On a Bounded Runtime Algorithm for the Homeomorphism Problem for Hyperbolic Manifolds



Joe Scull
Mansfield College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

2022

Acknowledgements

First and foremost I want to thank my husband SJ Spencer, at least 50 percent of the work of completing this PhD was in finding the willpower to write and rewrite once the "maths bit" was done. SJ gave me the emotional and existential support required to turn all the maths I'd worked out into documents other people can now read, so if you're reading this or my papers, then you have him to thank.

The other half of the work was actually working out the maths. In this endeavour I must first thank my supervisor Marc Lackenby. Marc would make a great supervisor to most mathematicians, but for me he was everything I needed him to be. When I had questions he gave me answers, when I'd run out of questions he gave me more, and when I'd decided I didn't want to answer any more questions, he gave me license to stop. Marc took the job of supervising me seriously, in a way many supervisors do not. For that and much more, he deserves my thanks.

I also want to thank the members of the topology group here at Oxford. Though there's much about this university that doesn't make sense to me, its ability to bring together a bunch of interesting and interested mathematicians to learn together was invaluable to me in finishing this degree. From the first moments learning everything I could from anyone who wanted to tell me, to the unlocking of a difficult proof stood at a whiteboard together, you were all invaluable.

In particular, I want to thank Jan Steinebrunner, a great friend and housemate, but more relevantly a great mathematician. He answered so many of my questions, and always did so with a genuine curiosity and excitement that was infectious. There was a week in the middle where I would go to his office every 2 hours and ask him a question, then spend the next 2 hours understanding what he'd told me, on and off for the whole working week, it was the most progress I made in a single week for my whole PhD.

I also want to acknowledge that the work in this thesis has appeared before in other forms, Chapters 4 and 5 in particular contain work from my two papers, which can be found on the arXiv. The work has been edited to fit the story, but especially in the case of proofs, which have already been peer-reviewed, they appear here exactly as they do in the papers.

Abstract

We give a bounded runtime solution to the homeomorphism problem for closed hyperbolic 3-manifolds. This is an algorithm which, given two triangulations of hyperbolic 3-manifolds by at most t tetrahedra, decides if they represent the same hyperbolic 3-manifold with runtime bounded by

$$2^{2^{t^{O(t)}}}.$$

Along the way we prove two supplementary results, both of which we also prove in higher dimensions. First, we prove that one can compare geometric triangulations of hyperbolic n-manifolds in time bounded by

$$2^{2^{nt^{O(nt^3)}}}, \quad n \geq 4, \qquad 2^{t^{O(t)}}, \quad n = 3.$$

Second, we show that the length R of a systole of a closed or finite hyperbolic n-manifold ($n \geq 3$) admitting a triangulation by t n-simplices can be bounded below by a function of n and t, namely

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}.$$

Contents

1 Introduction **1**

1.1 Introduction 1

1.2 Future Work 5

2 Background **7**

2.1 Triangulations 7

2.1.1 Simplicial Complexes 7

2.1.2 Pachner Moves 10

2.1.3 Semi-Ideal Triangulations 14

2.1.4 Cocycles and Representations 15

2.2 Hyperbolic Geometry 16

2.2.1 The Models of Hyperbolic Geometry 16

2.2.1.1 Riemannian Metrics 16

2.2.1.2 Isometries between Models 17

2.2.1.3 Distance Functions for the Models 18

2.2.2 Thick-Thin Decomposition 19

2.2.3 Mostow and Other Rigidities 21

2.2.4 Geometric Inequalities and Dimensional Differences 24

2.3 3-Manifolds 28

2.3.1 Geometric Decomposition 28

2.3.2 Algorithms for 3-manifolds 29

2.3.3 The Work of Kuperberg 31

2.4 Systems of Polynomials 34

2.4.1 An Algorithm for Finding Solutions 34

2.4.2 Poincare's Polyhedron Theorem 37

3 An Algorithm for Geometric Triangulations **40**
3.1 Kalelkar and Phanse's Work in a Simple Case 40
3.2 Bounding Simplex Volume 49
3.3 A Combinatorial Bound for Geometric Triangulations 53

4 Injectivity Radius **72**
4.1 Introduction 72
4.2 Diameter of Margulis Tubes 75
4.3 Bounded Cocycles 80
4.4 The Cusped Case 88

5 The Main Algorithm **95**
5.1 Finding a Geometric Quotient 95
5.2 Finding the Geometric Triangulation 104
5.3 Comparing Geometric Triangulations 120

Bibliography **121**

Chapter 1

Introduction

1.1 Introduction

A natural question in the study of manifolds is whether two manifolds from a given class $\mathcal{C}$ are homeomorphic. More precisely, we ask whether an algorithm exists which, given two triangulations and the promise that the triangulations are manifolds in $\mathcal{C}$, can decide whether or not they are homeomorphic. This is called the homeomorphism problem for $\mathcal{C}$.

In this thesis we shall focus on classes of orientable manifolds and from now all manifolds are assumed to be orientable without comment.

In dimensions 1 and 2, the homeomorphism problem for closed connected manifolds is solved quite simply, as there is only one such 1-manifold and closed connected 2-manifolds are classified by their Euler characteristic and their orientability. In dimensions 4 or above, a result of Markov [Mar58] states that the homeomorphism problem, even for closed connected manifolds, is unsolvable.

In dimension 3 the problem hits a sweet spot of being tractable while still being quite difficult. For example all known solutions to the homeomorphism problem for 3-manifolds use Thurston's Geometrisation Conjecture [Thu82], and thus require Perelman's resolution of the conjecture [Per02][Per03b][Per03a].

In [Sel95] a solution to the homeomorphism problem for closed connected oriented 3-manifolds is given as a corollary of Thurston's Geometrisation before Pereleman had proved it. A more accessible description of this algorithm is given in the book of Bessieres et al. [BBB+10], citing both a series of lectures by Jaco and Sela's work. In these algorithms, hyperbolic manifolds are dealt with using Sela's solution to the isomorphism problem for torsion free hyperbolic groups. Scott and Short also provide an alternative algorithm [SS14] which avoids Sela's work.

With the decidability of the homeomorphism problem settled, the next question is naturally whether we can give bounds on runtime. Partial results for certain subclasses of 3-manifolds had previously been given by Mijatovic [Mij03][Mij04] [Mij05a] [Mij05b]. Currently, the best known result is given by Kuperberg [Kup19] in which he proves that the homeomorphism problem for closed oriented 3-manifolds admits a bounded runtime solution, though the bound is partially unknown. That is, Kuperberg's solution has runtime bounded by a tower of exponentials in the number of tetrahedra, but the height of this tower is unknown. The question asked at the end of Kuperberg's paper is what height might such a tower have, whether for his or an improved algorithm?

A common thread in all the above approaches is that they deal with hyperbolic 3-manifolds using a bespoke algorithm. In Kuperberg's case for example, he algorithmically finds the geometric decomposition of the 3-manifold and then compares the (Seifert fibred or hyperbolic) pieces. Thus a solution to the homeomorphism problem for all closed 3-manifolds seems destined to require one for hyperbolic 3-manifolds.

In the following paper we present a solution to the homeomorphism problem for closed hyperbolic 3-manifolds with runtime bounded by a tower of exponentials of height 4.

Theorem 1.1.1. *If M_1 and M_2 are triangulated closed hyperbolic 3-manifolds, each triangulated by less than t tetrahedra, then we can decide in time bounded by*

$$2^{2^{t^{O(t)}}}$$

whether or not they are homeomorphic.

Note here that the algorithm which decides whether the manifolds are homeomorphic takes as input solely the combinatorial data of the triangulations. In particular it requires no information regarding the hyperbolic structure. This means that the algorithm we create here needs to be capable of finding hyperbolic structures as well as comparing them.

As part of the proof of Theorem 1.1.1 for example we prove the following result, adapted from Theorem 5.2.15, which gives a bounded runtime algorithm for finding the hyperbolic structure of a given hyperbolic 3-manifold M, in the form of a geometric triangulation of M. Finding the hyperbolic structure of a hyperbolic 3-manifold was already proven to be algorithmically possible by Manning [Man01] (Scott and Short's Paper also touches on this [SS14]) and of course Kuperberg's paper [Kup19] provides a bounded runtime algorithm. The following provides the best known bound on runtime for such an algorithm.

Theorem 1.1.2. *Given a triangulation of a closed hyperbolic 3-manifold M by t tetrahedra, there is an algorithm which produces a geometric triangulation of M consisting of less than*

$$2^{t^{O(t)}}$$

tetrahedra in time bounded by

$$2^{2^{t^{O(t)}}}.$$

Furthermore, this geometric triangulation has the property that all edges have length less than the injectivity radius of the manifold.

Theorem 1.1.2 requires the following two results which may be of independent interest.

Theorem 1.1.3. *Given a triangulation of a closed hyperbolic 3-manifold M by t tetrahedra, M admits a geometric triangulation where the number of tetrahedra in the geometric triangulation is bounded by*

$$2^{t^{O(t)}}.$$

Furthermore there is some subdivision of the triangulation of M and a surjective degree one simplicial map from this subdivision to our geometric triangulation. The size of the subdivision is bounded similarly.

Theorem 1.1.3 is adapted from Corollary 5.1.12 which proves a more general result for closed hyperbolic n-manifolds, and in it we show that the geometric triangulation is in fact a simplicial quotient of a subdivision of M. We shall define simplicial quotients later, but in particular the simplicial map mentioned in the above is non-degenerate on all simplices.

Finally, to find a simplicial quotient of a certain subdivision, we first need to be able to find this subdivision in bounded runtime. Such an algorithm exists in the works of Kuperberg [Kup19] and Kalelkar and Phanse [KP21].

The following is adapted from Theorem 3.1.14.

Theorem 1.1.4. *Suppose we have a closed 3-manifold with a triangulation $\mathcal{T}_1$ by t tetrahedra and a subdivision $\mathcal{T}_2$ of $\mathcal{T}_1$ by T tetrahedra. Then $\mathcal{T}_1$ and $\mathcal{T}_2$ are related by a sequence of Pachner moves of length bounded by*

$$O(tT).$$

In fact, we define composite moves (combinations of Pachner moves) which preserve subdivisions, and so we can list all subdivisions of a certain complexity by performing all sequences of these composite moves with the length of the sequence bounded as above.

On route to proving Theorem 1.1.1, we need some auxiliary results which are themselves of independent interest. Our overall algorithm takes a pair of triangulations, finds a geometric structure on each and then uses bounds on the geometry of those structures to compare them. Much of our work (see Chapter 5) is in finding a geometric triangulation, but we also need to be able to compare geometric triangulations and understand their geometry. These results, along with Theorem 1.1.4 make up our Chapters 3 and 4.

In fact, we can prove these auxiliary results in higher dimensions and do so. In Chapter 3, we improve upon the results of Kalelkar and Phanse to compare geometric triangulations of hyperbolic n-manifolds, recreating their results in dimension 3. We show that it is no longer necessary to know geometric information about the triangulations in advance.

Theorem 1.1.5. *Let M be a closed hyperbolic n-manifold, $n \geq 3$ and let $\mathcal{T}_1, \mathcal{T}_2$ be two geometric simplicial triangulations of M, consisting of t_1, t_2 tetrahedra respectively. Then, taking t to be the maximum of t_1, t_2, there is a function $f(t, n)$ which bounds the number of Pachner moves necessary to take $\mathcal{T}_1$ to $\mathcal{T}_2$.*

In particular

$$f(t, 3) \leq 2^{2^{O(t^2)}}$$

and

$$f(t, n) \leq 2^{2^{tn^{O(tn^3)}}}$$

for $n \geq 4$.

The other major result that we prove in higher dimensions but use only in 3 dimensions, and the focus of Chapter 4, is on the injectivity radius of hyperbolic manifolds. In fact, the proof we use for the n-dimensional finite volume result here is overpowered for use in the closed 3-dimensional case, and a slightly better bound can be found in dimension 3 using simpler methods.

Theorem 1.1.6. *Given a finite volume hyperbolic n-manifold ($n \geq 3$) with a simplicial triangulation by t semi-ideal n-simplices, let R be the length of a systole in M, then*

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}.$$

We prove the results of chapters 3 and 4 in higher dimensions, mostly to show that it's possible and because these results are of independent interest. However, as the following subsection suggests, there may well be a future use for having these auxiliary results in the strength they are here.

1.2 Future Work

One obvious next step for this work would be to extend this to an algorithm for the general case of closed 3-manifolds. Kuperberg already lays out methods which one might be able to refine and find explicit bounds for, just like how we have in this thesis for the closed hyperbolic case.

On one hand there is no reason to believe any of our results would be useful to the case of Seifert fibred spaces. On the other, the first obvious question would be the case of finite volume hyperbolic 3-manifolds, as these arise naturally in Thurston's geometric decomposition. When we prove results in this thesis on injectivity radius, we in fact prove results about systole length in order to extend them to the finite volume case. Furthermore, the generality of the methods used in this work suggest that the case of n-dimensional hyperbolic manifolds (either closed or finite volume) may be tractable. Again, if we don't require bounded runtime, the homeomorphism problem is already known to be decidable for this class of manifolds, via an application of a solution to the isomorphism problem for their fundamental groups [DGr08]. The author believes these methods could well be used to prove a similar result for higher dimensional manifolds, and those of finite volume. In particular, we claim that there are very few parts of our proof in which new techniques would be necessary to upgrade them to higher dimensions and cusped manifolds.

However, that does not mean that upgrading would be a simple task. As we demonstrate throughout this thesis, a lot of the work of using our methods is in the meticulous description of systems of polynomials. As one increases dimensions or adds the complexity of cusps, we lose some of the shortcuts we've used in the closed 3-dimensional case. This means that not only do the systems themselves get much bigger, but so does the work of describing them.

One example of a shortcut we've used which means extra work required in higher dimensions is that higher dimensions introduce the possibility of orbifolds. In higher dimensions Poincaré's Polyhedron Theorem (Section 2.4.2 doesn't guarantee a hyperbolic manifold in the same way it does in 3 dimensions, providing only a hyperbolic orbifold. This clearly presents problems as orbifolds are much harder to work with combinatorially. We believe these problems are assailable but do require a good amount of work.

It is because of this possible future work that we sometimes prove or cite results in more generality than is necessary, to pave the way for future work. The hope is that the reader can view this paper as a blueprint for how one might go on to solve the finite volume problem in all dimensions.

Chapter 2

Background

2.1 Triangulations

2.1.1 Simplicial Complexes

We start here with a definition due to Thurston [Thu97], who calls this concept a gluing, which is a generalisation of a triangulation.

Definition 2.1.1 (Gluing)**.** An n-dimensional (rectilinear) gluing consists of a finite set of n-simplices, a choice of pairs of codimension one faces such that each face appears in exactly one of the pairs and an affine identification map between the faces of each pair. We also sometimes call the quotient space derived from this process a gluing as well.

Definition 2.1.2 (Simplicial Complex)**.** A simplicial complex K is a locally finite collection of simplices satisfying the following two conditions:

- Any face (subsimplex) of a simplex in K is itself in K.
- The intersection of two simplices in K is either empty or a common face of both.

The union of all the simplices is called the realisation of K and is denoted by $|K|$.

Note that taking the double barycentric subdivision of a gluing gives a simplicial complex.

This means that if we are happy to do a double barycentric subdivision, we can simply assume that we are always given a manifold in the form of a finite simplicial complex. We now define some basic terms for dealing with simplicial complexes in the context of 3-manifolds.

Definition 2.1.3 (Star and Link)**.** The star $\mathrm{st}(\Delta)$ of a simplex Δ in a simplicial complex K is the smallest subcomplex of K containing all simplices which have Δ as a face. The link $\mathrm{lk}(\Delta)$ of a simplex Δ in a simplicial complex K consists of all simplices in $\mathrm{st}(\Delta)$ which do not contain Δ.

Definition 2.1.4 (Subdivision)**.** A subdivision of a simplicial complex K is a simplicial complex K' such that $|K| = |K'|$ and each simplex of K' is contained within a simplex of K.

Definition 2.1.5 (Simplicial Isomorphism)**.** Let $f : |K| \to |L|$ be a map, we say that f is simplicial if for all simplices Δ of K, $f\mid_\Delta$ is affine and $f(\Delta)$ is a simplex in L. A homeomorphism which is simplicial will be called a simplicial isomorphism and similarly we say two simplicial complexes are simplicially isomorphic if such a homeomorphism exists.

Definition 2.1.6 (Simplicial Quotient)**.** Given an n-dimensional gluing M, and a collection of simplicial identifications $\Delta \to \Delta'$ for Δ, Δ' n-simplices in M, consider the quotient space Q of all these identifications. If Q is itself an n-dimensional gluing then we call it a simplicial quotient of M.

Remark 2.1.7. We won't talk about simplicial quotients for a while, but will later need to be able to list degree one simplicial quotients of certain subdivisions of triangulations. Here degree one simply means that the signed number of preimages of any given simplex is exactly 1. This presupposes an orientation on both the quotient and original triangulation.

Definition 2.1.8 (Combinatorial Equivalence)**.** We declare two simplicial complexes to be combinatorially equivalent if there exist subdivisions of each which are simplicially isomorphic.

Definition 2.1.9 (Combinatorial n-Manifold)**.** A combinatorial n-manifold is a simplicial complex such that the link of every vertex is combinatorially equivalent to the boundary of the standard n-simplex.

Definition 2.1.10 (Triangulation)**.** A triangulation of an n-manifold M is simply a gluing paired with a homeomorphism to M. If the gluing is a simplicial complex we call the triangulation simplicial. If the gluing is a combinatorial manifold, then we call the triangulation combinatorial.

Definition 2.1.11 (Piecewise Linear)**.** A manifold equipped with a homeomorphism to a combinatorial manifold is called a PL manifold. A PL (or Piecewise Linear) homeomorphism is a homeomorphism between combinatorial n-manifolds which is simplicial with respect to some subdivision of each. We call two such manifolds PL equivalent, and sometimes call this homeomorphism a PL equivalence

Note that it is possible to define PL manifolds similarly to smooth manifolds - via a system of charts where the transition maps are in some sense piecewise linear. It turns out that this chart based definition is equivalent to the definition above so we omit it here (shown in chapter 3 of [HSoM67]).

With the above definition of a PL manifold in mind, we now state a theorem which will inspire a lot of what is to come in this thesis. A Pachner move is a combinatorial move that changes one simplicial complex to another, we shall see a formal definition in the next subsection.

Theorem 2.1.12 (Theorem 5.5 [Pac91])**.** *Two closed combinatorial n-manifolds are PL homeomorphic if and only if they are related by a sequence of Pachner moves and simplicial isomorphisms.*

This statement is exciting as it allows us to change from a manifold perspective (a system of charts to $\mathbb{R}^n$) to a combinatorial perspective (performing moves on simplicial complexes). However, it would appear to pose an immediate problem by introducing the requirement of a PL equivalence rather than simply a homeomorphism. Fortunately this can be solved in two different ways. Firstly, in 3 dimensions the concepts can be proven to be completely indistinguishable.

Theorem 2.1.13 (Theorems 1, 3 and 4 [Moi52])**.** *Every closed 3-manifold is triangulable, every simplicial triangulation of a closed 3-manifold is combinatorial and any two triangulations of a closed 3-manifold are PL homeomorphic.*

Thus, when talking about triangulated 3-manifolds, we only need to be able to discern whether triangulations are equivalent, a much easier task than searching for an explicit homeomorphism. Alternatively, as we will be dealing mostly with hyperbolic manifolds, we can take another route to deal with the fact that Pachner's theorem only applies to PL homeomorphic manifolds. Fortunately, all hyperbolic manifolds admit geometric triangulations (see section 2.2 for a precise definition) and any geometric triangulation is combinatorially equivalent to any other (see 2.1.20). We will see later (Chapter 5) that when given a non-geometric triangulation of a hyperbolic manifold we can find a geometric one and so we can apply Pachner's theorem to these geometric triangulations instead.

2.1.2 Pachner Moves

Pachner's methods are more geometric than combinatorial, so we won't use his methods explicitly in what follows, but his theorem points us towards the strategy of trying to find sequences of Pachner moves between manifolds.

Definition 2.1.14 (Pachner Move)**.** Given a combinatorial n-manifold K, and some combinatorial n-disc D which is a subcomplex both of K and of the boundary of the $n+1$-simplex Δ, then a Pachner move consists of replacing D by $\partial\Delta - D$.

A Pachner move is defined by the number of simplices which make up D, and so it makes sense to refer to a (1-4) move for example which replaces one 3-simplex with four 3-simplices.

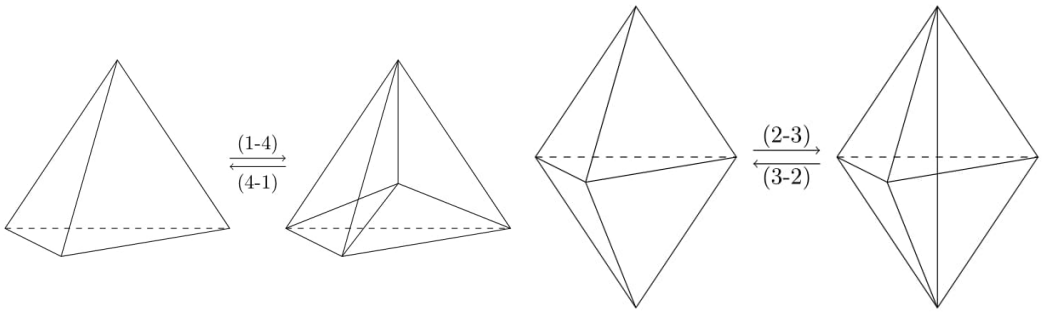


Figure 2.1: All four 3-dimensional Pachner moves

We shall return later to an interesting paper of Kalelkar and Phanse [KP21], but a major insight in that paper is the application of Theorem 2.1.17, due to Adiprasito and Benedetti [AB17] in conjunction with Theorem 2.1.18, due to Lickorish [Lic99]. This canny combination of ensuring shellability by using subdivisions and then working with a shellable n-simplex allows us to describe sequences of Pachner moves without really knowing much about the starting triangulations.

This single idea is very powerful, and Kalelkar and co-authors ([KP21], [KP20] and [KR20]) have used this and similar techniques to great effect in the past couple of years.

Definition 2.1.15 (Polytopal Subdivision)**.** A polytope P in $\mathbb{R}^n$ is the convex hull of finitely many points. A face F of P is either P itself or $P \cap H$ where H is some hyperplane such that it intersects P non-trivially and P lies entirely in one of the half spaces defined by H. A polytopal complex is a collection of polytopes in $\mathbb{R}^n$ which is closed under taking faces and such that two polytopes intersect in a common face or not at all. A polytopal subdivision C' of a polytopal complex C is a polytopal

complex which has the same underlying space as C and such that every face of C' lies in some face of C.

Definition 2.1.16 (Shellability)**.** A combinatorial n-ball is shellable if there exists an ordering of the n-simplices such that we can remove all but the final n-simplex in order, at each stage being left with a connected combinatorial n-ball.

Theorem 2.1.17 (Theorem A - [AB17])**.** *If C is any (polytopal) subdivision of a convex polytope, the second barycentric subdivision of C is shellable. If dim(C) = 3, already the first barycentric subdivision of C is shellable.*

On its face, it may not be clear why shellability is desirable, and hence why Adiprasito and Benedetti's theorem is so interesting. A simple lemma due to Lickorish provides a reason to search for shellability: shellable n-simplices are a small number of Pachner moves away from a much simpler subdivision, the cone.

Theorem 2.1.18 (Lemma 5.7 [Lic99])**.** *If X is a shellable combinatorial n-ball made up of k n-simplices, then X can be transformed into a cone on its boundary by a sequence of k Pachner moves.*

Proof. We use induction on k. If $k = 1$ then the (1-$(n+1)$) Pachner move suffices. Suppose the statement is true for all balls with fewer than k n-simplices, then number the n-simplices of X according to some shelling and consider the $k-1$ simplices labelled $2, \ldots, k$. By shellability, these form a combinatorial n-ball themselves and so by induction we can perform $k-1$ Pachner moves converting this ball to a cone on its boundary. At this point, the simplex labelled 1 shares a face with $j \leq n$ simplices of this ball, and performing a $((j+1)$ - $(n+1-j))$ move respectively leaves us with a cone on ∂X. □

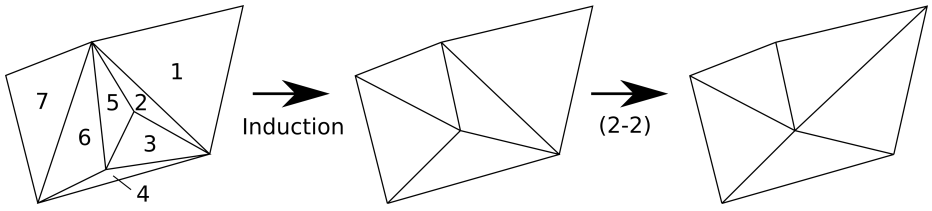


Figure 2.2: A particularly simple example of coning a shelling.

A combination of the above two results presents a strategy of coning off simplices to produce "coned subdivisions". Details, including a definition of a coned subdivision can be found in Section 3.1. The key is that moving between coned subdivisions

is substantially easier as a coned subdivision of an n-simplex is determined by the subdivision of its boundary, and so we can reduce the dimension of the problem and care only about subdivisions of 2-complexes, more on this later.

Another noted use of this strategy of coning simplices in order to perform Pachner Moves was in the work of Mijatovic. He proved in a series of papers that one could bound the number of Pachner moves between triangulations of certain subclasses of 3-manifolds. In particular his papers dealt with the 3-sphere [Mij03], Seifert-fibred spaces [Mij04] and knot complements [Mij05a] [Mij05b]. In these he used this idea of coning simplices and a process he called "the changing of cones" which is similar to the "coned subdivision" techniques we define in Definition 3.1.7.

The following lemma demonstrates how we link this bound on combinatorial complexity back to our search for algorithms. In fact it is a general theme that bounding combinatorial complexity allows for an algorithmic approach.

Lemma 2.1.19. *If there is a computable function $f(t) > t$ which bounds the number of Pachner moves required between two triangulations of a 3-manifold by at most t tetrahedra, then there is an algorithm which solves the homeomorphism problem and has runtime bounded on the order of $f(t)^{f(t)}$.*

Proof. Performing a Pachner move on a simplicial 3-complex introduces at most 3 new tetrahedra, hence the number of tetrahedra in a complex after $f(t)$ moves is at most $t + 3f(t)$. Furthermore, note the number of $((5-k)$-$k)$ moves possible at any stage is dependent on the number of $(k-1)$-simplices in the complex satisfying certain conditions. For example, a (3-2) move is possible only when there is an edge (a 1-simplex) such that its star consists of 3 tetrahedra. Thus the number of possible moves at any point is bounded by the sum of the number of k-simplices for all k, which is bounded by some constant multiple of the number of tetrahedra. Thus the number of possible sequences of Pachner moves is bounded by

$$O\left((t + 3f(t))^{f(t)}\right) = O(f(t)^{f(t)}$$

□

Clearly, this idea that a bound on combinatorial distance implies a bound on runtime of a comparison algorithm can be proven in more generality. One can also imagine similar statements and proofs regarding Tietze transformations of hyperbolic group presentations or Reidemeister moves for knot projections.

Hence, in some cases, instead of searching for an algorithm explicitly, we search for a bound on Pachner moves and then implicitly use the above algorithm. In fact,

our final algorithm will use a combination of the above idea and a less combinatorial approach.

A notable use of this approach, and especially relevant to our study are the previously mentioned results of Kalelkar and his coauthors Phanse and Raghunath. Note their result holds for spherical and Euclidean manifolds though we shall only use the hyperbolic result.

Theorem 2.1.20 (Kalelkar Phanse [KP21]). *Let M be a compact spherical, Euclidean or hyperbolic n-manifold with geometric simplicial triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ with t_1 and t_2 n-simplices respectively with lengths of edges bounded above by L and let $inj(M)$ be the injectivity radius of M. When M is spherical, we insist that $L \leq \frac{\pi}{2}$. Then $\mathcal{T}_1$ and $\mathcal{T}_2$ are related by $c(m,n)t_1t_2(t_1+t_2)$ Pachner moves where $c(m,n) = (2n+2)!^2(n+1)!^{3m}$ and m is an integer greater than $\mu log(\frac{2L}{inj(M)})$ where μ is as follows:*

1. *When M is Euclidean, $\mu = n+1$*
2. *When M is spherical, $\mu = 2n+1$*
3. *When M is hyperbolic, $\mu = ncosh^{n-1}(L)+1$*

We will later prove our own even more restricted version of Kalelkar and Phanse's result which improves our final algorithm in the 3-dimensional case. We also cite some other major results of theirs as, though we wont use their results explicitly, the work in their papers serves as inspiration for many of the ideas in this thesis.

A question one might ask when presented with such a result relating geometric triangulations, is whether all the Pachner moves preserve the geometric triangulation. If you follow the proof given in this thesis or in Kalelkar and Phanse's original paper, the Pachner moves necessarily distort the geometric structure, a hint at which can be seen in Figure 3.3. However, Kalelkar and Phanse followed their paper up with a short paper addressing this question.

Theorem 2.1.21 ([KP20]). *Let M be a closed hyperbolic 3-manifold with geometric simplicial triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ with t_1 and t_2 tetrahedra respectively. Then $\mathcal{T}_1$ and $\mathcal{T}_2$ are related by a sequence of geometric Pachner moves.*

Kalelkar also extends his work to the 3-dimensional cusped case in his paper with Raghunath.

Theorem 2.1.22 ([KR20]). *Let M be a complete orientable cusped hyperbolic 3-manifold. Let $\mathcal{T}_1$, $\mathcal{T}_2$ be geometric ideal triangulations of M with at most t_1, t_2 tetrahedra respectively and all dihedral angles at least θ_0. Let $t = t_1 + t_2$, then the number of Pachner moves needed to relate $\mathcal{T}_1$ and $\mathcal{T}_2$ is less than*

$$2.8 \times 10^{12} \cdot \frac{t^{11/2}}{sin(\theta_0)^{12m+27/2}}$$

Before we move on, the following is a fun observation for which the author has found no practical use, we'll first state the fact and prove it and then talk about what pathology this perhaps strange lemma points at.

Lemma 2.1.23. *The number of k-simplices in a simplicial n-manifold M made up of t n-simplices is bounded below by a function of n which is independent of t. That function is*

$$\binom{n+2}{k+1}$$

Proof. A given n-simplex Δ in M has $n+1$ adjacent n-simplices along its codimension one faces. Each of these intersect with Δ only in this face by definition of a simplicial complex, and thus they are all distinct. Thus there are at least $n+2$ n-simplices in any such triangulation. The simplest thing for these remaining $n+1$ n-simplices is for them all to be glued to one another, which is in fact what occurs in the standard triangulation of the n-sphere coming from the boundary of the $(n+1)$-simplex. Thus we see that the minimal number of k-simplices in a simplicial n-manifold is the number of k-simplices in this canonical triangulation of the n-sphere or equivalently in the $(n+1)$-simplex, which is simply $\binom{n+2}{k+1}$. □

This silly little result masks the fact that clearly the number of k-simplices is not independent of t, it's just that t itself is bounded below by this function of n. This is really a demonstration of the well understood idea that simplicial complexes can't be "too small".

2.1.3 Semi-Ideal Triangulations

We will also want to talk about triangulations of cusped (that is finite volume) hyperbolic manifolds. This is often done through the use of ideal triangulations, though whether it is possible to use Pachner moves to traverse between ideal triangulations remains an open problem. For this thesis, we instead use semi-ideal triangulations which were suggested for this purpose by Kuperberg [Kup19]. We'll see throughout

their many advantages, including that they allow barycentric subdivisions and have concrete vertices whose positions can be modelled.

Definition 2.1.24 (Semi-ideal n-simplex)**.** A semi-ideal n-simplex is one which has had exactly one of its vertices removed. We call this vertex the ideal vertex. Any simplex which has not had any of its vertices removed shall be called non-ideal.

Definition 2.1.25 (Semi-Ideal Triangulation)**.** Let M be a finite volume hyperbolic n-manifold, then a semi-ideal triangulation of M is a simplicial complex where we remove some of the vertices before defining a homeomorphism to M such that each simplex has at most one ideal vertex. Note that semi-ideal triangulations may include semi-ideal simplices, non-ideal simplices or both.

2.1.4 Cocycles and Representations

One way in which we shall be finding bounds on the geometry of our triangulated manifolds is by finding bounded representations of their fundamental groups. In fact it will be much easier to work with a slightly more general definition than a representation, the cocycle. Cocycles induce representations and so can give us all the information of a representation but also relate more directly to the triangulations on which they are defined.

Definition 2.1.26 (Simplicial Path)**.** A sequence of oriented 1-simplices in a triangulation such that the end vertex of each simplex is the start vertex of the next is called a simplicial path.

Definition 2.1.27. Let M be a triangulated n-manifold and Λ a group. A cocycle $\alpha \in C^1(M, \Lambda)$ is a map from the oriented 1-simplices of M to Λ which satisfies the following conditions

- For each 2-simplex of M, let a, b, c be the three edges of the simplex oriented such that a and c start at the same vertex, and b and c end at the same vertex, then $\alpha(a)\alpha(b) = \alpha(c)$.
- if $\bar{a}$ is the same edge as a with the opposite orientation then $\alpha(a)^{-1} = \alpha(\bar{a})$.

Remark 2.1.28. We can extend α to a map on all simplicial paths $\gamma = e_1 \dots e_n$ in M by defining

$$\alpha(\gamma) = \alpha(e_1) \dots \alpha(e_n).$$

Because a homotopy of simplicial paths can be realised as a sequence of homotopies across 2-simplices, then by the first bullet point above, α is invariant under homotopy fixing end points. This means that if we pick some vertex x_0 as our basepoint, then α induces a homomorphism $\alpha_* : \pi_1(M, x_0) \to \Lambda$.

Convention. In the case of a semi-ideal triangulation of a manifold M, a cocycle $\alpha \in C^1(M, \Lambda)$ shall take as inputs only the non-ideal edges of the triangulation. This is justified as the semi-ideal triangulation deformation retracts to its non-ideal simplices.

Remark 2.1.29. Let M be a closed hyperbolic manifold, $p : \tilde{M} \to M$ its universal cover, $\tilde{x_0}$ some lift of a vertex x_0 in M and y_0 some chosen basepoint in $\mathbb{H}^n$. Note that if we are given a cocycle $\alpha \in C^1(M, \Lambda)$ where $\Lambda =$Isom$(\mathbb{H}^n)$, then we can uniquely define a map F from $\tilde{M}$ to $\mathbb{H}^n$ by for each vertex $\tilde{x}$ in $\tilde{M}$ simply selecting a simplicial path $\tilde{\gamma}_{\tilde{x}}$ from $\tilde{x_0}$ to $\tilde{x}$ and defining

$$F(\tilde{x}) = \alpha(p(\tilde{\gamma}_{\tilde{x}}))(y_0)$$

then defining F on each simplex to be geodesic interpolation between the values of F at its vertices.

In this way it makes sense to say that a cocycle induces both a representation on the fundamental group and a map on the triangulation. In fact these are respectively the holonomy and developing map as defined in [Thu97]. Note that by definition the map restricts to at least a geodesic immersion on each simplex and sometimes even a geodesic embedding of each simplex. We will see later how this definition can be extended to semi-ideal triangulations.

2.2 Hyperbolic Geometry

2.2.1 The Models of Hyperbolic Geometry

The following basic information about the different models of hyperbolic space is compiled from notes by Marc Lackenby [Lac00] and Benedetti and Petronio's book [BP92], though these facts might be more easily referenced by careful use of Wikipedia.

2.2.1.1 Riemannian Metrics

Definition 2.2.1 (Hyperboloid Model)**.** In $\mathbb{R}^{n+1}$ consider the bilinear form of signature $(n, 1)$ defined by $q(x_1, \ldots, x_{n+1}) = x_1^2 + \ldots + x_n^2 - x_{n+1}^2$. Let

$$I_n := \{x \in \mathbb{R}^{n+1} \mid q(x) = -1, x_{n+1} > 0\}$$

The restriction of q to the tangent space of is positive definite, so induces a metric. This gives us a Riemannian manifold I_n which happens to be isometric to $\mathbb{H}^n$.

Definition 2.2.2 (The Poincaré Disc Model)**.** Let D^n denote the ball of (Euclidean) radius 1 in $\mathbb{R}^n$. Define a Riemannian metric $\langle v, w\rangle$ on the tangent space at each point $x \in D^n$ by

$$\langle v, w\rangle = \langle v, w\rangle_{D^n} \cdot \left(\frac{2}{1-|x|^2}\right)^2$$

Definition 2.2.3 (Upper Half Plane Model)**.** Consider U^n, the upper half space of $\mathbb{R}^n$ (all points with $x_n > 0$.) Define a metric $\langle v, w\rangle$ on the tangent space at a point $x \in U^n$ by

$$\langle v, w\rangle = \frac{\langle v, w\rangle}{x_n^2}$$

There is also the Klein Model of hyperbolic space which is a certain metric on the Euclidean 1-ball. We shall not make use of its metric, nor shall we ever be taking isometries between the Klein model and other models but we will need to understand its distance function and so will define this later. We will also make use of the following fact.

Lemma 2.2.4. *All geodesics in the Klein model are Euclidean.*

This means that we can isometrically embed any given n-simplex in the Klein model and it will be Euclidean. We use this later to define a grid of boxes which cut up an n-simplex, this is usually quite a tricky thing to do as hyperbolic space doesn't decompose into a grid in the same way Euclidean space does.

2.2.1.2 Isometries between Models

Lemma 2.2.5 (Hyperboloid and Disc)**.** *An isometry from the hyperboloid to the disc is given by*

$$(x_1, \ldots, x_{n+1}) \mapsto \frac{(x_1, \ldots, x_n)}{1 + x_{n+1}}$$

Its inverse is given by

$$(y_i, t) \mapsto \frac{(2y_i, 1 + \Sigma y_i^2 + t^2)}{1 - \Sigma y_i^2 - t^2}$$

Lemma 2.2.6 (Disc and Upper Half Plane)**.** *The disc and upper half plane are related by the self-inverse map $I : \mathbb{R}^n \to \mathbb{R}^n$ define which gives the desired isometry when restricted to either D^n or U^n.*

$$x \mapsto 2\frac{x + e_n}{|x + e_n|^2} - e_n$$

Corollary 2.2.7 (Upper Half Plane and Hyperboloid)**.** *Composing the above two isometries gives an isometry from the hyperboloid to the upper half plane as follows*

$$(x_1, \ldots, x_{n+1}) \mapsto 2\frac{(x_{n+1}+1)(x_1, \ldots, x_n + x_{n+1} + 1)}{(x_1^2 + \ldots x_{n+1}^2 + 2x_n x_{n+1} + 2x_n + 2x_{n+1} + 1)}$$

.

2.2.1.3 Distance Functions for the Models

We already have a definition of the Riemannian metric, but below we give a function which determines the distance between points.

Lemma 2.2.8 (Hyperboloid Model)**.** *The distance between two points x, y denoted by $(x_1, \ldots, x_{n+1})$ and $(y_1, \ldots, y_{n+1})$ respectively is given by $arcosh(Q(x, y))$ where Q is the bilinear map induced by the bilinear form q defined above.*

$$Q(x, y) = x_1 y_1 + \ldots x_n y_n - x_{n+1} y_{n+1}$$

Note that cosh is injective on the positive reals so $d(x, y) = d(z, w)$ if and only if $Q(x, y) = Q(z, w)$.

Lemma 2.2.9 (Upper Half Space)**.** *Let $\underline{x} := (x_1, \ldots, x_n), \underline{y} := (y_1, \ldots, y_n)$ lie in the upper half space model of hyperbolic space, then*

$$d(\underline{x}, \underline{y}) = arcosh\left(1 + \frac{\sum_{i=1}^{n}(x_i - y_i)^2}{2x_n y_n}\right)$$

$$Alternatively\ d(\underline{x}, \underline{y}) = 2ln\left(\frac{\sqrt{\sum_{i=1}^{n}(x_i - y_i)^2} + \sqrt{(\sum_{i=1}^{n-1}(x_i - y_i)^2) + (x_n + y_n)^2}}{2\sqrt{x_n y_n}}\right)$$

Lemma 2.2.10 (Klein Model)**.** *If the Euclidean distance to the origin of a point x in the Klein model is r, then the hyperbolic distance is*

$$\frac{1}{2} log\left(\frac{1+r}{1-r}\right) = arctanh(r)$$

Note that this means that radial segments of equal Euclidean length have greater hyperbolic length the further from the origin they lie.

Corollary 2.2.11. *The distance from a point to the vertical axis in upper half space model is given by*

$$arcosh\left(\frac{\|\underline{x}\|}{x_n}\right)$$

Proof. Let $\underline{x} := (x_1, \dots x_n)$ be a point in the upper half space model and note that the closest point projection of $\underline{x}$ to the vertical axis is $\underline{y} := (0, \dots, 0, \|\underline{x}\|_{Eucl})$ as both points lie on a half circle with centre on the boundary at infinity which intersects the vertical axis at a right angle.

Now a formula for distance in the upper half space model is given by

$$d(\underline{x}, \underline{y}) = arcosh\left(1 + \frac{\sum_{i=1}^{n}(x_i - y_i)^2}{2x_n y_n}\right)$$

So substituting in our values gives us

$$d(\underline{x}, \underline{y}) = arcosh\left(1 + \frac{\left(\sum_{i=1}^{n-1}(x_i)^2\right) + (x_n - \|\underline{x}\|)^2}{2x_n\|\underline{x}\|}\right)$$

$$= arcosh\left(1 + \frac{2\|\underline{x}\|^2 - 2x_n\|\underline{x}\|}{2x_n\|\underline{x}\|}\right) = arcosh\left(\frac{\|\underline{x}\|}{x_n}\right)$$

□

2.2.2 Thick-Thin Decomposition

We present here some standard results on the thick-thin decomposition, following the textbooks of Purcell and Benedetti-Petronio [Pur20] [BP92].

Remark 2.2.12. We note that the model of hyperbolic space we use will vary between sections. For the purpose of understanding cocycles and defining a system of polynomials, we will use the hyperboloid model because its isometry group is easily described using polynomials, in comparison with $PSL(2, \mathbb{C})$ which is defined as a quotient. In this section, and later sections we will use the Euclidean structure of horospheres and so we use the upper half space model because any given horosphere can be taken to a horizontal Euclidean plane by an isometry of $\mathbb{H}^n$.

Definition 2.2.13 (Injectivity Radius)**.** We define the injectivity radius at a point x in a hyperbolic manifold M to be

$$\text{injrad}(x) = \sup\{r \mid B(x, r) \text{ is an embedded } r\text{-ball in } M\}.$$

The injectivity radius of a manifold is then given by

$$\text{injrad}(M) = \inf\{\text{injrad}(x) \mid x \in M\}.$$

Remark 2.2.14. Note that in the case where M is closed, there is some $x \in M$ such that $\text{injrad}(x) = \text{injrad}(M)$.

We give the following simple compactness proof for completeness but we'll be leaving this style of proof to the reader in future.

Proof of Remark 2.2.14. Let I be the injectivity radius of M. The set A_n of points with injectivity radius $> I + \frac{1}{n}$ is an open set. Thus the set of all such A_n for $n \in \mathbb{N}$ is an open cover and as M is compact, it has a finite subcover. But this would mean that actually the injectivity radius of M is at least $I + \frac{1}{n}$ for some n but gives $I > I + \frac{1}{n}$, a contradiction. □

Definition 2.2.15. A systole of a hyperbolic manifold M is a closed geodesic γ in M of shortest length.

Definition 2.2.16. We define the ϵ-thin part $M_{(0,\epsilon)}$ of a hyperbolic manifold M to be

$$M_{(0,\epsilon)} = \{x \in M \mid \text{injrad}(x) < \frac{\epsilon}{2}\}.$$

Similarly the ϵ-thick part is defined to be

$$M_{[\epsilon,\infty)} = \{x \in M \mid \text{injrad}(x) \geq \frac{\epsilon}{2}\}.$$

Theorem 2.2.17 (Margulis' Lemma, D3.3 in Benedetti and Petronio [BP92])**.** *There exists a universal constant ϵ_n for each $n \geq 3$ such that the following holds.*

Let M be a complete oriented hyperbolic n-manifold (not necessarily compact or finite volume). The thin part $M_{(0,\epsilon_n)}$ is the union of pieces homeomorphic to one of the following types:

1. $\mathring{D}^{n-1} \times S^1$.
2. $V \times (0, \infty)$ *where V is a differentiable oriented $(n-1)$-manifold without boundary supporting a Euclidean structure.*

Furthermore,

- *The pieces are a positive distance from one another.*
- *The second case occurs only when the manifold is non-compact.*
- *When M is finite volume non-compact, the second case occurs and the manifolds V are closed.*

From now on, if we say thin part, rather than ϵ-thin part, then we mean the ϵ_n-thin part.

Remark 2.2.18 ([BP92])**.** For each component of the thin part there is a subgroup H of the fundamental group G of M and a representation $\rho : G \to Isom(\mathbb{H}^n)$ such that either:

- In case 1 of Theorem 2.2.17, $\rho(H)$ is generated by a single loxodromic element whose axis projects down to the core curve $\{0\} \times S^1$.
- In case 2 of Theorem 2.2.17 $\rho(H)$ consists entirely of parabolic isometries fixing ∞ and such that V is homeomorphic to $(\{0\} \times R^{n-1})/\rho(H)$.

Furthermore, the injectivity radius of a point p in a given component of the thin part is the same as the injectivity radius of any point in $\mathbb{H}^n/\rho(H)$ which is both a lift of p and lies in the thin part of $\mathbb{H}^n/\rho(H)$.

Definition 2.2.19. We shall refer to the two types of component of the thin part as tubes (also called Margulis tubes) and cusp neighbourhoods.

As it stands, Margulis' Lemma presents an algorithmic problem. If we ever want to use ϵ_n in an algorithm then we will have a constant which varies with n in a way we don't fully understand. Fortunately, this is an area of much research and lower bounds are known. Robert Meyerhoff gives a bound on ϵ_3 [Mey87]. For general n, a result of Ruth Kellerhals [Kel04] gives a lower bound for ϵ_n.

Theorem 2.2.20 (Meyerhoff [Mey87])**.** *With the definition of the thick-thin decomposition given in Theorem 2.2.17, the value of ϵ_3 is at least* 0.052 *.*

Theorem 2.2.21 (Kellerhalls [Kel04])**.** *The value of ϵ_n is bounded below by $\frac{1}{(6\pi)^n}$.*

In fact Kellerhals' paper gives a tighter and more technical lower bound than this. We shall use this coarser bound for simplicity of expression as doing so has no impact on our final bounds.

2.2.3 Mostow and Other Rigidities

Underlying everything we shall be working on in this thesis are the many theorems which establish the rigidity of hyperbolic manifolds. First among these of course is Mostow-Prasad Rigidity.

Theorem 2.2.22 (Mostow-Prasad Rigidity [Pra73])**.** *Let Γ, Γ' be two lattices in $Isom(\mathbb{H}^n)$ (where $n \geq 3$) and $\phi : \Gamma \to \Gamma'$ a group isomorphism. Then there exists an isometry g in $Isom(\mathbb{H}^n)$ such that*

$$g\gamma g^{-1} = \phi(\gamma) \quad \forall \gamma \in \Gamma$$

We state this here in the language of conjugate lattices for later use, but the more frequent formulation essentially states that hyperbolic manifolds are uniquely determined by their fundamental groups. This means that to check homeomorphism between two hyperbolic manifolds is equivalent to checking isomorphism between their fundamental groups. Note that the requirement that $n \geq 3$ is certainly necessary as there are many non-isometric hyperbolic structures on the genus 2 surface. This is explained in any introduction to Teichmuller space such as [Mar16] [FM11].

Remark 2.2.23. An interesting consequence of this is that we may have a sequence of hyperbolic manifolds where there is at first glance no reason that they should be homeomorphic. However, if the process used to get from one to other preserves fundamental group, then they must be homeomorphic. This will be relevant later when we build new manifolds from old in various ways, but can prove that their fundamental groups are the same such as in Theorem 5.2.11

Remark 2.2.24. A further corollary of this rigidity is not only that group isomorphisms imply homeomorphisms, but also that homeomorphisms imply isometries. The first implication is false in the case of spherical manifolds, where there are lens spaces of the same fundamental group which are not homeomorphic.

This second implication is not true for example in Euclidean manifolds for many reasons. Firstly one can scale a Euclidean manifold without changing its curvature due to the sectional curvature of a Euclidean manifold being 0, whereas if one simply applies a scalar to a hyperbolic metric it ceases to have curvature -1 as the curvature scales as well and thus ceases to be hyperbolic.

Even if we account for scaling by normalising the volume of Euclidean manifolds to be 1 problems persist. For example, one can create infinitely many distinct unit volume Euclidean structures on the 3-torus by identifying the sides of various different shapes of unit volume cuboids.

Where this is useful is that if we have a triangulation and find multiple ways of giving it a hyperbolic structure, then we know that the resulting hyperbolic manifolds are isometric.

Mostow-Prasad rigidity is particularly useful when we are finding hyperbolic structures for triangulations. This is not the only way we shall be finding hyperbolic structures. Another way we will be finding them is by looking for lattice representations.

Definition 2.2.25 (Lattice Representation)**.** We call a representation

$$\rho : G \to Isom(\mathbb{H}^n)$$

a lattice representation if its image is a lattice (a discrete co-finite volume subgroup).

We shall see how to find a collection of representations which includes a representation from each component of the variety of representations. In this context the following rigidity theorems are immensely useful as they guarantee that one of our representations is indeed a faithful lattice representation.

Definition 2.2.26. Let G be a group, $\Lambda = Isom(\mathbb{H}^n)$ and let $H := \{H_1, \dots H_q\}$ be a collection of subgroups of G. Define

$$\mathrm{Hom}_{\mathrm{par}}(G, H; \Lambda) := \{\rho \in \mathrm{Hom}(G, \Lambda) \mid \rho(\gamma) \text{ is parabolic } \forall \gamma \in H_j, \forall H_j \in H\}$$

If G is a subgroup of Λ, then we define

$$\mathrm{Hom}_{\mathrm{par}}(G; \Lambda) := \mathrm{Hom}_{\mathrm{par}}(G, H_{\max}; \Lambda)$$

where $H_{\max}$ is a collection of representatives of the conjugacy classes of maximal parabolic subgroups of G. Note that if G is a uniform lattice, then $Hom_{par}(G; \Lambda) = Hom(G; \Lambda)$.

Theorem 2.2.27 (Calabi-Weil Rigidity, Thms. 4.19 and 8.55 in [Kap00])**.** *Let G be a lattice in $Isom(\mathbb{H}^n)$ where $n \geq 3$. The identity representation $\rho_0 : G \to Isom(\mathbb{H}^n) =: \Lambda$ lies in a connected component of $Hom_{par}(G; \Lambda)$ made up entirely of conjugates of ρ_0.*

Note that there is a somewhat hidden requirement in this theorem, that is that in the finite volume case where parabolics can exist, it is required that we consider only representations which send parabolics to parabolics. The following remark details why this presents difficulties.

Remark 2.2.28. Given a representation of the fundamental group of a finite volume hyperbolic manifold into its group of isometries and trying to understand whether a given subgroup maps entirely onto parabolic elements varies in difficulty based on dimension.

For dimension 3, note that the only orientation preserving isometries of $\mathbb{R}^2$ with no fixed points are translations. Thus any parabolic isometry of $\mathbb{H}^3$ which appears in an orientation preserving lattice restricts to a translation on each horosphere based at its fixed point. Thus in this dimension, the product of any two parabolic isometries with the same fixed point at infinity restricts to the product of their translations on the same horospheres and is thus either trivial or again a parabolic isometry. This means

that if we know the generators of a subgroup map to parabolics with the same fixed point, then so does the whole subgroup, making it quite easy to check parabolicity in dimension 3.

In higher dimensions, unfortunately, as $\mathbb{R}^n$ admits more complicated isometries, a phenomenon can occur in which we have two isometries of $\mathbb{R}^n$ with no fixed points, but whose product is both non-trivial and has a fixed point. For example in $\mathbb{R}^3$ one can imagine one isometry which translates in one axis and rotates in the other 2, and another which simply inverts the translation. Together, these isometries combine to give a rotation and thus their corresponding parabolics would combine to give a elliptic isometry. Hence, in higher dimensions, knowing that the generators of a subgroup are parabolic is insufficient to know that the whole subgroup is parabolic and so algorithmically deciding whether a representation preserves parabolics is difficult. Note that Euclidean subgroups do have finite index translation subgroups, even effectively so (see [Oli80]). Nevertheless, finding such translation subgroups still presents quite a problem.

As dispiriting as the above remark may seem, there is a result due to Garland and Raghunathan which removes the requirement of preserving parabolicity exactly in the dimensions in which it poses problems.

Theorem 2.2.29 (Theorem 7.2 in [GR70])**.** *Let G be a lattice in $Isom(\mathbb{H}^n)$ for $n \geq 4$. The identity representation $\rho_0 : G \to Isom(\mathbb{H}^n) =: \Lambda$ lies in a connected component of $Hom(G; \Lambda)$ made up entirely of conjugates of ρ_0.*

2.2.4 Geometric Inequalities and Dimensional Differences

We want to bound runtime as a function of combinatorial data (number of n-simplices for example). We do this by relating both concepts to geometrical ones. However, we have no guarantee that the geometric property relating to runtime is the same as that coming from the combinatorics. Thus we will often need to relate different aspects of the geometry of a manifold. Even if one doesn't care about the combinatorics, this is already a major tool for understanding Riemannian manifolds and thus an active area of research. In this section we talk about some historic results in this area, and also outline how and why these results can vary with dimension.

Before we get into the meat of the chapter, we first quote an interesting result which will pop up a couple of times in this thesis, concerning the volume of the n-sphere.

Theorem 2.2.30 ([Hub82])**.** *The volume of the n-sphere achieves a maximal value at $n = 6$.*

This remarkable little theorem means that whenever we're talking about volumes of spheres, if we only care about bounding from above, then we can replace a function which varies with n by the constant $vol(S^6)$. With this little aside complete, we return to the study of hyperbolic manifolds.

Two key geometric properties, especially for understanding injectivity radius, are volume and diameter. In a 1978 paper, Gromov proves a fundamental result relating the two.

Theorem 2.2.31 (Gromov [Gro78])**.** *Let M be a closed n-manifold of negative curvature in the interval $[-1, 0)$, for $n \geq 8$. Then there is some constant C_n dependent only on dimension such that the following inequality relating volume and diameter holds.*

$$vol(M) \geq C_n(1 + diam(M))$$

For $n \geq 4$ he proves a weaker result where diameter is replaces by its cube root.

Here we see the first point at which hyperbolic 3-manifolds distinguish themselves from their higher dimensional kin. Not only did Gromov not prove such a theorem for $n = 3$, it is in fact not possible. To see why we first recall Thurston's hyperbolic Dehn surgery theorem.

Theorem 2.2.32 (Hyperbolic Dehn Surgery [Thu78])**.** *Let K be a hyperbolic knot, with hyperbolic complement in the 3-sphere M_K. For all but finitely many boundary slopes, Dehn filling along a given slope yields a closed hyperbolic 3-manifold. Furthermore, the volume of these manifolds is bounded above by the volume of M_K and there are slopes which give manifolds of arbitrarily small injectivity radius.*

From this we can clearly see that we have many infinite families of hyperbolic 3-manifolds with arbitrarily small injectivity radius but with volume uniformly bounded above. Yet again, we encounter a situation where we need to translate between different geometric properties, here between injectivity radius and diameter. There's certainly reason to hope that injectivity radius and diameter should be related in 3 dimensions. Firstly, they are related in higher dimensions, as the following theorem shall imply.

Theorem 2.2.33 ([Rez95])**.** *Let M be a compact Riemannian manifold of constant curvature -1. If $n = dim(M) \geq 4$, then*

$$injrad(M) \geq C_n(Vol(M))^{-(1+4/(n-3))}$$

where C_n is a constant dependent on n.

Combining Reznikov's work with the more general formulation of Gromov's theorem mentioned above gives the desired relation between injectivity radius and diameter for hyperbolic manifolds of dimension 4 and above. Fortunately, the injectivity radius of a closed hyperbolic 3-manifold is also correlated with its diameter. In fact, the diameter of a tube in the thin part of a hyperbolic manifold is inversely related to the length of its core curve, from this the relation follows. This relationship is surely known to experts, for example it is used in recent work of Futer, Purcell and Schleimer ([FPS19]), and follows from Margulis' Lemma. This statement about tubes actually holds in all dimensions greater than 2 and in this thesis we shall give a more explicit version of the following lemma, which relates the length of a systole (shortest closed geodesic) in a hyperbolic n-manifold to its diameter.

Lemma 2.2.34 (Theorem 4.2.5)**.** *The smaller the length of the core curve of a Margulis tube in a finite volume hyperbolic n-manifold ($n \geq 3$), the greater the diameter of the tube.*

The above lemma or variations on it are known as the collar lemma, see for example **??**, **??**. You can also prove this in dimension 2 but as our thesis focuses exclusively on dimensions 3 and above, so shall our collar lemmas. From this we can see that both Gromov's and Reznikov's inequalities can't hold in 3 dimensions as the following corollary demonstrates.

Corollary 2.2.35. *Let V be some constant such that there is a hyperbolic knot whose complement has volume V, then there is no bound on the diameters of closed hyperbolic 3-manifolds of volume V.*

Proof. Hyperbolic Dehn surgery gives us an infinite family of closed 3-manifolds with volume bounded above by V and with arbitrarily small injectivity radius. Now in a closed hyperbolic 3-manifold, the injectivity radius is realised by a ball based on a systole, that is the length of a systole is always twice the injectivity radius. Thus the length of the systoles in our infinite family are also arbitrarily small. Thus by the above lemma, the diameters of their Margulis tubes are arbitrarily large and hence the diameters of the manifolds containing them are arbitrarily large also. □

While we're noting the differences between the dimensions, let us ask the same question about how volume and diameter are related in 2 dimensions as another example of how dimension 2 acts very differently than the others. In this case there is once again no relation between diameter and volume, but it is much easier to show. The hyperbolic structure on a surface can be built up by gluing together hyperbolic pairs of pants along their boundaries (again see [Mar16] or [FM11]. However, there are hyperbolic structures on the pair of pants with arbitrary diameter and thus the same is true for any hyperbolic surface.

The area(or volume) of a hyperbolic surface on the other hand is uniquely determined by its Euler characteristic (this is the Gauss-Bonnet Theorem) and thus again is unrelated to diameter though for a completely different reason than in dimension 3.

This difference in dimension poses an important question. If we see that in certain dimensions there are no ways to relate different geometric properties, one might worry whether or not we can relate combinatorial data to geometric data. Fortunately we can relate the two and in one direction, there is an almost tautological argument for why the relationship exists.

Lemma 2.2.36. *For any constant T, any dimension n and any finite geometric property $G(M)$ of a hyperbolic n-manifold M triangulated by less than T n-simplices (such as the volume of a finite volume hyperbolic manifold or the injectivity radius of a closed hyperbolic manifold), there are functions f, F such that $f(T) < G(M) < F(T)$.*

Proof. As there are only finitely many simplicial complexes with less than T n-simplices, there are only finitely many such manifolds and so f can be taken to be simply

$$min\{G(M) \mid M \text{ is triangulated by } < T \ n\text{-simplices}\}$$

and F is defined similarly as the maximum over the same set.

□

This tells us that it is not whether, but rather how we can bound geometry as a function of combinatorics that is interesting. Just because the functions f, F described above exist, doesn't mean we know what they are or that they're even computable. What we shall be finding throughout this thesis are actual computable bounds on these kinds of relationships that help us understand how T, F, and f all interact.

While limiting our view to a finite collection of manifolds clearly bounds the geometry of those manifolds, this statement can more surprisingly be reversed. The

following states that if we consider only manifolds with certain geometry bounds, then there are only finitely many manifolds which meet these restrictions, this is known as Cheeger finiteness (see [Che70] or a more modern treatment in [Cha95]).

Theorem 2.2.37 (Cheeger Finiteness)**.** *Given real numbers $n, d, V, \Lambda > 0$, let $\mathcal{K}$ denote sectional curvature, then there exist only finite many diffeomorphism classes of compact n-dimensional Riemannian manifolds satisfying*

$$diam(M) \leq d, \qquad Vol(M) \geq V, \qquad |\mathcal{K}| \leq \Lambda.$$

The above examples demonstrate at least that this line of reasoning, the comparison of many different geometric and combinatorial properties is one which has seen much success in the past. Throughout this thesis we will build on these inequalities and find our own to allow us to understand the geometry and combinatorics of hyperbolic manifolds.

2.3 3-Manifolds

2.3.1 Geometric Decomposition

We assume a good amount of knowledge of the basics of 3-manifolds in this chapter. For a recap of the definitions of terms, one can read any introduction to 3-manifolds such as [Hem76].

Thurston's Geometrisation is the third in a trio of developments in the theory of 3-manifolds from the 20th century which together describe a very useful decomposition of closed orientable 3-manifolds.

The first of the three is the prime decomposition theorem, the existence and uniqueness of which were proven respectively by Kneser and Milnor.

Theorem 2.3.1 (Prime Decomposition, Kneser-Milnor [Kne29, Mil62])**.** *Let M be a closed, oriented 3-manifold which is not S^3, then M admits a unique finite connected sum decomposition up to reordering. The summands are closed oriented prime 3-manifolds and are unique up to oriented homeomorphism. Note that for uniqueness, summands are not allowed to be homeomorphic to S^3.*

The second result in our trio is due to Jaco, Shalen and Johannson, which shows how to take the prime pieces we have produced above and decompose them along tori.

Theorem 2.3.2 (JSJ Decomposition, Jaco-Shalen, Johannson [JS79, Joh79])**.** *Given a closed, prime, orientable 3-manifold M, there exists a (possibly empty) minimal collection of incompressible tori embedded in M such that when we cut along these tori, the resulting components are either Seifert fibred or atoroidal. This collection is unique up to isotopy.*

The first two results give us a way of cutting up 3-manifolds into simpler pieces. Our next step is trying to impose some structure on those simple pieces. The JSJ decomposition gives us a useful Seifert fibred structure on some of the pieces, but we need Thurston's Geometrisation to understand the remaining atoroidal pieces.

Theorem 2.3.3 (Thurston's Geometrisation Conjecture or Perelman's Theorem, [Thu82, Per02, Per03a, Per03b])**.** *Let M be a compact orientable prime 3-manifold with boundary, such that the boundary consists entirely of tori or is empty. If M is atoroidal then M is either Seifert fibred or its interior admits a finite volume hyperbolic structure.*

The consequence is thus that we can decompose a given compact orientable 3-manifold into pieces which are either Seifert fibred or hyperbolic. Thus one can gain an understanding of 3-manifolds by first understanding hyperbolic or Seifert fibred 3-manifolds, and then considering how they glue together. In this thesis we will mostly focus on hyperbolic 3-manifolds, using tools of hyperbolic geometry which continue to be useful in higher dimensions. First, let's look at how dimension 3 is special from an algorithmic perspective, and what results already exist for 3-manifolds.

2.3.2 Algorithms for 3-manifolds

First we give a definition of our core concern, the homeomorphism problem.

Definition 2.3.4 (Homeomorphism Problem)**.** We define the homeomorphism problem for a class of manifolds $\mathcal{M}$ to be the following question:

Is there an algorithm which takes as input two triangulations of manifolds in $\mathcal{M}$ and outputs whether or not they are homeomorphic?

We define the oriented homeomorphism problem on oriented manifolds similarly except we ask whether two oriented manifolds admit an orientation preserving homeomorphism.

Definition 2.3.5 (Recognition Problem)**.** We can also define a similar but slightly easier problem by fixing one of the manifolds in the homeomorphism problem. For a fixed manifold M we call this the M-recognition problem. Similar to above, we can consider the M-recognition problem for manifolds in $\mathcal{M}$.

Remark 2.3.6. Note that the algorithms we are talking about here when we specify a class of manifolds are something called promise problems or promise algorithms. These take as input a triangulation of the manifold, and they assume that the manifold belongs to the class $\mathcal{M}$. The fact that the manifolds are in $\mathcal{M}$ is not something we have to prove to the algorithm as part of the input.

It is well known that for n less than 3, the homeomorphism problem for closed n-manifolds is solved. For 2-manifolds, one need only check whether a triangulation is orientable or not, and then its Euler characteristic determines it.

In higher dimensions however, the problem quickly becomes intractable. Here are a few results that demonstrate this.

First, a result due to Markov. This is shown by a demonstration that the triviality problem for finitely presented groups (already known to be undecidable) can be reduced to the M^n-recognition problem for a particular simply n-manifold M^n.

Theorem 2.3.7 ([Mar58])**.** *For every natural number $n > 3$ one can create a closed n-manifold M^n such that the M^n-recognition problem is undecidable. Hence the homeomorphism problem for closed n-manifolds is undecidable.*

Corollary 2.3.8. *The homeomorphism problem for (not necessarily closed) n-manifolds is undecidable.*

In fact, for dimensions > 4, the manifold M^n can be taken to be the n-sphere.

Theorem 2.3.9 (A Theorem of S. P. Novikov in [Vol74])**.** *The n-sphere recognition problem for $n > 4$ is undecidable.*

Remark 2.3.10. We have seen that 2-sphere recognition is very simple and that 5-sphere recognition is impossible. Thus the interesting dimensions would appear to be 3 and 4. In fact this holds true as 4-sphere recognition remains an open problem while 3-sphere recognition is decidable, as we shall soon see, but requires a great deal of effort.

Theorem 2.3.11. *There is no algorithm which decides whether a given simplicial n-complex is a combinatorial manifold for $n > 5$.*

Proof. This follows simply from the fact that to tell whether a simplicial complex is a combinatorial manifold requires that we be able to tell whether the links of all its vertices are spheres. However, given a simplicial n-complex one can easily construct a simplicial $(n+1)$-complex which has a vertex whose link is the chosen n-complex, for example, this is true in the suspension of the n-complex. □

This isn't the end of the woes for higher dimensional combinatorial manifolds. One might guess that the problem described in the above proof comes from the seemingly pathological step of taking a suspension in order to complicate a link, and that by restricting to the world of manifolds we might avoid such pathologies. Unfortunately this is not the case. The following is a well known example of a triangulation of the 5-sphere which we shall see is non-combinatorial.

Theorem 2.3.12 ([Edw06])**.** *There exists a homology 3-sphere M with $M \ncong S^3$ such that the double suspension $\Sigma^2 M$ is homeomorphic to S^5.*

To see that this gives a non-combinatorial triangulation of S^5 consider the simplicial complex K given by taking the double suspension of any triangulation of the homology sphere M. The suspension points of K each have as their link a triangulation of the single suspension ΣM. Thus all we need to show is that ΣM is not a 4-sphere. This follows from the fact that if ΣM were a 4-sphere, then removing two points would leave a simply connected manifold, but removing the suspension points from our suspensions leaves us with a space which is homotopy equivalent to M, and by the Poincare conjecture, M cannot be simply connected. This is thus an example of a triangulated manifold where the triangulation is simplicial but not combinatorial.

The above serves to warn us that, especially in higher dimensions, there are many traps we could fall into. Our algorithms need to specifically use the properties of our manifolds, whether that be using their dimension or their geometry, a more general algorithm which ignores these properties is guaranteed to fail.

Fortunately, as we saw in Theorem 2.1.13, the 3-dimensional case is simpler from a simplicial viewpoint, and in fact many algorithms already exist to solve the homeomorphism problem for 3-manifolds.

2.3.3 The Work of Kuperberg

The work of this thesis is so influenced by the work of Kuperberg in his paper "Algorithmic Homeomorphism of 3-Manifolds as a Corollary of Geometrization" [Kup19] that it is important to give a short overview of the results of that paper. This has

many aims beyond the very necessary recognition of such a direct influence. Firstly, we hope that a reiteration of Kuperberg's work can help to illuminate the specific ways in which this thesis contributes to the field, by demonstrating where it is we improve upon or stray from the ideas Kuperberg lays out. Furthermore, Kuperberg's work lays out a path forward for the ideas discussed in this thesis, he opens the way to a project of finding an algorithm for the homeomorphism problem for closed 3-manifolds which has explicitly bounded runtime. This thesis hopes to make a major contribution to this project.

Kuperberg's paper achieves two goals, firstly he establishes that Perelman's proof of the Poincaré conjecture and of the wider Geometrisation conjecture was the final step necessary to prove the decidability of the homeomorphism problem. He does this by providing an algorithm using only mathematics that would have been available in the 70s when Thurston made his conjecture. Secondly he provides a modern algorithm, using technology developed in the intervening 40-50 years, which can solve the homeomorphism problem in bounded time.

In particular he proves the algorithm lies in the complexity class ER, or Elementary Recursive.

Definition 2.3.13. An algorithm is elementary recursive, or ER, if it runs in time

$$O\left(2^{2^{\cdot^{\cdot^{\cdot^{n}}}}}\right)$$

where n is the size of the input and the ellipsis denotes a tower of 2s of bounded height.

Remark 2.3.14. One could instead consider fixing the height of the tower in advance, for example considering algorithms which have runtime and solution size bounded by

$$2^{O(n)}.$$

Note, however, that if you have two such algorithms but we compose them (the output of one is the input of the other) then the complexity of this combined algorithm could end up having runtime as great as

$$2^{2^{O(n)}}.$$

For example, if an algorithm has runtime bounded by $2^{O(n)}$ and its output is a triangulation, this output could have size $2^{O(n)}$ also, and then plugging this into a new algorithm gives doubly exponential runtime.

Continuing these kinds of compositions causes the height of the tower in the runtime bound to multiply exponentially and thus any fixed bound can be easily outrun. However, this composition of algorithms doesn't pose any problems if we only care that there is some finite tower of exponentials bounding our runtime.

Kuperberg makes use of the feature of ER described in the above remark in his paper, for example one might have a situation in which one uses a solution to a system of polynomials to build another system of polynomials. If all you need to do is show that the overall algorithm is in ER, then you just need to know that both systems of polynomials can be constructed. If you need to know the height of the tower, then you need to understand the complexity of the first system of polynomials as well as exactly how the complexity of its solution affects the construction of the second system. This allows the following wide reaching results without requiring the technical detailed description of polynomial systems present in some sections of this thesis.

Theorem 2.3.15 (Theorem 8.1 in [Kup19])**.** *It is in ER to find and triangulate the prime summands of a closed, oriented 3-manifold M, to find and triangulate the JSJ components within each prime summand, to find their JSJ graph, and to recognize which JSJ components are Seifert-fibered and find their fibrations.*

Theorem 2.3.16 (Theorem 8.3 in [Kup19])**.** *If a compact, oriented 3-manifold N has a closed or cusped hyperbolic structure, then it is in ER to find a geometric triangulation and specify its geometric data with algebraic numbers. The isomorphism and automorphism problems are also both in ER.*

Core to Kuperberg's algorithm is that it finds a prime and JSJ decomposition of the given manifold as well as identifying the structures on the pieces, before then applying algorithms which are tailored to the Seifert fibred or hyperbolic cases componentwise. This is in contrast to the algorithms of Scott and Short [SS14] and Bessiéres et al [BBB+10] which avoid the JSJ decomposition entirely by appealing to results on Haken manifolds whenever the JSJ decomposition is non trivial. In this way Kuperberg's algorithm performs something quite satisfying in that it actually uses algorithmic versions of the decomposition theorems that we use theoretically all the time.

2.4 Systems of Polynomials

2.4.1 An Algorithm for Finding Solutions

What we present here is adapted from [Gri86] (which is a survey of results including those of [GV85], see also [GV88]) with a few corollaries which show how to adapt the theorem to our situation. This section shows how given a system of polynomial inequalities, one can find solutions which are bounded in terms of the size of the system.
Let a system of polynomial inequalities

$$f_1 > 0, \ldots, f_m > 0, f_{m+1} \geq 0, \ldots, f_\kappa \geq 0$$

be given, where the polynomials $f_i \in \mathbb{Q}[X_1, \ldots, X_N]$ satisfy the bounds

$$deg_{X_1,\ldots,X_N}(f_i) < d, \qquad l(f_i) < M, \qquad 1 \leq i \leq \kappa.$$

Here $deg_{X_1,\ldots,X_N}(f_i)$ is the maximum degree of a monomial in f_i, where the degree of a monomial is the sum of the exponents of the monomial. The length, or complexity l is defined on rational numbers by $l(\frac{p}{q}) = log_2(|pq|+2)$ where p, q are coprime integers and on polynomials it is defined as the maximum complexity among its coefficients.

Let $\alpha = (\alpha_1, \ldots, \alpha_N)$ be a solution to the system of inequalities where each α_i is an algebraic number. Then we can represent each α_i in terms of a primitive element, θ, of the field $\mathbb{Q}(\alpha_1, \ldots, \alpha_N) = \mathbb{Q}(\theta)$. We represent θ by providing an irreducible polynomial $\Phi(X) \in \mathbb{Q}(X)$ of which θ is a root and an interval $(\beta_1, \beta_2) \subseteq \mathbb{R}$ with endpoints in $\mathbb{Q}$ which determines θ among the roots of Φ. With θ defined, one has $\alpha_i = \sum_j \alpha_i^{(j)} \theta^j$ for $\alpha_i^{(j)} \in \mathbb{Q}$. Using this notation we can now state the theorem of Grigoriev.

Theorem 2.4.1 (Grigoriev and Vorobjov [Gri86])**.** *For a given system of inequalities such as those above, in each connected component of the set of real solutions to this system, there exists an algebraic solution $\alpha = (\alpha_1, \ldots, \alpha_N)$ represented as above for which the following are true:*

$$deg(\Phi) \leq (\kappa d)^{O(N)}; \qquad l(\Phi), l(\alpha_j^{(i)}), l(\beta_1), l(\beta_2) \leq M(\kappa d)^{O(N)}.$$

Furthermore, an algorithm exists which finds such a solution for each connected component in time polynomial in $M(\kappa d)^{N^2}$.

The following corollary follows the same method of proof as that of Lemma 8.9 in [Kup19], but for self-containedness, we provide it here.

Corollary 2.4.2. $\frac{1}{2^{M(\kappa d)^{O(N)}}} \leq |\theta| \leq 2^{M(\kappa d)^{O(N)}}$.

Proof. Note that we can scale Φ so that it is an integer polynomial without negating the statement that $l(\Phi) \leq M(\kappa d)^{O(N)}$. So we may assume that the coefficients, γ_i of Φ are integers, and hence the γ_i have size bounded by $2^{l(\Phi)} - 2$ Now note that as $\Phi(\theta) = 0$ we get that

$$-\gamma_{deg(\Phi)}\theta^{deg(\Phi)} = \gamma_0 + \ldots + \gamma_{deg(\Phi)-1}\theta^{deg(\Phi)-1}$$

and so either $|\theta| \leq 1$ and the upper bound holds automatically or

$$|\theta| = \left|\frac{\gamma_0 + \ldots + \gamma_{deg(\Phi)-1}\theta^{deg(\Phi)-1}}{\theta^{deg(\Phi)-1}\gamma_{deg(\Phi)}}\right|$$

$$\leq \sum_{i=0}^{deg(\Phi)-1} \frac{|\gamma_i|}{|\theta^{deg(\Phi)-1-i}||\gamma_{deg(\Phi)}|} \leq \sum_{i=0}^{deg(\Phi)-1} |\gamma_i| \ .$$

The final inequality holds as $|\theta|$ and $|\gamma_{deg(\Phi)}|$ are greater than or equal to 1. Note that θ^{-1} is a root of

$$\Psi(x) = x^{deg(\Phi)}\gamma_0 + \ldots + \gamma_{deg(\Phi)} = x^{deg(\Phi)}\Phi(x^{-1}).$$

Thus once again either $|\theta^{-1}| \leq 1$ and the lower bound holds automatically or

$$|\theta^{-1}| \leq \sum_{i=0}^{deg(\Phi)-1} |\gamma_i|.$$

Applying the bounds on the γ_i gives

$$|\theta| \leq \sum_{i=0}^{deg(\Phi)-1} |\gamma_i| \leq \sum_{i=0}^{deg(\Phi)-1} 2^{l(\Phi)} \leq deg(\Phi)2^{l(\Phi)}$$

and similarly,

$$|\theta| \geq \frac{1}{deg(\Phi)2^{l(\Phi)}}.$$

Finally, Theorem 2.4.1 tells us that

$$deg(\Phi)2^{l(\Phi)} \leq (\kappa d)^{O(N)}2^{M(\kappa d)^{O(N)}} \leq 2^{M(\kappa d)^{O(N)}}$$

and the statement follows. □

From this, we can derive a statement about the size of the solutions themselves.

Corollary 2.4.3. *Let α be as in Theorem 2.4.1, then for each i, if $\alpha_i \neq 0$*

$$\frac{1}{2^{M((\kappa+2N)d)^{O(N)}}} \leq |\alpha_i| \leq 2^{M(\kappa d)^{O(N)}}.$$

Proof. First see that we can achieve an upper bound on $|\alpha_i|$ by applying the triangle inequality in the following manner:

$$|\alpha_i| \leq |\sum_{j=0}^{deg(\Phi)} \alpha_i^{(j)}\theta^j| \leq \sum_{j=0}^{deg(\Phi)} |\alpha_i^{(j)}||\theta^j| \leq \sum_{j=0}^{deg(\Phi)} 2^{M(\kappa d)^{O(N)}} 2^{M(\kappa d)^{O(N)}}$$

$$\leq ((\kappa d)^{O(N)} + 1)2^{M(\kappa d)^{O(N)}} 2^{M(\kappa d)^{O(N)}} \leq 2^{M(\kappa d)^{O(N)}}.$$

Note that a lower bound on $|\alpha_i|$ is equivalent to an upper bound on $|\alpha_i^{-1}|$, so we can modify our system of polynomial inequalities by adding some β_i such that $\alpha_i\beta_i = 1$ for all i. This is represented by two inequalities and one new variable for each of the original variables. Now applying the theorem to this new system of inequalities gives an upper bound for $|\beta_i|$ as above except κ is replaced by $\kappa + 2N$, as $O(2N)$ and $O(N)$ are the same. Taking reciprocals of both sides gives us the lower bound in the statement. □

We should note that the theorem requires that the system be given as polynomial inequalities. When we create our system, we shall allow for equalities as well, which can be replaced by a pair of inequalities. Because of big O notation, doubling the number of polynomials will not affect the final bound and so we shall ignore this distinction from here on.

An obvious question in response to such a theorem is whether it is sharp. In the context of hyperbolic manifolds, we do not make any claims that taking the information of a hyperbolic manifold and transforming it into a system of polynomials should give sharp results, as in some sense the transformation "forgets" the hyperbolicity of the input. However, Corollary 2.4.3 is sharp as theorem about polynomial systems, the following remark demonstrates that 2^{2^N} can be very easily attained as the size of a solution.

Remark 2.4.4. Consider the following system of polynomials:

$$X_1 = 4, X_2 = X_1^2, X_3 = X_2^2, \ldots X_N = X_{N-1}^2$$

Then there is only one solution given by $X_i = 2^{2^i}$. Thus we see that it is not possible to improve the bound in Corollary 2.4.3 to anything better than doubly exponential in N. As one might guess, this dependence on N is exactly where the worst algorithmic bounds come from and so for our purposes the bound of this corollary is the best we can do.

2.4.2 Poincare's Polyhedron Theorem

In this section we shall first state Poincaré's Polyhedron Theorem and then show how it gives us a means of checking whether an assignment of hyperbolic structure to the n-simplices of a gluing induces a hyperbolic structure on the gluing itself.

Though we shall not state the theorem in full generality, it is possible to state this for different geometries, types of building block and with different resulting objects. For the full statement, as well as a brief history of the Poincaré Polyhedron theorem, the reader is directed to [EP94], from which this entire section is adapted.

Definition 2.4.5 (Model n-Simplices)**.** Let M be a triangulated n-manifold consisting of t n-simplices. Let $\mathcal{M}$ be a set of t geodesic n-simplices in the upper half space model of $\mathbb{H}^n$ such that each is identified with a distinct simplex in M by a simplicial map. We call such a collection a set of model n-simplices for the triangulation of M.

Definition 2.4.6 (Dual Graph of a Triangulation)**.** We define the dual graph $\Lambda(M)$ of a triangulation to be the graph whose vertices are the n-simplices of the triangulation, and whose edges are codimension one faces along which two simplices meet.

Definition 2.4.7. Consider some spanning tree Λ' of $\Lambda(M)$ and suppose $\mathcal{M}$ is such that the face pairings in M can be realised by isometries. Then define a polyhedron Y by gluing the n-simplices in $\mathcal{M}$ only along faces corresponding to the edges of Λ'. If we also pick some base n-simplex Δ_0 in $\Lambda(M)$ then we can define a map $D : Y \to \mathbb{H}^n$ by first defining D to be the identity on the tetrahedron Δ_0 and then using Λ' to define inductively on the other tetrahedra. That is to say that an edge of Λ' connecting Δ_0 to some Δ_1 defines a unique embedding of Δ_1 into $\mathbb{H}^n$ by choosing the embedding which agrees with that of Δ_0 on the shared face and which ensure the interiors of their embeddings don't overlap, repeating this over all edges defines the map D. We will see that $D(Y)$ becomes a fundamental polyhedron for a lattice in $\mathrm{Isom}(\mathbb{H}^n)$ after applying the Poincaré Polyhedron Theorem.

Theorem 2.4.8 (Poincaré Polyhedron Theorem)**.** *Let M, $\mathcal{M}, \Lambda(M)$ be as in Definitions 2.4.5 and 2.4.6 and suppose the following hold:*

- *(Edge Equations) For any two n-simplices in M which intersect in an edge, the corresponding edges in $\mathcal{M}$ are of equal length.*
- *(Angle Equations) For any co-dimension 2 face of the triangulation, let $\alpha_1, \ldots \alpha_m$ be the dihedral angles at the corresponding face in each of the n-simplices of $\mathcal{M}$. Then $\sum \alpha_i = 2\pi$.*

- *(Orientation) There is an orientation of $\mathbb{H}^n$ with respect to which all the identifications in $\mathcal{M}$ are orientation preserving.*

Then M admits a hyperbolic structure such that when we restrict to any specific n-simplex the structure is isometric to that given in $\mathcal{M}$. Furthermore, M can be described as the quotient of the action of some uniform lattice $\Gamma \leq Isom^+(\mathbb{H}^n)$ where Γ has the following properties.

- *If we pick some spanning tree Λ' of $\Lambda(M)$ and define Y, D as in Definition 2.4.7, then D is injective on the interior of Y and has image a fundamental domain for the lattice Γ.*

- *For each edge e of $\Lambda(M)$ not in Λ' there are two faces F_1, F_2 in Y corresponding to this edge, and there is a unique orientation preserving isometry g_e (possibly $g_e = id.$) of $\mathbb{H}^3$ which takes $D(F_1)$ to $D(F_2)$. The set of all the g_e generates the lattice Γ.*

Definition 2.4.9. If a triangulation admits a set of model n-simplices satsifying the hypotheses of the above theorem, we say this triangulation admits the structure of a geometric triangulation, or simply admits a geometric/hyperbolic structure.

Remark 2.4.10. Suppose M satisfy the hypotheses of the above theorem. If we consider the n-simplices in $\mathcal{M}$ all embedded in $\mathbb{H}^n$, then we can define a unique orientation preserving isometry h_E of $\mathbb{H}^n$ for each oriented edge E of $\Lambda(M)$ (also thought of as an ordered pair of faces $(F \subseteq \Delta, F' \subseteq \Delta')$) by requiring that it takes F to F'. Note that as our tetrahedra satisfy the orientation requirement, this isometry will have the property that $h_E(\Delta)$ and Δ' have disjoint interiors.

The face pairings g_e of the polyhedron $D(Y)$ defined above can then be described as a product of the face pairings h_{E_i} where the E_i make up a loop in $\Lambda(M)$ consisting of e and the unique path in Λ' between its endpoints.

Remark 2.4.11. The above theorem also holds in more generality serving as a means of constructing geometric orbifolds from geometric building blocks of many sorts. For our purposes however, we want the result to be a hyperbolic manifold rather than an orbifold, this is equivalent to requiring the lattice be torsion free. Fortunately, in dimension 3 we get the extra information that the lattice is torsion free as torsion elements are ruled out by a combination of the angle and orientation equations. Indeed, the orientation restriction ensures that all elements of the lattice are orientation preserving, ensuring that the faces cannot contain the singular points that a reflection

would create. Furthermore, the angle equations ensure that interior points of edges cannot be singular points, thus the only points that could be singular are vertices. However, orientation preserving point stabilisers in $Isom^+(\mathbb{H}^3)$(elements of $SO(3)$) are rotations about some axis, and so this axis has to form part of the singular locus of the manifold thus guaranteeing that the singular set cannot be 0-dimensional, if it exists. Thus the only conclusion is that the singular locus is empty and the geometric orbifold is in fact a geometric manifold, as desired.

In this thesis whenever we apply Poincaré's Polyhedron Theorem we will either work with 3-manifolds, where the above discussion already confirms that the triangulation is a hyperbolic manifold, or we will deal with higher dimensional manifolds where the triangulation is already confirmed to admit the structure of a geometric triangulation of a hyperbolic manifold. Thus, we will not consider further the possibility of orbifolds in this thesis.

We finish by linking this result back into the results on systems of polynomials.. Define Σ to be a system of equations where the variables are given by the vertices of the tetrahedra in $\mathcal{M}$ and the equations are the edge equations, angle equations and equations which are equivalent to the orientation hypothesis. We will soon construct such a Σ more explicitly and ensure the equations are represented by polynomials but for now we simply state the following less precise corollary.

Corollary 2.4.12. *The triangulation of M admits a hyperbolic structure if and only if Σ admits a solution. Furthermore, each solution to Σ encodes a hyperbolic structure of the triangulation M.*

It's worthwhile here to reiterate that all different hyperbolic structures on a triangulation coming from all these different solutions give the same hyperbolic structure on the manifold, this follows from Mostow Rigidity.

Chapter 3

An Algorithm for Geometric Triangulations

3.1 Kalelkar and Phanse's Work in a Simple Case

This section attempts to demonstrate what kind of techniques one might use to prove the results that we shall later simply cite from Kalelkar and Phanse's paper. We do this by restricting our scope dramatically in comparison to their work, and focusing solely on the hyperbolic 3-dimensional case. We use slightly different techniques to theirs though very much inspired by their work, our focus on 3 dimensions is mostly for ease of reading but we also later in this section achieve a small improvement on the bounds in the 3-dimensional case when compared to [KP21].

For now, however, we shall simply be attempting to show that if one can control the geometry of geometric triangulations, then one can find a sequence of Pachner moves relating two geometric triangulations of the same hyperbolic 3-manifold. Furthermore, that sequence has length which only grows polynomially in the number of tetrahedra in each triangulation.

A key idea both here and in Kalelkar and his coauthors' work [KP21] [KP20] [KR20] is that generically, two tetrahedra embedded geodesically in a 3-manifold don't intersect one another too many times. It is only when the geodesic tetrahedra are embedded so as to be long, thin and in such a way that they "wrap around" the manifold multiple times, that they are able to intersect one another in ways that cause problems. We can counteract this by preventing such pathological tetrahedra from being included in our triangulations.

We also acknowledge the work of Mijatovic, [Mij03] [Mij03] [Mij04] [Mij05a] [Mij05b] whose descriptions of simple methods for changing triangulations via Pachner moves are the building blocks of our method here.

Theorem 3.1.1. *Let M be a closed hyperbolic 3-manifold and let $\mathcal{T}_1$, $\mathcal{T}_2$ be two geometric simplicial triangulations of M, consisting of t_1 , t_2 tetrahedra respectively. Then there is a function*

$$f(t_1, t_2, L, V) = 1000(t_1^2 + t_2^2) \cdot v(2L)/V$$

which bounds the number of Pachner moves necessary to take $\mathcal{T}_1$ to $\mathcal{T}_2$. Here L is the maximum length of an edge in either triangulation, V is the minimum volume of a tetrahedron in either triangulations, and $v(r)$ represents the volume of a ball of radius r in hyperbolic 3-space.

We prove this at the end of this section. The idea is that we simply provide a method for constructing a sequence of Pachner moves between two given triangulations of the same manifold. We take a common subdivision of both triangulations, call it $\mathcal{T}_3$ and construct a sequence of Pachner moves from $\mathcal{T}_1$ to $\mathcal{T}_3$. That the two triangulations have a common subdivision is provided by the fact that, by Theorem 2.1.13 ([Moi52]), the two triangulations are PL equivalent and thus combinatorially equivalent. However, we want a bounded sequence of moves and thus require a "bounded" subdivision. For this purpose we utilise that the triangulations are geometric.

We fix a copy of the hyperbolic manifold M which they triangulate and embed both triangulations at once. If we then intersect these triangulations, in most situations we get a polytopal subdivision. This only fails when two tetrahedra fail to intersect in a polyhedron but have non-trivial intersection, such as when two faces overlap but the interiors are disjoint. Cases such as these can however be managed by simply moving the vertices some infinitesimal amount, and thus we may assume that after some wiggling of vertices, our triangulations intersect in a polytopal complex which is then a polytopal subdivision of each of $\mathcal{T}_1$ and $\mathcal{T}_2$. We then use Adiprasito and Benedetti's Theorem 2.1.17 to see that the first barycentric subdivision of this polytopal subdivision is a simplicial subdivision of each triangulation which has the property that it induces a shellable subdivision on each of the original tetrahedra in our triangulations. We name this triangulation $\mathcal{T}_3$ and set about bounding its size. Once the size is bounded, we shall then find a path from $\mathcal{T}_1$ to $\mathcal{T}_3$using the shellability provided by Theorem 2.1.17.

As a first step in understanding the size of $\mathcal{T}_3$, we consider that it is made by barycentric subdivision of a polytopal subdivision, so we need to know how complex these polytopes (or polyhedra) can be.

Lemma 3.1.2. *The triangulated polyhedron formed by taking the intersection of two hyperbolic tetrahedra whose interiors intersect in $\mathbb{H}^3$ and then taking its barycentric subdivision is made up of ≤ 112 tetrahedra.*

Proof. Each face, F, of the polyhedron is formed by the intersection of one of the tetrahedra, T_1, with a face, $F_2 \subseteq T_2$, of the other tetrahedron. The intersection of T_1 with the plane containing F_2 is either a triangle or a square, depending on how the plane separates the vertices of T_1. So F is the intersection of the triangle F_2 with either a triangle or a square. In particular as each edge of F comes from a unique edge of the two polygons that form it, F has at most 7 edges. Similarly, as each tetrahedron has 4 faces, the polyhedron can only have at most 8 faces. Now, the barycentric subdivision of a polygon with 7 sides has 14 simplices, and so the number of simplices in the barycentric subdivision of our polyhedron is at most $8 \cdot 14 = 112$. □

Lemma 3.1.3. *The maximum number of polyhedra formed by intersecting a single 3-simplex T_1 of $\mathcal{T}_1$ with the triangulation $\mathcal{T}_2$ is bounded by $v(2L)/V$ where L, $v(2L), V$ are as defined in Theorem 3.1.1.*

Proof. Let T_1 be a simplex in $\mathcal{T}_1$. We want to bound the number of components of $T_1 \cap \mathcal{T}_2$. Let $\tilde{T_1}$ be any lift of T_1, then we only need to bound the number of lifts of tetrahedra in $\mathcal{T}_2$ which intersect $\tilde{T_1}$ as each such lift intersects in exactly one component as both tetrahedra are convex. Let p be some point in the interior of $\tilde{T_1}$ and let $\tilde{T_2}$ be a lift of some tetrahedron $T_2 \in \mathcal{T}_2$ which intersects $\tilde{T_1}$, then the distance from p to $\tilde{T_2}$ is bounded by the diameter of $\tilde{T_1}$ and so $\tilde{T_2}$ lies in $B(p, diam(\tilde{T_1}) + diam(\tilde{T_2}))$. Both these diameters are certainly less than L and so every such $\tilde{T_2}$ lies in the ball $B(p, 2L)$ and so in particular, the volume taken up by all these lifts must not exceed the volume of the ball. Thus the number of lifts of T_2 which intersect $\tilde{T_1}$ is bounded by

$$v(2L)/V.$$

□

Combining these two lemmas allows us to find a bound on the size of $\mathcal{T}_3$.

Corollary 3.1.4. *Let M be a closed hyperbolic 3-manifold and let $\mathcal{T}_1$, $\mathcal{T}_2$ be two geometric simplicial triangulations of M, consisting of t_1 , t_2 tetrahedra respectively. Let $L, V, v(2L)$ be the length and volume bounds defined in Theorem 3.1.1, then the following is a bound on the size (number of tetrahedra) of the common subdivision $\mathcal{T}_3$*

formed by intersecting the two triangulations and taking the barycentric subdivision of the resulting polytopal complex:

$$112 \cdot t_1 v(2L)/V.$$

Proof. The number of polyhedra in the polytopal complex formed by intersecting the two triangulations is bounded by $t_1 v(2L)/V$, this can be seen by applying Lemma 3.1.3 to each of the t_1 tetrahedra. Lemma 3.1.2 then tells us that each polyhedron barycentrically subdivided contributes less than 112 tetrahedra, as required. □

With the size of $\mathcal{T}_3$ bounded we now begin to describe a sequence of moves which allow one to traverse from a given triangulation to its subdivision. Our first step is to change to a simpler subdivision, which we call $\mathcal{T}_1'$. The subdivision $\mathcal{T}_3$ induces a subdivision on each tetrahedron Δ of $\mathcal{T}_1$. We now take the induced subdivision on $\partial\Delta$ and simply cone to this boundary and use that instead of the original subdivision, repeating this for each tetrahedron in $\mathcal{T}_1$ gives a subdivision of $\mathcal{T}_1$, and it is this subdivision which we call $\mathcal{T}_1'$. This subdivision is quite close to $\mathcal{T}_3$ in terms of Pachner moves, as Lemma 3.1.6 shows.

Remark 3.1.5. $\mathcal{T}_3$ is defined such that it induces shellable subdivisions on each tetrahedron of $\mathcal{T}_1$. This is a necessary condition for the use of our coning arguments. However, if a subdivision doesn't satisfy this, then its barycentric subdivision does.

Lemma 3.1.6. *Let B be a bound on the number of tetrahedra in $\mathcal{T}_3$. Then B is also a bound on the number of moves from $\mathcal{T}_3$ to the subdivision $\mathcal{T}_1'$ defined above.*

Proof. This is simply an application of Lickorish's lemma, our Theorem 2.1.18. The triangulations induced on a single tetrahedron of $\mathcal{T}_1$ by $\mathcal{T}_3$ and T_1' are a shellable triangulation and a cone respectively. This is because by construction $\mathcal{T}_3$ satisfies the premises of Adiprasito and Benedetti's Theorem 2.1.17. Thus if the triangulation induced by $\mathcal{T}_3$ comprises k tetrahedra, then it takes k moves to make it agree with $\mathcal{T}_1'$. As the sum of all these ks is by definition bounded by B, the result follows. □

Clearly, the fact that this subdivision is composed of cones is instrumental in the ease of the above proof. We formalise this concept now.

Definition 3.1.7 (Coned Subdivisions and Preserving the Cone Structure)**.** Coned subdivisions are defined to be those subdivisions of a simplicial n-complex which are formed by subdividing the $(n-1)$-skeleton and then taking the subdivision on each n-simplex to be a cone to its boundary. Moves which preserve the property of being a coned subdivision shall be said to preserve the cone structure.

Remark 3.1.8. A coned subdivision of a triangulated 3-manifold is uniquely determined by the subdivision it induces on the 2-skeleton of the original triangulation. Thus it makes sense to refer to **the** coned subdivision which induces a certain subdivision of the 2-skeleton.

Definition 3.1.9. Note that we can get from $\mathcal{T}_1$ to a very simple coned subdivision by performing t_1 (1-4) Pachner moves, one for each tetrahedron of $\mathcal{T}_1$. We shall call this the first coned subdivision of $\mathcal{T}_1$.

Thus we have two coned subdivisions of $\mathcal{T}_1$, its first coned subdivision and $\mathcal{T}_1'$. That we should be able to traverse between these two coned subdivisions is not a given, but we should be encouraged by the following observations. Firstly, as coned subdivisions are defined by their subdivision on the 2-skeleton, we have in some sense restricted our problem from a 3-dimensional object to a 2-dimensional one. Secondly, note that for any 2-simplex in $\mathcal{T}_1$ its subdivision induced by $\mathcal{T}_1'$ is "suspended", which is to say that the suspension of the simplicial 2-complex embeds in our subdivision. This is because, being a coned subdivision, the subdivided 2-simplex has cones on either side. The following shows why this is so important.

Lemma 3.1.10. *Let $\Delta_1, \ldots \Delta_k$ be 2-simplices arranged so that it is possible to perform a 2-dimensional Pachner move, (k= 1, 2 or 3 depending on the move). Then this move can be realised in the suspension $S(\Delta_1 \cup \ldots \cup \Delta_k)$ by performing two 3-dimensional Pachner moves. We shall refer to these composite moves as two-dimensional moves.*

Proof.

The (1-3) Move

Perform a (1-4) move on one of the tetrahedra, and then perform a (2-3) move incorporating the other tetrahedron.

The (2-2) Move

Perform a (2-3) move on the tetrahedra on one side of the suspension, then a (3-2) move incorporating the tetrahedra on the other side.

The (3-1) Move

Perform a (3-2) move on the tetrahedra on one side of the suspension, then a (4-1) move incorporating the tetrahedra on the other side. Note that this is simply the reverse of the (1-3) case.

□

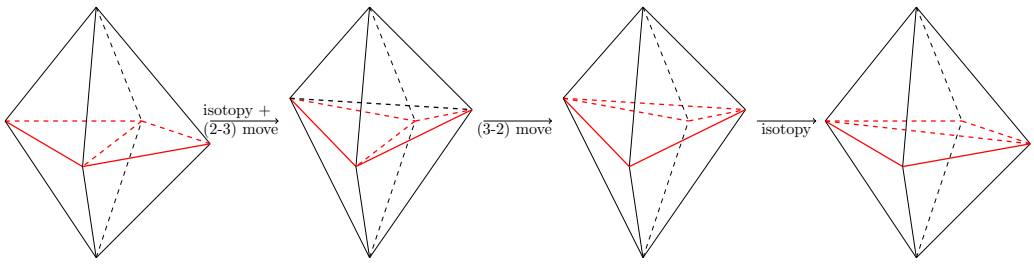


Figure 3.1: Performing the (2-2) move in suspension.

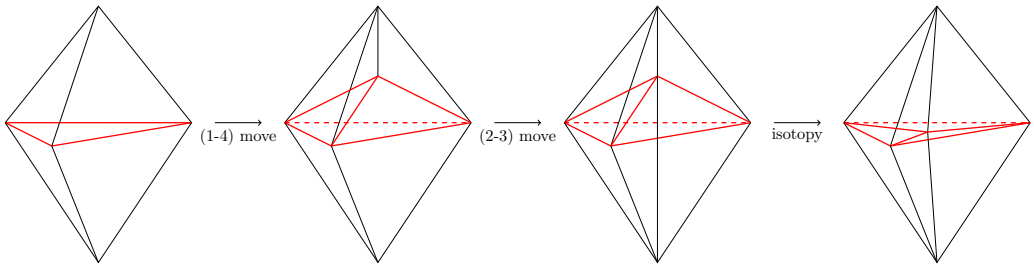


Figure 3.2: Performing the (1-3) move in suspension. Note the (3-1) move can be done by simply performing these moves in reverse.

This allows us to take the subdivisions of 2-simplices of $\mathcal{T}_1$ and apply 2-dimensional Pachner moves to them, this is almost enough to completely change the subdivision of the 2-skeleton, in fact we can use these two-dimensional moves to change the subdivision of each 2-simplex so that it too is a cone. We call the resulting subdivision a twice coned subdivision.

Definition 3.1.11 (k-Times Coned Subdivision)**.** A k-times coned subdivisions are defined to be those subdivisions of a simplicial n-complex which are formed by subdividing the $(n-k)$-skeleton and inductively taking the subdivision on each $(n-k+i)$-simplex to be a cone to its boundary. Again if a sequence of Pachner moves takes one k-times coned subdivision to another, we say that sequence preserves the cone structure.

The only problem is that Pachner moves on a subdivided 2-simplex cannot change the subdivision of the boundary. Fortunately, changing the subdivision of a 1-complex simply consists of inserting vertices into edges.

Lemma 3.1.12. *Let $\mathcal{T}$ be any twice coned subdivision of $\mathcal{T}_1$, and e an edge of $\mathcal{T}$ which lies in the 1-skeleton of $\mathcal{T}_1$, then we can insert a vertex into e while preserving the cone structure. Let k be the number of tetrahedra in the star of e, then this process takes k Pachner moves.*

Proof. Consider the star of a given edge e, $st(e)$ is made up of k tetrahedra labelled $T_1, \dots T_k$ arranged anticlockwise around e.
Let us also label the vertices of e (which are shared by all the T_i) by u and v and label by w_i the vertex shared by T_i and T_{i+1} (mod k). Performing a (1-4) move on T_1 creates a new vertex, x, and 4 tetrahedra, only one of which shares a face with T_2 - the tetrahedron with vertices (u, v, x, w_1). The goal is to make x the new vertex bisecting e.
We can now perform a (2-3) move on the pair formed by T_2 and the tetrahedron (u, v, x, w_1), which gives us 3 tetrahedra, each containing x, one of which shares a face with T_3. We can keep performing (2-3) moves like this, each time connecting x to a new w_i until x is connected in this way to w_{k-1}. At this point, the vertex x is connected to all the vertices of $st(e)$ but we still have the original edge e, which we want to eliminate. The three tetrahedra T_k, (u, v, x, w_{k-1}) and (u, v, w_{k-1}, w_k) surround e and a (3-2) move on these three removes e. Thus we are finished after k Pachner moves.

Note that the w_i are coning vertices making $\mathcal{T}$ a coned subdivision, and the process described above ensures exactly that each of these vertices now also connects to the new vertex x, thus this entire composite move preserves the cone structure of this twice coned subdivision. Note also that the new edges we have created have the same number of tetrahedra in their stars as the original edge did, because their stars are isomorphic. □

We have now established both two-dimensional moves and vertex adding moves, and described them as combinations of Pachner moves. It is now relatively simple to describe the path from the first coned subdivision to the $\mathcal{T}_1'$.

Lemma 3.1.13. *Let $\mathcal{T}_1$ be a triangulation consisting of t tetrahedra, and let $\mathcal{T}$ be a coned subdivision of $\mathcal{T}_1$ consisting of T tetrahedra and such that the subdivision induced by $\mathcal{T}$ on any 2-simplex of $\mathcal{T}_1$ is shellable. Note that this last requirement is always met if we construct $\mathcal{T}$ by taking the barycentric subdivision of a polytopal subdivision and then converting to a cone structure.*

We can get from this first coned subdivision of $\mathcal{T}_1$ to $\mathcal{T}$ through a sequence of combinatorial moves which preserve the cone structure. The number of vertex adding moves required is bounded by $3T$ and the number of two-dimensional moves is bounded above by $T + 2t$.

Proof. As vertex adding moves and two-dimensional Pachner moves preserve the cone structure, and coned subdivisions are uniquely determined by their restriction to the

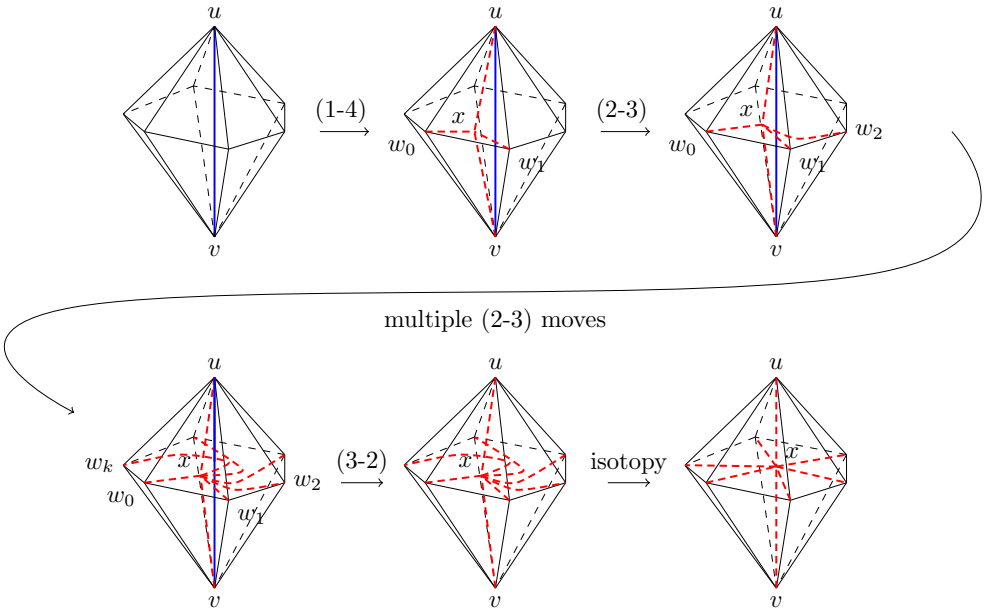


Figure 3.3: A sequence of Pachner moves inserting a vertex into an edge of our triangulation.

2-skeleton of $\mathcal{T}_1$, we only need to show that we can change subdivision on the 2-skeleton of $\mathcal{T}_1$ to that of $\mathcal{T}$ using only vertex adding moves and two-dimensional moves. Firstly, we extend the first coned subdivision of $\mathcal{T}_1$ to a twice coned subdivision. This consists entirely of performing exactly $2t$ two-dimensional (1-3) moves on the 2-simplices of $\mathcal{T}_1$. We then perform vertex adding moves until the triangulation agrees with $\mathcal{T}$ on the 1-skeleton using Lemma 3.1.12. The number of necessary vertex adding moves is bounded above by the number of vertices in $\mathcal{T}$ which lie in the 1-skeleton of $\mathcal{T}_1$. As each 3-simplex has at least one vertex on a cone point, this is bounded above by $3T$.

We then have a triangulation which agrees on the 1-skeleton with $\mathcal{T}$, and in fact when we restrict to the 2-skeleton of $\mathcal{T}_1$, our triangulation is a coned subdivision of the 2-skeleton of $\mathcal{T}_1$.

This means that if we look at the subdivision of each 2-simplex of $\mathcal{T}_1$ induced by our subdivision so far, compared to the subdivision induced by $\mathcal{T}$ the first is a cone to its boundary, and the second is a shellable triangulation of the 2-simplex which agrees with the first on its boundary. Thus using Lemma 2.1.18 we can perform two-dimensional Pachner moves after which the triangulation agrees with $\mathcal{T}'$ on this 2-simplex. Now the total number of such moves we need to make over all 2-simplices is bounded above by the total number of 2-simplices in $\mathcal{T}$ which lie in the 2-skeleton of $\mathcal{T}_1$, as $\mathcal{T}$ is a coned subdivision, each 3-simplex of $\mathcal{T}$ has exactly one such face and

so we perform at most T such moves.

Thus, after performing the above process on each 2-simplex, we have transformed the first coned subdivision of $\mathcal{T}_1$ to $\mathcal{T}$ in $\leq 4T$ vertex adding moves and $\leq T + 2t$ two-dimensional moves. □

The results of this section so far have proven an underlying fact, that a subdivision of a triangulation satisfying the simple condition that it induces a shellable triangulation on each simplex is related to the original triangulation by a sequence of Pachner moves. In fact the length of this sequence is almost linear in the number of tetrahedra in the subdivision. While we seek to use this fact to prove the relation between two geometric triangulations, it is useful on its own.

Theorem 3.1.14. *Given a triangulation $\mathcal{T}_1$ of a closed 3-manifold consisting of t_1 tetrahedra, and $\mathcal{T}_3$ a subdivision consisting of T tetrahedra. The number of Pachner moves required to change $\mathcal{T}_1$ to $\mathcal{T}_3$ is bounded by*

$$6Tt_1 + 9t_1 + 2T$$

Proof. Our first step is to move to the twice coned subdivision of $\mathcal{T}_1$, which takes $5t_1$ moves, t_1 to cone the 3-simplices and 2 moves to cone each of the $2t_1$ 2-simplices. We can then apply Lemma 3.1.13. This gives that we need to perform $3T$ vertex adding moves and $T + 2t_1$ two-dimensional moves. Each of these vertex adding moves consists of, by Lemma 3.1.12 at most $2t_1$ Pachner moves, as the star of any edge in the simple twice coned subdivision of $\mathcal{T}_1$ contains at most $2t_1$ simplices, as the original edge contained only t_1 at most and the coning process splits each of these in half. Thus the total Pachner move contribution from vertex adding moves is $6Tt_1$. Each two-dimensional move contributes 2 Pachner moves and so the total number contributed by such moves is $2T + 4t_1$. Combining the three numbers gives a bound on the total sequence length of

$$5t_1 + 6Tt_1 + 2T + 4t_1 = 6Tt_1 + 9t_1 + 2T$$

Note that for constant t_1 this depends linearly on T. □

We now have all the tools we need to prove our version of Kalelkar and Phanse's result.

Proof of Theorem 3.1.1. We simply use the bound on the length of the sequence of Pachner moves taking $\mathcal{T}_1$ to $\mathcal{T}_3$ and then by symmetry we have a sequence from $\mathcal{T}_2$ to

$\mathcal{T}_3$ and thus a sequence all the way from $\mathcal{T}_1$ to $\mathcal{T}_2$. Firstly, the number of tetrahedra in $\mathcal{T}_1, \mathcal{T}_2$ is t_1, t_2 respectively, and the number in $\mathcal{T}_3$ is $112 \cdot t_1 v(2L)/V$.

Thus using Theorem 3.1.14 we see that the number of Pachner moves required to get from $\mathcal{T}_1$ to $\mathcal{T}_3$ is bounded by

$$6 \cdot 112 t_1^2 v(2L)/V + 9t_1 + 2112 \cdot t_1 v(2L)/V \leq 1000 t_1^2 v(2L)/V$$

Repeating this for $\mathcal{T}_2$ and adding them together gives the result.

□

3.2 Bounding Simplex Volume

In the section after this one, we shall be constructing a system of polynomials with solutions which, by Poincare's Polyhedron Theorem, correspond to a geometric structure on a triangulation. Through this we shall be able to bound geometric properties of the triangulation and use those bounds to replace the constants L and V used in the previous section. We will do this by expressing the geometric properties of the simplices involved as polynomial functions of the positions of the vertices of these simplices.

Simplex volume is not a polynomial function of its vertices. We will see later that this simple fact doesn't always stop us. Edge length is not polynomial in the positions of the edge's vertices either, and yet Theorem 3.3.13 provides a system of polynomials which gives us a variable corresponding to edge length. To do this we bound the exponential of the length, hence its presence in the following lemma, but more on this in the next section.

However, in the case of volume, the function is simply too complex for the cheap tricks we use in the edge length case. Because of this, we opt for a coarser lower bound by measuring the volume of a subset of the simplex. To do this we shall use bounds on some obscure geometric properties of the simplex, relating to how they are embedded in the Poincare ball model of hyperbolic space. These properties are chosen, not because they are relevant in general, but because they are easy to find bounds on using polynomial systems and because they allow us to bound simplex volume specifically. Note also that though we only need this relationship in dimension 3, the argument works in n dimensions, and so we prove the more general case.

This process of turning bounds on obscure but easy to measure geometric properties into bounds on relevant but hard to measure geometric properties is key throughout this thesis.

As we have done before (see Lemma 2.2.36), we shall start with a somewhat tautological bound which we will later show has a well understood relationship to t, n.

Lemma 3.2.1. *There exists a function $B(t,n)$ such that for any geodesic triangulation, $\mathcal{S}$, of a closed hyperbolic n-manifold by t n-simplices, there exists a geodesic triangulation, $\mathcal{T}$ such that the following hold:*

- *$\mathcal{S}$ and $\mathcal{T}$ are simplicially isomorphic*
- *Let l be the length of an edge of $\mathcal{T}$, then the following inequality holds.*

$$B(t,n) \geq e^l \geq l \geq \frac{1}{B(t,n)}$$

- *Pick any vertex v_0 of an n-simplex of $\mathcal{T}$ and isometrically embed the n-simplex in the Poincaré ball model such that the chosen vertex is at the origin. Now measure the Euclidean distance from each other vertex $v_0, \ldots, v_n$ in this embedding to its opposing face, label these distances $h_0, \ldots, h_n$ respectively, then we have*

$$h_i \geq \frac{1}{B(t,n)}, \quad \forall i.$$

Proof. For each t, n there are only finitely many simplicial complexes made up of t n-simplices, and hence only finitely many homeomorphism classes of n-manifolds triangulable by t n-simplices up to simplicial isomorphism. Only finitely many of these admit a hyperbolic structure, so if we fix one structure for each such triangulation, then we have a finite number of hyperbolic structures, we can thus take minima and maxima over these and define $B(t,n)$ accordingly.

Note that when one embeds an n-simplex in the Poincaré ball model such that a certain vertex v_i is at the origin, the distance h_i is invariant under all isometries which fix v_i. This is because they are elliptics fixing the origin and thus elements of $SO(n)$, and thus they preserve Euclidean distances. □

In what follows $\mathcal{T}$ will be a geometric triangulation satisfying the bounds in the above lemma, and T will be an n-simplex in $\mathcal{T}$. Our plan is to use the above to bound the volume of T from below. We consider T as embedded in the Poincaré ball model of hyperbolic space, with one of its vertices v_0 at the origin. We then take a ball B around v_0 with radius equal to the distance between v_0 and its opposing face such that the interior of the ball doesn't intersect the opposing face of T. By the above

lemma, this means that the Euclidean radius of the ball (corresponding to h_0 in the above theorem) is at least $1/B(t,n)$ and as hyperbolic distance from the origin in the Poincaré ball model is always greater than Euclidean distance, the ball has hyperbolic radius r also bounded below by $1/B(t,n)$. Thus the volume of T is bounded below by

$$Vol(B \cap T) = Vol(B) \cdot Vol(B \cap T)/Vol(B).$$

Remark 3.2.2. As the interior of B doesn't touch the opposing face of T the fraction of the ball which lies in T, $Vol(B \cap T)/Vol(B)$, is the same as the fraction of the boundary of B which lies in T, $Vol_{\partial B}(\partial B \cap T)/Vol_{\partial B}(\partial B)$. Here $Vol_{\partial B}$ is the volume induced by considering ∂B as a subspace of $\mathbb{R}^n$. This measuring of what fraction of the ball lies in the n-simplex is a parallel of measuring the angle at a certain vertex in a triangle in 2 dimensions. Unfortunately, in higher dimensions the relationship between angles and volumes is much more complicated.

Lemma 3.2.3. *The value of $Vol_{\partial B}(\partial B \cap T)/Vol_{\partial B}(\partial B)$ is bounded below by*

$$\frac{1}{n!}\frac{1}{B(t,n)^{2n}}\frac{1}{Vol(S^6)}.$$

Proof. By convexity $\partial B \cap T$ forms a spherical $(n-1)$-simplex in ∂B. Projection onto the Euclidean $(n-1)$-simplex Δ spanned by the same vertices is a distance-decreasing map and hence reduces volume. Furthermore, the volume of ∂B is bounded above. Indeed, B is a ball of radius at most 1, so ∂B has smaller volume than the n-sphere, which has its volume bounded above by the 6-sphere, as we saw in Theorem 2.2.30. Thus we get that the following inequality holds:

$$Vol_{\partial B}(\partial B \cap T)/Vol_{\partial B}(\partial B) \geq Vol(\Delta)/Vol(S^6).$$

Now note that a standard formula for the volume of a Euclidean n-simplex is given by $\frac{1}{n!}h_1'h_2'\dots h_n'$ where, in the case of Δ, h_i' is the distance between a given vertex, w_i of Δ to the opposing face F_i' of Δ. Now w_i is a point on ∂B and B has radius bounded below by $1/B(t,n)$, furthermore the point v_i lies at euclidean distance < 1 from v_0 as the Poincare ball has euclidean radius 1. Thus, the point $(1/B(t,n))v_i$ lies inside B and thus lies between v_0 and w_i on the edge connecting them. At the same time F_i' lies in the face of T opposite v_i, which lies in some plane P_i, and clearly $1/B(t,n)$ is closer to this plane than w_i is, and so

$$h_i' = d(w_i, F_i') \geq d(w_i, P_i) \geq d((1/B(t,n))v_i, P_i) = (1/B(t,n))d(v_i, P_i) = (1/B(t,n))\, h_i.$$

Thus we get a relationship between h_i' and h_i, but we know h_i to bounded below by $1/B(t,n)$ itself, and thus $h_i' \geq 1/B(t,n)^2$ and so plugging this into the above formula gives that

$$Vol(\Delta) \geq \frac{1}{n!} frac1\big(B(t,n)^2\big)^n.$$

Combining this bound with the earlier inequality gives the stated result. □

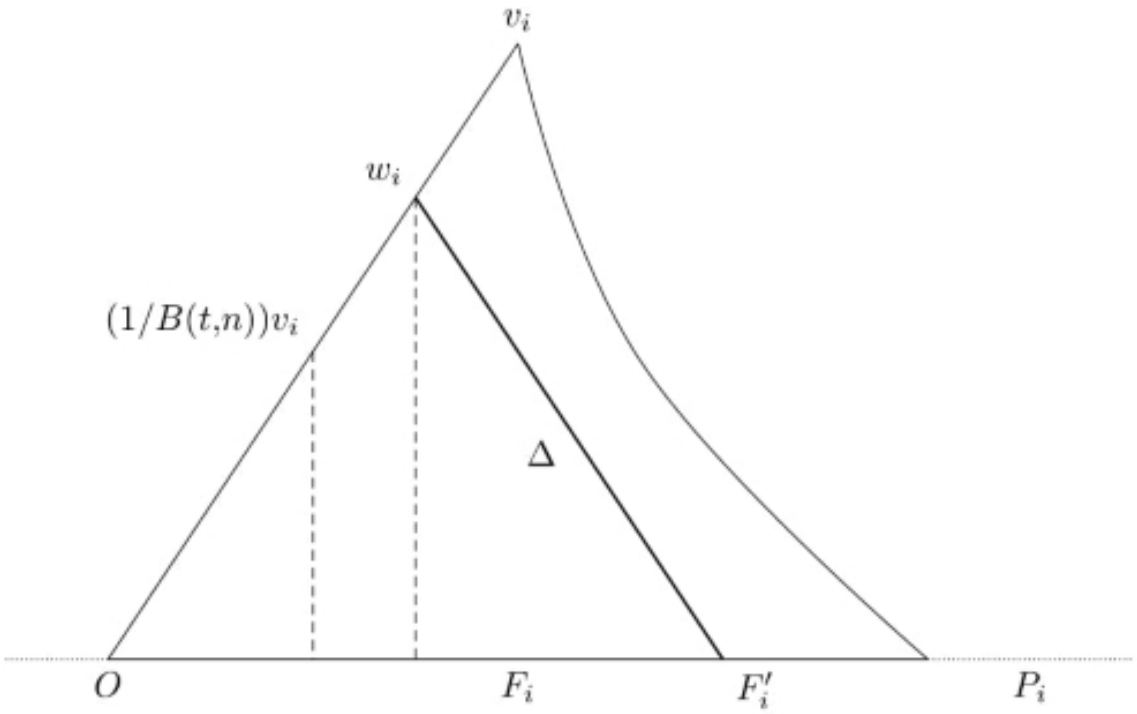


Figure 3.4: A 2-dimensional representation of the distances described above

We now have a lower bound on what fraction of the ball intersects the n-simplex T, so we only need to bound the volume of the ball itself.

Lemma 3.2.4. *The volume of a hyperbolic ball of radius r is bounded below by the volume of a Euclidean ball of radius r.*

Proof. One can calculate the volume of an n-ball in either geometry by integrating the surface area of the concentric $(n-1)$-spheres that make up the ball. The formulae for the surface areas $S_H(r), S_E(r)$ of hyperbolic and Euclidean $(n-1)$-spheres are:

$$S_H = \Omega_n sinh^{n-1}(r), \quad S_E = \Omega_n r^{n-1}, \quad \text{where } \Omega_n \text{ is a function of n.}$$

Thus, that the volume of a hyperbolic ball is bigger than that of a Euclidean ball is a consequence of the fact that $sinh(x) > x$. □

Proposition 3.2.5. *The volume of an n-simplex in $\mathcal{T}$ is bounded below by*

$$\frac{1}{n!}\frac{1}{B(t,n)^{2n}}\frac{1}{Vol(S^6)}\frac{\pi^{n/2}}{\Gamma(n/2+1)}\frac{1}{B(t,n)^n} \geq \frac{1}{(nB(t,n))^{3n}}$$

Proof. The volume of $B \cap T$ is just the product of the volume of B and the proportion of B lying in T. Theorem 3.2.3 gives us the latter, and Theorem 3.2.4 gives us that a lower bound on the former is the volume a Euclidean n-ball, which we know to have radius bounded below by $1/B(t,n)$. Thus using the formula $\frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)}r^n$ for the volume of a ball of radius r and noting that both $n!$ and $\Gamma(\frac{n}{2}+1)$ are bounded above by n^n we get the above bound. Here $\Gamma(n/2+1)$ is the Gamma function (a relative of the factorial function) evaluated at $n/2+1$. □

3.3 A Combinatorial Bound for Geometric Triangulations

With a lower bound on the volume of a simplex in our geometric triangulation comes a lower bound on the volume of the manifold, simply by multiplying by t. From this and the "tautological bound" found in Lemma 3.2.1, we can replace the result of Theorem 3.1.1 with one which depends entirely on t,n where t is the maximum of the number of n-simplices in either triangulation.

However, this requires the use of the function $B(t,n)$ defined in the previous section, and we have no understanding of how this function behaves. Thus the aim of the section is to find an easily understood bound on $B(t,n)$ and through that get not only a bounded runtime algorithm for comparing geometric triangulations but a bounded runtime algorithm for finding such a geometric structure on a triangulation in the first place.

In our discussion of the Poincaré Polyhedron Theorem (Section 2.4.2), we saw that when given the hyperbolic structure of the n-simplices involved in the triangulation, we only need to check that the gluing satisfies the edge, angle and orientation equations and then we know that the manifold admits a hyperbolic structure. We shall now construct a system of polynomials which encode these equations. To do this we start with variables coming from the vertices of model simplices. Consider each n-simplex as embedded in the upper half space model of hyperbolic space, then the hyperbolic structure of each n-simplex is determined by the coordinates of its vertices. Thus the n values which make up the position vector of a given vertex are the n variables corresponding to that vertex. As we are looking to find a bound on

edge length as well as the "heights" h_i defined in Lemma 3.2.1, we shall also introduce variables corresponding to their values.

Our reason for constructing such a system of polynomials whose variables include the edge lengths and the h_i's is that we can then apply the theorem of Grigoriev and Vorobjov from Section 2.4.1 to get quantifiable bounds both on the runtime of the algorithm finding a solution and on the size of the variables themselves.

Remark 3.3.1. Throughout this section we summarise the contribution of each lemma to the complexity of our system of polynomials at the end of each proof for use later. Recall that these measures of complexity are as follows

$$\begin{aligned} N &= \text{number of variables} \\ \kappa &= \text{number of polynomials} \\ d &= \text{maximum degree of polynomials} \\ M &= \text{maximum complexity of coefficients} \end{aligned}$$

Furthermore, note that with the variables and polynomials, we are counting, and so each new contribution increases the count. However, with degree and coefficients we are taking a maximum, so only contributions that increase the maximum affect the overall complexity.

Readers may be questioning at this point how we are going to express the edge and angle equations in a system of polynomials with rational coefficients. Indeed the length of an edge as a function of its vertices involves both a logarithm and a quotient. Furthermore the angles need to sum up to 2π which is a transcendental number and hence presents an immediate problem. Much of this section is focused on solving this problem of translating these non-polynomial equations into polynomial equivalents. Once these are sorted out, all that is left is a simple counting exercise. We start by noting that we can clean up our equations to separate out quotients and square roots.

Remark 3.3.2. Note that if we pay attention to how it affects numbers of polynomials/variables we can accept equations which involve both quotients and square roots of polynomials. For example if P, Q are polynomials, then the formula

$$\frac{P}{\sqrt{Q}} = 1$$

would appear to cause us some problems. However this can be fixed by introducing a new variable $v_{\sqrt{Q}}$ such that

$$(v_{\sqrt{Q}})^2 = Q, \qquad v_{\sqrt{Q}} \geq 0.$$

We also fix the fact that it is a quotient by multiplying through by the denominator to get that the original equation is equivalent to the polynomial

$$P - v_{\sqrt{Q}} = 0.$$

Remark 3.3.3. Sometimes, there are problems that cannot be solved with a single system of polynomials. When we are defining our system of polynomials, we may need to make an arbitrary choice from a possibility space of bounded size. For example it might be useful for our algorithm to know in which quadrant of the complex plane a complex number lies. Our response to a situation like this is to "make a guess", and encode a new system for each possible choice in which the variable must satisfy polynomials corresponding to the choice we made. In this case this is a complex number $x + iy$ satisfying something like $x \geq 0, y \geq 0$. In this example, this guess would lead to four new systems, one for each possible quadrant. Note that if the system without the guess had a solution then, at least one of our new systems must do too and any solution to any of the four new systems corresponds to a solution of the old one. At this point we can just run our algorithm as many times as there are possible choices (in this example four) and guarantee that we find a solution in one of the runs. Thus there are two factors in determining the complexity of our algorithm, the first is how long Grigoriev and Vorobjov's algorithm takes to run, which is determined by the size of our polynomial systems, the second is how many times we run it, which is determined by how many guesses we have made, and how many choices there are in each guess. We shall attempt to keep track of both these complexities throughout the section.

Remark 3.3.4. The first and most important place that this technique comes in is when we express the angle equations. As noted above, the angle equations naturally introduce a transcendental number, which poses algorithmic problems and certainly doesn't fit into our system of polynomials with rational coefficients. To solve this, we borrow an idea from Thurston's gluing equations (see [Thu78]or [Thu97]). We ask not that the sum of the angles θ_j add up to 2π, but rather that the product of the exponentials $e^{i\theta_j}$ comes to 1. These are not equivalent requirements, however, as the product coming to 1 is realised whenever the θ_j sum to any multiple of 2π including 0. This is where the need for a choice comes in. Before we can describe this choice, we note that we don't yet have a way of expressing the θ_j in terms of our existing variables, the position vectors of the vertices of our n-simplices. We shall show later how to define a variable which corresponds to $cos(\theta_j)$, and in doing so, also one which

corresponds to $sin(\theta_j)$. We thus have a pair of variables which denotes the position in the complex plane of each $e^{i\theta_j}$. It is using these variables that we can make guesses which control which multiple of 2π we are dealing with.

Remark 3.3.5. The following is not the exact method we follow, ours will be an approximation, however it is hopefully an easier to follow approach which is evocative of the true method detailed in the proof of Lemma 3.3.6.

Suppose the number of angles round a given codimension 2 face is k, then we pick some margin of error ϵ so that $k\epsilon$ is less than 2π. We then split the circle $S^1 \subseteq \mathbb{C}$ into equal segments of length $\leq \epsilon$ and for each angle θ_j round the given codimension 2 face, we choose a segment in which it lies. These segments each have minima and maxima, which means that the total angle sum also has a minimum and a maximum. We only consider segment choices that are such that the minimum angle sum is less than 2π and the maximum angle sum is greater than 2π. Note that because of our choice of ϵ, the only multiple of 2π which lies in the range of our minimum/maximum angle sum is 2π itself. Thus, the combination of this choice of segments and the equation $\prod e^{i\theta_j} = 1$ guarantees that the angles have to sum exactly to 2π.

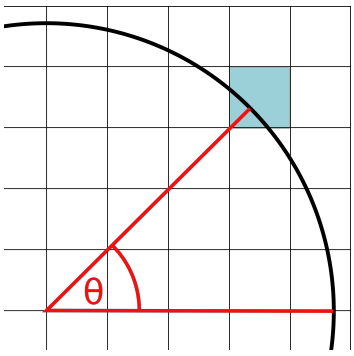


Figure 3.5: We can restrict the value of θ by requiring that the corresponding point on S^1 lies in a certain grid square.

So now we need to convert this idea into something which is actually encodable as polynomials. Remember again that we are assuming that we have a system of polynomials containing variables which correspond to the values $cos(\theta_j), sin(\theta_j)$, for every dihedral angle at every codimension 2 face of every n-simplex.

Lemma 3.3.6. *Let M be a hyperbolic n-manifold triangulated by t n-simplices modelled in the upper half space model of $\mathbb{H}^n$ and suppose we have variables which we know correspond to $cos(\theta)$ and $sin(\theta)$, for each θ a dihedral angle of our model n-simplices. For each codimension 2 face of the triangulation, we can list a finite set of choices*

of restrictions on the angles round that edge, and having made such a choice we can encode the angle equations as polynomials. At least one sequence of such choices for all codimension 2 faces of the triangulation corresponds to the actual angles one could measure in some hyperbolic structure on the manifold.

Proof. Firstly, we reiterate that half of the job shall be ensuring that the angle sum is some multiple of 2π, this is achieved by asking that $\Pi e^{i\theta} = 1$. This is polynomial in $cos(\theta), sin(\theta)$ as $e^{i\theta} = cos(\theta) + isin(\theta)$. In fact as $\Pi e^{i\theta} = 1$ is actually a statement about complex numbers, we get two polynomials, one setting the real part equal to 1, the other setting the imaginary part equal to 0, both have degree bounded by $t\binom{n+1}{2}$, as this is the total number of dihedral angles.

The choice of restrictions we now make is that for each θ, we set the pair $(cos(\theta), sin(\theta))$ to lie in a box of the form $[j\epsilon/18, (j+1)\epsilon/18] \times [j'\epsilon/18, (j'+1)\epsilon/18]$ where $\epsilon = 1/t\binom{n+1}{2}$ and j, j' are integers in $[-18/\epsilon, 18/\epsilon)$. Knowing these bounds tells us that $e^{i\theta}$ lies in a square in $\mathbb{C}$ of edge length $\epsilon/18$ and thus on a subarc of the circle of length $< \epsilon/6 < 2\pi\epsilon$. Thus, by the remark above, we know that for any given choice of these boxes, the total angle sum lies in a window of size less than 2π. Thus if we pick our choices so that 2π lies in the window of possible angle sums then it is the only possible multiple of 2π which can be in the window, and the product equations in the first paragraph give us that the angle sum must be a product of 2π, and so together we recreate the equation $\sum \theta = 2\pi$.

As noted in Remark 3.3.3, we have had to introduce an element of choice here to recreate the angle equations. We now track how many possible choices, and hence how many parallel systems of polynomials we now have to work with. The total number of possible boxes is $(36/\epsilon)^2 = 6^4 t^2 (\binom{n+1}{2})^2$ and we have to pick one for each of the $t\binom{n+1}{2}$ angles, so the possible number of distinct choices is at most $(6^4 t^2 (\binom{n+1}{2})^2)^{t\binom{n+1}{2}}$. Note that this number is certainly not optimised, but we shall see that it doesn't affect the runtime of an algorithm for deciding whether a triangulation admits a geometric structure. Furthermore, many of these choices may be nonsensical, for example angles can't be negative, and must be smaller than π and so we only consider choices where the boxes result in angles which lie in $(0, \pi)$. This reduces our total number of choices and thus reduces our eventual runtime.

For use later we note that the total number of polynomials we would need here would be the 2 product equations per codimension 2 face of the triangulation and the 4 equations per dihedral angle, defining the box in which it lies. In total, this is less than $6t\binom{n+1}{2}$ polynomials, and no new variables. The degree of each is less

than $t\binom{n+1}{2}$, and the most complex coefficient would be $\frac{18/\epsilon-1}{18/\epsilon}$ which has complexity $< log_2((18t\binom{n+1}{2})^2+1) < 36t(n+1)^2$. □

From this point on we assume when making our system of polynomials that we have already made one such choice of squares in which all our angles lie.

Thus, to complete the angle equations, all that remains is to define variables corresponding to the $sin(\theta)$ $cos(\theta)$ used above and to define polynomials which relate them to the vertices so that they truly are the angles to which they correspond.

Unfortunately, though our embedding in the upper half space model will be useful for our edge equations later and for many of the proofs in this thesis, one thing which is quite difficult in the upper half space model is understanding angles. This is because most hyperbolic geodesics and geodesic planes in the upper half space model are (from a Euclidean perspective) curved and so aren't easily described using linear equations. Fortunately, the Poincaré ball model is such that any two hyperplanes through the origin are subsets of Euclidean planes through the origin and thus we can easily find the angle between them by comparing the defining equations of each plane or equivalently by comparing normal vectors to each plane. Let's first show that we can add polynomials and variables that correspond to taking our n-simplices from one model to the other, in fact as we only need this to measure angles we need only introduce variables for the normal vectors to the faces of our n-simplex.

Lemma 3.3.7. *For each n-simplex in the upper half space model and each vertex of the n-simplex, we can embed the n-simplex in the Poincaré ball model such that that vertex is at the origin and then take unit normal vectors to each of the faces of the simplex which contain the chosen vertex. We can introduce variables corresponding to these normal vectors while contributing the following to the complexity of our system*

$$\begin{aligned} N &\leq 3tn^2(n+1) \\ \kappa &\leq 3tn^2(n+1) \\ d &\leq n+3 \\ M &\leq 3 \end{aligned}$$

Proof. Suppose the n-simplex is given by the vertices $\mathbf{v}_0, \ldots, \mathbf{v}_n$ with $\mathbf{v}_i = (v_{i_1}, \ldots, v_{i_n})$. First we take $\mathbf{v}_0$ to $(0, \ldots, 0, 1)$. The isometry ϕ defined by

$$(x_1, \ldots, x_n) \mapsto \tfrac{1}{v_{0_n}}(x_1 - v_{0_1}, \ldots, x_{n-1} - v_{0_{n-1}}, x_n)$$

takes $\mathbf{v}_0$ to $(0,\dots,0,1)$. We then map to the Poincaré ball model. The upper half space and Poincaré ball models are related by the following isometry, I, which is in fact an involution of $\mathbb{R}^n \cup \{\infty\}$ which takes the subspace corresponding to each model isometrically onto the other:

$$(x_1,\dots,x_n) = \mathbf{x} \mapsto 2\frac{\mathbf{x}+e_n}{\langle \mathbf{x}+e_n, \mathbf{x}+e_n\rangle} - e_n \qquad \text{where } e_n = (0,\dots,0,1)$$

Let us introduce new variables and polynomials which correspond to the images of these vectors, $I\circ\phi(\mathbf{v}_i) = \mathbf{v}'_i = (v'_{i_1},\dots v'_{i_n})$. The equation $v'_{i_j} = (I\circ\phi(\mathbf{v}_i))_j$ is not yet polynomial, but can be made so by multiplying through by the denominators present in the formulae for both I and ϕ, doing so yields a polynomial with degree at most 4 and with largest coefficient 4.

If we reorder the vertices of each n-simplex so that a different vertex becomes $\mathbf{v}_0$ then we can define ϕ exactly as above but on the new reordered simplex. If for each vertex we pick a reordering which labels that vertex as $\mathbf{v_0}$ and then define such a ϕ for each reordering, then as each new vertex in the image is made up of n variables we are introducing $tn^2(n+1)$ new polynomials defining as many new variables.

Because ϕ maps $\mathbf{v}_0$ to e_n and I maps e_n to the origin in the Poincaré ball model, all geodesics through $\mathbf{v}_0$ are mapped to straight lines by $I\circ\phi$. So all faces which contain $\mathbf{v}'_0$ lie in Euclidean hyperplanes and so it makes sense to calculate normal vectors for them. A codimension one face is defined by which vertex it doesn't contain, without loss of generality we shall consider the faces not containing $\mathbf{v}'_1$ and $\mathbf{v}'_n$ respectively.

A normal vector to the planes containing these codimension 1 faces can be derived by considering the $(n-1)\times n$ matrices

$$A = \begin{pmatrix} \mathbf{v}'_1 \\ \vdots \\ \mathbf{v}'_{n-1} \end{pmatrix}, B = \begin{pmatrix} \mathbf{v}'_2 \\ \vdots \\ \mathbf{v}'_n \end{pmatrix}$$

We define normal vectors $\mathbf{a} = (a_1,\dots a_n), \mathbf{b} = (b_1,\dots,b_n)$ by defining a_i (resp. b_i) to be the $(n-1)\times(n-1)$ minor of A (resp. B) formed by removing the ith column. Then $\mathbf{a},\mathbf{b}$ are orthogonal to their respective planes. Defining a normal vector consists of n new variables for each of the n faces of each of the $3t$ embeddings of n-simplices, and thus introduces $tn^2(n+1)$ new variables and polynomials. As they are defined as determinants of an $(n-1)\times(n-1)$ matrix, the degree of the polynomials is bounded by $n-1$ and the coefficients are all 1.

Now we can define $\mathbf{a}' := \mathbf{a}/(\mathbf{a}\cdot\mathbf{a})$ a unit vector in the same direction as $\mathbf{a}$, and by multiplying through by $\mathbf{a}\cdot\mathbf{a}$ the defining equations become polynomial of degree 3 and coefficients all 1. Again to do this we need $tn^2(n+1)$ new polynomials and variables. □

Lemma 3.3.8. *Adding the angle equations to our system of equations contributes the following to our system of polynomials*

$$\begin{aligned} N &\leq 3t\binom{n+1}{2} \\ \kappa &\leq 10t\binom{n+1}{2} \\ d &\leq 2n + t\binom{n+1}{2} \\ M &\leq 36t(n+1)^2 \end{aligned}$$

Proof. Given any pair of codimension 1 faces of one of our simplices, choose some embedding of our simplex such that both lie in Euclidean planes. At this point we have unit normal vectors to both and so the angle between the planes is given by $\theta = cos^{-1}(\mathbf{a}' \cdot \mathbf{b}')$. Thus $cos(\theta) = \mathbf{a}' \cdot \mathbf{b}'$ and $sin(\theta) = \sqrt{1 - (\mathbf{a}'_j \cdot \mathbf{b}'_j)^2}$. Combining these means that for each θ_j round an angle

$$e^{i\theta_j} = \mathbf{a}'_j \cdot \mathbf{b}'_j + i\sqrt{1 - (\mathbf{a}'_j \cdot \mathbf{b}'_j)^2}.$$

As noted in Remark 3.3.2 we can introduce new variables to deal with the square root. Indeed if θ_j is one of the angles round some codimension 2 face of the triangulation, then we can define the following variables and polynomials

$$C_j = \mathbf{a}'_j \cdot \mathbf{b}'_j, \quad D_j^2 = 1 - (\mathbf{a}'_j \cdot \mathbf{b}'_j)^2, \qquad D_j > 0.$$

And so the angle equation can be expressed as

$$\prod_{j=1}^{k} (C_j + iD_j) = 1$$

In the above, k corresponds to the number of n-simplices meeting at a codimension-2 face. Separating the real and imaginary parts of the equation above gives us a pair of polynomials with maximum coefficient equal to 1 and with degree equal to k, which is bounded above by $t\binom{n+1}{2}$. We have one such polynomial for each codimension 2 face of the triangulation, thus the number of polynomials contributed by just the angle equations themselves is at most $t\binom{n+1}{2}$.

We also introduced the new variables and polynomials to define C_j, D_j, which adds $2t\binom{n+1}{2}$ new variables and $3t\binom{n+1}{2}$ new polynomials of degree 2n and with coefficients at most 2.

As noted earlier, the contribution of requiring the angles lie within their chosen squares is at most $6t\binom{n+1}{2}$ polynomials, no new variables. The requirement that a given angle lie in the square is given by linear polynomials so their degree contribution is 1 with coefficients of complexity at most $36t(n+1)^2$. □

We now move on to the edge equations, which are slightly simpler in that they don't require us to make any choices in advance. However, we do still have the problem that the equations are not polynomial in the vertices, measuring the distance between two points remains a logarithm. This is solved however by introducing a new variable corresponding not the edge length itself, but to the exponential of the edge length. This is because if one bounds e^l above, then the logarithm of this bound is a bound on l itself. Furthermore, if we can bound below the variable $e^l - 1$, then we have a lower bound on l as either $1 < l$ or $1/2(e^l - 1) < l$.

Lemma 3.3.9. *Given a simplicial triangulation $\mathcal{T}$ of a closed hyperbolic n-manifold by t n-simplices, the edge equations can be expressed as a system of polynomials which introduces the following to the complexity of the system of equation*

$$\begin{aligned} N &\leq 4t\binom{n+1}{2} \\ \kappa &\leq (t\binom{n+1}{2})^2 + 7t\binom{n}{2} \\ d &\leq 2 \\ M &\leq 3 \end{aligned}$$

Proof. As noted in Section 2.2.1, in the upper half plane model of hyperbolic space, the distance function is defined by

$$d_{\mathbb{H}^n}((x_1,\ldots,x_n),(y_1,\ldots,y_n) = \\ log\left(\frac{\sqrt{(y_1-x_1)^2+\ldots+(y_{n-1}-x_{n-1})^2+(y_n-x_n)^2}+\sqrt{(y_1-x_1)^2+\ldots+(y_{n-1}-x_{n-1})^2+(y_n+x_n)^2}}{2\sqrt{y_n x_n}}\right)$$

So then setting variables D, D' equal to $e^l - 1, e^{l'} - 1$ the edge equation $l = l'$ is equivalent to $D = D'$. We do that in the following way, defining for each length l 4 variables, A, B, C, D as follows:

$$
\begin{aligned}
A^2 &= (y_1 - x_1)^2 + \ldots + (y_{n-1} - x_{n-1})^2 + (y_n - x_n)^2, && A > 0,\\
B^2 &= (y_1 - x_1)^2 + \ldots + (y_{n-1} - x_{n-1})^2 + (y_n + x_n)^2, && B > 0,\\
C^2 &= 4y_n x_n, && C > 0.\\
CD &= A + B - C
\end{aligned}
$$

If we define A', B', C', D' similarly for l', then D represents $e^l - 1$ and we can express $e^l = e^{l'}$ as

$$D = D'$$

So, to express the edge equations as polynomials, we introduce 4 new variables and 7 new polynomials for each edge of each model n-simplex, which amounts to $4t\binom{n+1}{2}$ new variables and $7t\binom{n+1}{2}$ new polynomials. Add to this the number of edge equations, which is bounded above by (number of edges in the model)^2, which amounts to $(t\binom{n+1}{2})^2$. As we can see above, the maximum degree among these polynomials is 2 and the biggest coefficient is 4 which has complexity 3. □

Now with the edge equations and the angle equations dealt with, we simply need to ensure that all n-simplices are similarly oriented in the upper half space. This is because once we know where an isometry sends a codimension 1 face, there are only two options for that isometry, one of which is orientation reversing. Thus because the edge equations holding tells us we can glue two n-simplices along their faces, the orientation equation determines uniquely the isometry gluing one n-simplex to the other.

Lemma 3.3.10. *The orientation equations can be included in our system of polynomials and contribute the following:*

$$
\begin{aligned}
N &\leq (n+1)^2 t\\
\kappa &\leq 2(n+1)^2 t\\
d &\leq 2\\
M &\leq 1
\end{aligned}
$$

In the process we also define variables which correspond to the Euclidean distances $h_1, \ldots, h_n$ defined earlier. We do this for all possible choices of which vertex v_0 is sent to the origin. We shall later use these to bound h_0 as well.

Proof. For every face of every n-simplex we have $n+1$ embeddings, the polynomials and variables for which we defined earlier when calculating angles. In n of these embeddings the image of the face is Euclidean, so to calculate the distance from this face to its opposing vertex, we define a new variable which is just the dot product of the unit normal vector of the face and the vertex. The modulus of this variable is the Euclidean distance between the vertex and the plane containing the face, which is a lower bound on the Euclidean distance between the vertex and the face. Once we have these variables, a polynomial of the form $A > 0$ or $A < 0$ for each variable according to the desired orientation of the n-simplex guarantees that our simplices are correctly oriented and nondegenerate.

This corresponds to two equations and one polynomial per embedding of each face of each n-simplex, so $2(n+1)^2t$ polynomials and $(n+1)^2t$ variables. The polynomials defining them have degree 2 and coefficient complexity 1. □

In the above proofs we have built up a system of polynomials, the complexity of which he have tracked, where a solution corresponds to a geometric structure on the n-simplices of M which all satisfy the edge, angle and orientation equations. We now summarise the size of that system of polynomials.

Lemma 3.3.11. *For each choice described in Lemma 3.3.6, we can build a system of polynomials which encodes the edge equations, the angle equations and the orientation hypotheses for a given triangulation consisting of t n-simplices. This system also has variables corresponding to exponentials of edge lengths and to a generating set for the lattice described in Theorem 2.4.8. The complexity of this system is defined by the following for some universal constant C:*

$$\kappa \leq Ct^2n^4$$
$$N \leq Ctn^3$$
$$d \leq Ctn^2$$
$$M \leq Ctn^2$$

Proof. Simply adding or taking maxima of the complexity numbers in the lemmas of this section gives the following bounds on complexity.

$$\kappa \leq (t\binom{n+1}{2})^2 + 7t\binom{n}{2} + 3tn^2(n+1) + 8t\binom{n+1}{2} + 2(n+1)^2t \tag{3.1}$$

$$N \leq 4t\binom{n+1}{2} + 3tn^2(n+1) + 3t\binom{n+1}{2} + (n+1)^2t \tag{3.2}$$

$$d \leq 2n + t\binom{n+1}{2} \tag{3.3}$$

$$M \leq log_2((4t\binom{n+1}{2})^2) + 2 \tag{3.4}$$

$$\tag{3.5}$$

We then note that there is some (not particularly large) constant C such that each of these is bounded as required. □

We can now find a more explicit bound to replace $B(t, n)$ in Lemma 3.2.1. This improved version of that lemma follows from combining the two main results of Section 2.4, Poincaré's Polyhedron theorem and Grigoriev and Vorobjov's algorithm for solutions of systems of polynomials and applying them to the system of polynomials developed in this section.

Theorem 3.3.12. *For any geodesic triangulation, $\mathcal{S}$, of a closed hyperbolic n-manifold by t n-simplices, there exists a geodesic triangulation, $\mathcal{T}$ such that the following hold:*

- *$\mathcal{S}$ and $\mathcal{T}$ are simplicially isomorphic*
- *Let l be the length of an edge of $\mathcal{T}$, then the following inequality holds.*

$$2^{2^{O(t^2n^4)}} \geq 2^{(tn)^{O(tn^3)}} \geq e^l \geq l \geq 2^{-(tn)^{O(tn^3)}} \geq 2^{-2^{O(t^2n^4)}}$$

- *Pick any vertex v_0 of an n-simplex of $\mathcal{T}$ and isometrically embed the n-simplex in the Poincaré ball model such that the chosen vertex is at the origin. Now measure the Euclidean distance from each vertex $v_0, \ldots, v_n$ in this embedding to its opposing face, label these distances $h_0, \ldots, h_n$ respectively, then the following inequality holds for any choice of n-simplex in $\mathcal{T}$ and any choice of embedding.*

$$h_i \geq 2^{-(tn)^{O(tn^3)}} \geq 2^{-2^{O(t^2n^4)}}$$

Proof. We start with any triangulation $\mathcal{T}$ simplicially isomorphic to $\mathcal{S}$ and attempt to impose a geometric structure on it. Given some choice of squares in which our angles lie, we can set up a system of polynomials which has a solution if and only

if the n-manifold defined by the triangulation $\mathcal{T}$ admits a hyperbolic structure such that the dihedral angles satisfy the constraints imposed by the choice of squares. This system satisfies the following bounds on its complexity:

$$N \leq 4t\binom{n+1}{2} + 3tn^2(n+1) + 3t\binom{n+1}{2} + (n+1)^2t$$
$$\kappa \leq (t\binom{n+1}{2})^2 + 7t\binom{n}{2} + 3tn^2(n+1) + 8t\binom{n+1}{2} + 2(n+1)^2t$$
$$d \leq 2n + t\binom{n+1}{2}$$
$$M \leq log_2((4t\binom{n+1}{2})^2) + 2$$

We know that the triangulation admits a geometric structure because $\mathcal{S}$ is a geodesic triangulation, so we know there is some choice of squares which yields a solution, consider the system of polynomials resulting from one such choice of squares. If we then apply Corollary 2.4.3, we see that as solution(s) $\alpha = (\alpha_1, \dots \alpha_n)$ to the system exists, then there is a solution which satisfies the following bounds on α_i

$$\frac{1}{2^{M((\kappa+2n)d)^{O(N)}}} \leq |\alpha_i| \leq 2^{M(\kappa d)^{O(N)}}$$

so substituting in our values for N, κ, d, M and simplifying big O notation gives:

$$2^{-2^{O(t^2n^4)}} \leq 2^{-(tn)^{O(tn^3)}} \leq |\alpha_i| \leq 2^{(tn)^{O(tn^3)}} \leq 2^{2^{O(t^2n^4)}}$$

Now, as we have variables representing $e^l - 1$ for every edge length so using the fact that $min\{1, \frac{1}{2}e^l - 1\} \leq l \leq e^l$ and that multiplying by one half and adding one are actions that disappear after simplifying big O notation, we get the following

$$2^{-(tn)^{O(tn^3)}} \leq min\{1, \frac{1}{2}e^l - 1\} \leq l \leq e^l \leq 2^{(tn)^{O(tn^3)}}$$

The same argument holds for the h_i, for $i \geq 1$, except that as the h_i are themselves variables, we don't need to translate to or from exponentials.

This leaves only the case of h_0. Note that we have bounds on the Euclidean distance between each vertex and its opposing face in all embeddings except when the vertex is sent to the origin. Note further that in the ball model, hyperbolic distance is always greater than Euclidean distance. So there is a ball of hyperbolic radius $r \geq 2^{-2^{O(t^2n^4)}}$ at the origin which doesn't touch the face opposite the origin. Now this ball has Euclidean radius $R = tanh(\frac{1}{2}r) \geq min\{\frac{1}{2}, \frac{1}{4}r\} \geq 2^{-(tn)^{O(tn^3)}}$, and thus we have our lower bound for h_0. □

If we instead apply Theorem 2.4.1 to the system of polynomials built up in this chapter, in particular that there is an algorithm to find a solution in bounded runtime, we see that if there is a hyperbolic structure making a certain triangulation geometric, then we can find one in bounded runtime. Because the runtime of the algorithm to find a positive result is bounded, we also get an algorithm for the negative result, either a triangulation admits a hyperbolic structure and we find it within the allotted runtime, or we run out of time and can conclude that no such structure exists.

Theorem 3.3.13. *Let M be a 3-manifold triangulated by T tetrahedra. Then if M admits the structure of geometric triangulation of a hyperbolic 3-manifold as described in Section 2.4.2, we can find such a structure in time bounded by*

$$T^{O(T^2)}.$$

Hence, we can decide whether or not a triangulation admits such a structure in the same runtime.

Proof. If we run the algorithm of Theorem 2.4.1 on all the systems corresponding to all the choices described in Lemma 3.3.6, then we get a solution for at least one of the choices if and only if M is a hyperbolic manifold and the triangulation admits a geometric structure. This follows from Corollary 2.4.12 and the fact that the list of choices exhausts all possibilities for the dihedral angles of the tetrahedra. The runtime of each algorithm and the number of times we must run it are both bounded by

$$(T)^{O(T^2)}.$$

This follows from applying Theorem 2.4.1 to the polynomial system constructed in this section, using the bounds on the size of the system given in Lemma 3.3.11 and from the count of choices in the proof of Lemma 3.3.6. Thus this is also a bound for the runtime of running Grigoriev's algorithm over all choices, after simplifying big O notation.

That we can then decide whether or not a triangulation admits a hyperbolic structure follows from the fact that if a structure exists it will be found by the algorithm within the given runtime bound, and thus if the algorithm reaches its runtime bound without finding a solution, then none exist. □

Before we use this to prove the main result of this chapter, we first prove some simple corollaries which will be of great use to us later.

Firstly we can simply apply this recognition algorithm to simplicial quotients of a triangulation also.

Corollary 3.3.14. *Given a triangulated 3-manifold M (by T tetrahedra) and a collection of simplicial identifications of the form $\Delta \to \Delta'$ for Δ, Δ' tetrahedra in M, we can decide whether those identifications define a simplicial quotient Q, whether that quotient is indeed a manifold and whether it admits a hyperbolic structure, all in time bounded by*

$$T^{O(T^2)}.$$

Proof. That we can find a hyperbolic structure for a triangulated 3-manifold in time bounded as above is given by Corollary 3.3.13. That we can check in polynomial time whether a 3-dimensional gluing is a 3-manifold is one of the 'simple algorithms' described in [FM97]. Thus it only remains to show that we can check whether a collection of identifications defines a simplicial quotient and that we can do so in time bounded as above.

This can indeed be done. Note that the number of possible identifications between two tetrahedra is $|A_4| = 12$ and the number of possible pairs of tetrahedra is $\binom{T}{2}$ and the number of possible equivalence classes of tetrahedra after identification is $< T$ and so the size of the data defining the quotient is polynomial in T. First we check, for each equivalence class, whether the identifications induce any self identifications of tetrahedra in the quotient and then check that no three tetrahedra now meet at a face in the quotient. Both these steps are polynomial in the size of the data defining the quotient and hence polynomial in T. □

The obvious next step is then to consider the situation where you know that a triangulation admits a simplicial quotient which has a hyperbolic structure. Can you find this hyperbolic structure? The answer is another simple corollary.

Corollary 3.3.15. *We can list all possible oriented geometrically triangulated simplicial quotients of a given triangulation of a closed 3-manifold consisting of T tetrahedra in time bounded by*

$$T^{O(T^2)}.$$

Proof. As shown above, the number of possible identifications between two tetrahedra is $|A_4| = 12$ and the number of possible pairs of tetrahedra is $\binom{T}{2}$. Thus there are less than

$$2^{12\binom{T}{2}}$$

possible combinations of identifications. For each we use Corollary 3.3.15 to check whether it gives a geometrically triangulated simplicial quotient in time bounded by

$$T^{O(T^2)}.$$

Thus the time taken to run this algorithm for all of the possible combinations of all the possible identifications is

$$2^{12\binom{T}{2}} T^{O(T^2)}$$

which simplifies to

$$T^{O(T^2)}.$$

At this point, as checking whether a triangulated manifold is oriented is decidable in polynomial time [FM97], we can also find the oriented sublist in the same time bound. □

As mentioned earlier, we will later also need to understand not just simplicial quotients but those which are degree one (the total signed number of preimages of each tetrahedron is 1).

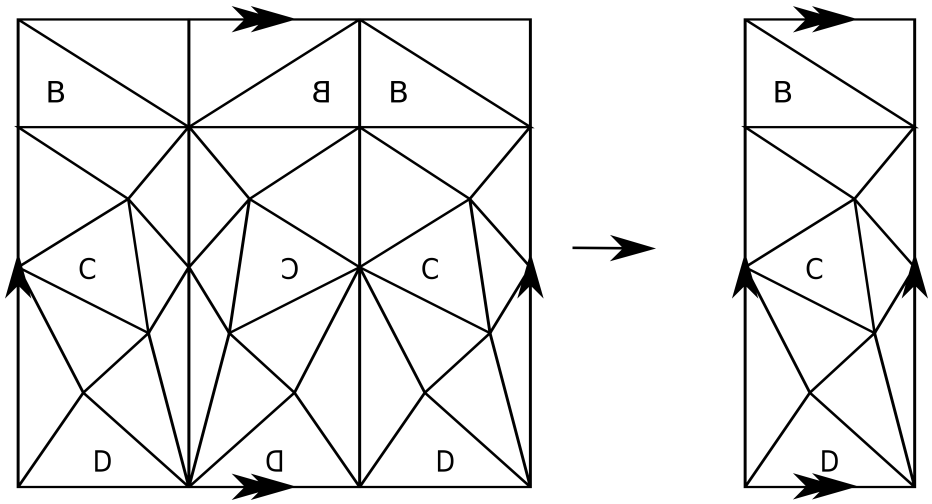


Figure 3.6: An example of a degree one simplicial quotient map from the torus to itself. The faces labelled with letters indicate how faces are identified to make the letters line up.

Lemma 3.3.16. *Given a simplicial quotient $\phi : M \to Q$ of oriented manifolds, we can find the degree of ϕ in polynomial time.*

Proof. As shown above, a simplicial quotient $M \to Q$ can be defined by a collection of simplicial identifications between tetrahedra in M. Also note that these collections are such that any induced self identifications are trivial. Thus if we have an orientation on M, and we pick some tetrahedron Δ in M, then we can consider the collection of tetrahedra which are identified with Δ and ask whether that identification is orientation preserving or reversing in each case. The degree of the map is then the number of tetrahedra glued in an orientation preserving manner (including Δ itself) minus the number glued in an orientation reversing manner. This can be checked in time polynomial in the number of tetrahedra in M. □

We summarise these corollaries in the following theorem.

Theorem 3.3.17. *Given a closed 3-manifold M triangulated by t tetrahedra and a bound T on the size of the subdivisions (number of tetrahedra) we are considering, we can list all oriented geometrically triangulated simplicial quotients of subdivisions of M where the quotient map is degree one and the subdivision consists of at most $T >> t$ tetrahedra in time bounded by*

$$T^{O(T^2)}.$$

Proof. Note that by Theorem 3.1.14, we can get from our triangulation of M to any subdivision of size bounded by T using composite moves consisting of $O(T^2)$ Pachner moves. We noted in the proof of Lemma 2.1.19 that the number of possible Pachner moves at any stage with this setup is linear in T, and so the total number of sequences and hence of subdivisions is at most $T^{O(T^2)}$. Applying the listing process of Corollary 3.3.15, gives us a list of geometric simplicial quotients in time $T^{O(T^2)}$ and then refining this list to one consisting of only degree one quotients is polynomial in T for each quotient. Hence, the result follows by simplifying big O notation. □

Theorem 3.3.18. *Let M be a closed hyperbolic n-manifold, $n \geq 3$ and let $\mathcal{T}_1, \mathcal{T}_2$ be two geometric simplicial triangulations of M, consisting of t_1, t_2 tetrahedra respectively. Then, taking t to be the maximum of t_1, t_2, there is a function $f(t, n)$ which bounds the number of Pachner moves necessary to take $\mathcal{T}_1$ to $\mathcal{T}_2$.*

In particular

$$f(t, 3) \leq 2^{2^{O(t^2)}}$$

and

$$f(t,n) \leq 2^{2^{tn^{O(tn^3)}}}$$

for $n > 4$.

Proof. The idea of this proof is to plug our geometry bounds into Kalelkar and Phanse's Theorem. In dimension three however, we instead use Theorem 3.1.1 which improves slightly on the final result. Starting with the 3D case, $f(t,3)$ is bounded above by

$$2000(t^2) \cdot v(2L)/V$$

Recall that L, V are the edge length upper bound and tetrahedron volume lower bounds respectively.

Firstly, using the bound from Theorem 3.3.12 in Proposition 3.2.5 tells us that V is at least

$$1/2^{2^{O(t^2)}}.$$

Furthermore, the volume of a ball in hyperbolic space is bounded above by e^L up to multiplying by constants. This we showed to also be bounded above by

$$2^{2^{O(t^2)}}.$$

Thus substituting these bounds into the bound from Theorem 3.1.1 and simplifying big O notation gives the following bound on Pachner moves in the 3D case

$$2^{2^{O(t^2)}}.$$

Now, moving up to n dimensions. It is possible to prove this relation using some complicated conversions between injectivity radius, volume and diameter. However, for brevity of proof we quote our bound on injectivity radius from the next section and plug that into Kalelkar and Phanse's result along with the bounds again from Theorem 3.3.12. This gives the following bound on the integer m in Theorem 2.1.20

$$m \leq (ncosh^{n-1}(L)+1)log(2L/inj(M) \leq ne^{nO(L)}log(2L/inj(M))$$

Now, by Theorem 4.4.8 and Theorem 3.3.12 we get both

$$\frac{2L}{inj(M)} \leq 2^{tn^{O(tn^3)}}, \qquad e^{nO(L)} \leq \left(e^{(}L)\right)^{O}(n) \leq 2^{tn^{O(tn^3)}}.$$

Combining these gives the following bound on m

$$m \leq 2^{tn^{O(tn^3)}}.$$

Substituting this into the main bound on Pachner Moves from Theorem 2.1.20 gives

$$(2n+2)!(n+1)!^{2^{tn^{O(tn^3)}}} \leq 2^{2^{tn^{O(tn^3)}}}$$

as required. □

Chapter 4

Injectivity Radius

4.1 Introduction

Our next step, having established how to traverse between geometric triangulations will be to work out how to find a geometric triangulation when given a non-geometric one. This, however, requires an understanding of the relationship between the number of tetrahedra in a simplicial triangulation of a hyperbolic manifold (a combinatorial property) and the geometric complexity of a possible geometric triangulation (a geometric property). The way we'll do this is by first attempting to understand the geometry of the manifold, which we can then use to construct bounded complexity geometric triangulations.

This chapter therefore shall focus on the injectivity radius of a hyperbolic manifold, how this relates to diameter, and how we can relate diameter to simplicial complexity.

Remark 4.1.1. The methods we use are very similar to the previous chapter, we use systems of polynomials to find bounds on variables which eventually correspond to geometric properties of the manifold. The difference, however, is that instead of searching for a geometric structure on the triangulation using the equations from Poincaré's Polyhedron theorem, we'll be searching for a hyperbolic structure for the manifold which comes from a representation of the fundamental group into the isometries of $\mathbb{H}^n$. This representation is induced by a cocycle, and that cocycle will be found as variables in our system of polynomials. In this sense, though the same tool provides the bound, the geometry involved is quite different and requires quite some explanation.

This chapter will consist of two major results, firstly we relate diameter to injectivity radius, then we relate diameter to combinatorial complexity. In this chapter

we actually focus on systole length over injectivity radius. This is because in finite volume hyperbolic manifolds, the injectivity radius of a point can be arbitrarily small as one travels further towards a cusp. However, the length of a smallest closed geodesic exists and is positive for both closed and finite volume hyperbolic manifolds of dimension ≥ 3.

For hyperbolic manifolds of dimension ≥ 3, Margulis' Lemma and the resulting thick-thin decomposition tell us that having a sufficiently short systole is indicative of the existence of Margulis tubes, whose diameters grow as the length of their core curves decrease. In this chapter we give an explicit lower bound on how the diameter of these tubes relate to the length of the core curve, and further show how this depends on n.

Theorem 4.2.5. *Suppose that M is a finite volume hyperbolic n-manifold ($n \geq 3$) with systole(s) of length $R \leq 2\epsilon_n$, where ϵ_n is the Margulis constant in dimension n. Then the distance from a systole to the boundary of the Margulis tube containing it is bounded below by*

$$\frac{1}{n} log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4).$$

A less explicit version of this theorem is surely known to experts. In dimensions ≥ 4 Theorem 2.2.33 which we cited earlier due to Reznikov [Rez95] shows that for a closed hyperbolic manifold a small injectivity radius implies that the manifold has both a large diameter and volume, but how this depends on dimension isn't known and the theorem doesn't hold in dimension 3. Note that one can reprove this relation for $n \geq 8$ using the above Theorem 4.2.5 and Gromov's result relating diameter and volume (Theorem 2.2.31).

Belolipetsky and Thomson [BT11] also prove a similar result for some infinite families of hyperbolic manifolds whose injectivity radii tend to 0. It is worth noting that their result, which builds on earlier work of Agol [Ago06], demonstrates the existence of hyperbolic manifolds with arbitrarily small systoles and thus proves that there can be no universal lower bound on systole length.

Remark 4.1.2. In this chapter will need that the edges in our triangulations are not loops, that is their vertices are distinct. As has been noted before in this thesis, if we are given a non-simplicial triangulation we can barycentrically subdivide to make our triangulations simplicial. In particular if a triangulation is simplicial, its edges aren't loops. Unfortunately this barycentric subdivision increases the size of the triangulation from t n-simplices to roughly tn^n simplices. However, triangulations

often naturally occur which already meet this requirement. Because of this, we will sometimes provide two bounds in our theorem, one for triangulations which have no loops, and one for triangulations which may have loops.

In the second main theorem of this chapter we indeed show that triangulation complexity can be used to bound from below the length of a systole.

Theorem 4.1.3. *Given a closed (or finite volume) hyperbolic n-manifold ($n \geq 3$) triangulated by t n-simplices (possibly semi-ideal), the length R of a systole of M is bounded below by a function of t and n, in particular*

$$R \geq \frac{1}{2^{(nt)^{O(n^{n+5}t)}}}.$$

This improves if we can assume the triangulation has no loop edges (all edges have distinct vertices) to

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}$$

Remark 4.1.4. Note that for a non-orientable hyperbolic n-manifold, we can apply Theorem 4.1.3 to its orientable double cover and use that to bound the length of systoles in the original manifold. In doing so we double the number of tetrahedra and halve the bound on systole length. This difference disappears when we simplify big O notation however.

Because of this, we shall from this point on assume all our manifolds to be orientable, and thus all our lattices shall lie in $\mathrm{Isom}^+(\mathbb{H}^n)$.

Remark 4.1.5. Another thing to note is that the definition of semi-ideal used here and defined in Section 2.1.3 is non-standard (here semi-ideal simplices are required to have exactly one ideal vertex and no more). Semi-ideal simplices can be defined more generally to have possibly more than one ideal vertex, but throughout this paper we shall deal only with semi-ideal simplices with exactly one ideal vertex. This is because a geodesic semi-ideal simplex with only one ideal vertex can be modelled in the upper half space model of hyperbolic space so that the ideal vertex is at the point at infinity. This has the benefit that all semi-ideal faces are then vertical, and so can be easily understood simply by knowing their non-ideal vertices.

We don't lose much generality by making this restriction either, as if we are given an ideal triangulation or a semi-ideal triangulation using the more general definition, its barycentric subdivision is a semi-ideal triangulation using the more restrictive definition we have given.

We've spoken already about how these statements interact with the higher dimensional results of Reznikov and Gromov, but if we restrict to 3 dimensions, more again is known. Futer, Purcell and Schleimer [FPS19] prove a 3-dimensional version of Theorem 4.2.5 which they prove to be sharp up to the additive constants and which is still linear in $\log(1/R)$. Kalelkar and Raghunath [KR20] also provide a result linking the number of tetrahedra in a triangulation of a cusped hyperbolic 3-manifold M to systole length. Their result gives much better bounds in terms of t but requires that the original triangulation be geometric and also depends on the geometry of the triangulation, which we've already seen is quite a restriction. A closed 3-dimensional version of Theorem 4.1.3 follows from a combination of Theorem 4.2.5 (or indeed [FPS19]) and another result relating combinatorial and geometric properties of hyperbolic 3-manifolds, due to Matthew White [Whi00a] (see also [Whi00b] [Whi02]). White's work was never published, but a stronger version of his theorem was proven by Agol and Liu [AL10].

Theorem (Theorem 5.9, [Whi00b])**.** *There is an explicit constant $K > 0$ such that if M is a closed, connected, hyperbolic 3-manifold, and*

$$P = \langle x_1, \ldots, x_n \mid r_1, \ldots, r_m \rangle$$

is a presentation of its fundamental group, then $diam(M) < K \cdot l(P)$, where

$$l(P) = \sum_{i=1}^{m} l(r_i)$$

and $l(r_i)$ is the word length of a given relator.

The 3-dimensional version of Theorem 4.1.3 that follows from this combination has a tighter bound on systole length of the form

$$R \geq \frac{1}{2^{O(t)}}.$$

We give a proof of this in the final chapter, Corollary 5.1.4.

4.2 Diameter of Margulis Tubes

In this section we shall describe how to bound the length of a systole from below as a function of the diameter of the Margulis tube containing it.

The tool we will use to do this is the thick-thin decomposition, which we described in Section 2.2.2.

Remark 4.2.1. The model of hyperbolic space we shall use in this section is different from that which we will use in the next. For the purpose of understanding cocycles and defining a system of polynomials, we will use the hyperboloid model because its isometry group is simply described by polynomials. In this section, we will use the Euclidean structure of horospheres and so we use the upper half space model where any given horosphere can be taken to a horizontal Euclidean plane by an isometry of $\mathbb{H}^n$.

We begin by recalling a quick lemma about closest point projection to the vertical axis in the upper half space model.

Theorem 2.2.11. *The distance from a point $\underline{x}$ to the vertical axis through the origin in the upper half space model is given by*

$$arcosh\left(\frac{\|\underline{x}\|}{x_n}\right)$$

where $\|\underline{x}\|$ is the standard Euclidean norm.

Statements about the geodesic between 0 and ∞ are not as specific as they may seem. In fact, they allow us to understand Margulis tubes because the action of the isometry group of hyperbolic n-space on the boundary of hyperbolic space is 2-transitive and hence we may assume that our systole lifts to any geodesic of our choosing.

Recall that the translation length of a loxodromic isometry is precisely the distance it moves any point which lies on its invariant geodesic axis.

Remark 4.2.2. When considering the injectivity radius of a point p in $\mathbb{H}^n/\langle\phi\rangle$ for ϕ a loxodromic isometry with axis $0, \infty$, it is not sufficient to consider simply the distance between p and $\phi(p)$. Though these two numbers do coincide for p on the axis consider ϕ to be the isometry of $\mathbb{H}^3$ of the form $x \mapsto (1+\delta)\cdot Ax$ where A is a rotation by π in the x, y coordinates. Consider the point p to be $(D, 0, 1)$, for $D >> 0$ then its image under ϕ is $(-(1+\delta)D, 0, (1+\delta))$ and the distance between the two is greater than the distance from p to $(0, 0, 1)$ which is approximately $2log(D) >> 0$. However, if we apply ϕ twice, its image is $((1+\delta)^2 D, 0, (1+\delta)^2)$ which has a much smaller distance to p, in particular as δ tends to 0, so does distance to p whereas $2log(D)$ is independent of δ.

In a more general setting, A may not be of order 2, it may be infinite order for all we know and so we will have to work out how many times we need to apply ϕ until we can see this dependence on δ, but the same trick can be applied. Note that δ here

is intrinsically linked to the translation length of ϕ along its axis, and hence to the length of the systole in $\mathbb{H}^n/\langle\phi\rangle$.

Definition 4.2.3. Note that up to swapping 0 and ∞, all loxodromics fixing the $(0,\infty)$ axis are of the form $(1+\delta)Ax$ where $\delta > 0$ and A only acts on the first $n-1$ coordinates as an orthogonal map. Call this A the orthogonal map.

In the following we generalise the above example by finding a power k such that when the orthogonal map is taken to this power, it is very close to the identity, and thus when the loxodromic is taken to this power, the distance it translates a given point is predominantly dependent on the translation distance of ϕ and the power k.

Lemma 4.2.4. *Let ϕ be a loxodromic isometry in $Isom^+(\mathbb{H}^n)$ acting on the upper half space model with fixed points at 0 and ∞. Suppose the translation length of ϕ along its axis is R. If $\underline{x}$ is such that $\frac{\|\underline{x}\|}{x_n} \leq e^D$, then for each $1 > a > 0$, we can find an integer $0 < k \leq \left(\frac{4e^D}{a}\right)^{n-1}$ s.t. $d\left(\phi^k(\underline{x}), e^{kR}(\underline{x})\right) < a$.*

Proof. A loxodromic isometry ϕ of translation length R acting on the upper half space model of hyperbolic space fixing 0 and ∞ is a composition of a scaling $\underline{x} \mapsto e^R\underline{x}$ and an orthogonal map $A \in SO(n-1)$ acting only on the first $n-1$ variables. Thus if $\underline{x} = (x_1, \dots x_n)$ then

$$d(\phi^k(\underline{x}), e^{kR}\underline{x}) = d(A^k \cdot e^{kR}(\underline{x}), e^{kR}\underline{x}) = d(A^k\underline{x}, \underline{x})$$

We can apply a further isometry such that $\underline{x}, A^k\underline{x}$ both lie on the plane $x_n = 1$ and by Lemma 2.2.11 it is still true that $\frac{\|\underline{x}\|}{x_n} \leq e^D$ and hence if we let π be the projection map onto the first $n-1$ factors we get

$$e^{2D} \geq \|\underline{x}\|^2 = \|\pi(\underline{x})\|^2 + 1 \qquad \text{and so} \qquad \sqrt{e^{2D}-1} \geq \|\pi(\underline{x})\|.$$

In particular, if we restrict to the hyperplane $\mathbb{R}^{n-1} \times \{1\}$ the $\frac{a}{2}$-ball around each of the translates $A^k\underline{x}$ lies within the $2e^D$-ball in $\mathbb{R}^{n-1} \times \{1\}$. Indeed, let $\underline{x}'$ lie in such an $\frac{a}{2}$-ball, then $(\frac{a}{2})^2 < a < 1$ and A being orthogonal gives $\|\pi(A^k\underline{x})\| = \|\pi(\underline{x})\|$ and thus the following inequalities hold:

$$\|\pi(\underline{x}')\|^2 \leq \left(\|\pi(A^k\underline{x})\| + \frac{a}{2}\right)^2 \leq \|\pi(\underline{x})\|^2 + a\|\pi(\underline{x})\| + \left(\frac{a}{2}\right)^2$$

$$< e^{2D} - 1 + a\sqrt{e^{2D}-1} + 1 < 2e^{2D}.$$

Thus $\|\pi(\underline{x}')\| < \sqrt{2}e^D < 2e^D$.

Now let $vol(B(r))$ denote the volume of the r-ball in Euclidean $(n-1)$-space and define

$$f(a) = \frac{vol(B(2e^D))}{vol(B(\frac{a}{2}))}.$$

If $k > f(a)$ then there must be some pair of points $A^i\underline{x}, A^j\underline{x}$ with $i < j$ such that their $\frac{a}{2}$-balls overlap, otherwise all the balls are disjoint but the sum of their volumes is greater than the volume of the $2e^D$-ball containing them.

Thus the result follows as

$$\frac{vol(B(2e^D))}{vol(B(\frac{a}{2}))} = \frac{(2e^D)^{n-1}}{\left(\frac{a}{2}\right)^{n-1}} \frac{vol(B(1))}{vol(B(1))} = \left(\frac{(4e^D)}{a}\right)^{n-1}$$

and if the neighbouring balls of $A^i\underline{x}, A^j\underline{x}$ overlap then so do the balls round $\underline{x}, A^{j-i}\underline{x}$ and so $A^{j-i}\underline{x}$ is within Euclidean (and hence also hyperbolic) distance a of $\underline{x}$ as required. □

This understanding of the distance that a given isometry translates points in its neighbourhood is directly related to the injectivity radius of points in the neighbourhood of a systole. If we can use this to show that all points a certain distance away from the systole have injectivity radius bounded above by the Margulis constant, then we have shown that this neighbourhood is part of the Margulis tube and hence isometrically embeds into the original manifold. We formalise this below.

Theorem 4.2.5. *Suppose that M is a finite volume hyperbolic n-manifold ($n \geq 3$) with systole(s) of length $R \leq 2\epsilon_n$, where ϵ_n is the Margulis constant in dimension n. Then the distance from a systole to the boundary of the Margulis tube containing it is bounded below by*

$$\frac{1}{n} log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4).$$

Proof. Let γ be a systole of M. We consider $N = N(\gamma, D)$ in an attempt to bound below the maximum value of D such that this neighbourhood lies in the thin part of M.

By definition γ is a geodesic and so the corresponding isometry of hyperbolic space in the deck group of M is a loxodromic with geodesic axis $\tilde{\gamma}$ which projects down to γ under the covering map. By conjugation this can be modelled in the upper half space model of hyperbolic space by the isometry

$$\phi : \underline{v} \mapsto A \cdot e^R \underline{v}$$

with geodesic axis the geodesic from 0 to ∞ where R is the translation length of ϕ and A is an orthogonal map on the first $n-1$ variables.

Now let $\underline{x} := (x_1, \ldots, x_n)$ be a point in the upper half space model for $\mathbb{H}^n$ within distance D from the axis of ϕ, so in particular $\frac{\|\underline{x}\|}{x_n} < cosh(D) < e^D$ and let $\underline{x}' := (x_1, \ldots, x_{n-1}, e^{kR}x_n)$ be the vertical translate of $\underline{x}$ by hyperbolic distance kR. Then

$$d(\underline{x}, \phi^k(\underline{x})) \leq d(\underline{x}, \underline{x}') + d(\underline{x}', e^{kR}\underline{x}) + d(e^{kR}\underline{x}, \phi^k(\underline{x})).$$

Now $d(\underline{x}, \underline{x}')$ is just the length of the vertical geodesic from x to x'. This is $log(e^{kR}x_n) - log(x_n) = log(e^k R) = kR$.

Furthermore, the distance from $\underline{x}'$ to $e^{kR}\underline{x}$ is bounded above by their distance in the horosphere, which is Euclidean up to scaling. So if $\pi(x)$ represents the projection to the first $n-1$ factors, the following holds

$$d(\underline{x}', e^{kR}(\underline{x})) \leq \frac{1}{e^{kR}x_n}\|e^{kR}\pi(x) - \pi(x)\| = \frac{(e^{kR}-1)\|\pi(x)\|}{e^{kR}x_n} \leq (e^{kR}-1)(e^D).$$

So by Lemma 4.2.4

$$d(x, \phi(x)) \leq kR + (e^{kR}-1)e^D + a.$$

At this point we note that the following statement follows by simple rearrangment and taking logarithms

$$D < \frac{1}{n}log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4) \quad \Leftrightarrow \quad R\left(\frac{4e^D}{\epsilon_n}\right)^n < 1 \tag{4.1}$$

Thus, if we set $a = \epsilon_n$ and require

$$D < \frac{1}{n}log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4)$$

then for $k < \left(\frac{4e^D}{a}\right)^{n-1}$ it is both true that

$$kR \leq R\left(\frac{4e^D}{\epsilon_n}\right)^{n-1} < R\left(\frac{4e^D}{\epsilon_n}\right)^n < 1$$

and

$$kR(2e^D+1) < \left(\frac{4e^D}{\epsilon_n}\right)^{n-1} R(4e^D) < R\left(\frac{4e^D}{\epsilon_n}\right)^n \epsilon_n < \epsilon_n$$

Combining this with the fact that $e^x - 1 < 2x$ on the interval $(0,1)$ gives that

$$d(\underline{x}, \phi(\underline{x})) \leq kR + (e^{kR}-1)e^D + a \leq (1+2e^D)kR + \epsilon_n < 2\epsilon_n$$

and hence if $\underline{x}$ is within distance D of the axis of ϕ then $\underline{x}$ maps to the thin part of M. Thus the thick-thin decomposition implies that the diameter of the Margulis tube is at least

$$\frac{1}{n} log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4).$$

□

4.3 Bounded Cocycles

With a link between diameter and injectivity radius established, we now search for a way to bound the diameter of a hyperbolic manifold as a function of its number of simplices. We will do this by finding an edge length bound as we did when we were searching for a geometric triangulation.

However, in this case the edge length bound we will find will not be for a geometric triangulation, in fact it won't be for a triangulation at all. In this section we take our topological triangulation and pull each simplex in turn into a straight or geodesic simplex, meaning that a lot of the structure is preserved, all our simplices are still paired along their faces as they were before, but we introduce degeneracy. It is not hard to find an example of two topologically embedded 2-simplices in $\mathbb{R}^2$ for example where if one keeps the vertices fixed and straightens the edges, the resulting triangles overlap. This overlapping means that we no longer have a triangulation, and lose much of the structure that this provides.

However, it can be made clear that this straightening process is also homotopic to doing nothing at all and so we know that the natural map from this original triangulation to the straightened 'triangulation' is a homotopy equivalence. We can use this to show that the map is also surjective, and hence we have created for ourselves a cover of our manifold by geometric pieces, whose geometry and number we can control. With this information in hand, bounding the geometry of the entire manifold is simple. All we need to show therefore, is how the geometry of the pieces relates to the number of simplices in the starting triangulation.

We note that this process of straightening was used quite notably by Gromov. See chapter C of [BP92] for a description of its application in Gromov's alternate proof of Mostow Rigidity.

We will start by defining and bounding this straightening homotopy equivalence, but first some notation.

Let M be an n-manifold admitting a finite volume hyperbolic structure equipped with a triangulation (semi-ideal if the manifold is not closed) and let $X = G\backslash\mathbb{H}^n$

be a hyperbolic n-manifold homeomorphic to M, so $G \leq Isom^+(\mathbb{H}^n)$. Let f be a homeomorphism $f : M \to X$, then f induces a G-equivariant homeomorphism $\tilde{f} : \tilde{M} \to \mathbb{H}^n$ because the action of G on $\tilde{M}$ is defined by following the chain of isomorphisms

$$G \cong \pi_1(X) \xleftarrow{f_*} \pi_1(M) \cong Deck(\tilde{M}).$$

Definition 4.3.1. Let $\tilde{M}$ be the universal cover of M and let ϕ be the isomorphism $\pi_1(M, x_0) \to Deck(\tilde{M})$ corresponding to some lift $\tilde{x_0}$ of x_0, then for each $[\gamma] \in \pi_1(M, x_0)$ we define $g_\gamma := \phi([\gamma])$.

Lemma 4.3.2. *Let M, X, G be as in the preceeding paragraph, x_0 a vertex. In particular M is a triangulated n-manifold. Let α be a cocycle such that α induces $\alpha_* : \pi_1(M, x_0) \to Isom^+(\mathbb{H}^n)$, a representation of $\pi_1(M)$ with $Im(\alpha_*) = G$. Then we can define a map $\tilde{F} : \tilde{M} \to \mathbb{H}^n$ such that for each (possibly semi-ideal) n-simplex Δ in $\tilde{M}$, $\tilde{F}(\Delta)$ is a geodesic n-simplex (it is the convex hull of its vertices) and $\tilde{F}$ induces a homotopy equivalence $F : M \to X$. Furthermore, the map which F induces on fundamental groups is exactly α_*.*

Proof. Let $\tilde{x_0} \in \tilde{M}$ be some lift of x_0 and take $\tilde{F}(\tilde{x_0}) = \tilde{y_0}$ some arbitrary point in $\mathbb{H}^n$. Now let $\tilde{x}$ be any other non-ideal vertex in the 0-skeleton of $\tilde{M}$ and $\tilde{\gamma}$ a simplicial path from $\tilde{x_0}$ to $\tilde{x}$. Let γ be the projection of $\tilde{\gamma}$ to M, then define $\tilde{F}(\tilde{x}) = \alpha(\gamma) \cdot \tilde{y_0}$. For now, assume that $\tilde{F}$ is defined and G-equivariant on the vertices (including the non-ideal ones), we shall define $\tilde{F}$ on the rest of the simplices dimension by dimension using induction. For each dimension k, starting from dimension 1, pick one representative of each G-orbit, and define a homeomorphism on the representative k-simplex by picking some homeomorphism which extends the homeomorphism on its boundary. We then define $\tilde{F}$ on all the translates of our simplex by $\tilde{F}(g(x)) = \alpha_*(g)\tilde{F}(x)$ for all x in our chosen representative. Once we've defined $\tilde{F}$ on all G-translates of all k-simplices, we increment dimension and do the same again.

First consider the closed case, $\tilde{F}$ is already defined on the 0-skeleton, so we just need to show that it is G-equivariant. By construction, $\tilde{F}$ is G-equivariant if and only if the map on the 0-skeleton is. Let $\tilde{x}$ be as above, then $\tilde{x} = g_\gamma(\tilde{x_0})$. Take $g_{\gamma'}$ some other element of Deck$(\tilde{M})$, then

$$\tilde{F}(g_{\gamma'}(\tilde{x})) = \tilde{F}(g_{\gamma'} g_\gamma(\tilde{x_0})) = \alpha(\gamma')\alpha(\gamma)\tilde{F}(x_0) = \alpha(\gamma')\tilde{F}(g_\gamma(\tilde{x_0})) = \alpha(\gamma')\tilde{F}(\tilde{x}).$$

Hence, $\tilde{F}$ is G-equivariant.

Note here that if we take some homeomorphism F' from M to X, then this lifts to a map $\tilde{F}'$ from $\tilde{M}$ to $\tilde{X}$ and by point-pushing G-equivariantly, we can assume that

this map agrees with $\tilde{F}$ on the vertices. Because of this we can define a straight line homotopy on each simplex of $\tilde{M}$ which commutes with the G-action and takes $\tilde{F}$ to $\tilde{F}'$. Thus F and F' are homotopic, and as F' is a homeomorphism, F must be a homotopy equivalence.

In the cusped case, we can use the same argument but we first have to define $\tilde{F}$ on the ideal vertices in $\tilde{M}$ and show it to be G-equivariant on them too. Note here that we are in fact defining $\tilde{F}$ by defining a map from $\tilde{M} \cup \{\text{ideal vertices}\}$ to $\mathbb{H}^n \cup \partial\mathbb{H}^n$ and then taking its restriction to $\tilde{M}$. However, we will abuse notation to refer to both maps as $\tilde{F}$ throughout. We number the cusps and take a spanning tree of the 1-skeleton of the link of the ideal vertex v_i corresponding to the ith cusp and add some edge which ends at v_i. We also choose a path δ in the 1-skeleton of M from our basepoint to this tree such that the union is still a tree, call this union Γ_i. Any edge in the link of the ideal vertex which is outside of the spanning tree gives a loop σ by adjoining the unique geodesic in Γ_i linking its endpoints. We can make σ be based at x_0 by adjoining our path δ at both ends. As σ is homotopic into the ideal vertex, $\alpha(\sigma)$ is thus a parabolic element of G, fixing some point p in $\partial\mathbb{H}^n$.

If $\tilde{v_i}$ is the lift of v_i which lies in the same lift of Γ_i as $\tilde{x_0}$, then we define $\tilde{F}(\tilde{v_i}) = p$. Each lift $\tilde{v_i}$ of v_i lies in a lift of Γ_i so if we take some path γ from $\tilde{x_0}$ to this new lift of Γ_i, then we define $\tilde{F}(\tilde{v_i}) = \alpha(\gamma)p$. From here, we only need to check two things, first that $\tilde{F}$ is well defined, and second that it is G-equivariant. For the well definedness, note that for a lift $\tilde{v_i}$, there are many lifts of Γ_i to which it might belong and thus distinct choices, γ, γ' for the path defined above. However all such pairs of lifts of Γ_i can be connected by a path which lies entirely in the link of the lift and maps down to some loop δ which is homotopic into v_i. Thus $\alpha(\delta)$ corresponds to a parabolic element fixing p and so as $\gamma' = \gamma\delta$, we get that

$$\alpha(\gamma')p = \alpha(\gamma)\alpha(\delta)p = \alpha(\gamma)p$$

so $\tilde{F}$ is well defined. For the G-equivariance, note that if g is some deck transformation, $\tilde{x_0}$ and $g(\tilde{x_0})$ are linked by a path γ with $\alpha(\gamma) = g$. Now, as $g(\tilde{v_i})$ lies in the same lift of Γ as $g(\tilde{x_0})$ we have that

$$\tilde{F}(g\tilde{v_i}) = \alpha(\gamma)p = \alpha(\gamma)\tilde{F}(\tilde{v_i}) = g\tilde{F}(\tilde{v_i})$$

and hence for any $h \in G$ and any ideal vertex $\tilde{v} = g\tilde{v_i} \in \tilde{M}$, the following holds

$$\tilde{F}(h\tilde{v}) = \tilde{F}(hg\tilde{v_i}) = hg\tilde{F}(\tilde{v_i}) = h\tilde{F}(\tilde{v}).$$

□

For most of this section we will restrict to the closed case and deal with the cusped case separately in the following section. In this proof, however, we dealt with the cusped case as well to avoid too much restatement of notation and ideas. We now bring in the fact that this homotopy equivalence is necessarily a surjection. This fact is surely well known, but the author couldn't find a reference. The proof below mimics many similar proofs in Chapter 3 of [Hat01].

Lemma 4.3.3. *A homotopy equivalence $f : M \to N$ between closed orientable n-manifolds is a surjection.*

Proof. Note that as f is a homotopy equivalence, it induces an isomorphism on nth homology groups, sending a fundamental class $[M]$ to a fundamental class $[N]$.
Now suppose f is not surjective, and let $y \in N$ be outside the image of f. Then consider the natural isomorphism $H_n(N) \cong H_n(N, N-y)$ then for any representative of $[M]$, $f[M]$ lies entirely in $N - y$ and so $f[M]$ is trivial in $H_n(N, N-y)$ and thus trivial in $H_n(N)$ contradicting that f is a homotopy equivalence. □

Remark 4.3.4. It is also possible (see Lemma 4.4.1) to prove that the homotopy equivalence defined in Lemma 4.3.2 is surjective in the non-compact case. That requires a more nuanced argument as homotopy equivalences between non-compact manifolds can easily not be surjections, for example consider a standard map from $\mathbb{R}^n$ to itself which maps everything to the point.

We have now seen how we can describe a homotopy equivalence when given a cocycle, and even hinted at how we can use the image of this homotopy equivalence to understand the geometry of the manifold. However, when we find cocycles as the solutions of a system of polynomials, these cocycles could be trivial or induce non faithful representations, so to use the results of Section 2.4.1 we need some connected component of the solution set to consist of cocycles which all induce faithful lattice representations (see Definition 2.2.25). We shall prove the existence of this connected component in the closed case using Calabi-Weil Rigidity (Theorem 2.2.27.

Lemma 4.3.5. *Let M be a closed hyperbolic n-manifold (where $n \geq 3$) with a triangulation by t n-simplices, such that each edge has distinct vertices. There is a system of polynomial inequalities as in Theorem 2.4.1 such that the following all hold:*

- *The solution set consists of cocycles in $C^1(M, SO(n,1))$.*
- *The solution set has a component consisting entirely of cocycles each of which induces a faithful lattice representation.*

- *The system consists of* $\leq (n+1)^5 t$ *polynomials, in* $\leq (n+1)^4 t$ *variables with degree* ≤ 2 *and coefficients all* ± 1.

Proof. For each edge e of the triangulation of M (of which there are less than $\binom{n+1}{2}t$), define two sets of $(n+1)^2$ variables corresponding to the entries of a matrix in $SO(n,1)$, these variables correspond to the images of both orientations of e under some cocycle α. Then, for each dimension 2 face of the triangulation (of which there are less than $\binom{n+1}{3}t$) we get a relation which can be expressed as $(n+1)^2$ polynomials (for the $(n+1)^2$ variables defining the matrix) which are satisfied if and only if the matrices corresponding to the first two edges of the face multiply to give the matrix corresponding to the third edge, these being the face relations of the cocycle. We also include polynomials which ensure that opposite orientations of an edge are sent to inverse elements, this consists of $(n+1)^2$ polynomials for each edge, and these polynomials have coefficients ± 1 and are quadratic.

Thus we have a system of $\leq (n+1)^5 t$ polynomials in $\leq (n+1)^4 t$ variables with coefficients always 1 or -1 and which are at most quadratic. Furthermore, solutions to this system correspond to maps from the oriented 1-simplices of M to $SO(n,1)$ which satisfy the face relations and orientation reversing relations, thus by definition solutions correspond to 1-cocycles.

As our manifold M admits a hyperbolic structure, we know that it admits a faithful lattice representation and thus there exists a cocycle which induces such a representation. Now any continuous deformation of this cocycle induces a continuous deformation of the induced representation. Thus Calabi-Weil Rigidity implies that any component containing a cocycle which induces a faithful lattice representation consists entirely of such cocycles. □

At this point we could immediately apply the results of Section 2.4.1 and get bounds on the size of the matrices in our cocycle and attempt to use that to find bounds on the geometry of M. Instead, we're going to continue to employ the trick of expanding the system of polynomials so that the geometric quantities we want to bound appear among the variables and thus the results of Section 2.4.1 immediately give us the desired bounds. Note that when we do this expansion, we don't materially change what cocycles can be in the solution set.

It is in the following lemma that it is necessary that our triangulation has no loops (edges where both vertices are the same). This is so that any cocycle can be altered (by multiplying wherever necessary the value on all edges round a vertex by a sufficiently small isometry) to take non-zero values on every edge while inducing

the same representation. Because of this, we don't miss any representations by only considering those cocycles which take non-identity values on all edges.

Lemma 4.3.6. *We can expand the system of polynomials above to include variables corresponding to the lengths of the edges in the image of the homotopy equivalence F defined in Lemma 4.3.2. We also add the requirement that each edge be non-zero length.*

The solutions of this new system are in one to one correspondence with the subset of the old solution set consisting of all cocycles which take non-zero values on all edges.

Proof. In the proof of Lemma 4.3.2 we didn't specify a model for the universal cover of either manifold, though both have universal cover $\mathbb{H}^n$. Here we shall take the hyperboloid model as our model and we shall take $(0, \dots, 0, 1)$ as our basepoint. Now $SO(n, 1)$ acts on the hyperboloid by matrix multiplication. Thus the map on universal covers $\tilde{F} : (\tilde{M}, \tilde{x_0}) \to (\mathbb{H}^n, (0, \dots, 0, 1))$ is defined by sending the point at the end of the path γ in $\tilde{M}$ to the point $A_\gamma(0, \dots, 0, 1)^T$ where $\alpha(\gamma) =: A_\gamma \in SO(n, 1)$. Thus the vertices of the triangulation of the universal cover are given by $A(0, \dots, 0, 1)^T$ where A is A_γ for some γ.

As there are at most $\binom{n+1}{2}t$ edges in M we can find a lift $\tilde{e}$ of each edge e of M such that both vertices $\underline{x}, \underline{y}$ of $\tilde{e}$ lie within $\binom{n+1}{2}t$ edges of the basepoint. Picking paths γ_1, γ_2 of length $\leq \binom{n+1}{2}t$ to each vertex of $\tilde{e}$, yields group elements $A_{\gamma_1}, A_{\gamma_2}$ which can be expressed as $(n+1)^2$ polynomials of degree at most $\binom{n+1}{2}t$ in the original variables (each A_γ is a product of A_{e_i} with e_i edges of the triangulation).

We build this up edge by edge, at each point we define at most one new vertex for each new edge as a set of $n+1$ new variables. In total we need $\leq (n+1)\binom{n+1}{2}t$ new variables and polynomials defining these vertices. The length of the edge between two such vertices $\underline{x}$ and $\underline{y}$ where $\underline{x} = (x_0, \dots, x_n)$ and $\underline{y} = (y_0, \dots, y_n)$ can then be calculated as

$$l(e) = d(\underline{x}, \underline{y}) = arcosh(-x_0y_0 - \dots - x_{n-1}y_{n-1} + x_ny_n).$$

This is not a polynomial and so we instead define a variable for each edge as

$$C = cosh(l(e)) - 1 = cosh(d(\underline{x}, \underline{y})) - 1$$

which is certainly a polynomial, the number of variables C is the number of edges which is $\leq \binom{n+1}{2}t$. We then require that $C > 0$, and thus that $l(e) > 0$.

In total we have added less than $(n+2)\binom{n+1}{2}t$ new variables and $(n+2)\binom{n+1}{2}t$ polynomials with degree bounded by $\binom{n+1}{2}t$ and coefficient complexity still at 1. $\square$

Combining the above two proofs gives a polynomial system which has size described (in the language of Section 2.4.1) by

- $N \leq (n+1)^4 t + (n+2)\binom{n+1}{2} t \leq (n+2)^4 t$
- $\kappa \leq (n+1)^5 t + (n+2)\binom{n+1}{2} t \leq (n+2)^5 t$
- $d \leq \binom{n+1}{2} t \leq (n+1)^2 t$
- $M \leq 2$

Thus, as the results of Section 2.4.1 bound the size of non-zero entries of solutions the following theorem is a simple corollary of Corollary 2.4.3 and Lemmas 4.3.2, 4.3.5, and 4.3.6.

Theorem 4.3.7. *Let M be a closed hyperbolic n-manifold triangulated such that each edge has distinct vertices, there exists a homotopy equivalence F as described in Lemma 4.3.2 such that the image of each edge has length $l(e)$ bounded in the following way:*

$$cosh(l(e)) - 1 \leq 2^{2((n+2)^7 t^2)^{O((n+2)^4 t)}}$$

This can be simplified, due to big O notation to the following bound:

$$cosh(l(e)) \leq 2^{(nt)^{O(n^4 t)}}$$

Which can be used to show that

$$0 < l(e) \leq (nt)^{O(n^4 t)}.$$

Proof. Lemmas 4.3.5 and 4.3.6 provide us with a system of polynomials, the solution set of which has a connected component, all elements of which induce faithful lattice representations. Mostow Rigidity tells us each of these lattices defines a hyperbolic manifold X as in the setup of Lemma 4.3.2. Thus taking a bounded solution given by Corollary 2.4.3 means that all the non-zero entries of this solution are bounded as in the statement of the Theorem. In particular, the entries $C = cosh(l(e) - 1)$ corresponding to edge lengths are required to be greater than zero and so the upper bounds on $cosh(l(e)$ and hence on $l(e)$ follow. Furthermore as $cosh(l(e))$ is then greater than 1, and $l(e)$ is non-negative, it must be positive. □

Theorem 4.3.8. *Given a closed hyperbolic n-manifold ($n \geq 3$) triangulated by t n-simplices such that each edge has distinct vertices, let R be the length of a systole in M, then*

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}.$$

Proof. Let X be the hyperbolic manifold defined in Lemma 4.3.2, then by Theorem 4.3.7 we know that X admits a cover by t immersed hyperbolic n-simplices with edge lengths bounded above by some bound $B(t)$ with

$$B(t) \leq (nt)^{O(n^4 t)}.$$

Pick some $\epsilon > \frac{1}{(6\pi)^n}$ so that $M_{(0,\epsilon)}$ satisfies the conclusion of Theorem 2.2.17. This is possible as the bound in Theorem 2.2.21 is not strict.

If the thin part $M_{(0,\epsilon)}$ is empty, then $R \geq \frac{2}{(6\pi)^n}$ and so the conclusion holds. Suppose the thin part is non empty, then Theorem 4.2.5 applies. Consider a systole γ. This curve has length R and lies in the centre of one of the tubes of the thin part. We know from the above that the diameter of this tube is at least

$$\frac{1}{n} log(1/R) + log(\frac{\epsilon_n}{4}).$$

However, we know that M is covered by t n-simplices of bounded edge length, and thus of bounded diameter, and so the diameter of M is at most

$$tB(t) \leq t(nt)^{O(n^4 t)}$$

and hence

$$\frac{1}{n} log(1/R) + log(\frac{\epsilon_n}{4}) \leq t(nt)^{O(n^4 t)}.$$

Simplifying big O notation further and applying Theorem 2.2.21

$$log\left(\frac{1}{R}\right) \leq (nt)^{O(n^4 t)} + nlog(4(6\pi)^n).$$

Taking exponentials on both sides gives

$$\frac{1}{R} \leq e^{(nt)^{O(n^4 t)}} (24\pi)^{n^2}$$

which simplifies, after rearrangement, to

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}.$$

□

The following is then a direct corollary of Theorem 4.3.8 and Remark 4.1.2.

Theorem 4.3.9. *Given a closed hyperbolic n-manifold ($n \geq 3$) triangulated by t n-simplices such that each edge has distinct vertices, let R be the length of a systole in M, then*

$$R \geq \frac{1}{2^{(nt)^{O(n^{n+5} t)}}}.$$

4.4 The Cusped Case

With some small but technical adaptations, the arguments from the closed case hold in the cusped (finite volume) case where M is given by a finite semi-ideal triangulation and M admits a finite volume hyperbolic structure.

Lemma 4.4.1. *The homotopy equivalence F defined in Lemma 4.3.2 is surjective also in the finite-volume case.*

Proof. In the proof of Lemma 4.3.2 we showed that the homotopy equivalence $F : M \to X$ is homotopic to a homeomorphism F', and in fact the homotopy fixes ideal vertices, so the entire homotopy extends to the space we get when we add the ideal vertices back in (replace all ideal simplices with non-ideal ones). Call these pseudomanifolds M', X'. Then the nth homology of each of M' is $\mathbb{Z}$ with fundamental class $[M']$ given by the triangulation formed by replacing the ideal vertices with concrete ones. First note that as the homotopy to F' fixes ideal vertices and F' is a homeomorphism, F has to surject onto the ideal vertices. Thus, we only need to show it is surjective on non-ideal vertices x in X', whose neighbourhoods are balls which thus have $H_n(X', X' - x) \cong \mathbb{Z}$. Now note that $[M']$ generates both $H_n(M')$ and $H_n(M', M' - x)$ for x a non-ideal point, and so, as $F[M']$ is homotopic and thus homologous to $F'[M']$, we know the same is true for $F[M']$ in $H_n(X')$ and $H_n(X', X' - x)$ for x non-ideal point. Thus, if F misses a non-ideal point x, then $F[M']$ lies in $X' - x$ and so is trivial in $H_n(X', X' - x)$, which is a contradiction as F and F' are homotopic but $F'[M']$ generates $H_n(X', X' - x) \cong \mathbb{Z}$ so is non-trivial. □

In order to use Calabi-Weil in the finite volume case we need control over parabolic isometries which can be tricky, especially in higher dimensions where parabolic subgroups are not necessarily translation groups. We solve this problem in higher dimensions by using a stronger statement of rigidity, due to Garland and Raghunathan instead, which we repeat here.

Theorem 2.2.29. *Let G be a lattice in $Isom(\mathbb{H}^n)$ for $n \geq 4$. The identity representation $\rho_0 : G \to Isom(\mathbb{H}^n) =: \Lambda$ lies in a connected component of $Hom(G; \Lambda)$ made up entirely of conjugates of ρ_0.*

Lemma 4.4.2. *Let M be a finite volume hyperbolic n-manifold with a semi-ideal triangulation such that all non-ideal edges have distinct vertices, $n \geq 4$. There is a system of polynomial inequalities as in Theorem 2.4.1 such that the following all hold:*

- *The solution set consists of cocycles in* $C^1(M, SO(n,1))$.
- *The solution set has a component consisting entirely of cocycles which induce faithful lattice representations.*
- *The system consists of* $< (n+1)^5 t$ *polynomials in* $< (n+1)^4 t$ *variables, with degree bounded by* $(n+1)^2 t$ *and all coefficients* ± 1.

Furthermore, this systems contains variables for the lengths of all non-ideal edges.

Proof. The statements and proofs of Lemmas 4.3.5 and 4.3.6 hold in the cusped case (for $n \geq 4$) if we replace mention of edges and simplices by non-ideal edges and non-ideal simplices. The only other difference being that instead of applying Calabi-Weil Rigidity, we apply Theorem 2.2.29. □

In the 3-dimensional case, Garland and Raghunathan's result no longer holds, and Calabi-Weil rigidity requires that we have control over where parabolics are sent. For this we consider $Isom^+(\mathbb{H}^3) \cong PSL(2,\mathbb{C})$ instead, while sticking with the hyperboloid model. Note that $PSL(2,\mathbb{C})$ is not usually the group of isometries of this model, but it can be made to act on the hyperboloid model with an action that we describe in the proof of Lemma 4.4.5. The reason behind this is that $PSL(2,\mathbb{C})$ offers a remarkably easy method of verifying parabolicity, namely, an element g in $PSL(2,\mathbb{C})$ is parabolic if and only if for $A \in SL(2,\mathbb{C})$ a matrix representative of g, trace(A) ± 2, or equivalently, its trace squares to 4.

There is one small hitch with looking at representations into $PSL(2,\mathbb{C})$. To apply the results of section 2.4.1 we need to be looking at elements of a variety. Fortunately, the following tells us we can equivalently look at representations into $SL(2,\mathbb{C})$.

Theorem 4.4.3 (Corollary 2.3 [Cul86])**.** *If a discrete subgroup* G *of* $PSL(2,\mathbb{C})$ *has no 2-torsion then it lifts to* $SL(2,\mathbb{C})$*. That is there is a homomorphism*

$$G \to SL(2,\mathbb{C})$$

such that the composition with the natural projection

$$SL(2,\mathbb{C}) \to PSL(2,\mathbb{C})$$

is the identity on G.

Combining this with the fact that connected components of $Hom(G, SL(2,\mathbb{C}))$ map down to connected components of $Hom(G, PSL(2,\mathbb{C}))$ tells us that if we apply the results of Section 2.4.1 to a system such as those created so far but with images in $SL(2,\mathbb{C})$, then we can find a connected component of the solution set which contains a solution inducing a faithful lattice representation, and thus by Theorem 4.4 the whole component consists of such representations.

Because of the focus on parabolic elements, which correspond to loops which are freely homotopic to ideal vertices, it will be useful to be able to talk about the parabolics corresponding to each individual ideal vertex with the following definition.

Definition 4.4.4 (Cusp Fundamental Group)**.** We define the fundamental group of a given cusp in M to be fundamental group of the star of a chosen ideal vertex. Note that choosing a path from our chosen basepoint for M to our chosen basepoint for the star of an ideal vertex defines an embedding of the cusp fundamental group into the fundamental group of M. All such choices define conjugate embeddings.

We can now develop a system of polynomials which suits the final case of cusped hyperbolic 3-manifolds.

Lemma 4.4.5. *Let M be a finite volume hyperbolic 3-manifold triangulated by t tetrahedra with fundamental group G, such that all non-ideal edges have distinct vertices. There is a system of polynomial inequalities as in Theorem 2.4.1 such that the following all hold:*

- *The solution set consists of cocycles in $C^1(M, SL(2,\mathbb{C}))$ which induce representations in $Hom_{par}(G, SL(2,\mathbb{C}))$.*
- *The solution set has a component consisting entirely of cocycles each of which induces a faithful lattice representation.*
- *The system has each of its measures of complexity N, κ, d, M (as in Section 2.4.1) bounded above by a constant multiple of t.*

Furthermore, this systems contains variables for the lengths of all non-ideal edges in the image of the homotopy equivalence defined in Lemma 4.3.2.

Proof. Outside of the handling of parabolics, the proofs here are very similar to the $SO(n,1)$ case, so we will only outline where the proofs differ from those given for Lemmas 4.3.5 and 4.3.6.

To get a system of polynomials which has cocycles in $C^1(M, SL(2,\mathbb{C}))$ as its solutions (not yet including variables for edge lengths and not necessarily inducing representations in $Hom_{par}(G, SL(2,\mathbb{C}))$, the same proof as for $SO(n,1)$ holds. We get some constant multiple of t as a bound for the number of polynomials and variables, and the polynomials have degree at most quadratic with coefficients ± 1.

For edge lengths, we can again define variables for the end points of lifts of edges as well as their lengths in the hyperboloid model. The numbers of polynomials and variables added is bounded by a multiple of t and the degree and coefficient complexity of the polynomials are both bounded by a constant. This uses the following action of $SL(2,\mathbb{C})$ on the hyperboloid model.

One can represent the point (x, y, z, t) in the hyperboloid model by the matrix

$$X = \begin{pmatrix} t+z & x-iy \\ x+iy & t-z \end{pmatrix}$$

and then $A \in SL(2,\mathbb{C})$ acts as follows (A^* denotes conjugate transpose)

$$X \mapsto AXA^*.$$

Note that taking the determinant gives the standard quadratic form used to define the hyperboloid and the action preserves determinant.

Finally the big difference in the cusped case is that we need to check that parabolic subgroups of $\pi_1(M)$ are indeed sent to parabolics in $PSL(2,\mathbb{C})$ for each representation induced by a cocycle solution to our system. In three dimensions it is enough to check that generators of the cusp fundamental group are sent to parabolics which fix the same point at infinity, as parabolics in $Isom^+(\mathbb{H}^3)$ restrict to translations on horospheres about their fixed point, and so products of such parabolics are also parabolics, as non-trivial products of translations are non-trivial translations. Note this is not true in higher dimensions where parabolics don't have to restrict to translations and so two parabolics with the same fixed point at infinity can have a non-trivial elliptic as their product.

To do this we take generating sets of each cusp fundamental group as described in the proof of Lemma 4.3.2. That is we pick a spanning tree of the link of each ideal vertex, then the remaining edges define loops in the link, each of which can be homotoped arbitrarily close to the ideal vertex and thus maps to a parabolic, and which together generate the fundamental group of the cusp. As there are less than $6t$ total edges, we see that there are less than $6t$ such generating curves, each of which is a product of less than $6t$ edges. So for each curve we define polynomials which

require that the square of the trace of the matrix associated to each such curve be 4. As this matrix is a product, the degree of these polynomials could be as much as $6t$. We also introduce new variables which correspond to the shared fixed points of the generating set of a cusp, and polynomials requiring that they be fixed.

At all times however, all measures of complexity of this system are bounded by a multiple of t.

Thus we have a system of polynomials whose solutions correspond to cocycles in $C^1(M, SL(2,\mathbb{C}))$ and also to the lengths of edges in the image of the homotopy equivalence. We also know that parabolic generators (and hence the entirety of each parabolic subgroup) are mapped to parabolic elements in the induced representation into $PSL(2,\mathbb{C})$. We know that there is a solution to this system inducing a faithful lattice representation and the entire connected component of the solution set containing this solution induces representations into $PSL(2,\mathbb{C})$ which all lie in the same component of $Hom_{par}(\pi_1(M), PSL(2,\mathbb{C}))$ and hence, by Calabi-Weil Rigidity these representations are all themselves faithful lattices. □

As in the closed case, the following lemma follows by applying Theorem 2.4.1 to the systems of polynomials defined in Lemmas 4.4.2 and 4.4.5 and simplifying big O notation.

Lemma 4.4.6. *Let M be a finite volume hyperbolic manifold equipped with a semi-ideal triangulation by t semi-ideal n-simplices such that all non-ideal edges have distinct vertices. Then there exists a homotopy equivalence F as described in Lemma 4.3.2 such that the image of each non-ideal edge has length $l(e)$ bounded in the following way:*

$$0 < l(e) \leq (nt)^{O(n^4 t)}.$$

Theorem 4.4.7. *Given a finite volume hyperbolic n-manifold ($n \geq 3$) triangulated by t semi-ideal n-simplices such that all non-ideal edges have distinct vertices , let R be the length of a systole in M, then*

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}.$$

Proof. Throughout, we let $B(t)$ be a bound on the edge length of the image of non-ideal edges under the homotopy equivalence F. Let $\bar{X}$ be the image in X under F of the non-ideal simplices in M. For each simplex in M, the diameter of the image of its non-ideal part is bounded by the maximum length of its non-ideal edges, hence by

$B(t)$, and thus the diameter of $\bar{X}$ is bounded by $tB(t)$. Note also that, as the links of ideal vertices consist entirely of non-ideal simplices, the images of links of ideal vertices lie in $\bar{X}$.

We aim to show that points in X either lie in a bounded neighbourhood of $\bar{X}$, or they lie in the subset of the thin part made up of cusp neighbourhoods.

Note that the complement of $\bar{X}$ is covered by the images of the stars of the ideal vertices of M. Consider one such vertex v_i, and its star $st(v_i)$, let S_i denote the image of $st(v_i)$ in X and L_i denote the image of $lk(v_i)$. Then S_i is a union of geodesic semi-ideal n-simplices. Take a collection of lifts (one per simplex) of these semi-ideal n-simplices in $\mathbb{H}^n$ (the upper half space model) such that the collection is connected, and shares the same point on the boundary. We can perform an isometry of $\mathbb{H}^n$ taking this shared point to the point at infinity. We shall call this collection T_i, and note that we still have a canonical map from T_i into X.

Consider the map ϕ_d which takes each point p in the upper half space model and translates it a hyperbolic distance d along the unique (vertical) geodesic through p and the point at infinity. Then this map scales distances in $\mathbb{H}^n$ by

$$\frac{1}{e^d}.$$

Let ψ_d be the map on S_i which is induced by $\phi_d \mid_{T_i}$, then any edge L_i has its length scaled by the scale factor above. Now as $L_i \subseteq \bar{X}$, its diameter is bounded by $tB(t)$. Thus setting

$$d_0 = log\frac{tB(t)}{\epsilon_n},$$

we know that for all $d \geq d_0$, the image of L_i under ψ_d has diameter bounded above in the following way

$$Diam(\psi_d(L_i)) < \frac{1}{e^{log\frac{t(B(t))}{\epsilon}}} t(B(t)) = \epsilon_n.$$

Thus as the union of all $\psi_d(L_i)$ covers S_i, every point $p \in S_i$ lies in some $\psi_d(L_i)$. If $d \leq d_0$, then p lies in a d_0 neighbourhood of $\bar{X}$. If $d > d_0$, then there is some curve of length $< 2\epsilon_n$ based at p which lies entirely in $\psi_d(L_i)$, and hence p , and in fact all of $\psi_d(S_i)$ lies in the thin part of X. As $\psi_d(S_i)$ is a neighbourhood of an ideal vertex, it must lie in one of the cusp neighbourhoods of the thin part.

The thin part of a manifold decomposes as a disjoint union of cusp neighbourhoods and tubes which are disjoint. Furthermore, there are no geodesics in the cusp neighbourhoods as translation towards the end of the cusp decreases the length of a

curve. Thus either all systoles lie in the thick part and have length bounded below by $2\epsilon_n$ or they lie in the thin part and are disjoint from all the cusp neighbourhoods.

Each systole must then lie in the d_0-neighbourhood of $\bar{X}$. This means that the entire 1-skeleton of $\bar{X}$ lies within distance

$$diam(\bar{X}) + d_0 < tB(t) + log\frac{tB(t)}{\epsilon_n}$$

of each systole. By using Theorem 4.4.6 as a bound for $B(t)$, using Theorem 2.2.21 as a bound for ϵ_n and simplifying big O notation, we get that $\bar{X}$ lies within distance

$$(nt)^{O(n^4t)}.$$

of the systole. This is therefore also an upper bound for the distance from each systole to the boundary of the tube containing it. Indeed if $\bar{X}$ lies in this tube, then a generating set for $\pi_1(X)$ lies in this subspace which has fundamental group $\mathbb{Z}$, and so $\pi_1(X)$ is a subgroup of $\mathbb{Z}$, a contradiction as $\pi_1(X)$ is the fundamental group of a finite volume hyperbolic manifold and so cannot be $\mathbb{Z}$.

Thus by Theorem 4.2.5

$$\frac{1}{n}log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4) < (nt)^{O(n^4t)}$$

and hence we can rearrange this (simplifying big O notation) to

$$log\left(\frac{1}{R}\right) \leq (nt)^{O(n^4t)} + nlog(n(6\pi)^n)$$

and thus simplifying again

$$R \geq \frac{1}{2^{(nt)^{O(n^4t)}}}.$$

□

Theorem 4.4.8. *Given a finite volume hyperbolic n-manifold ($n \geq 3$) triangulated by t semi-ideal n-simplices, let R be the length of a systole in M, then*

$$R \geq \frac{1}{2^{(nt)^{O(n^{n+5}t)}}}.$$

Proof. It again follows from Remark 4.1.2, that the requirement that all non-ideal edges have distinct vertices is satisfied when the triangulation is simplicial, and thus performing a barycentric subdivision allows us to remove that requirement from Theorem 4.4.8, at the cost of increasing the size of the triangulation from t to $O(n^n t)$. □

Chapter 5

The Main Algorithm

5.1 Finding a Geometric Quotient

We shall recall the results of Chapter 4 in a form which will be useful for this section. Here we shall start our description of the algorithm which shall solve the homeomorphism problem. This section demonstrates the theoretical existence of a geometric triangulation which we find as a quotient of a subdivision of a given topological triangulation. We do this by non-algorithmically constructing such a triangulation. It is the nature of the construction that allows us to bound the complexity (in terms of both geometry and combinatorics) of the triangulation. Having this bound will later allow us to find this triangulation again using an algorithm which we will describe in the later sections of this chapter.

Throughout this section, as elsewhere in the thesis, we prove results in as general a case as doesn't make substantial changes to the proofs. In this section our proofs extend naturally to n-dimensional closed hyperbolic manifolds.

Theorem 5.1.1 (Lemmas 4.3.2, 4.3.3, 4.4.1 and Theorem 4.3.7)**.** *Let M be a closed hyperbolic n-manifold ($n \geq 3$) triangulated by t n-simplices such that all edges have distinct vertices. Let X be a hyperbolic manifold homeomorphic to M, then there is a homotopy equivalence*

$$F : M \to X$$

such that the image of each n-simplex in M is a geodesically immersed hyperbolic n-simplex in X. Furthermore F is surjective and the length $l(e)$ of an edge in the image of F is bounded as follows.

$$0 < l(e) \leq (nt)^{O(n^4 t)}.$$

Recall also the following two theorems which give control over the length of systoles and hence over the injectivity radius in closed hyperbolic n-manifolds.

Theorem 4.1.3. *Given a closed hyperbolic n-manifold ($n \geq 3$) triangulated by t n-simplices such that all non-ideal edges have distinct vertices. The length R of a systole (shortest closed geodesic) of M is bounded below by a function of t and n, in particular*

$$R \geq \frac{1}{2^{(nt)^{O(n^4 t)}}}.$$

Theorem 4.2.5. *Suppose that M is a finite volume hyperbolic n-manifold ($n \geq 3$) with systole(s) of length $R \leq 2\epsilon_n$, where ϵ_n is the Margulis constant in dimension n. Then the distance from a systole to the boundary of the Margulis tube containing it is bounded below by*

$$\frac{1}{n} log\left(\frac{1}{R}\right) + log(\epsilon_n) - log(4).$$

The first theorem tells us that our original triangulation of M admits a degree one map $M \to M$ (in fact a homotopy equivalence) such that its image consists of geodesic n-simplices which cover the manifold. We now attempt to use what the second and third theorems tell us about the injectivity radius of the manifold to form a subdivision of each of these geodesic simplices so that the subdivided cover has the property that any two simplices can only intersect in one convex component.

For closed hyperbolic manifolds, this is as simple as subdividing so that the n-simplices have diameter less than the injectivity radius of the manifold. For this we use the above bound on injectivity radius. However, in three dimensions, a better bound exists by a combination of the following theorem of White (see also the stronger result of Agol and Liu [AL10]) and Theorem 4.2.5.

Theorem 5.1.2 (Theorem 5.9, White [Whi00a])**.** *There is an explicit constant $K > 0$ such that if M is a closed, connected, hyperbolic 3-manifold, and*

$$P = \langle x_1, \ldots, x_n \mid r_1, \ldots, r_m \rangle$$

is a presentation of its fundamental group, then $diam(M) < K(l(P))$, where

$$l(P) = \sum_{i=1}^{m} l(r_i)$$

and $l(r_i)$ is the word length of a given relator.

Now, given a simplicial complex, we can pick some spanning tree of its 1-skeleton, then a generator set is given by the remaining edges of the 1-skeleton, and a relator set, each of which is a word of length at most 3, is given by the 2-simplices. This leads to the following corollary.

Corollary 5.1.3. *Let M be a closed, connected, hyperbolic 3-manifold triangulated by t tetrahedra. The number of 2-simplices is exactly $2t$ and so $l(P) \leq 6t$ where P is the presentation described above. Hence, $diam(M) < 6Kt < O(t)$.*

Corollary 5.1.4. *Let M be a closed, connected, hyperbolic 3-manifold triangulated by t tetrahedra. The injectivity radius of M is bounded below by*

$$e^{-O(t)}$$

Proof. Either the injectivity radius of M is bigger than ϵ_3 in which case the result follows, or Theorem 4.2.5 applies and by Corollary 5.1.3

$$log(\frac{1}{R}) < O(t)$$

from which the statement follows. □

Remark 5.1.5. Let $F : \Delta \to X$ be a geodesic immersion of a hyperbolic tetrahedron Δ and $\tilde{F} : \Delta \to \mathbb{H}^n$ be a lift of F. We note that for such a geodesic immersion to self intersect, the diameter of the lift of the tetrahedron must be greater than the length of a systole, or twice the injectivity radius of the hyperbolic manifold X.

Our first step is to use this remark to find a subdivision of each simplex in M (subdivisions may not agree on the boundary and so this might not be a subdivision of the whole triangulation) such that each of the new simplices embeds into X. The next two lemmas describe how to perform this subdivision.

Lemma 5.1.6. *Let Δ be a hyperbolic k-simplex with $k \leq n$ and edge lengths bounded above by $L > 0$. Suppose we have a subdivision of one of its faces F into K hyperbolic polyhedra, each of which has diameter bounded above by $\frac{k-1}{n}D$ for some $D > 0$ and has at most $2(k-1)$ faces. Then we can subdivide the k-simplex itself into $K\frac{Ln}{D}$ polyhedra, each of diameter bounded above by $\frac{k}{n}D$ with at most $2k$ faces.*

Proof. Take an embedding of Δ in the Klein model of hyperbolic k-space such that the vertex v opposite F is at the origin. Next, cone the polyhedral decomposition of F to the vertex v.

Now, we insert parallel scaled copies of F throughout the k-simplex such that each lies in the $\frac{D}{n}$-neighbourhood of the last. We do this by ensuring that, for every vertex of F, and each corresponding edge coning to v, the length of the edge segment e_i between two copies F_1 and F_2 along this (hyperbolic geodesic) edge is less than $\frac{D}{n}$. This ensures that every point on F_1 is with a $\frac{D}{n}$-neighbourhood of F_2 as for any point the radial arc from F_1 to F_2 is both shorter in Euclidean length and has its vertices closer to the origin than the maximal length e_i, thus its hyperbolic length is also less than that of e_i, due to how the metric in the Klein model scales down as we go further from the origin.

This gives us a polyhedral subdivision into $K\frac{Ln}{D}$ polyhedra each within at most $2k$ faces (each of which has (2k−2) faces and so on). To see that the diameter of each polyhedron is bounded above by $\frac{kD}{n}$, note that any two points in the polyhedron can be connected by a radial arc, followed by a path in a scaled copy of one of the polyhedra in F, thus this path has total length bounded above by the lengths of the two subpaths

$$\frac{D}{n} + \frac{(k-1)D}{n} = \frac{kD}{n}$$

□

Lemma 5.1.7. *The number of n-simplices required to triangulate a convex polyhedron formed by the intersection of J halfspaces bounded by hyperplanes in $\mathbb{R}^n$ is bounded above by the following*

$$\left(2^J\right)^{n+1}$$

In particular if $J = 2(n+1)$, then the number of n-simplices required is bounded by

$$\left(2^{2n+2}\right)^{n+1} \leq 2^{7n^2}$$

Proof. We can triangulate the polyhedron by barycentrically subdividing it. Note that each n-simplex Δ of the barycentric subdivision is uniquely determined by a sequence $\sigma_0 \subseteq \sigma_1 \subseteq \ldots \sigma_n$ where each σ_i is some i-dimensional face of the polyhedron where Δ is the unique n-simplex which intersects the interior of each face of the sequence. The number of faces of a given dimension is certainly less than 2^J as this is the size of the power set of the hyperplanes and every face is an intersection of hyperplanes. Thus the total number of sequences of faces is bounded above by by

$$\left(2^J\right)^{n+1}$$

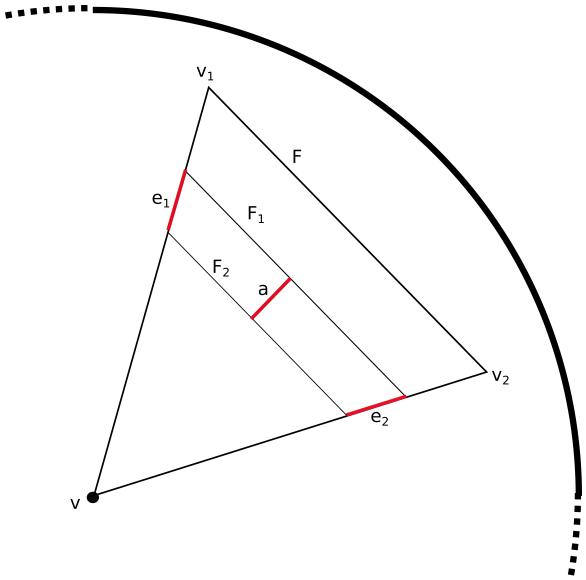


Figure 5.1: Let a be a radial segment between two translates F_1, F_2 of F and let e_i denote the segment of the edge v-v_i lying between F_1, F_2 for v_i the vertices of F. If v_i is the further vertex from v, then the angle made by e_i with the translates F_1, F_2 is further from perpendicular than the angle made by a, thus it has has both greater Euclidean length and is further from the origin than a. Hence the hyperbolic length of e_i is greater than the hyperbolic length of a, so if both segments have length bounded above by D/n, that is also a bound for the length of a.

□

We now combine the above two lemmas to see how it relates to our original triangulation.

Lemma 5.1.8. *Given $(M, \mathcal{T})$ a triangulated n-manifold homeomorphic to a hyperbolic n-manifold X, $F : M \to X$ the homotopy equivalence defined in Theorem 5.1.1 and some desired bound $D < \frac{1}{n}$ on the diameter of n-simplices, we can find a subdivision of each simplex of $(M, \mathcal{T})$ such that F restricted to each simplex has image a geodesically immersed hyperbolic n-simplex of diameter less than D.*

Furthermore, the union of these subdivisions consists of less than the following number of n-simplices:

$$(nt)^{O(n^4 t)} \frac{1}{D^n}.$$

Proof. For each hyperbolic n-simplex with edge lengths bounded by $L > 0$, we first inductively perform the process described in Lemma 5.1.6 starting with the base case of a hyperbolic 1-simplex. Clearly we can get a subdivision of a 1-simplex of length at most L into "polyhedra" of diameter less than $\frac{D}{n}$ which consists of $\frac{Ln}{D}$ polyhedra. Thus

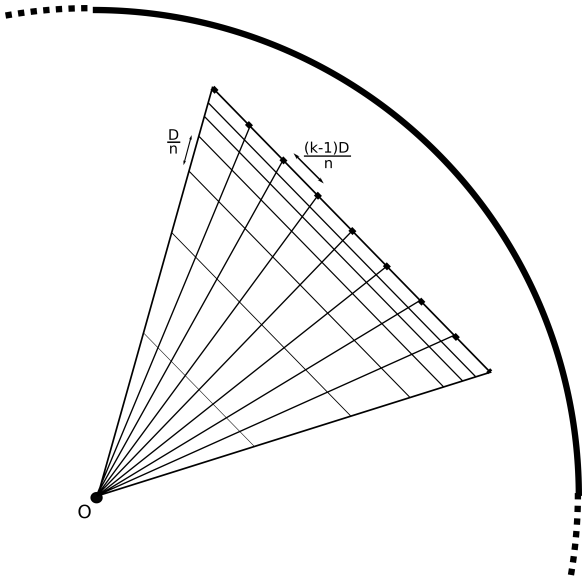


Figure 5.2: Note that the radial segments lie a similar euclidean distance apart, while the translates of the exterior face bunch up as we go further from the origin. This is because in the Klein model, as distance from the origin increases smaller and smaller euclidean segments can measure the same hyperbolic distance.

inductive application of Lemma 5.1.6 gives that the n-simplex admits a polyhedral decomposition into $\left(\frac{Ln}{D}\right)^n$ polyhedra each bounded by at most $2n$ hyperplanes. By Lemma 5.1.7 this means each polyhedron can be triangulated by at most 2^{2n^2+2n} hyperbolic n-simplices. Hence the total number of hyperbolic n-simplices required to triangulate a hyperbolic n-simplex of edge length less than $L > 0$ into n-simplices of diameter less than D is bounded above by

$$2^{2n^2+2n}\left(\frac{Ln}{D}\right)^n.$$

Thus, for F as in Theorem 5.1.1 we can subdivide a lift of $F(\Delta)$ as in Lemma 5.1.6 for each Δ in the triangulation of M and the preimage of this subdivision is a subdivision of Δ with the required property. Thus substituting in our known bounds on the edge lengths of the lift of $F(\Delta)$, summing over all t n-simplices and simplifying big O notation gives the desired bound:

$$t2^{2n^2+2n}\left(\frac{(nt)^{O(n^4t)}n}{D}\right)^n \leq (nt)^{O(n^4t)}\frac{1}{D^n}.$$

□

This theorem then combines with our bound on injectivity radius to give the following corollary.

Corollary 5.1.9. *Given $(M, \mathcal{T})$ a triangulated n-manifold homeomorphic to a hyperbolic n-manifold X, $F : M \to X$ the homotopy equivalence defined in Theorem 5.1.1, we can find a subdivision of each simplex of $(M, \mathcal{T})$ such that F restricted to each simplex has image a geodesically embedded hyperbolic tetrahedron of diameter less than half the injectivity radius of the manifold or even $1/c$ times the injectivity radius for some integer $c \geq 1$.*

Furthermore, in the n-dimensional case, the union of these subdivisions consists of less than the following number of n-simplices:

$$c^n 2^{(nt)^{O(n^4 t)}}.$$

In the 3-dimensional case, the bound improves to

$$c^n t^{O(t)}$$

Proof. This is a simple corollary of Theorem 4.1.3 and Lemma 5.1.8 if we set $D = R/2c$ and simplify big O notation. In the 3-dimensional case, we instead apply Corollary 5.1.4. That the n-simplices are embedded follows from Remark 5.1.5.

□

So if we set $c = 2$ in the above we now have a collection of embedded hyperbolic n-simplices which cover our hyperbolic manifold X and can intersect pairwise in at most one component because their diameter is bounded by half the injectivity radius of the manifold. There are at most T such n-simplices with $T \leq 2^{nt^{O(n^4 t)}}$ if $n \geq 4$ or $T \leq t^{O(t)}$ if $n = 3$.

Our next task is to triangulate the intersections of these tetrahedra to give a true geometric triangulation of X, the preimage of this triangulation is then a triangulation of M which has the geometric triangulation of X as its simplicial quotient.

Lemma 5.1.10. *Suppose we are given T geodesic n-simplices cover a closed hyperbolic n-manifold X such that the n-simplices have diameter bounded by half the injectivity radius of the manifold. Then there is a geometric triangulation of X such that the interior of each geometric n-simplex of the triangulation is either contained within or disjoint from the interior of each simplex of the cover. This triangulation has size bounded by*

$$T 2^{Tan^2}$$

for some constant a.

Proof. We build up this triangulation starting with a single geometric n-simplex and by adding one of our T geometric n-simplices (which we shall call the "old" n-simplices) at a time and subdividing so that the union so far is geometrically triangulated. Consider the first two n-simplices, they intersect in one component due to the requirement on injectivity radius. Now consider this pair in the universal cover of hyperbolic space and extend each of their faces to the hyperbolic hyperplanes which contain them, then we have $2n+2$ such hyperplanes, and choosing half spaces and taking their intersection yields at most 2^{2n+2} convex polyhedra, a subset of which tiles the union of two n-simplices. In the Klein ball model for the universal cover these polyhedra are indeed Euclidean and so we can geometrically triangulate each polyhedron by at most

$$2^{7n^2}$$

n-simplices by Lemma 5.1.7 and thus we get a geometric triangulation of the union of the two n-simplices by

$$2^{2n+2}2^{7n^2} < 2^{a(n^2)}$$

for some constant a. If we then add a third tetrahedron, we need to intersect it with all these $2^{a(n^2)}$ tetrahedra and each intersection yields at most $2^{a(n^2)}$ tetrahedra itself, making $\leq \left(2^{a(n^2)}\right)^2$ n-simplices so far. In the end we are left with a geometric triangulation of X by at most

$$\left(2^{a(n^2)}\right)^T = 2^{Tan^2}$$

tetrahedra.

Because new tetrahedra either lie entirely in old tetrahedra or their complements this triangulation is such that the preimage of it under F is a subdivision of the original triangulation of M, and the map F is simplicial when considered as a map from this subdivision to the newly subdivided X. Each of the 2^{Tan^2} new n-simplices can lie in at most T old n-simplices and so has at most T preimages in M, hence the subdivision of M is by at most $T2^{Tan^2}$ n-simplices.

□

Remark 5.1.11. Note that by Corollary 5.1.9 the edge lengths of the tetrahedra in this final geometric triangulation are less than the injectivity radius of the manifold.

The following summarises the results of the section.

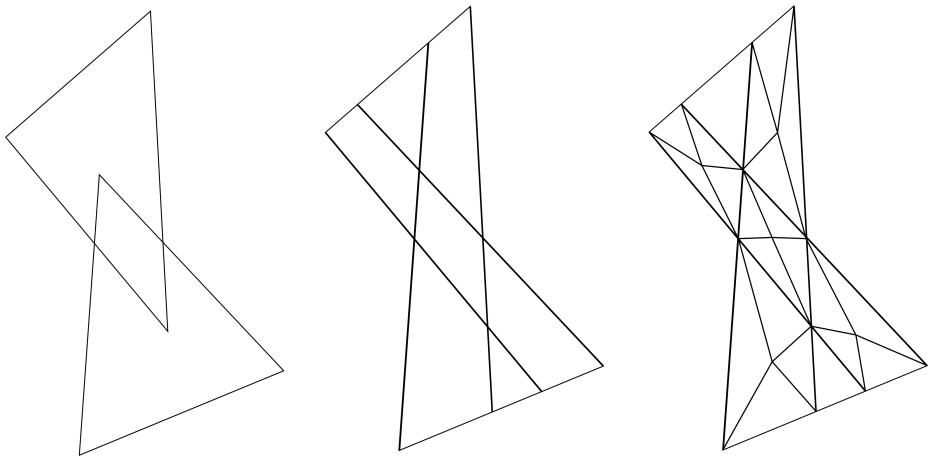


Figure 5.3: The intersection of two simplices can create non-convex (and hence hard to triangulate) polyhedra. Intersecting half spaces instead unfortunately creates more polyhedra, but ensures they are all convex (and thus easier to triangulate).

Corollary 5.1.12. *Suppose we are given a manifold M triangulated by t n-simplices such that M admits a hyperbolic structure. Then M admits a geometric triangulation where the number of tetrahedra in the geometric triangulation is bounded by*

$$2^{2^{(nt)^{O(n^4 t)}}}$$

and furthermore this triangulation is a degree one simplicial quotient of some subdivision of the triangulation of M where the subdivision has its number of tetrahedra bounded by the same bound.

In dimension 3, the bound can be improved to

$$2^{t^{O(t)}}$$

Proof. This follows from using the value of T from Corollary 5.1.9 (setting $c = 2$) in Lemma 5.1.10 to get a geometric triangulation of X and simplifying big O notation. The subdivisions of M then comes from taking the preimage of this geometric triangulation under the homotopy equivalence F from Corollary 5.1.9. This is a degree one simplicial quotient as F is a homotopy equivalence with image a triangulation of X. □

We now have bounds on both the combinatorial complexity (number of tetrahedra in the subdivision of which our manifold is a simplicial quotient) and the geometric complexity (edge lengths are bounded above by half the injectivity radius). We shall use the first bound in Section 5.2 to build a list on which our manifold must appear.

In Section 5.3 the bound on geometric complexity will allow us to navigate from one geometric triangulation to another in bounded time.

Remark 5.1.13. The results of this section have been to take a topological triangulation and find from it a geometric triangulation. Not only that, we have found a geometric triangulation with complexity bounded in multiple ways, and have given a step by step guide to finding it and finding how it relates to our original triangulation. One could be forgiven for thinking that the next step here is simply to describe how an algorithm performs these steps in bounded time. However, the steps we've described in this section are not immediately computable, and our algorithm shall take a different path.

We will instead attempt to use the combinatorial and geometric bounds on the complexity of the geometric triangulation which we have just shown to exist. These bounds allow us to formulate a finite list of possible triangulations, by listing all simplicial quotients of all subdivisions of our original triangulation within a certain complexity range. We will then try to find geometric triangulation from among that list, and find at least one which is a geometric triangulation of the original manifold.

Remark 5.1.14. On the topic of whether the steps in this section are algorithmic, we note that we use the existence of a homeomorphism from our triangulation to a geometric manifold. Even describing such a homeomorphism is computationally complicated, but note that if we were able to find such a homeomorphism and prove that it was indeed a homeomorphism in bounded time, then we'd have an algorithm to detect whether any given triangulation was a hyperbolic manifold, we spoke earlier about how such an algorithm would allow us to detect spheres of all dimensions, which is known to be undecidable (see Theorem 2.3.9).

However, it might be possible to produce an algorithm which takes as input a triangulation, and returns as output a map to a geometric manifold such that the map is guaranteed to be a homeomorphism if the original triangulation is homeomorphic to a hyperbolic manifold. In this case the algorithm doesn't provide proof that the map is a homeomorphism and thus doesn't provide proof that the original manifold is hyperbolic so doesn't run into undecidability problems.

5.2 Finding the Geometric Triangulation

We now return to the 3-dimensional case. So far we have seen that Grigoriev's Theorem (Theorem 2.4.1) gives a "bounded" solution for every component of the

solution set and that one of the components contains a faithful lattice representation. Together these imply the existence of a "bounded" faithful discrete representation. However, a subtle point that we have thus far obscured is that we gave (and know of) no general method for finding which of the solutions provided by Grigoriev's algorithm was our desired faithful discrete representation. This poses yet another road block to the idea of forming an algorithm from the steps of the proof in Section 5.1.

Thus we need another route to get to the geometric triangulation, and we will use the bounds on its complexity that we found in earlier sections. First for ease of reading we shall recap the major results of our previous sections that culminate in our algorithm. These are summarised in Theorem 5.2.1.

In Section 5.1 we saw that given a manifold M triangulated by t tetrahedra, there exists a geometric triangulation of M which is a degree one simplicial quotient of a subdivision of the original triangulation of M consisting of at most

$$2^{t^{O(t)}}$$

tetrahedra. This is exactly Corollary 5.1.12.

The results of Section 3.1, in particular Theorem 3.1.14 then tell us that we can perform a sequence of composite moves consisting of at most

$$2^{t^{O(t)}}$$

Pachner moves to take us from the original triangulation of M to the subdivision which admits the geometric quotient described above. All sequences of the form described in Section 3.1 end in subdivisions of the original triangulation of M. So if we perform all such sequences of composite moves which have their length in Pachner moves bounded in this way we get a list of subdivisions, at least one of which admits a geometric triangulation as a simplicial quotient. We can even discard any subdivisions whose size is bigger than our subdivision size bound. As the number of possible moves at each step is on the order of at most

$$2^{t^{O(t)}}$$

the total length of the list as well as the runtime of an algorithm producing it is bounded by

$$2^{2^{t^{O(t)}}}.$$

Given one subdivision N of the original triangulation of M from the list described above, we showed in Theorem 3.3.17 that we can list all oriented gluings which can be

found as a simplicial quotient (where the quotient map is degree one) of N and which admit a structure of a geometric triangulation. In fact, for each such item of the list, the information of a geometric triangulation includes the vertices of a triangulation of a fundamental polyhedron in $\mathbb{H}^3$ and matrices in $PSL(2,\mathbb{C})$ which perform the face pairings. This list is found in time exponential in the number of tetrahedra in the subdivision :

$$2^{2^{t^{O(t)}}}.$$

These gluings may or not be homeomorphic to the original manifold M but they do admit degree one maps from N and thus from M. Now performing this for each subdivision on the list has run time bounded by

$$\left(2^{2^{t^{O(t)}}}\right)^2$$

but this once again simplifies to

$$2^{2^{t^{O(t)}}}.$$

The above is thus a bound on the number of possible candidates in our list as well as a bound on the runtime of the algorithm creating the list.

We summarise this in the following Theorem.

Theorem 5.2.1. *Given an oriented closed hyperbolic 3-manifold M triangulated by t tetrahedra, we can construct a list of geometrically triangulated oriented closed hyperbolic 3-manifolds X_i paired with degree one simplicial maps from some subdivision of the original triangulation of M to X_i.*

This list is constructed in time bounded by

$$2^{2^{t^{O(t)}}}$$

and this is also a bound on the length of the list. Each triangulation has its number of tetrahedra bounded by

$$2^{t^{O(t)}}.$$

Furthermore, the list contains a geometric triangulation which is simplicially isomorphic to the geometric triangulation defined in Corollary 5.1.12.

Definition 5.2.2 (Candidate Manifolds)**.** We shall call the manifolds in the list created in the above theorem candidate manifolds, as they are so far the best candidates to be a true geometric triangulation of our original manifold M and at least one of them actually is such a triangulation.

As we know that all these manifolds are degree one simplicial quotients of subdivisions of M, all their fundamental groups admit surjections from $\pi_1(M)$. Note that as hyperbolic 3-manifold groups are residually finite, they are Hopfian, that is they don't admit self-surjections with non-trivial kernel (see [Mal40] or chapter 4 of [LS77]). Thus, if some Y in our list also admits π_1 surjections to the rest of the list and some X in our list is homeomorphic to M, then there exists a chain of surjections:

$$\pi_1(M) \twoheadrightarrow \pi_1(Y) \twoheadrightarrow \pi_1(X) \cong \pi_1(M).$$

The induced surjection $\pi_1(M) \twoheadrightarrow \pi_1(M)$ is an isomorphism because $\pi_1(M)$ is Hopfian and hence $\pi_1(M) \cong \pi_1(Y)$. Thus, by Mostow Rigidity, M and Y are homeomorphic. This surjection property characterises which candidates are true geometric triangulations of M.

Our plan moving forward is to identify such surjections by understanding the images of a given simplicial generating set of $\pi_1(M)$. For each candidate manifold X in our list the quotient map $M \to X$ is simplicial on some subdivision, so a simplicial generating set of $\pi_1(M)$ maps via this subdivision to a generating set of $\pi_1(X)$, each element of which is simplicial. However the algorithm which provides us with the hyperbolic structures on our list entries doesn't provide us with a simplicial generating set, it provides us with a set of generators for a lattice in $Isom^+(\mathbb{H}^3)$ or equivalently a generating set for the deck group of M.

We now show how to relate these two notions algorithmically, for this we slip back into the n-dimensional setting.

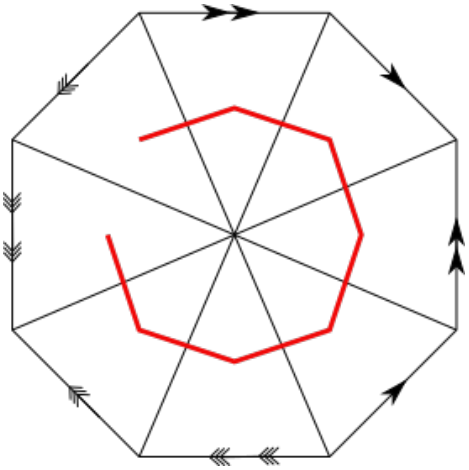
First, we recall some notions from Section 2.4.2. In that section we worked with a model for the geometry of the tetrahedra in a 3-dimensional triangulation but here we shall work purely with the combinatorial data of the n-dimensional triangulation of a hyperbolic manifold X.

Definition 5.2.3 (Face Pairings)**.** Recall that given some spanning tree Λ' of the dual graph $\Lambda(X)$ to a triangulation of a hyperbolic n-manifold X, we can define a polyhedron Y formed by ungluing the n-simplices which triangulate X and then regluing only along the (codimension one) faces which correspond to edges of Λ'.

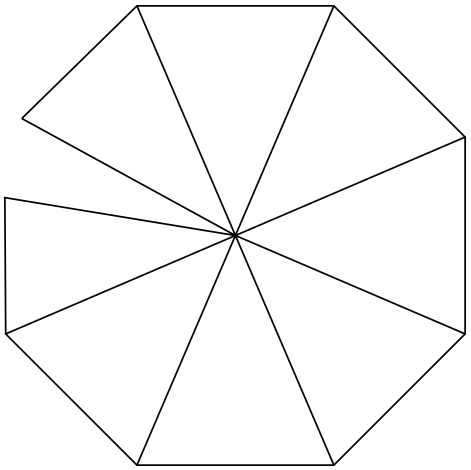
There is then a natural quotient map $\pi_Y : Y \to X$ which reglues along all faces to recover X. As Y is simply connected, given some choice of basepoint, π_Y lifts to a map $\tilde{\pi}_Y : Y \to \tilde{X}$ with image a fundamental domain containing this basepoint for the deck group action on $\tilde{X}$.

As X is a manifold, any translate of $\tilde{\pi_Y}(Y)$ can be reached by a path through codimension one faces and thus the deck group is generated by deck transformations which identify the faces of $\tilde{\pi_Y}(Y)$, we call these generators the face pairing generators or simply the face pairings of Y.

Note that these face pairings correspond to edges of the dual graph $\Lambda(X)$ which are not in Λ'. Indeed, if γ is a loop in $\Lambda(X)$ which consist of a path in Λ' and an edge $e \notin \Lambda'$ such that e connects the faces F and F', then γ lifts to a path which connects $\tilde{\pi_Y}(Y)$ to its translate under the deck transformation which identifies F with F'.



(a) A certain choice of spanning tree Λ' of the dual graph $\Lambda(X)$ where X is the genus 2 surface.



(b) The polyhedron Y associated to the spanning tree Λ' defined in the first figure.

Figure 5.4

We now seek to relate the face pairings to the simplicial loops in X. We do this by taking a subdivision of X which contains both the original triangulation and its dual graph.

Definition 5.2.4 (Coning to the Boundary)**.** Given a k-simplex Δ, a subdivision of its boundary and a point p in its interior, we can naturally associate Δ with

$$C(\partial\Delta) = \partial\Delta \times [0,1]/(x,1) \sim (y,1) \forall x, y$$

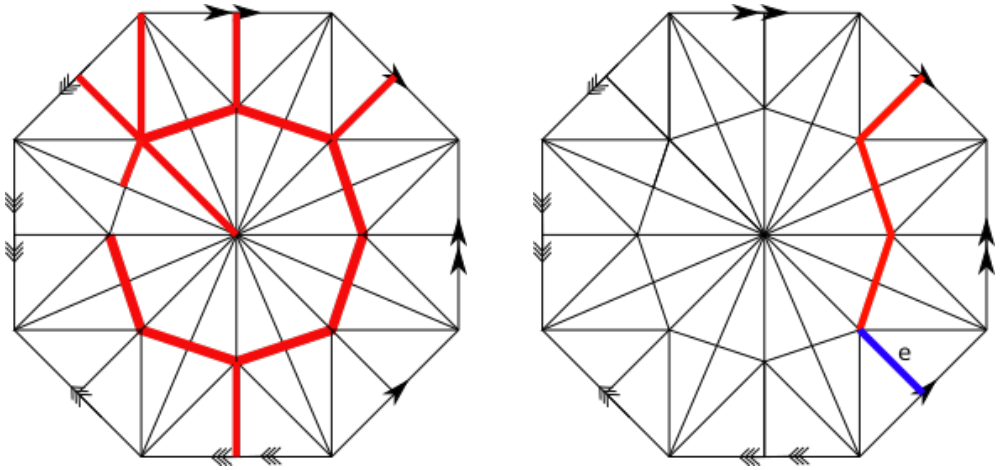
by identifying each line segment between p and some boundary point x with the interval $x \times [0,1]$. $C(\partial\Delta)$ has a natural structure as a triangulation, and we call this the triangulation formed by coning to the boundary of Δ.

Definition 5.2.5 (X' - a Partial Barycenteric Subdivision)**.** Define the subdivision X' of X as the triangulation formed by coning all $(n-1)$-simplices to their boundaries

and then coning all n-simplices to their now subdivided boundaries. This can be thought of as a partial barycentric subdivision. Note that the number of n-simplices in this subdivision is on the order of n^2T, where T is the number of n-simplices in the triangulation of X.

Note that $\Lambda(X)$ doesn't simplicially embed in X', but its barycentric subdivision naturally embeds as the graph spanned by the coning vertices. From this point on, we shall abuse notation to ignore this subdividing step and simply say that a loop or subgraph in $\Lambda(X)$ embeds if its barycentric subdivision does.

Definition 5.2.6 (Γ - a Spanning Tree for the 1-skeleton of X')**.** A spanning tree Λ' of $\Lambda(X)$ also embeds in X' and we can extend this to a spanning tree Γ of the entire 1-skeleton of X' by adding only edges which contain the coning vertices of the n-simplices. We can perform this subdivision also on the polyhedron Y to get a polyhedron Y' and the quotient map π_Y defined above induces a simplicial quotient map $\pi_{Y'} : Y' \to X'$. Note that by construction, the preimages $\Lambda'_{Y'} \subseteq \Gamma_{Y'}$ of $\Lambda' \subseteq \Gamma$ in Y' are both still trees.



(a) The spanning tree Γ extends Λ', note that it embeds also in Y without being disconnected.

(b) The face pairing loop associated to the "single arrowhead" face pairing.

Figure 5.5

Remark 5.2.7. Note that the face pairing loop γ corresponding to a given face pairing of Y defined in Definition 5.2.3 embeds in X' as the concatenation of an embedded path in Γ and one edge not in Γ. The set of all simplicial loops meeting this description (an embedded path in Γ and an edge e not in Γ connecting its endpoints)

shall be called the simplicial generating set, and such a curve is called the simplicial generator associated to e. By construction the face pairing loops corresponding to the face pairings of Y form a subset of the simplicial generating set of X'.

Lemma 5.2.8. *These face pairing loops generate the fundamental group of X and in fact the other generators can be expressed as at most T-fold products of these face pairings where T is the number of n-simplices in the given triangulation of X. We can find these products algorithmically in time bounded by*

$$(nT)^{O(nT)}.$$

Proof. Let e be some edge in Y' not in $\Gamma_{Y'}$, but with at least one endpoint w in $\Gamma_{Y'}$, note that all edges of X' not in Γ lift to such an edge. Let v denote the other endpoint of e and let u be the unique vertex of $\Gamma_{Y'}$ such that $\pi_{Y'}(u) = \pi_{Y'}(v)$ in X'. The path γ formed by concatenating a path in $\Gamma_{Y'}$ from u to w with e projects to the loop $\pi_{Y'}(\gamma)$ in X' which is exactly the simplicial generator associated to the edge $\pi_{Y'}(e)$.

If $\tilde{X}'$ denotes the universal cover of X', then recall that $\tilde{\pi}_{Y'} : Y' \to \tilde{X}'$ is the lift of the quotient map $\pi_{Y'} : Y' \to X'$ and it maps Y' to a fundamental polyhedron for the deck group action on $\tilde{X}'$. Consider $\tilde{u}, \tilde{v} = \tilde{\pi}_{Y'}(u), \tilde{\pi}_{Y'}(v)$ respectively, then the deck transformation mapping $\tilde{u}$ to $\tilde{v}$ is exactly the deck transformation corresponding to the loop $\pi_{Y'}(\gamma)$. Thus understanding how $\pi_{Y'}(\gamma)$ can be expressed as a product of face pairing loops is equivalent to understanding what sequence of face pairings maps $\tilde{u}$ to $\tilde{v}$.

Now as X is a manifold, there is a sequence of distinct translates $\tilde{\pi}_Y(Y) = Y_1, \ldots, Y_k$ of $\tilde{\pi}_Y(Y)$, all of which contain $\tilde{v}$, each of which intersects the next along a codimension one face such that the deck transformation taking Y_1 to Y_k maps $\tilde{u}$ to $\tilde{v}$. Define F_i to be the face of Y_1 such that its translate in Y_i connects Y_i to Y_{i+1}. Denote by E_i the face of Y_1 to which each F_i is paired and let g_{F_i} denote the face pairing deck transformation which maps E_i to F_i. The deck transformation taking Y_1 to Y_k is then exactly

$$g_{F_1} \cdot \ldots \cdot g_{F_k}$$

Note that as there are $\leq (n+1)T$ translates of Y_1 around any one vertex, $k \leq (n+1)T$.

Now as each translate contains $\tilde{v}$

$$\tilde{v} \in g_{F_1} \cdot \ldots \cdot g_{F_{k-i}}(Y_1) = Y_{k-i}$$

for all i, and so we can see that

$$g_{F_{k-i}}^{-1} \cdot \ldots \cdot g_{F_1}^{-1}(\tilde{v}) =: w_i$$

is a vertex of Y_1. Furthermore as

$$g_{F_1} \cdot \ldots \cdot g_{F_k}(\tilde{u}) = \tilde{v}$$

we can see that

$$w_i = g_{F_{k-i}}^{-1} \cdot \ldots \cdot g_{F_1}^{-1}(\tilde{v}) = g_{F_{k-i+1}} \cdot \ldots \cdot g_{F_k}(\tilde{u}).$$

Thus if we set $w_0 = \tilde{u}$, then each w_i is a vertex of Y_1 related to the next in the sequence by a face pairing deck transformation. Thus to find a sequence of deck transformations taking $\tilde{u}$ to $\tilde{v}$ we can simply consider all sequences of $\leq (n+1)T$ face pairings, and build up a sequence of vertices of Y_1

$$w_i = g_{F_{k-i+1}} \cdot \ldots \cdot g_{F_k}(\tilde{u})$$

for each and search for such sequences which take $\tilde{u}$ to $\tilde{v}$. Note that as X is in fact a hyperbolic manifold, there is no deck transformation which fixes a point, and thus the fact that a deck transformation takes $\tilde{u}$ to $\tilde{v}$ uniquely determines it, so any such sequence of face pairings must give the chosen deck transformation.

As both the total length of the sequence of pairings and the number of choices of pairing at each stage is bounded above by $(n+1)T$, and we know such a sequence always exists, we can find one in time $(nT)^{O(nT)}$. We do so for each of the edges not in $\Lambda_{P'}$ and thus we do this on the order of n^2T times, and the result follows by simplifying big O notation. □

We remark here some technicalities concerning the above proof.

Remark 5.2.9. There are two natural ways of thinking of fundamental polyhedra. Often with a geometric manifold M which admits a triangulation, we consider a fundamental polyhedron P of its associated lattice action such that P is triangulated by lifts of the n-simplices in M. In this case we have a triangulation of P and can represent M as P glued along some face pairings as we do in Figure 5.4(a). In this case you have many faces of n-simplices which are interior to the polyhedron and thus are paired by the identity map. Furthermore, the face pairings of exterior faces of the polyhedron correspond to non-trivial lattice elements.

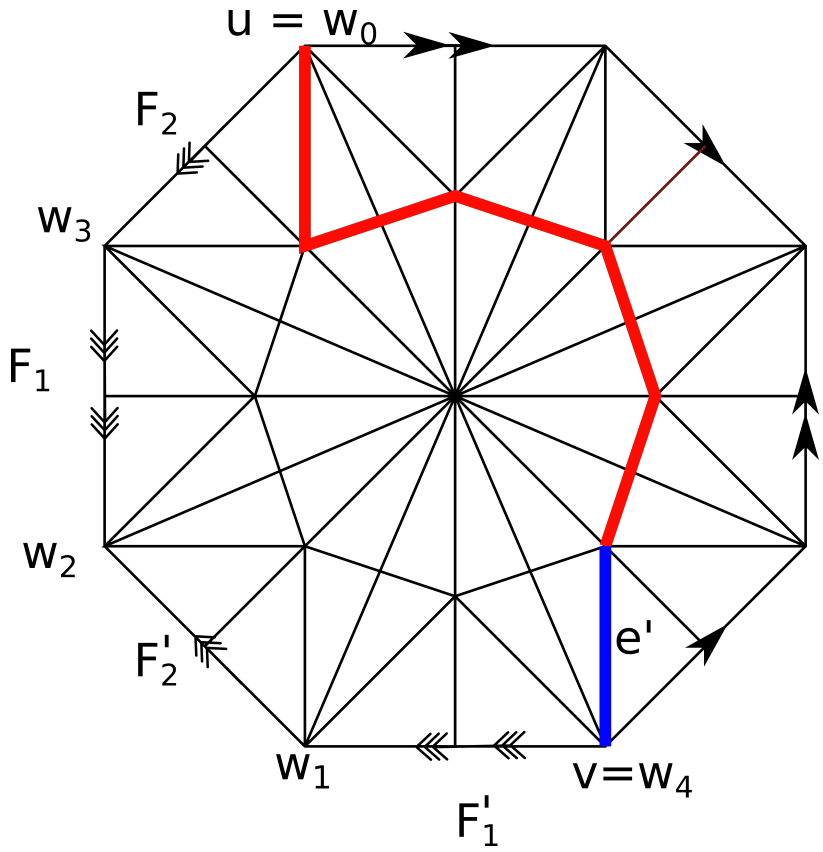


Figure 5.6: The simplicial loop corresponding to the edge e' is not a face pairing loop and its corresponding deck transformation can be written as the product of face pairings $g_{F_1'}g_{F_2}g_{F_1}g_{F_2'}$. (Here $g_{F_i'}$ maps F_i onto F_i' and $g_{F_i} = g_{F_i'}^{-1}$).

However, if we only have the combinatorial data, we don't know which faces will be interior and exterior to our polyhedron once its embedded. Once we have chosen a spanning tree for the dual graph to our triangulation, only those faces corresponding to edges of the spanning tree can be known in advance to be identified by the identity map. We see this combinatorial polyhedron in figure 5.4(b).

Remark 5.2.10. If we only know the combinatorial data, then we are actually working with the face pairings of Y rather than the exterior face pairings of its image Y_1 under the map $\tilde{\pi}_Y$ which maps it into $\mathbb{H}^n$. Thus we only actually know where $g_{F_{k-i}}$ sends the vertices of E_i and no further information, so if at any point $g_{F_{k-i}}$ is not defined on w_i, we simply discard that sequence. A further technical quirk is that as we only know the combinatorial data coming from Y, there are some face pairings of Y which are not face pairings of its image Y_1 (because they occur on the interior of Y_1), but will still occur in our sequences, this occurs when the map $\tilde{\pi}_Y$ identifies a pair of faces E, F by simply embedding each in the same place as the other, in this case

the corresponding deck transformation taking E to F would simply be the identity.

These technicalities do not preclude the following two facts. Firstly, the existence of a sequence of face pairing deck transformations as defined above implies the existence of a sequence (still of length bounded above by $(n+1)T$) of face pairings of Y which eventually relates u to v. Furthermore, given a sequence of face pairings of Y relating u to v, the corresponding sequence of face pairing deck transformations (some of which may be the identity) translates $\tilde{u}$ to $\tilde{v}$. This means that whether we have any geometric info or not, we can still generate the entire simplicial generating set using only face pairings.

We now return to the 3-dimensional case. Looking for surjections between groups is not in general a simple task. It is certainly not a good idea to check through all possible maps from one group to the other and check whether they are surjective homomorphisms. We circumnavigate this by picking a specific 'homomorphism' in advance to check. Let X_1, X_2 be candidate manifolds. If X_1 is indeed homeomorphic to M and $f_{X_1} : M \to X_1$ is the quotient map then $f_{X_1*} : \pi_1(M) \to \pi_1(X_1)$, the induced map on fundamental groups, is an isomorphism and so admits an inverse and so $f_{X_2*} \circ {f_{X_1*}}^{-1} : \pi_1(X_1) \to \pi_1(X_2)$ is a surjective homomorphism. Thus to find an X_i which is homeomorphic to M we don't need to check over all possible maps, we only need to check whether the above map exists and is a surjection for all targets $\pi_1(X_j)$. Existence is important here as when f_{X_1*} is not an isomorphism then the above need not even be a well defined map.

Theorem 5.2.11. *Let X_1, X_2 be a pair of candidate manifolds for M. Suppose X_1, X_2 are triangulated by at most*

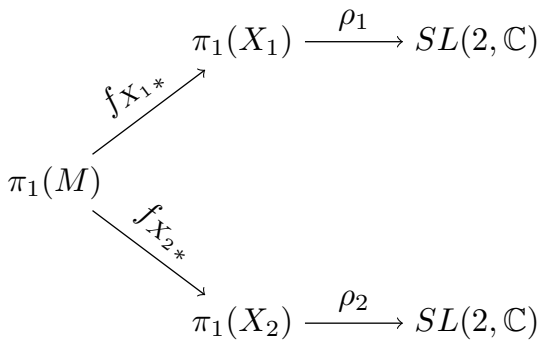
$$T < 2^{2^{t^{O(t)}}},$$

then there is an algorithm which searches for a surjective homomorphism from $\pi_1(X_1)$ to $\pi_1(X_2)$. If X_1 is homeomorphic to M, then it is guaranteed to find such a homomorphism, and the homomorphism is indeed the F defined above. Otherwise, the algorithm may or may not find a surjective homomorphism, even when one does exist.

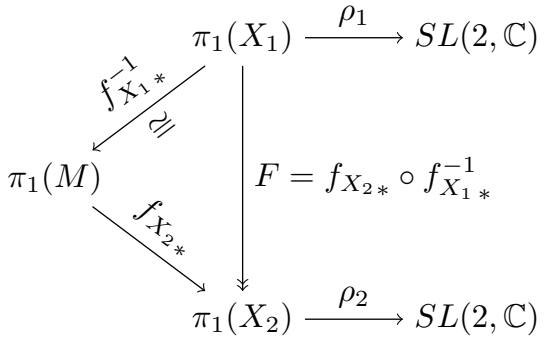
The runtime of this algorithm is bounded above by

$$T^{O(T^2)}.$$

Proof. Let $S := \{g_1, \dots g_n\}$ be a simplicial generating set for $\pi_1(M)$ as in definition 5.2.7 for some spanning tree of the 1-skeleton of M. Each g_i then consists of at most $6t$ edges. As each X_i is a quotient of a subdivision by at most T simplices, the image



(a)



(b)

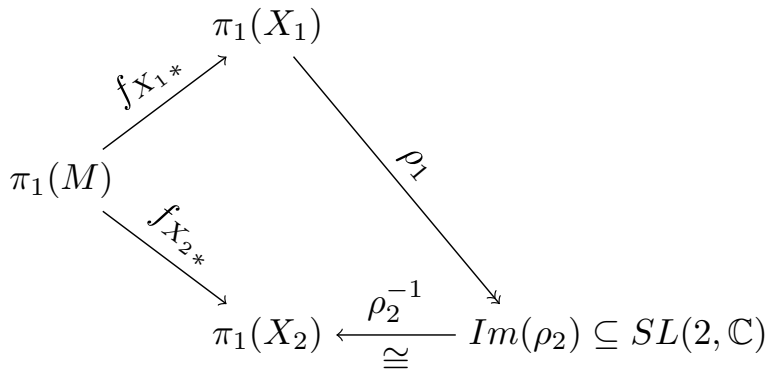
Figure 5.7: Figure a shows the setup before we know anything about X_1,X_2 other than that they are each candidate manifolds. Figure b shows that if X_1 is homeomorphic to M then F is a well defined surjective homomorphism.

of each g_i in $\pi_1(X_i)$ can be expressed as a simplicial loop with at most $36tT$ edges, thus it is at most a $36tT$-fold product of the simplicial generators of each $\pi_1(X_i)$. If X_1 is indeed homeomorphic to M then F is a well defined surjective homomorphism, and it takes the image $f_{X_1*}(g_i)$ to $f_{X_2*}(g_i)$.

We take as a starting point the very same system of polynomials that we used as input for the algorithm which found us the geometric structure on X_2.

To this system we add polynomials with variables the real and imaginary parts of entries of complex matrices $H_1, \ldots H_k$ corresponding to the simplicial generators of $\pi_1(X_1)$. As $\pi_1(X_1)$ admits a presentation $\pi_1(X_1) = \langle h_1, \ldots, h_n \mid r_1, \ldots r_n \rangle$, we also add polynomials corresponding to these relations, so for example, if $r_1 = h_3 h_2 h_1^{-1}$, then we'd add polynomials which encode the requirement that $H_3 H_2 H_1^{-1} = \mathrm{Id}$.

Then if we also require that these matrices satisfy that their determinants be 1, we know that a solution to a system gives both a geometric structure on X_2 and a representation of $\pi_1(X_1)$ into $SL(2,\mathbb{C})$.



(a)

Figure 5.8: This shows how we get a surjective homomorphism whenever there is a solution to the combined system of polynomials.

If for a faithful representation ρ_2 of $\pi_1(X_2)$ into $SL(2,\mathbb{C})$, there is a representation ρ_1 of $\pi_1(X_1)$ where the image of each $f_{X_1*}(g_j)$ is $\rho_2(f_{X_2*}(g_j))$, then this encodes a surjective homomorphism $\rho_2^{-1} \circ \rho_1 : \pi_1(X_1) \to \pi_1(X_2)$. Thus by combining the system of polynomials which we used to find a hyperbolic structure on X_2 (and which has variables corresponding to the lattice generators) with the system for representations of $\pi_1(X_1)$ as we did above, we can then add polynomials which relate the two systems. As noted above, both $\rho_1 f_{X_1*}(g_j)$ and $\rho_2(f_{X_2*}(g_j))$ are at most $36tT$-fold products of the images of the simplicial generators which are themselves at most T-fold products of the lattice generators, so the eight (two per complex variable of the complex matrices) polynomials required to encode the equation

$$\rho_1(f_{X_1*}(g_j)) = \rho_2(f_{X_2*}(g_j))$$

are of degree $36tT^2$. Adding these polynomials makes a system that has a solution if and only if X_2 has a hyperbolic structure (which we already know it to have) such that $\pi_1(X_1)$ admits a representation which surjects onto $\rho_2(\pi_1(X_2))$.

Thus running the algorithm from Theorem 2.4.1 gives us a solution every time when $X_1 \cong M$ and any time it gives a solution, we know at least that $\pi_1(X_1)$ surjects onto $\pi_1(X_2)$. To find the runtime we just need to understand the complexity of the system of polynomials. The original system for finding the hyperbolic structure on

X_2 has complexity described as below, for some constant C:

$$\kappa \leq CT^2$$
$$N \leq CT$$
$$d \leq CT$$
$$M \leq CT$$

The new system adds variables corresponding to the entries of matrices corresponding to each oriented 1-simplex in the triangulation of X_1, hence the number of new variables is bounded by some constant multiple of T. We add polynomials for the relations, which come from the 2-simplices of X_2 and thus have their number bounded by a constant multiple of T. The other polynomials come from the generators of $\pi_1(M)$ and thus have their number bounded by a constant multiple of $t << T$. The complexity of coefficients in the added polynomials is ± 1 and the degree of the polynomials is bounded by $36tT$ which is on the order of T. Thus adding the new variables and polynomials has no impact on the complexity of the system up to altering the constant C. Thus an algorithm can find whether F exists and is surjective in time bounded by

$$T^{O(T^2)}$$

just as in Corollary 3.3.13. □

Corollary 5.2.12. *We can check whether a manifold X_1 in our list of candidate manifolds is homeomorphic to M in time bounded by*

$$2^{2^{t^{O(t)}}}$$

Proof. As noted above, running the algorithm of Theorem 5.2.11 gives us a solution every time when $X_1 \cong M$, regardless of what X_2 is, and any time it gives us a solution, we know at least that $\pi_1(X_1)$ surjects onto $\pi_1(X_2)$. Thus, we know that running the algorithm above on each pair will return at least one candidate manifold which π_1-surjects onto all the others in time bounded by

$$T^{O(T^2)}.$$

Substituting the known bound on T given by Corollary 5.1.12 gives that for each pair this algorithm has a runtime bound of the form

$$\left(2^{t^{O(t)}}\right)^{2^{t^{O(t)}}} \approx 2^{2^{t^{O(t)}}}.$$

Now we saw above that the list of candidates has size bounded by the same bound. Thus we can check for all such surjections in time bounded by

$$\left(2^{2^{t^{O(t)}}}\right)^2 \approx 2^{2^{t^{O(t)}}}.$$

We now repeat why it is that if all the above surjections exist this X_1 is homeomorphic to M. First note that on account of being hyperbolic manifolds, each element of the list is residually finite and hence Hopfian, this means that if two manifolds are such that the fundamental group of each admits a surjection onto the other, then these surjections are isomorphisms. We also know that the fundamental group of every manifold in the list admits a surjection from $\pi_1(M)$ and that one element of the list is homeomorphic to M and thus that element, call it Y_1, has this same property. Thus X_1 and Y_1 are both such that their fundamental groups surject onto one another and are thus isomorphic. Now given that the manifolds are both hyperbolic, this isomorphism induces a homeomorphism by Mostow Rigidity. Thus X_1 is homeomorphic to M, as required.

□

We are now ready to bring all of this together into an algorithm which lists geometric triangulations of M. However, in Section 5.1, we showed the existence of not only a geometric triangulation, but one with edge lengths bounded by injectivity radius, the following lemma will allow us to recover this property in our geometric triangulations also.

Lemma 5.2.13. *Let M be a 3-manifold triangulated by t tetrahedra and let $\mathcal{T}$ be an alternative triangulation by $T < 2^{t^{O(t)}}$ tetrahedra. We can check if $\mathcal{T}$ admits the structure of a geometric triangulation of a hyperbolic 3-manifold with edge lengths bounded above by half of the injectivity radius. If it does admit such a structure we can find it in time bounded by*

$$2^{2^{t^{O(t)}}}.$$

Proof. Note that we are simply reproving Corollary 3.3.13, with the extra requirement of this injectivity radius bound and with a better knowledge of how T and t interact.

Note that as M admits a triangulation by t tetrahedra, then if it admits a hyperbolic structure, its injectivity radius will be bounded below by

$$\frac{1}{t^{O(t)}}$$

by Corollary 5.1.4. Thus, all we need to do is add equations into the system of polynomials of Lemma 3.3.11 which require that all the edge length variables be less

than half our known lower bound on injectivity radius. This introduces less than $6T$ new polynomials with coefficients of complexity bounded by $log_2(t^{O(t)})$. Thus the new system has size bounded as follows for some universal constant C' (compare with the size of the system in Lemma 3.3.11).

$$\begin{aligned} \kappa &\leq C'T^2 \\ N &\leq C'T \\ d &\leq C'T \\ M &\leq C'T \end{aligned}$$

and thus it follows by the same proof as Corollary 3.3.13 and our known bounds on T that the runtime of the algorithm is

$$C'T(C'T^2C'T)^{(C'T)^2} \leq T^{O(T^2)} \leq 2^{2^{t^{O(t)}}}.$$

□

Remark 5.2.14. Note that here we could just as easily have ensured edge length was less than any multiple $1/c$ of injectivity radius, for some integer $c \geq 1$, as we did in Corollary 5.1.9. The runtime here would be bounded by

$$c2^{2^{t^{O(t)}}}$$

as c would only affect the complexity M of the coefficients.

We can now prove the following theorem.

Theorem 5.2.15. *Given a triangulation of a hyperbolic 3-manifold M by t tetrahedra, then there is an algorithm which produces a geometric triangulations of M consisting of less than*

$$2^{t^{O(t)}}$$

tetrahedra in time bounded by

$$2^{2^{t^{O(t)}}}.$$

Furthermore, this geometric triangulation is simplicially isomorphic to the triangulation defined in Lemma 5.1.10. Hence, this geometric triangulation has the property that all edge lengths are bounded above by the injectivity radius of the manifold.

Proof. An algorithm producing triangulation first simply produces the list of candidate manifolds as described at the start of this section. It then performs the algorithm of Corollary 5.2.12 for each candidate manifold. The runtime of this algorithm and the length of the list of candidate manifolds is bounded by

$$2^{2^{t^{O(t)}}}$$

and thus this is a bound for the whole algorithm after simplifying big O notation.

The reason that we know one of our list must be simplicially isomorphic to the triangulation of Lemma 5.1.10 is that we have listed all possible subdivisions and all possible simplicial quotients within the given complexity bound, and have only discarded elements of the list if they don't admit geometric structures or aren't homeomorphic to M. Thus we know that for at least one element of the list, the algorithm of Lemma 5.2.13 returns a positive result, a geometric triangulation with edge lengths bounded by the injectivity radius of the manifold. □

As an aside, it is interesting to note here that though the above theorem provides a geometric triangulation which is simplicially isomorphic to the one derived in Lemma 5.1.10 and which shares the same property of edge lengths being less than injectivity radius, the geometric structures may still be different between the two. In this way, the geometric triangulation of Section 5.1 remains entirely theoretical.

Remark 5.2.16. Note that we have a bounded runtime algorithm to find a hyperbolic structure on a hyperbolic 3-manifold, given as a geometric triangulation. This doesn't however lead to an algorithm for recognising whether or not a 3-manifold admits a hyperbolic structure. This is because if we handed our algorithm a non-hyperbolic 3-manifold, then it may well return a geometric triangulation which is some degree one simplicial quotient of some subdivision of our manifold, but as the manifold is not hyperbolic, we can't use Mostow Rigidity to prove the two triangulations are homeomorphic. For example, in higher dimensions one could take a hyperbolic manifold and connect sum it with a homotopy n-sphere. This resulting manifold is not hyperbolic but admits a degree one map onto the original hyperbolic manifold so our algorithm could produce a geometric triangulation of the original hyperbolic manifold as a quotient of our new non-hyperbolic manifold. If this occurs, then it can't differentiate between the original hyperbolic manifold and the one with the homotopy sphere attached and would say they were the same manifold. In this way, the algorithm is capable of providing false positives when fed non hyperbolic manifolds and thus can't be used to recognise hyperbolicity. Hence, this safeguards against one of the bigger worries about generalising these proofs to higher dimensions.

5.3 Comparing Geometric Triangulations

Given two closed hyperbolic 3-manifolds M, N we have so far managed to produce bounded complexity geometric triangulations of these manifolds. One might think that we would now apply the results of Chapter 3 to compare them. However, as noted in that chapter, the advantages of our methods are that they don't require any further knowledge about the geometry of the triangulations involved. This is indeed true, and makes the results of Chapter 3 much easier to apply in general. Quite remarkably, however, we find ourselves in a situation where we know almost excessive amounts about the geometry of the triangulations we are comparing, and so here we can use the fact that Kalelkar and Phanse's result has a faster runtime in terms of the number of tetrahedra. This is precisely because the results of Chapter 4 tell us that the both of our geometric triangulations satisfy the condition that their edges are all shorter than the injectivity radius of the manifolds they triangulate. We now restate Kalelkar and Phanse's theorem and show why this relation is important.

Theorem 5.3.1 (Kalelkar Phanse '21 [KP21])**.** *Let M be a closed hyperbolic 3-manifold with geometric triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ by t_1, t_2 tetrahedra respectively. Let L be an upper bound on edge length and $inj(M)$ be the injectivity radius of M.*

$\mathcal{T}_1$ and $\mathcal{T}_2$ are related by $\leq f(t_1, t_2, L, inj(M))$ Pachner moves which do not remove common vertices. The function f is defined as follows

$$f(t_1, t_2, L, inj(M)) = 32(24^{4+3m})t_1t_2(t_1 + t_2)$$

where m is defined as a natural number greater than

$$(2cosh^2(L) + 1)ln(L/inj(M))$$

Corollary 5.3.2. *Let M be a closed hyperbolic 3-manifold with geometric triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ by $< T$ tetrahedra each and such that the lengths of the edges of tetrahedra in each triangulation are strictly less than the injectivity radius of the manifold. Then the number of Pachner moves required to get from one to the other is bounded by*

$$O(T^3).$$

Proof. Note that when $L \leq inj(M)$, the lower bound on m becomes negative, as $ln(x)$ is negative for all $x < 1$. Thus m can be taken to be 0 and so the result follows. □

Theorem 1.1.1. *If M_1 and M_2 are triangulated closed hyperbolic 3-manifolds, each triangulated by less than t tetrahedra, then we can decide in time bounded by*

$$2^{2^{t^{O(t)}}}$$

whether or not they are homeomorphic.

Proof. Theorem 5.2.15 provides for each manifold M_1, M_2 a geometric triangulation with edge lengths bounded by the injectivity radius, doing so in time bounded by

$$2^{2^{t^{O(t)}}}.$$

We also know that the number of tetrahedra in each of the triangulations is bounded above by

$$2^{t^{O(t)}}.$$

Thus by Corollary 5.3.2 the number of Pachner moves required to take $\mathcal{T}_1$ to $\mathcal{T}_2$ is bounded above by

$$\left(2^{t^{O(t)}}\right)^3 \approx 2^{t^{O(t)}}.$$

As the number of moves at each stage is polynomial in the number of tetrahedra, the number of possible sequences of this length is bounded by

$$2^{2^{t^{O(t)}}}$$

thus it is also a bound for the runtime of the algorithm which checks the output of all the sequences to see if any of them is simplicially isomorphic to $\mathcal{T}_2$.

There are two possibilities, either the algorithm returns a match, in which case the manifolds are homeomorphic or it doesn't. If there is no match then there is no sequence of Pachner moves of length less than the given bound relating the two geometric triangulations of M_1 and M_2 which contradicts Corollary 5.3.2 and thus the two manifolds must be distinct. □

Bibliography

[AB17] K. Adiprasito and B. Benedetti. Subdivisions, shellability, and collapsibility of products. *Combinatorica*, **37**(1), (2017), pp. 1-30.

[Ago06] Ian Agol. Systoles of Hyperbolic 4-Manifolds Preprint, arXiv:math/0612290.

[AL10] I. Agol, Yi Liu. Presentation length and Simon's Conjecture *J. Amer. Math. Soc*, **25**(1), 2021, pp151-187.

[Bas92] A. Basmaijan. Generalisinmg the Hyperbolic Collar Lemma *Bulletin o the American Mathematical Society* **27**(1), (1992), pp. 154-158.

[BT11] M. Belolipetsky, S. Thomson. Systoles of Hyperbolic Manifolds. *Algebr. Geom. Topol.* **11**(3), (2011), pp. 1455–1469.

[BP92] R. Benedetti and C. Petronio. *Lectures on Hyperbolic Geometry.* Springer Berlin Heidelberg, (1992).

[BBB+10] Laurent Bessières, Gérard Besson, Michel Boileau, Sylvain Maillot, and Joan Porti. *Geometrisation of 3-Manifolds.* European Mathematical Society Publishing House, (2010).

[Cha95] I. Chavel. *Riemannian Geometry : A Modern Introduction.* Cambridge Univ. Press (1995), Section IX.8

[Che70] J. Cheeger. Finiteness Theorems for Riemannian Manifolds. *American Journal of Mathematics* **92**(1), (1970), pp. 61-74

[Cul86] M. Culler. Lifting Representations to Covering Groups. *Advances in Mathematics* **59**, (1986), pp.64-70.

[DGr08] F. Dahmani, D. Groves. The isomorphism problem for toral relatively hyperbolic groups. *Publ. Math. Inst. Hautes Etudes Sci.*, **107**, (2008), pp. 211–290.

[Edw06] J. D. Edwards. Suspensions of Homology Spheres Preprint, arXiv:math/0610573.

[EP94] Epstein D. B. A., Petronio C. An Exposition of Poincaré's Polyhedron Theorem *Enseign. Math.* (2), 40(1-2):113–170, 1994.

[FM11] B. Farb and D. Margalit. *A Primer on Mapping Class Groups* Princeton University Press (2011).

[FM97] Fomenko, A.T., Matveev, S.V. *Algorithmic and Computer Methods for Three-Manifolds* Springer (1997).

[FPS19] D. Futer, J. Purcell and S. Schleimer. Effective distance between nested Margulis tubes *Trans. Amer. Math. Soc* **372**(6) (2019), pp.4211-4237

[GR70] H. Garland and M. S. Raghunathan Fundamental Domains for Lattices in (R-)rank 1 Semisimple Lie Groups. *Annals of Mathematics* Second Series, **92**(2) (1970), pp. 279-326.

[Gri86] D. Grigoriev. *Computational complexity in polynomial algebra.* In *Proc. Intern. Congress of Mathematicians*, **2**, Berkeley, 1986, pp. 1452-1460.

[GV85] D. Grigoriev and N. Vorobjov. Finding Real Solutions of Polynomial Inequalities in Subexponential Time. *Soviet Math. Dokl.* **32** (1985), pp.316-320.

[GV88] D. Grigoriev and N. Vorobjov. Solving Systems of Polynomial Inequalities in Subexponential Time. *Journal of Symbolic Computation* **5**(1) (1988), pp.37-64.

[Gro78] M. Gromov. Manifolds of Negative Curvature. *Journal of Differential Geometry* **13** (1978), pp.223-230.

[Hat01] A. Hatcher. *Algebraic Topology* Cambridge University Press, (2001).

[Hem76] J. Hempel *3-Manifolds* Princeton University Press (1976).

[Hub82] G. Huber. Gamma Function Derivation of n-Sphere Volumes *The American Mathematical Monthly* **89**(5), (1982), pp.301-302.

[HSoM67] J.F.P. Hudson, J.L. Shaneson, and Chicago (Ill.). University. Dept. of Mathematics. *Piecewise Linear Topology.* Number v. 1 in Mathematics lecture note series. Department of Mathematics, University of Chicago, (1967).

[JS79] W. Jaco and P. Shalen. Seifert fibered spaces in 3-manifolds. *Memoirs of the American Mathematical Society*, **21**(220), (1979).

[Joh79] K. Johannson. *Homotopy Equivalences of 3-Manifolds with Boundaries* Lecture Notes in Mathematics Springer Berlin. (1979).

[KP20] Geometric Bistellar Moves Relate Geometric Triangulations *Topology and its Applications* **285**, (2020) pp. 107390-107397.

[KP21] T. Kalelkar, A. Phanse. An Upper Bound on Pachner Moves Relating Geometric Triangulations, *Discrete Comput Geom* **66**, (2021), pp. 809-830

[KR20] T. Kalelkar, S. Raghunath. Bounds on Pachner moves and systoles of cusped 3-manifolds, Preprint, Accepted in Journal of Algebraic and Geometric Topology, (2020), arXiv:2007.02781.

[Kap00] M. Kapovich. *Hyperbolic Manifolds and Discrete Groups.* Birkhäuser Boston, 2010 (first published 2000).

[Kee73] L. Keen Collars on Riemann Surfaces. *Discontinuous Groups and Riemann Surfaces (AM-79), Volume 79: Proceedings of the 1973 Conference at the University of Maryland* edited by Leon Greenberg, Princeton: Princeton University Press, (2016), pp. 263-268.

[Kel04] R. Kellerhalls. On the Structure of Hyperbolic Manifolds. *Isr. J. Math*, **143** (2004), pp. 361-379.

[Kne29] H. Kneser. Geschlossene flächen in dreidimensionalen mannigfaltigkeiten. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, **38**, (1929), pp. 248-259.

[Kup19] G. Kuperberg. Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization. *Pacific Journal of Mathematics* **301**(1) (2019), pp. 189-241.

[Lac00] M. Lackenby. *Hyperbolic Manifolds* Lecture notes available at https://people.maths.ox.ac.uk/lackenby/

[Lic99] W. B. R. Lickorish. Simplicial moves on complexes and manifolds. In Proceedings of the Kirbyfest (Berkeley, CA, 1998), *Geom. Topol. Monogr.*, **2**. Geom. Topol. Publ. Coventry, (1999), pp. 299-320.

[LS77] R. C. Lyndon, P. E. Schupp. *Combinatorial Group Theory* Springer Berling/Heidelberg, (1977).

[Mal40] A. I. Malcev, On isomorphic matrix representations of infinite groups. *Mat. Sb.* **8**,(1940), pp. 405-422.

[Man01] J. Manning, Algorithmic detection and description of hyperbolic structures on closed 3–manifolds with solvable word problem. *Geom. Topol.* **6**(1), (2002), pp. 1-26.

[Mar58] A. A. Markov. The insolubility of the problem of homeomorphy. *Dokl. Akad. Nauk SSSR* **121**, (1958), pp. 218–220.

[Mar16] B. Martelli *An Introduction to Geometric Topology* Self-Published (2016).

[Mij03] A. Mijatovic. Simplifying triangulations of the 3-sphere *Pacific Journal of Mathematics* **208**(2), (2003), pp.291-324.

[Mij04] A. Mijatovic. Triangulations of Seifert fibred manifolds *Mathematische Annalen* **330**(2), (2003), pp.235-273.

[Mij05a] A. Mijatovic. Triangulations of fibre-free Haken 3-manifolds *Pacific Journal of Mathematics* **219**(1), (2005), pp.139-186

[Mij05b] A. Mijatovic. Simplicial structures of knot-complements *Mathematical Research Letters* **12**(5-6), (2005), pp. 843-856.

[Mil62] J. Milnor. A unique decomposition theorem for 3-manifolds. *American Journal of Mathematics*, **84**(1), (1962), pp. 1-7.

[Mey87] R. Meyerhoff. A Lower Bound for the Volume of Hyperbolic 3-Manifolds *Canadian Journal of Mathematics*, **39**(5) (1987), pp.1038-1056.

[Moi52] E. Moise. Affine structures in 3-manifolds: V. the triangulation theorem and hauptvermutung. *Annals of Mathematics*, **56**(1), (1952), pp.96–114.

[Oli80] R. Oliver. On Bieberbach's Analysis of Discrete Euclidean Groups *Proc. Amer. Math. Soc.*, **80**, (1980), pp. 15-21

[Pac91] U. Pachner. P.L. Homeomorphic Manifolds are Equivalent by Elementary Shellings *European Journal of Combinatorics* **12**(2), (1991), pp. 129-145.

[Per02] G. Perelman. The entropy formula for the Ricci flow and its geometric applications. arxiv:math/0211159, (2002).

[Per03a] G. Perelman. Finite extinction time for the solutions to the Ricci flow on certain three-manifolds. arxiv:math/0307245, (2003).

[Per03b] G.Perelman. Ricci flow with surgery on three-manifolds. arxiv:math/0303109, (2003).

[Pra73] G. Prasad. Strong Rigidity of $\mathbb{Q}$-rank 1 Lattices *Inventiones math.* **21** (1973) , pp. 255 - 286.

[Pur20] J. Purcell. *Hyperbolic Knot Theory.* arXiv:2002.12652, (2020).

[Rez95] A. Reznikov The Volume and the Injectivity Radius of a Hyperbolic Manifold *Topology* **34**(2) (1995) pp. 477-479.

[SS14] P. Scott, H. Short. The homeomorphism problem for closed 3-manifolds. *Algebr. Geom. Topol.* **14**(4), (2014), pp. 2431-2444.

[Sel95] Z. Sela. The isomorphism problem for hyperbolic groups. I. *Ann. of Math. (2)* **141**(2) (1995), pp. 217–283.

[Thu78] W. Thurston. *The geometry and topology of three-manifolds.* Princeton Univ. Math. Dept. (1978) (Available at http://msri.org/publications/books/gt3m/)

[Thu82] W. Thurston. Three dimensional manifolds, kleinian groups and hyperbolic geometry. *Bulletin of the American Mathematical Society* **6**(3), (1982), pp.357-382.

[Thu97] W. Thurston (ed. by Silvio Levy). *Three-Dimensional Geometry and Topology Volume One.* Princeton Mathematical Series, (1997).

[Vol74] I. A. Volodin et al. The Problem of Discriminating Algorithmically the Standard Three-Dimensional Sphere *Russian Mathematical Surveys* **29**(5), (1974).

[Whi00a] M. E. White. *Some Bounds for Closed Hyperbolic 3-manifolds* PhD Thesis (2000).

[Whi00b] M. E. White. A Diameter Bound for Closed, Hyperbolic 3-Manifolds. Preprint, (2000) arXiv:math/0104192 .

[Whi02] M. E. White. Injectivity Radius and Fundamental Group of Hyperbolic 3-Manifolds. *Communications in Analysis and Geometry* **10**(2) (2002), p. 377-395.