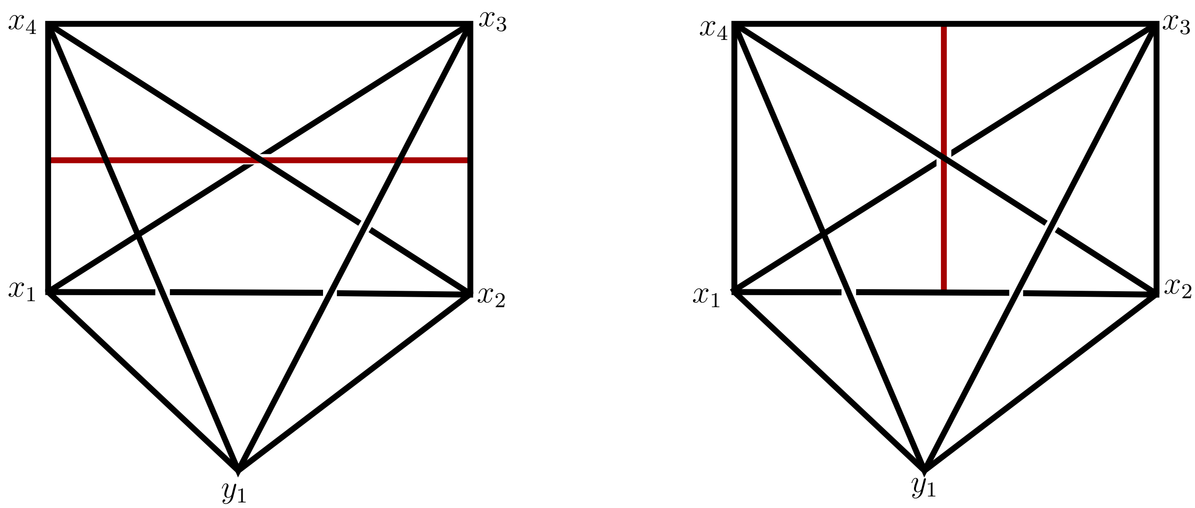

Figure 23: Horizontal or vertical axis.

Plugging in $z = 0$ gives w_2, so $w_2 = 1 - w_1/w_0 + w_1$. Therefore $w_2 = 1 - 1/u_0 + u_1$. We then get

$$z_0 = u_0 u_1/(1 - 1/u_0 + u_1)$$
$$z_1 = u_1(1/u_0 - u_1)/(1/u_0 - 1)$$
$$z_2 = 1 - 1/u_0 + u_1.$$

If the symmetry axis is instead vertical, then it is represented by the Mobius transformation

$$z \mapsto \frac{-z + (1 - w_1)(w_0 - 1) + 1}{-z + 1}.$$

Plugging in 0, we get $w_2 = (1 - w_1)(w_0 - 1) + 1$. And

$$z_0 = w_0/w_2$$
$$z_1 = (1 - w_2)w_1/(w_1 - w_2)$$
$$z_2 = w_2$$
$$z_3 = w_0(w_1 - w_2)/(w_1(w_0 - w_2)).$$

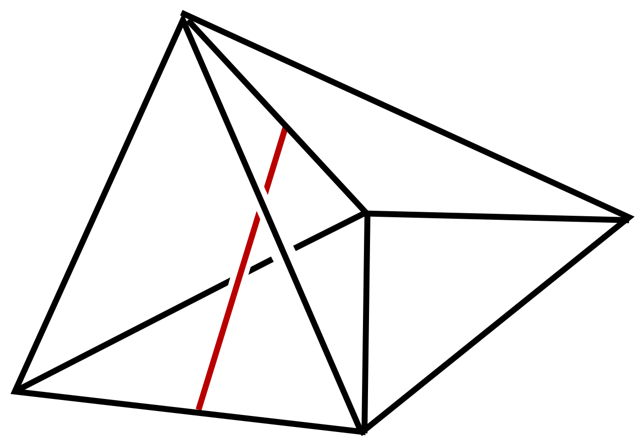

Figure 24: A two-to-three move is impossible because of the symmetry, which is order two rotation around the red axis.

5.5 The three-to-six move

Suppose that two distinct tetrahedra share a face, one of the tetrahedra has exactly one non-trivial symmetry, the other has only the trivial symmetry. View them as in figure 24. Then a $2 \to 3$ move through the shared face is impossible, because it is not invariant with respect to the symmetry.

Because of the symmetry, a copy of the tetrahedron on the right is also glued to the left face of the tetrahedron with the symmetry, as in figure 25. We view this union as a polyhedron with an order two symmetry.

If we want to modify the triangulation here, we should re-triangulate the union of these three tetrahedra in a way which is invariant with respect to the symmetry. Since we cannot do a $2 \to 3$ move, this is the next best thing. We can do this re-triangulation in two ways. Option 1: the new tetrahedra are

$$(u_0, u_1, v_0, w_0), (u_0, u_1, v_1, w_1), (u_0, u_1, w_0, w_1),$$
$$(u_0, v_1, w_0, w_1), (u_1, v_0, w_0, w_1), (v_0, v_1, w_0, w_1).$$

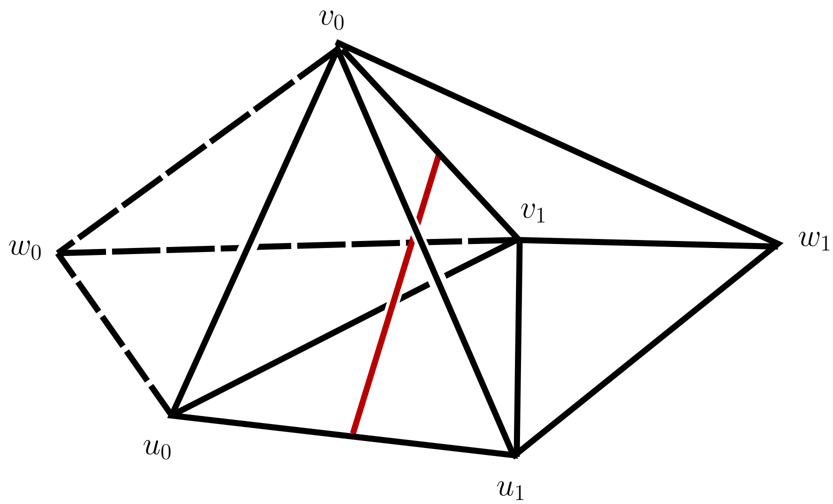

Figure 25: The union of the three tetrahedra is a polyhedron with an order two symmetry.

Option 2: the new tetrahedra are

$$(u_0, u_1, v_0, w_1), (u_0, u_1, v_1, w_0), (u_0, u_1, w_0, w_1),$$
$$(u_1, v_1, w_0, w_1), (u_0, v_0, w_0, w_1), (v_0, v_1, w_0, w_1).$$

These two re-triangulations are not really different combinatorially. The important difference arises when we consider the hyperbolic geometry of the tetrahedra. If, from the perspective of the figure, the geodesic (u_0, w_1) lies beneath the geodesic (u_1, w_0), then we should choose option 1. If (u_0, w_1) lies above (u_1, w_0) then we should choose option 2. If neither is true, in which case they intersect, then we can choose either option 1 or 2, although we accept that (u_0, u_1, w_0, w_1) will be flat.

Since the move turns three tetrahedra into six, we call it the $3 \to 6$ move. Because of the symmetry, we do not take all six tetrahedra, but instead four of them. For instance, in option 1, (u_0, u_1, v_0, w_0) is mapped to (u_0, u_1, v_1, w_1), and (u_0, v_1, w_0, w_1) is mapped to (u_1, v_0, w_0, w_1), so we only need to take one from each of those pairs. We also note that (u_0, u_1, w_0, w_1) and (v_0, v_1, w_0, w_1) each inherit the symmetry.

As with the $2 \to 3$ move, there are some geometric conditions which must be satisfied in order to be able to do the $3 \to 6$ move. Namely, we want the union of the three tetrahedra to be a convex polyhedron. This will be satisfied if and only if the dihedral angles at the edges $(u_1, v_)), (u_1, v_1)$, and (v_0, v_1) sum to less than or equal to π. Note that (v_0, v_1) has three tetrahedra contributing angle to it, while (u_1, v_0) and (u_1, v_1) just have two.

The task of determining the geometries of the new tetrahedra is made much easier when we realize that the $3 \to 6$ move can be accomplished by a series of "illegal" $2 \to 3$ moves. Suppose we are doing option 1. Forget the symmetry of the middle tetrahedron is there, and do a $2 \to 3$ move through the face (u_1, v_0, v_1). We now have tetrahedra

$$(u_0, u_1, v_0, w_1), (u_0, u_1, v_1, w_1), (u_0, v_0, v_1, w_1),$$

as well as (u_0, v_0, v_1, w_0), which did not change. Now do a $2 \to 3$ move through face (u_0, v_0, v_1). The triangulation now consists of

$$\begin{aligned}(u_0, u_1, v_0, w_1), (u_0, u_1, v_1, w_1), (u_0, v_0, w_0, w_1),& \\ (u_0, v_1, w_0, w_1), (v_0, v_1, w_0, w_1).&\end{aligned}$$

Since (u_0, w_1) lies beneath (u_1, w_0), we can do a $2 \to 3$ move through the face (u_0, v_0, w_1). This results in the desired re-triangulation by six tetrahedra. We just have to remove two tetrahedra and add the symmetry to two others, as discussed above. Since the $2 \to 3$ move already correctly changes the edge parameters of the tetrahedra, no other work is required.

5.6 Cancelling flat tetrahedra

When two tetrahedra are arranged around a valence 2 edge as in figure 26, we call their union a pillow. We can collapse the two tetrahedra onto each other, which we call a cancellation move or a $2 \to 0$ move. We discussed the version of this move without symmetries in section 2.5. We have other versions for each group of combinatorial symmetries of the pillow.

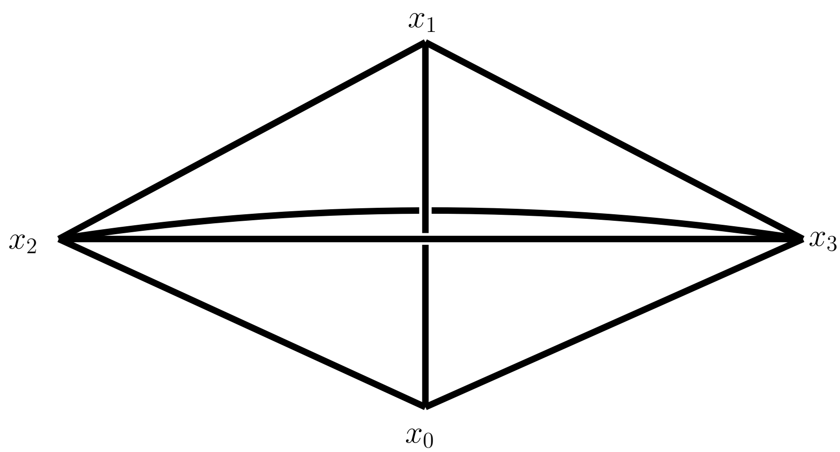

Figure 26: A pillow with labelled vertices.

The non-trivial symmetries are:

$$\varphi_1 : x_0 \mapsto x_1,\ x_1 \mapsto x_0,\ x_2 \mapsto x_2,\ x_3 \mapsto x_3$$

$$\varphi_2 : x_0 \mapsto x_1,\ x_1 \mapsto x_0,\ x_2 \mapsto x_3,\ x_3 \mapsto x_2$$

$$\varphi_3 : x_0 \mapsto x_0,\ x_1 \mapsto x_1,\ x_2 \mapsto x_3,\ x_3 \mapsto x_2.$$

The subgroups of the full symmetry group are H_1 = trivial group, $H_2 = \langle\varphi_1\rangle$, $H_3 = \langle\varphi_2\rangle$, $H_4 = \langle\varphi_3\rangle$, H_5 = full group. Let the ith case be the cancellation move corresponding to the subgroup H_i.

When we are looking for places in an orbifold triangulation to do a cancellation move, we might not "see" a full pillow. This is the usual situation when we are working with triangulated polyhedra with symmetries. Figures 27, 28, 29, 30 show what we will see in each case, before we do the cancellation move.

In case 2, we have a single tetrahedron with two faces glued to themselves in the indicated ways. We think of the two face gluing rotation axes as being a single rotation axis, and if we rotate the tetrahedron around that axis, we get the other tetrahedron of the pillow. When we do this cancellation move, the horizontal edge gets collapsed onto the rotation axis, making the two faces adjacent to that edge glued to each other.

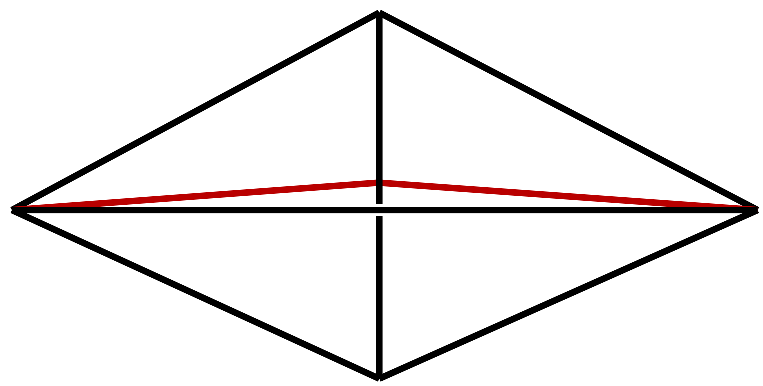

Figure 27: Case two of the cancellation move.

In case 3, we see both tetrahedra of the pillow from the start. They have an order two symmetry which is rotation around the red axis which extends through both of them, intersecting the vertical edge. To do the move, we just collapse the pillow.

In case 4, we only see one tetrahedron. The two faces adjacent to the vertical edge are glued to each other by an order two rotation around the vertical edge. When we collapse this tetrahedron, the two faces adjacent to the horizontal edge get glued to themselves, and the vertical edge becomes the rotation axis for these gluing maps.

In case 5, we again only see one tetrahedron. Because of its symmetry, we can actually suppose the two back faces are glued to each other as in case 4. So this is similar to case 4, except that, of the two faces adjacent to the horizontal edge, only one of them is glued to some other tetrahedron. The other face is not glued to anything, because of the symmetry.

For a hyperbolic orbifold triangulation, a pillow must be flat, and any tetrahedron or pair of tetrahedra we cancel will necessarily be flat. In the non-geometric setting, of course we do not talk about a tetrahedron being flat. There is an important obstruction to doing a cancellation move, specific to orbifolds. The horizontal edge(s), opposite the "inside" edge (which is vertical), must have label 1. This is because in some of these cases we push the horizontal edge into the singular locus. If the horizontal edge is already part of the singular locus, i.e. its label is greater than 1, then this illegally changes the orbifold structure.

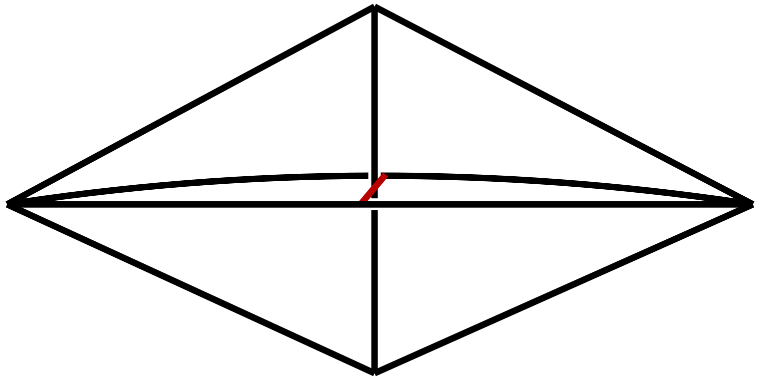

Figure 28: Case three of the cancellation move.

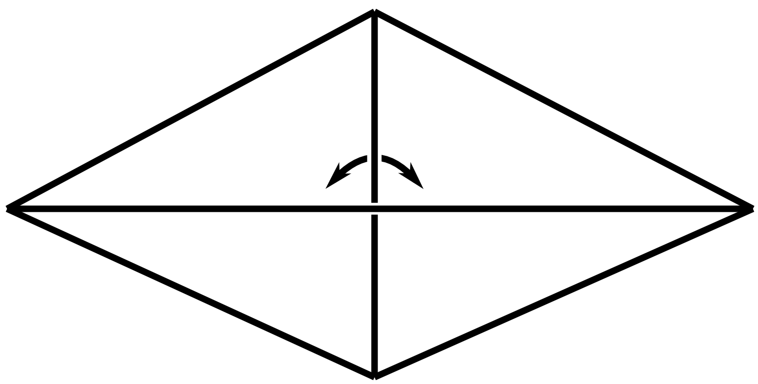

Figure 29: Case four of the cancellation move.

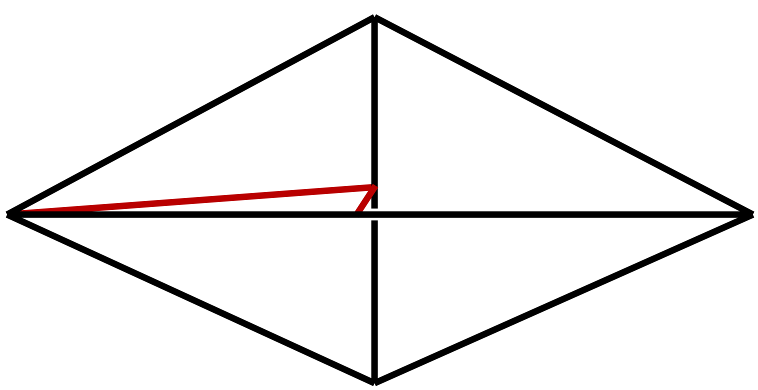

Figure 30: Case five of the cancellation move.

5.7 Re-triangulating a cube

Four regular tetrahedra, each glued to a fifth regular tetrahedron along each of its four faces, union together to form a cube. The edges of the inner tetrahedron are diagonals of the square faces of the cube. There is exactly one other triangulation of the cube, which is obtained by choosing a different inner tetrahedron: the one whose edges are the other diagonals of the cube's faces.

Suppose that the inner tetrahedron is labelled with the full symmetry group or the order 4 subgroup. Only one of the four regular tetrahedra around it will explicitly "be there", because of these symmetries. It could be that the face shared by this tetrahedron and the inner tetrahedron is concave. In that case, we would like to do a $2 \to 3$ move through this face, but unfortunately it is not possible, because of the symmetries. Similarly, we cannot do a $3 \to 6$ move. Since we know these tetrahedra correspond to a specific triangulation of the cube, we might hope that changing to the other triangulation of the cube will fix the concavity issue. That is exactly what this section's orbifold Pachner move does.

As noted above, when we change the triangulation of the cube, we swap diagonals in each face of the cube. This means we are changing the triangulation of the boundary of the cube. This is bad, because other tetrahedra which are part of this triangulation could be glued to the boundary of the cube with respect to its previous triangulation. To correct this, we must introduce a flat tetrahedron. Figure 31 shows the re-triangulation. The green edge is the "top" edge of the introduced flat tetrahedron.

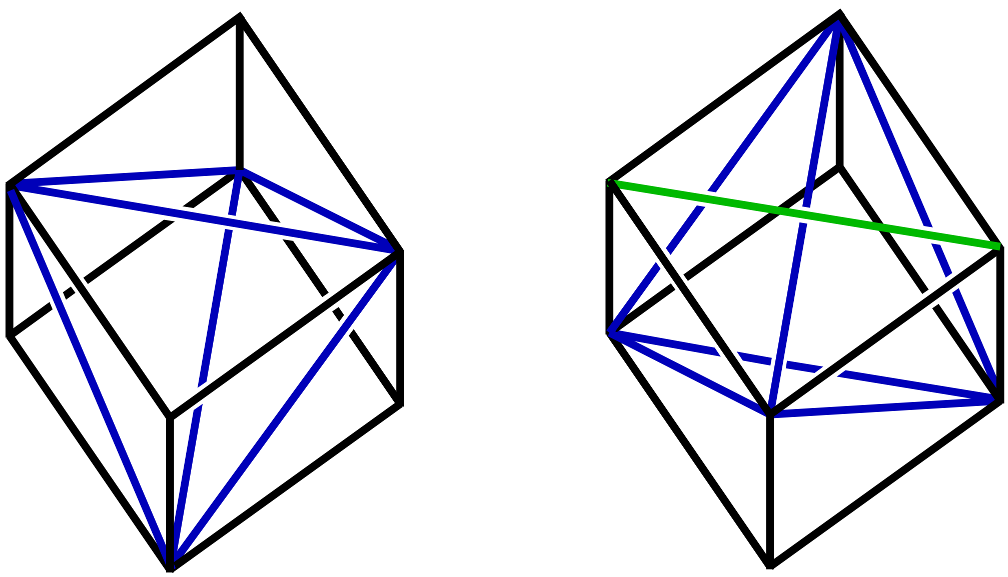

Figure 31: How to re-triangulate a cube.

6.0 Canonize part two: finding the canonical re-triangulation

Once we have succeeded with part one, we move on to *canonize part two*. We check if the proto-canonical triangulation obtained from part one is actually canonical. If so, then we have succeeded, so we exit. If not, then our proto-canonical triangulation is a sub-division of the canonical decomposition. We would like to turn our proto-canonical triangulation into the canonical re-triangulation. To accomplish this, we adapt to orbifolds SnapPy's canonical re-triangulation routine for manifolds.

As in SnapPy, we split canonize part two into two steps.

1. Step 1. Replace the given proto-canonical triangulation with the triangulation which has the following description. In each cell of the canonical decomposition, introduce a finite vertex. The boundary of the cell is triangulated by the given proto-canonical triangulation. Cone that boundary triangulation to the finite vertex. Do this for each cell.
2. Step 2. Each tetrahedron of this triangulation has exactly one finite vertex, and the face opposite it is a triangle in the triangulation of the face of the polyhedron induced by the original proto-canonical triangulation. For each tetrahedron, do a $2 \to 3$ move through that face. Unless it is glued to a flat tetrahedron, in which case do a *special* $4 \to 4$ *move*. Then do all possible cancellation moves.

To illustrate these two steps, consider again the cube example from Section 4.4. We start with the proto-canonical triangulation of figure 11. It is a sub-division of a cube, which gives a triangulation of the boundary of the cube. For step 1 of canonize part two, we want to turn our starting triangulation into the triangulation which is the boundary triangulation coned to the barycenter. If we cone the boundary triangulation to the barycenter, we get 12 tetrahedra. They all belong to the same orbit under the action of the symmetry group, so we only take one of them. Hence, the desired orbifold triangulation of step 1 is given in figure 32.

We now do step 2, meaning we do a $2 \to 3$ move through the face opposite the finite

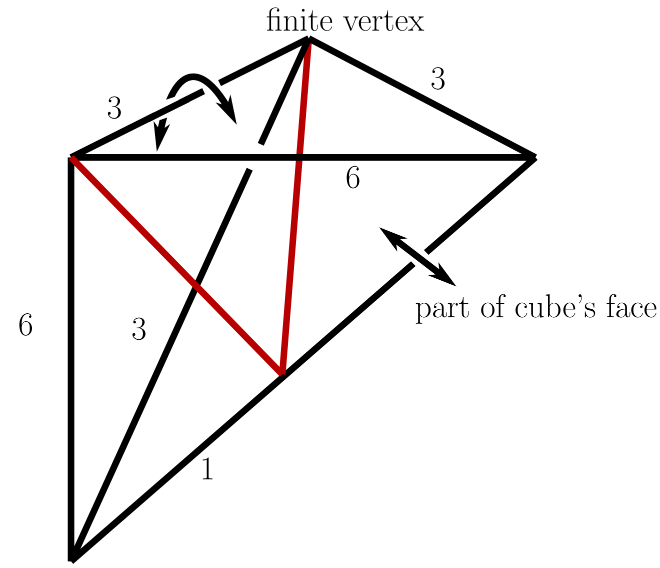

Figure 32: What we get after coning the cube's boundary triangulation.

vertex. The result is figure 33. The tetrahedron on the right can be cancelled, and the result is the desired canonical re-triangulation as in figure 14.

To accomplish step 1 is more complicated than its short description above makes it sound. While the desired triangulation which is the end goal of step 1 has a simple description in terms of the canonical decomposition, it is not straightforward to find it, given only an arbitrary proto-canonical triangulation.

To accomplish 1, we need new orbifold Pachner moves, which we describe in the next three sections. The remaining sections describe in detail how to do steps 1 and 2.

Warning. When we do canonize part two, we mostly forget about the geometric structure of the orbifold. This is because we introduce finite vertices, and keeping track of the geometry of hyperbolic tetrahedra which are not fully ideal is not as straightforward as with ideal tetrahedra. SnapPy does the same thing in its version of canonize part two. We do, however, keep track of which tetrahedra were flat in the original proto-canonical triangulation.

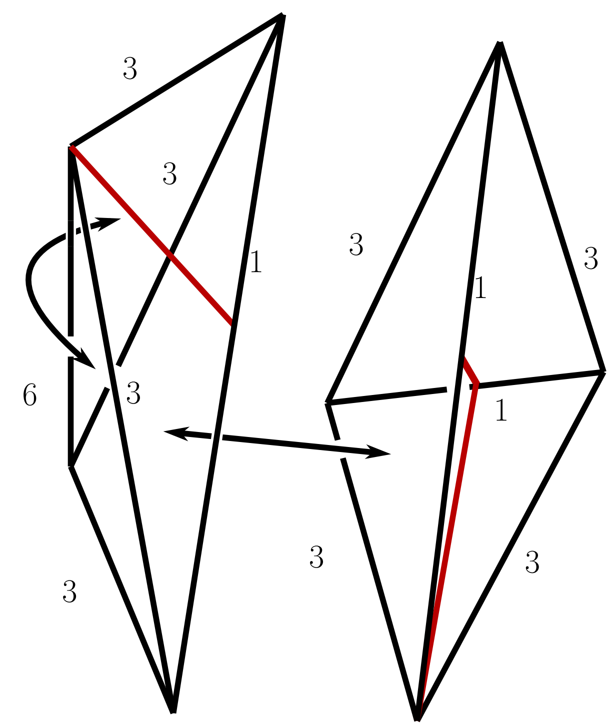

Figure 33: After the two-to-three move.

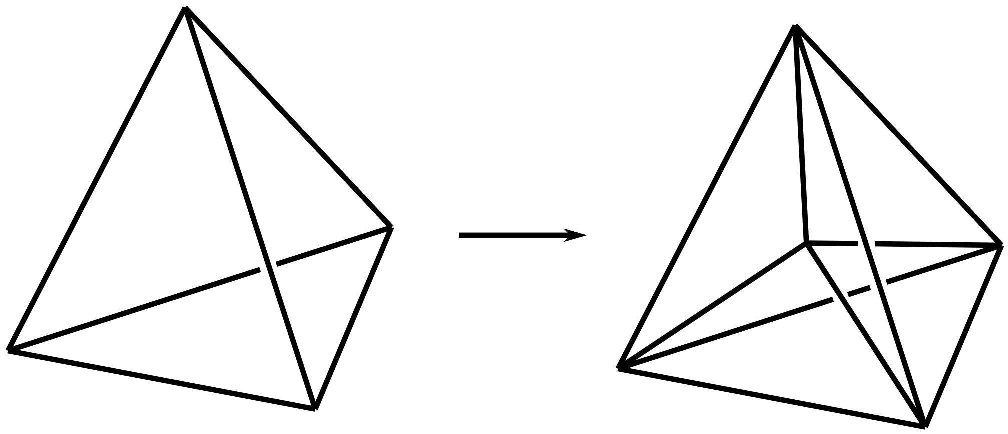

Figure 34: The one-to-four move with no symmetries.

6.1 The one-to-four and four-to-one moves

The $1 \to 4$ move splits 1 tetrahedron into 4 by introducing a point in the interior and coning each face to that point. In the setting of a hyperbolic orbifold triangulation, the new vertex must be finite, which means that after the $1 \to 4$ move we view the triangulation as just an orbifold triangulation, i.e. we forget about the geometric structure. The $4 \to 1$ move is the reverse of the $1 \to 4$ move. We do not use either move in canonize part one, but we do use the $1 \to 4$ move in canonize part two.

There is an orbifold version for each subgroup of symmetries of a tetrahedron. In figure 34, we see the $1 \to 4$ move in the case that the tetrahedron is only labelled with the trivial symmetry group. Figure 35 shows the case when the tetrahedron is labelled with the full symmetry group. Note that, in that case, we only take one of the resulting four tetrahedra.

6.2 Special cancellation

Let T be a tetrahedron in the orbifold triangulation with a vertex v whose opposite face f is glued to itself. Given that setup, there are two cases in which we can do a *special*

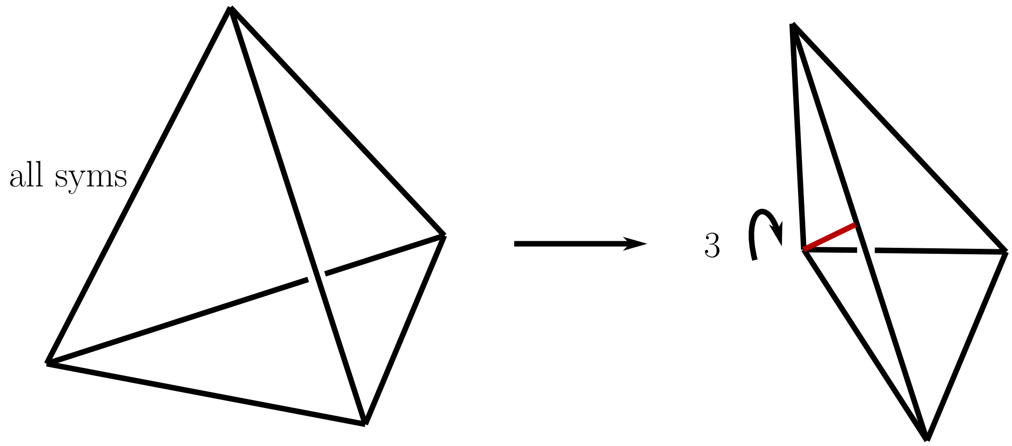

Figure 35: The one-to-four move with all symmetries.

cancellation move. In either case, the idea is to collapse T onto f.

Case 1. Suppose that the link of v is a sphere, the symmetry group of T is trivial, and all faces of T which are not f are glued to tetrahedra different from T. Then the *special cancellation* of T collapses T onto f, pushing v into the rotation axis of the face gluing map of f. See figure 36. The old face f gets sub-divided by faces of the green, blue, and purple tetrahedra. The rotation which glued f to itself now serves to glue the two faces of the green and blue tetrahedra to each other (with the edge between them now getting labelled 2) and to glue the face of the purple tetrahedron to itself.

Case 2. Now suppose that the link of v is the 2-orbifold whose underlying space is S^2 and whose singular locus consists of two points which are fixed points of order 3 rotations. Such an orbifold is called an "order 3 football". Suppose that the symmetry group of T is the order 3 group fixing v (hence the axis of this symmetry group intersects the link of v in one of its singular points). Finally, suppose if e is an edge of T containing v, then the edge label of e is 1. Then, as with case 1, the *special cancellation* of T collapses T onto f, pushing v into the rotation axis of the face gluing map of f. To understand this case, we may again use figure 36. Because of the symmetry group of T, the green, blue, and purple tetrahedra are all copies of each other. We can assume that of the three, we only have the purple one in the triangulation. Then, as in the figure, v collapses into the rotation axis of f, and the face of

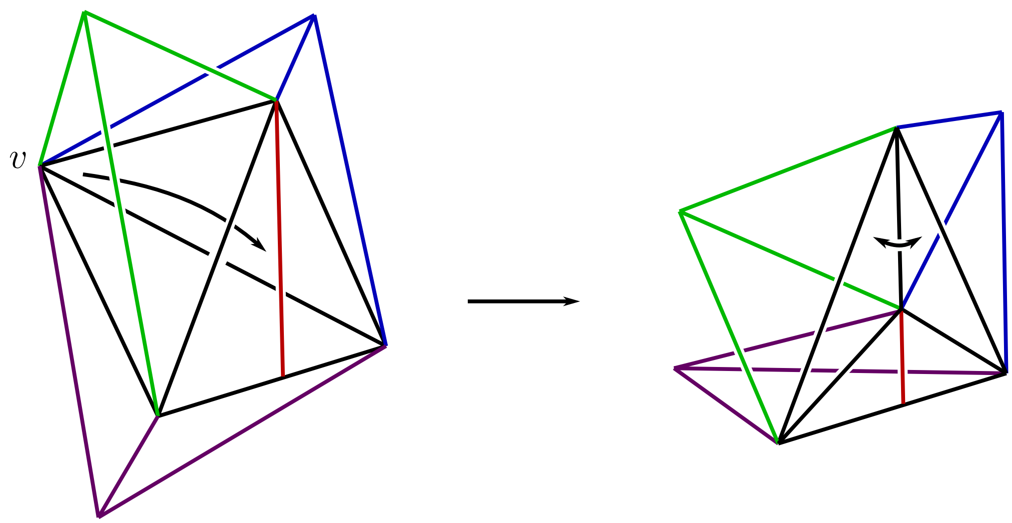

Figure 36: A special cancellation.

the purple tetrahedron which moves into f gets glued to itself. The two edges of the purple tetrahedron which contain v and are black get labelled 2 after the special cancellation.

If neither the assumptions of case 1 nor 2 are satisfied, then a special cancellation is not possible. The assumptions are important for this move to be valid. By "valid", we mean that the move does not change the orbifold. For example, it would not be valid to change the singular locus. Because v is spherical in case 1, it is not a singular point, so we can safely push it into the singular locus in f. In case 2, v lies in the interior of an edge of the singular locus, corresponding to the order 3 rotation of T. A vertex of this singular locus edge is at the center of f, and the special cancellation just pulls v into that vertex. We leave it to the reader to fill in other details about the validity of this move.

6.3 The stellar edge move

Let e be a quotient edge of the orbifold triangulation. Suppose that all the tetrahedra containing e either have trivial symmetry group or just the order 2 group mapping e to itself.

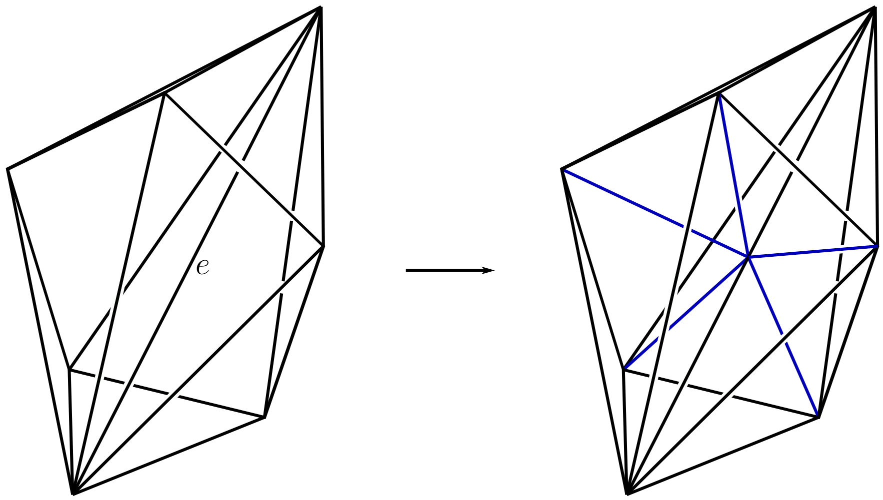

Figure 37: The stellar edge move.

Suppose further that if any face containing e is glued to itself, then its face gluing map takes e to itself. Then the set of all tetrahedra containing e forms a triangulated bipyramid with e running through the center. The bipyramid could have non-trivial symmetries, in which case not all of the tetrahedra are actually "there" in the orbifold triangulation. The *stellar edge move* of e puts a finite vertex in the interior of e and cones the boundary of the bipyramid to that point. See figure 37. If any of the tetrahedra containing e have a non-trivial symmetry preserving e, then the bipyramid has a non-trivial symmetry flipping e. All such symmetries have a unique fixed point, which is where we put the finite vertex in this case. As usual, if the bipyramid has symmetries then we just take one tetrahedron from each orbit.

6.4 Step 1

In general we do not know what the polyhedra of the canonical decomposition look like. All we have is the proto-canonical triangulation, which we know is some arbitrary sub-division of a polyhedral decomposition, whatever it is. To do step 1, we need to get a triangulation which is defined in terms of this polyhedral decomposition, without really knowing what this polyhedral decomposition is. We want to come up with some sequence of triangulation changing moves to apply to the proto-canonical triangulation which we somehow believe *must* lead us to the triangulation of step 1.

The function we use to accomplish step 1 is *ConePolyhedron*. Every time it is called, it tries to cone the triangulated boundary of some polyhedron to its barycenter. To keep track of which tetrahedra are part of a polyhedron we have already coned and which are not, we use the flags "coned" or "un-coned". Before the first application of ConePolyhedron, we set all flags to "un-coned". Any new tetrahedron we create will be set to "coned", even if it is not part of the fully coned triangulation of a polyhedron. We apply ConePolyhedron repeatedly until it returns False, at which point we know we have completed step 1.

The ConePolyhedron function is described in pseudo-code in Algorithm 1. We now describe the functions which it relies on.

The first function which ConePolyhedron calls is *InsertFiniteVertex*. See Algorithm 2. To start coning a polyhedron, we need to insert a finite vertex at a point which must be a barycenter of some polyhedron. Since we do not really "see" the polyhedra, only the proto-canonical triangulation, we need to take some care. For instance, in the $O5$ example whose proto-canonical triangulation is in figure 11, the barycenter of the cube is contained in the tetrahedron which has all symmetries, and not in the other tetrahedron. To insert a finite vertex in the tetrahedron with all symmetries, we do a $1 \to 4$ move on it. In general, the barycenter of the triangulated polyhedron could be contained in the interior of a tetrahedron, the interior of a face, or the interior of an edge. If it is contained in the interior of a face of a tetrahedron T, then InsertFiniteVertex inserts a finite vertex anywhere by doing a $1 \to 4$ move to any tetrahedron in this polyhedron. Then we collapse this vertex onto the rotation axis later on. If the barycenter is contained in the interior of an edge, we insert a finite vertex there with a stellar edge move.

The next function is *ExpandConedRegion*. See Algorithm 3. The idea is that we have a triangulation with one finite vertex of a polyhedron, and we want to make this triangulation closer to the goal triangulation of step 1. We look for a tetrahedron T containing the finite vertex v, whose face opposite v, call it f, is transparent and glued to some other un-coned tetrahedron T'. Then we expand the coned region by doing a $2 \to 3$ move through f. If T

```
def ConePolyhedron(Orbifold Triangulation orb):
    Result: Modify orb by coning a single polyhedron. Return True if we succeed,
            False if we cannot because all polyhedra are already coned.
    if InsertFiniteVertex(orb) = False then
        /* All polyhedra are already coned.                                   */
        return False;
    end
    while ExpandConedRegion(orb) = True do
        pass;
    end
    AttemptSpecialCancellation(orb);
    while AttemptCancellation(orb) = True do
        pass;
    end
    if VerifyConedRegion(orb) = False then
        /* We tried and failed to cone a polyhedron.  We are not sure if
           this can happen.                                                    */
        raise Exception("Failed to cone a polyhedron");
    end
    /* If we have gotten to this point, then we succeeded in coning a
       polyhedron.                                                             */
    return True;
```

Algorithm 1: This function tries to cone a polyhedron.

def `InsertFiniteVertex(`*Orbifold Triangulation orb*`):`

Result: Modify orb by inserting a finite vertex in an un-coned tetrahedron which we know must be a barycenter of some polyhedron. If we succeed, return True, otherwise all tetrahedra are already coned and we return False.

Make a list of any un-coned tetrahedra with more than 2 symmetries. If there are any, choose one of them with the most symmetries, do a $1 \rightarrow 4$ move on it, mark the newly created tetrahedra as coned, then return True;

Look for a transparent edge belonging to un-coned tetrahedra which has label > 1. If there is one, do a stellar edge move to it, mark the newly created tetrahedra as coned, then return True;

Look for a transparent edge belonging to un-coned tetrahedra and with more than one rotation axis intersecting it. If there is one, do a stellar edge move to it, mark the newly created tetrahedra as coned, then return True;

Look for any tetrahedra with 2 symmetries. If there are any, pick one, do a $1 \rightarrow 4$ move on it, mark the newly created tetahedra as coned, then return True;

Do a $1 \rightarrow 4$ move to any un-coned tetrahedron, mark the newly created tetrahedra as coned, then return True;

return False;

Algorithm 2: This function inserts a finite vertex.

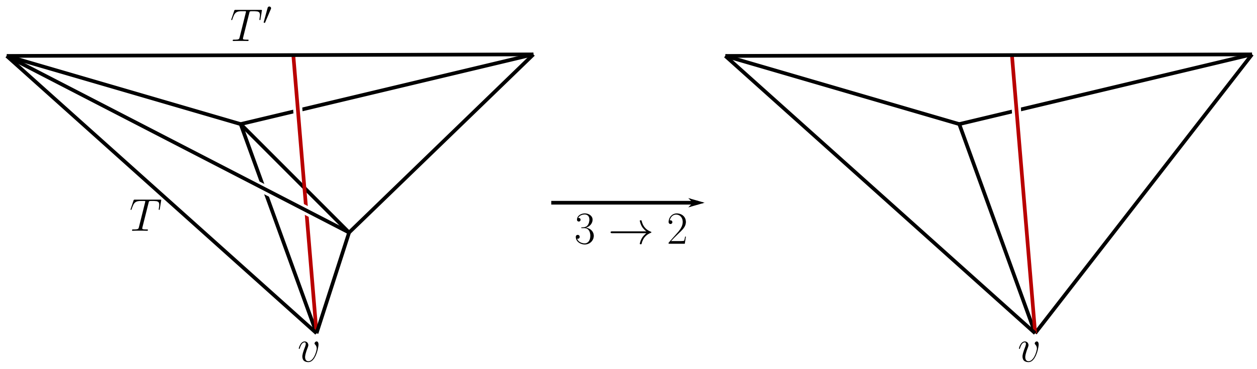

Figure 38: When a two-to-three move is not possible, we can do a three-to-two move instead.

or T' has the wrong symmetries, then of course we cannot do the $2 \to 3$ move. For our situation, the only bad setup of symmetries which can occur is when T has no non-trivial symmetries and T' has the order 2 group of symmetries as in figure 38. Note that a face of T is glued to itself by rotation around the red axis, and that red axis extends into T' as the axis of its symmetry group. In this case, we do a $3 \to 2$ move "through f" instead of a $2 \to 3$ move.

The next function used by ConePolyhedron is *AttemptSpecialCancellation*. This function looks for a transparent face opposite the finite vertex which is glued to itself, and does a special cancellation move on it. This collapses the finite vertex onto the rotation axis inside the face. This can only occur in the case that the barycenter is contained inside a face.

After this, we use *AttemptCancellation* repeatedly until all possible cancellation moves are done. They are necessary for removing transparent edges from inside the polyhedron.

Then we check that we succeeded in coning the polyhedron with *VerifyConedRegion*.

6.5 Step 2

After we have applied ConePolyhedron as many times as possible, step 1 is complete so we move on to step 2. First, we apply *EliminateOpaqueFace* repeatedly until it returns

```
def ExpandConedRegion(Orbifold Triangulation orb):
    Result: Try to modify orb by expanding the coned region. This means making
            the triangulation closer to the goal triangulation of step 1. If we can
            make such a modification, return True. Otherwise return False.
    Look for any tetrahedron T which is coned and has a transparent face f which is
     glued to some tetrahedron which is not coned.;
    if there exists such a T and f then
        if a 2 → 3 move through f is possible then
            Do the 2 → 3 move, mark the new tetrahedra as coned, return True;
        end
        if a 3 → 2 move through f is possible then
            Do the 3 → 2 move, mark the new tetrahedra as coned, return True;
        end
    end
    /* If we have not yet returned, then the coned region has already
       been expanded as much as possible.                              */
    return False;
```

Algorithm 3: This function expands a coned region.

False. See Algorithm 4. Then we apply AttemptCancellation repeatedly until it returns False. Each application of EliminateOpaqueFace does a $2 \to 3$ or a special $4 \to 4$ move through an opaque face, i.e. a face of the triangulated boundary of a polyhedron. Put differently, a face is opaque if it is opposite the finite vertex of a coned polyhedron. The special $4 \to 4$ move is needed when the polyhedron's face is glued to itself in a special way, which we now describe.

def `EliminateOpaqueFace`(*Orbifold Triangulation orb*):

> **Result:** Do a $2 \to 3$ or a special $4 \to 4$ on an opaque face. Return True if we succeed, False if not, in which case all opaque faces were already eliminated.
>
> Look for any tetrahedron T which is coned and not flat.
>
> **if** *such a T exists* **then**
>
> > Let f be the face opposite the finite vertex of T and T' the tetrahedron glued to T along f;
> >
> > `/*` f `is opaque, meaning it belongs to the boundary of the polyhedron */`
> >
> > **if** *T' is not flat* **then**
> >
> > > Do a $2 \to 3$ move through f, return True ;
> >
> > **end**
> >
> > **if** *T is flat* **then**
> >
> > > Do a special $4 \to 4$ move through f, return True ;
> >
> > **end**
> >
> > Mark the newly created tetrahedra as un-coned.;
> >
> > `/* So that we skip these tetrahedra on the next application of EliminateOpaqueFace */`
>
> **end**
>
> return False ;

Algorithm 4: This function tries to eliminate one opaque face.

If a face of a polyhedron is glued to itself, then the gluing map is π rotation around some axis contained in the face. This axis must connect a vertex to another vertex, a vertex to

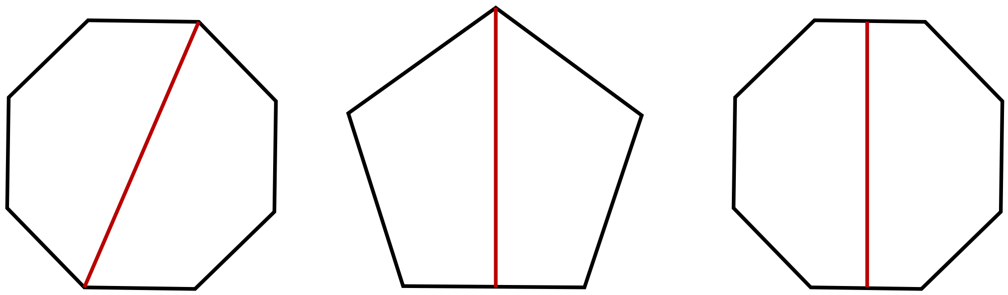

Figure 39: The three ways a face can be glued to itself.

an edge, or an edge to another edge. See figure 39.

Each face is triangulated, and the triangular faces should be glued to each other in a way which expresses how the polygonal face is glued to itself. However, in the third case this is only possible if there is a flat tetrahedron. This is an admissible flat tetrahedron, which we allow in proto-canonical triangulations. See figure 40. Therefore we do a special $4 \to 4$ move there, instead of a $2 \to 3$ move.

Assuming that step 1 succeeded, this step 2 algorithm should not fail. The $2 \to 3$ and special $4 \to 4$ moves connect finite vertices of adjacent polyhedra of the canonical decomposition with an edge, and the cancellation moves remove extra tetrahedra created by these moves which are not part of the canonical re-triangulation. The result is the canonical re-triangulation, so canonize part 2 is complete.

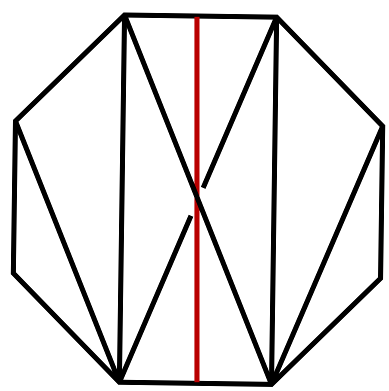

Figure 40: A triangulated polygonal face glued to itself.

7.0 A census of orbifolds commensurable to the figure eight knot complement

Two orbifolds are *commensurable* if they have a common finite sheeted cover. This is related to the notion of commensurability in group theory. Recall that two subgroups H_1, H_2 of some group G are *subgroup-commensurable* if $[H_1 : H_1 \cap H_2] < \infty$ and $[H_2 : H_1 \cap H_2] < \infty$. For finite volume hyperbolic 3-orbifolds $Q_1 = \mathbb{H}^3/\Gamma_1$ and $Q_2 = \mathbb{H}^3/\Gamma_2$, by Mostow-Prasad rigidity Q_1 is commensurable to Q_2 if and only if there exists a $g \in \mathrm{Isom}(\mathbb{H}^3)$ such that $g\Gamma_1 g^{-1}$ is subgroup-commensurable to Γ_2.

We define the *oriented commensurability category* $\mathcal{C}_3$ of the figure-eight knot complement M in the following way. Its objects are the orientation preserving isometry classes of oriented hyperbolic orbifolds commensurable to M. Its morphisms are the orientation preserving covering maps, defined up to pre- and/or post-composition with an isometry. In joint work, we have completely described $\mathcal{C}_3^{\leq 2v_0}$, the sub-category whose objects have hyperbolic volume at most $2v_0 = \mathrm{vol}(M)$, where $v_0 \approx 1.01494$ is the volume of a hyperbolic regular ideal tetrahedron. This project is in collaboration with Jason DeBlois, Anuradha Ekanayake, Tyler Gaona, Arshia Gharagozlou, and Priyadip Mondal. In this Chapter, we briefly describe some of the motivation and methods. In particular, Sym is heavily used to compute orbifold canonical decompositions.

7.1 Why study this commensurability class?

Understanding the commensurability class of M is one way to better understand M itself, which is a manifold of great importance both historically and at present. For instance:

1. Riley showed the existence of a hyperbolic structure on M at a time when few 3-manifolds were known to be hyperbolic [37].
2. M is the orientable non-compact hyperbolic 3-manifold of least volume [11].
3. M is the only arithmetic knot complement [36].

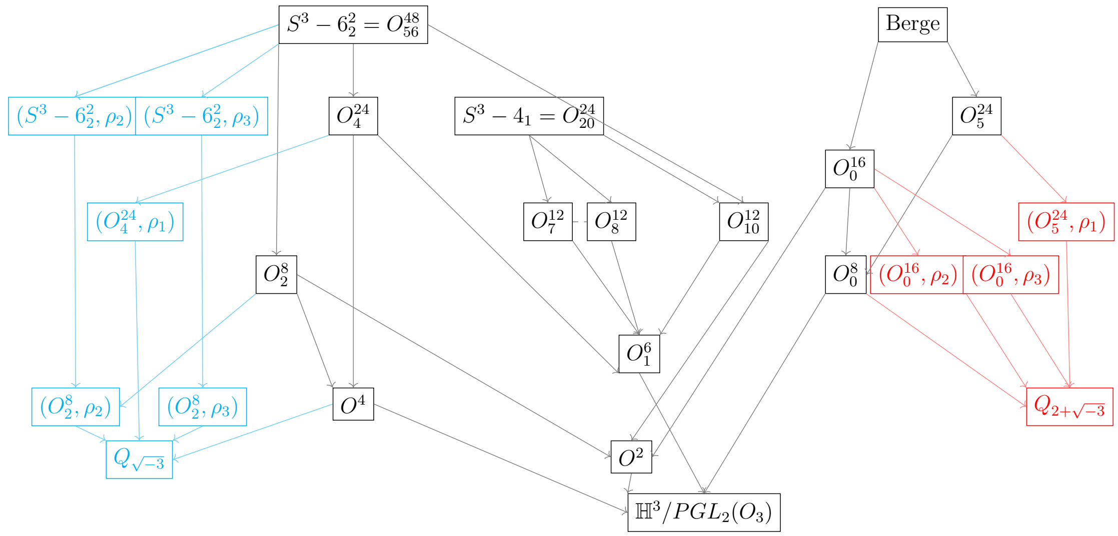

Figure 41: The covering lattice beneath three important manifolds.

Furthermore, this commensurability class contains the *Tetrahedral* manifolds, which are those hyperbolic manifolds admitting a triangulation by all regular ideal tetrahedra. A census of Tetrahedral manifolds was created by Fominykh, Garoufalidis, Goerner, Tarkaev, and Vesnin in [16]. Every orientable Tetrahedral manifold has volume $\geq 2v_0$, so our census is complementary to theirs, and one of our main goals was to sync up the two. As an example application, see figure 41, which shows the entire lattice of orientable orbifolds covered by at least one of M, $S^3 - 6^2_2$, or the Berge manifold.

Another motivation is to gain a better understanding of the collection of all low-volume cusped hyperbolic 3-orbifolds. It is well known that $\mathcal{C}_3$ contains minimal volume elements from many classes. As mentioned above, it contains the minimal volume orientable cusped hyperbolic 3-manifold, namely M itself. It also contains the minimal volume orientable one and two-cusped orbifolds [2], [3]. Tyler Gaona has independently shown that the minimal volume orbifold with one rigid and one smooth cusp also belongs to $\mathcal{C}_3$ (to appear in [18]). We can also recover the non-orientable orbifolds commensurable to M up to volume v_0, which includes other minimal volume examples such as the Gieseking manifold [1].

There is a sharp contrast between the commensurability class of an arithmetic orbifold and that of a non-arithmetic orbifold (see Machlachlan and Reid for an introduction to arithmetic hyperbolic 3-orbifolds [26]). The *commensurator* of a group $\Gamma \subset \mathrm{PSL}(2, \mathbb{C})$ is defined to be

$$\mathrm{Comm}(\Gamma) = \{g \in \mathrm{PSL}(2, \mathbb{C}) : [\Gamma : g\Gamma g^{-1} \cap \Gamma] < \infty\}.$$

Margulis showed that $\mathrm{Comm}(\Gamma)$ is discrete if and only if Γ is not arithmetic [27] (see also Zimmer [44]). Therefore if Γ is not arithmetic, then $\mathbb{H}^3/\,\mathrm{Comm}(\Gamma)$ is the unique minimal orbifold in the commensurability class of $\mathbb{H}^3/\Gamma$, where a *minimal orbifold* is an orbifold which does not cover any orbifold apart from itself. In [20], it was shown that in the non-arithmetic case geometric methods can be used to compute the commensurator, and hence determine if two non-arithmetic hyperbolic orbifolds are commensurable. On the other hand, there are many minimal orbifolds in the commensurability class of an arithmetic orbifold, which makes $\mathcal{C}_3$ very large and complicated.

7.2 Summary of the strategy

Our strategy is to enumerate finite covers of a particular well-known minimal orbifold in $\mathcal{C}_3$, which we call O^1, then to use canonical decompositions and number-theoretic methods to determine the rest of $\mathcal{C}_3^{\leq 2v_0}$. As we are working with orientable orbifolds, when we say "minimal orbifold" now, we mean an orbifold which covers no orientable orbifold but itself.

We now describe O^1. Denote the ring of integers of the quadratic number field $\mathbb{Q}(\sqrt{-3})$ as O_3. Then

$$\mathrm{GL}(2, O_3) = \{\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) : a, b, c, d \in O_3, ad - bc \in O_3^*\},$$

and $\mathrm{PGL}(2, O_3)$ is defined to be $\mathrm{GL}(2, O_3)/\{\left(\begin{smallmatrix} x & 0 \\ 0 & x \end{smallmatrix}\right) : x \in O_3^*\}$. Then $\mathrm{PGL}(2, O_3)$ is a discrete group of orientation-preserving isometries of $\mathbb{H}^3$, and we let O^1 be the quotient orbifold $\mathbb{H}^3/\,\mathrm{PGL}(2, O_3)$.

There is a geometric interpretation of $\mathrm{PGL}(2, O_3)$. As is classically known, there is a tiling $\mathcal{T}_3$ of the upper half space model of $\mathbb{H}^3$ by regular ideal tetrahedra for which the set of ideal vertices is $\mathbb{Q}(\sqrt{-3}) \cup \{\infty\}$. Then $\mathrm{PGL}(2, O_3)$ is exactly the group of orientation-preserving symmetries of this tiling (see Hatcher [22] for more details).

The tiling $\mathcal{T}_3$ projects to an orbifold triangulation of O^1 which we have mentioned before, in figure 4. It is canonical. With a computer search similar to the algorithm in [16], we construct finite covers of O^1 up to degree 48 using a sub-division of this triangulation. We then use the lifted triangulations to determine some covering maps. The category whose objects are these triangulated covers of O^1 and whose morphisms are triangulation-preserving maps we call $\mathcal{C}_{\text{main}}$. We label each object of $\mathcal{C}_{\text{main}}$ in the form O^n_k, which corresponds to the kth degree n cover. The covers of a fixed degree are ordered using the lexicographic order on *destination sequences*, where the destination sequence of O^n_k is a list of integers which encodes its triangulation lifted from O^1, similar to Burton's destination sequences [10]. If there is a single degree n cover then we omit a subscript and just call it O^n (e.g. O^1 rather than O^1_1). The remaining work is to extrapolate the data of $\mathcal{C}_3^{\leq 2v_0}$ from the data of $\mathcal{C}_{\text{main}}$.

The finite covers of O^1 correspond to finite index subgroups of $\mathrm{PGL}(2, O_3)$, hence $\mathcal{T}_3$ again descends to orbifold triangulations of these covers. However, these orbifold triangulations might not be canonical. Finding their canonical decompositions helps us understand $\mathcal{C}_3^{\leq 2v_0}$ in the following ways.

1. Two covers of O^1 could be isometric, but not isomorphic as objects of $\mathcal{C}_{\text{main}}$. This means that to determine which orbifolds are isometric, we need to do more than just generate triangulated covers by computer search as described above.
2. Similarly, the self-isometries of one of our covers might not preserve the lifted triangulation, so we need the canonical decomposition to determine the full isometry group. It is nice to know the full isometry group, but we also need it for more pressing reasons. It can be seen with arithmetic methods that every orbifold commensurable to M is regularly covered by some cover of O^1. Hence to find every object of $\mathcal{C}_3^{\leq 2v_0}$, our strategy is to find its regular cover as an object of $\mathcal{C}_{\text{main}}$ and the appropriate isometries to quotient by.

The arithmetic methods used in the second part above are an application of Borel's classification of minimal orbifolds in arithmetic commensurability classes [8], to understand minimal orbifolds in $\mathcal{C}_3^{\leq 2v_0}$. In general, there is a minimal orbifold Q_π of $\mathcal{C}_3$ for each prime or product of distinct prime ideals of O_3, which is indicated here by a generator π. From this perspective it turns out there is a canonical choice of regular cover $\tilde{Q}_\pi \to Q_\pi$ with $\tilde{Q}_\pi$ in

$\mathcal{C}_{\text{main}}$. We can then use what we know about the data of the lattice of covers above $\tilde{Q}_\pi$ to extrapolate the data about the lattice of covers above Q_π, completing the census.

To explore the two points listed above, Sym has a function which can compute all combinatorial maps from one orbifold triangulation to another. Then, to determine if two covers of O^1 are isometric, we just need to find their canonical decompositions, then use this function to see if they each have a combinatorial map to each other. In that case, they cover each other, so must be isometric. Similarly, this function explicitly finds all orientation-preserving self-isometries of an orbifold by finding all combinatorial maps from its canonical decomposition to itself.

As a sanity check that Sym is working correctly when applied to this project, we note the following.

- For any fixed prime π, covering space theory tells us that each cover of $\tilde{Q}_\pi$ will either have its triangulation-preserving symmetry group of index two in the full isometry group or will be non-triangulation preserving isometric to exactly one other cover of $\tilde{Q}_\pi$. This is exactly what Sym shows us.
- For $Q_{\sqrt{-3}}$ and Q_2, the two smallest Q_π, we used polyhedral decompositions from the literature to describe orbifold fundamental group presentations to feed into GAP's [19] LowIndexSubgroups routine, thereby independently recovering the structure of these orbifolds' lattices of covers. The lattice structure we generate this way matches the one we get using Sym and covering space theory.

Bibliography

[1] Colin C. Adams. The noncompact hyperbolic 3-manifold of minimal volume. *Proc. Amer. Math. Soc.*, 100(4):601–606, 1987.

[2] Colin C. Adams. Noncompact hyperbolic 3-orbifolds of small volume. In *Topology '90 (Columbus, OH, 1990)*, volume 1 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 1–15. de Gruyter, Berlin, 1992.

[3] Colin C. Adams. Volumes of hyperbolic 3-orbifolds with multiple cusps. *Indiana Univ. Math. J.*, 41(1):149–172, 1992.

[4] Alejandro Adem, Johann Leida, and Yongbin Ruan. *Orbifolds and stringy topology*, volume 171 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2007.

[5] Hirotaka Akiyoshi. Finiteness of polyhedral decompositions of cusped hyperbolic manifolds obtained by the Epstein-Penner's method. *Proc. Amer. Math. Soc.*, 129(8):2431–2439, 2001.

[6] Riccardo Benedetti and Carlo Petronio. *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin, 1992.

[7] Michel Boileau, Sylvain Maillot, and Joan Porti. *Three-dimensional orbifolds and their geometric structures*, volume 15 of *Panoramas et Synthèses [Panoramas and Syntheses]*. Société Mathématique de France, Paris, 2003.

[8] A. Borel. Commensurability classes and volumes of hyperbolic 3-manifolds. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)*, 8(1):1–33, 1981.

[9] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[10] Benjamin A. Burton. The Pachner graph and the simplification of 3-sphere triangulations. In *Computational geometry (SCG'11)*, pages 153–162. ACM, New York, 2011.

[11] Chun Cao and G. Robert Meyerhoff. The orientable cusped hyperbolic 3-manifolds of minimum volume. *Invent. Math.*, 146(3):451–478, 2001.

[12] Daryl Cooper, Craig D. Hodgson, and Steven P. Kerckhoff. *Three-dimensional orbifolds and cone-manifolds*, volume 5 of *MSJ Memoirs*. Mathematical Society of Japan, Tokyo, 2000. With a postface by Sadayoshi Kojima.

[13] Marc Culler, Nathan M. Dunfield, Matthias Goerner, and Jeffrey R. Weeks. SnapPy, a computer program for studying the geometry and topology of 3-manifolds. Available at `http://snappy.computop.org`.

[14] William D. Dunbar and G. Robert Meyerhoff. Volumes of hyperbolic 3-orbifolds. *Indiana Univ. Math. J.*, 43(2):611–637, 1994.

[15] D. B. A. Epstein and R. C. Penner. Euclidean decompositions of noncompact hyperbolic manifolds. *J. Differential Geom.*, 27(1):67–80, 1988.

[16] Evgeny Fominykh, Stavros Garoufalidis, Matthias Goerner, Vladimir Tarkaev, and Andrei Vesnin. A census of tetrahedral hyperbolic manifolds. *Exp. Math.*, 25(4):466–481, 2016.

[17] Roberto Frigerio and Carlo Petronio. Construction and recognition of hyperbolic 3-manifolds with geodesic boundary. *Trans. Amer. Math. Soc.*, 356(8):3243–3282, 2004.

[18] Tyler Gaona. *On hyperbolic 3-orbifolds of small volume.* PhD thesis, The University of Pittsburgh, 2022.

[19] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.11.1*, 2021.

[20] Oliver Goodman, Damian Heard, and Craig Hodgson. Commensurators of cusped hyperbolic manifolds. *Experiment. Math.*, 17(3):283–306, 2008.

[21] C. McA. Gordon and J. Luecke. Knots are determined by their complements. *J. Amer. Math. Soc.*, 2(2):371–415, 1989.

[22] Allen Hatcher. Hyperbolic structures of arithmetic type on some link complements. *J. London Math. Soc. (2)*, 27(2):345–355, 1983.

[23] Damian Heard. *Computation of hyperbolic structures on 3-dimensional orbifolds*. PhD thesis, The University of Melbourne, 2005.

[24] Sadayoshi Kojima. Polyhedral decomposition of hyperbolic 3-manifolds with totally geodesic boundary. In *Aspects of low-dimensional manifolds*, volume 20 of *Adv. Stud. Pure Math.*, pages 93–112. Kinokuniya, Tokyo, 1992.

[25] W. B. R. Lickorish. Simplicial moves on complexes and manifolds. In *Proceedings of the Kirbyfest (Berkeley, CA, 1998)*, volume 2 of *Geom. Topol. Monogr.*, pages 299–320. Geom. Topol. Publ., Coventry, 1999.

[26] Colin Maclachlan and Alan W. Reid. *The arithmetic of hyperbolic 3-manifolds*, volume 219 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2003.

[27] G. A. Margulis. *Discrete subgroups of semisimple Lie groups*, volume 17 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991.

[28] Bruno Martelli. An introduction to geometric topology. preprint, 2016, arxiv:1610.02592.

[29] Sergei Matveev. *Algorithmic topology and classification of 3-manifolds*, volume 9 of *Algorithms and Computation in Mathematics*. Springer, Berlin, second edition, 2007.

[30] Aleksandar Mijatović. Simplifying triangulations of S^3. *Pacific J. Math.*, 208(2):291–324, 2003.

[31] I. Moerdijk and D. A. Pronk. Orbifolds, sheaves and groupoids. *K-Theory*, 12(1):3–21, 1997.

[32] G. D. Mostow. *Strong rigidity of locally symmetric spaces.* Annals of Mathematics Studies, No. 78. Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1973.

[33] Udo Pachner. P.L. homeomorphic manifolds are equivalent by elementary shellings. *European J. Combin.*, 12(2):129–145, 1991.

[34] Gopal Prasad. Strong rigidity of $\mathbf{Q}$-rank 1 lattices. *Invent. Math.*, 21:255–286, 1973.

[35] John G. Ratcliffe. *Foundations of hyperbolic manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.

[36] Alan W. Reid. Arithmeticity of knot complements. *J. London Math. Soc. (2)*, 43(1):171–184, 1991.

[37] Robert Riley. Discrete parabolic representations of link groups. *Mathematika*, 22(2):141–150, 1975.

[38] I. Satake. On a generalization of the notion of manifold. *Proc. Nat. Acad. Sci. U.S.A.*, 42:359–363, 1956.

[39] Ichirô Satake. The Gauss-Bonnet theorem for V-manifolds. *J. Math. Soc. Japan*, 9:464–492, 1957.

[40] William P. Thurston. The geometry and topology of 3-manifolds. Mimeographed lecture notes., 1979.

[41] William P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.

[42] William P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy.

[43] Jeffrey R. Weeks. Convex hulls and isometries of cusped hyperbolic 3-manifolds. *Topology Appl.*, 52(2):127–149, 1993.

[44] Robert J. Zimmer. *Ergodic theory and semisimple groups*, volume 81 of *Monographs in Mathematics*. Birkhäuser Verlag, Basel, 1984.