

Three-dimensional Orbifolds and Cone-Manifolds

Daryl Cooper, Craig Hodgson, Steve Kerckhoff

Preface

The aim of this memoir is to give an introduction to the statement and main ideas in the proof of the “Orbifold Theorem” announced by Thurston in late 1981 ([83], [81]). The Orbifold Theorem shows the existence of geometric structures on many 3-dimensional orbifolds, and on 3-manifolds with a kind of topological symmetry.

In July 1998, fifteen lectures on the Orbifold Theorem were presented by the authors at a Regional Workshop in Tokyo. We would like to take this opportunity of thanking the Mathematical Society of Japan who organized and supported this workshop. In particular we thank and express deep gratitude to Sadayoshi Kojima whose idea this meeting was and who provided the stimulus required to produce this manuscript in a timely fashion. His tremendous organizational skills helped to ensure the success of the meeting.

The first six lectures were of an expository nature designed to meet the needs of graduate students. It is the content of these first six lectures, somewhat expanded, that form the bulk of this memoir. The content of the final nine lectures where we discussed the proof of the orbifold theorem have been summarized, and a detailed account of this proof will appear elsewhere.

It is our intention that this memoir should provide a reasonably self-contained text suitable for a wide range of graduate students and researchers in other areas, and to provide various tools that may be useful in other contexts. For this reason we have included many examples and exercises.

We develop the basic properties of orbifolds and cone-manifolds. In particular many ideas from smooth differential geometry are extended to the setting of cone-manifolds. We have also included an outline of a proof of the orbifold theorem. There is a short account of Gromov’s theory of limits

of metric spaces (as re-interpreted using ϵ -approximations by Thurston), and a discussion of deformations of hyperbolic structures.

This memoir should provide the background necessary to understand the proof of the orbifold theorem which will appear elsewhere. The prerequisites include some acquaintance with hyperbolic geometry, differential geometry and 3-manifold topology.

We have borrowed heavily (in the sense of cut and paste, as well as other ways) from Hodgson's thesis [43] and also from his notes on the orbifold theorem [44].

The Orbifold Theorem is due to Thurston. It is one part of a major program to geometrize topology in dimensions two and three. Unfortunately it was also one of the least documented results. Thurston outlined his proof in graduate courses at Princeton in 1982 and again in 1984. Two of the authors attended these lectures. Our original intent was to write out in detail Thurston's proof. In the course of doing this, we found it easier to develop a somewhat different proof. However our proof is closely based on Thurston's. Our main contributions are the generalized Bieberbach-Soul theorem for non-compact Euclidean cone-manifolds in all dimensions, and the more combinatorial approach to the collapsing case. Our treatment of the Euclidean/spherical transition when there are vertices is also very different to Thurston's treatment. We thank Bill Thurston for creating and sharing his wonderful ideas. His presence at the Tokyo workshop helped ensure its success.

We thank Penny Wightwick for her assistance in preparing the diagrams.

Contents

Preface	i
Introduction	1
1 Geometric Structures	5
1.1 Geometry of surfaces	6
1.2 Geometry of 3-manifolds	7
1.3 Thurston's eight geometries	8
1.4 Developing map and holonomy	9
1.5 Evidence for the Geometrization Conjecture	11
1.6 Geometric structures on 3-manifolds with symmetry	14
1.7 Some 3-manifolds with symmetry	15
1.8 3-manifolds as branched coverings	18
2 Orbifolds	21
2.1 Orbifold definitions	21
2.2 Local structure	24
2.3 Orbifold coverings	27
2.4 Orbifold Euler characteristic	29
2.5 Geometric structures on orbifolds	30
2.6 Some geometric 3-orbifolds	34
2.7 Orbifold fibrations	37
2.8 Orbifold Seifert fibre spaces	39
2.9 Suborbifolds	46
2.10 Spherical decomposition for orbifolds	47
2.11 Euclidean decomposition for orbifolds	48
2.12 Graph orbifolds	49
2.13 The Orbifold Theorem	50

3	Cone-Manifolds	53
3.1	Definitions	53
3.2	Local structure	59
3.3	Standard cone neighbourhoods	59
3.4	Geodesics	61
3.5	Exponential map	64
3.6	Dirichlet domains	66
3.7	Area and volume of cone-manifolds	71
4	Two-dimensional Cone-Manifolds	75
4.1	Developing map and holonomy	75
4.2	Two-dimensional spherical cone-manifolds	76
4.3	Two-dimensional euclidean cone-manifolds	79
4.4	Euclidean examples with large cone angles	83
4.5	Spaces of cone-manifold structures	84
4.6	Two-dimensional hyperbolic cone-manifolds	84
5	Deformations of Hyperbolic Structures	89
5.1	Introduction	89
5.2	Deformations and degenerations of surfaces	90
5.3	General deformation theory	92
5.4	Deforming hyperbolic cone-manifolds	94
5.5	Representation spaces	97
5.6	Hyperbolic Dehn filling	100
5.7	Dehn surgery on the figure eight knot	102
6	Limits of Metrics Spaces	107
6.1	ϵ -approximations	108
6.2	Limits with basepoints	110
6.3	Gromov's compactness theorem	113
6.4	Limits of hyperbolic cone-manifolds	114
6.5	Bilipschitz convergence	117
6.6	Convergence of holonomy	118
7	Proof of the Orbifold Theorem	119
7.1	Topological preliminaries	119
7.2	Deforming hyperbolic structures	122
7.3	Controlling degenerations	128
7.4	Euclidean to spherical transition	137
7.5	Analysis of the thin part	140

CONTENTS

v

7.6 Outline of the Collapsing Theorem	144
Background on the Orbifold Theorem	151
References	155
Index	163

Introduction

The theory of manifolds of dimension three is very different from that of other dimensions. On the one hand we do not even have a conjectural list of all 3-manifolds. On the other hand, if Thurston's Geometrization Conjecture is true, then we have a very good structure theory.

The topology of compact surfaces is well understood. There is a well known topological classification theorem, based on a short list of easily computable topological invariants: orientability, number of boundary components and Euler characteristic. For closed surfaces (compact with no boundary) the fundamental group is a complete invariant. The geometry of surfaces is also well understood. Every closed surface admits a metric of constant curvature. Those with curvature $+1$ are called spherical, or elliptic, and comprise the sphere and projective plane. Those with curvature 0 are Euclidean and comprise the torus and Klein bottle. The remainder all admit a metric of curvature -1 and are called hyperbolic. The Gauss-Bonnet theorem relates the topology and geometry

$$\int_F K dA = 2\pi\chi(F)$$

where K is the curvature of a metric on the closed surface F of Euler characteristic $\chi(F)$. In particular this implies that the sign of a constant curvature metric is determined by the sign of the Euler characteristic. However in the Euclidean and hyperbolic cases, there are many constant curvature metrics on a given surface. These metrics are parametrized by a point in a Teichmüller space.

The topology of 3-dimensional manifolds is far more complex. At the time of writing there is no complete list of closed 3-manifolds and no *proven* complete set of topological invariants. However if Thurston's Geometrization Conjecture were true, then we would know a complete set of topological invariants. In particular for irreducible atoroidal 3-manifolds, with

the exception of lens spaces, the fundamental group is a complete invariant. However this group, on its own, does not provide a practical method of identifying a 3-manifold. On the other hand, once the geometric structure has been found then there are geometrical invariants which can be practically calculated and completely determine the manifold.

A *geometric structure* on a manifold is a complete, locally homogeneous Riemannian metric: every two points have isometric neighbourhoods. The universal cover of such a manifold is a *homogeneous space* and is thus the quotient of a Lie group by a compact subgroup. In dimension two it is a classical result that every surface admits a geometric structure. There are eight geometries needed for compact 3-manifolds. The connected sum of two geometric three manifolds is usually not geometric. However the Geometrization Conjecture states that every closed 3-manifold can be decomposed (in a way to be described) into geometric pieces.

The first step in the decomposition of orientable 3-manifolds is into irreducible pieces by cutting along *essential* 2-spheres and capping off the resulting boundaries by attaching 3-balls. This theory was worked out by Kneser and refined by Milnor. For 3-dimensional manifolds the irreducible pieces obtained are unique. The corresponding statement in higher dimensions is false. Some important classes of 3-manifolds which were studied early on include the following:

- The quotient of the 3-sphere by a finite group of isometries acting freely (a spherical space form). These include the lens spaces (quotients of the round 3-sphere by a cyclic group of isometries) which provide the only known examples of distinct irreducible, atoroidal 3-manifolds with the same fundamental group. The famous Poincaré homology 3-sphere is the quotient of the 3-sphere by the binary icosahedral group (the double cover in $SU(2)$ of the icosahedral subgroup of $SO(3)$.)
- The Seifert fibre spaces. These are compact 3-manifolds which can be foliated by circles and were classified by Seifert. A special case is a *circle bundle* over a closed surface F . If F and the total space M are both orientable this bundle is determined by its Euler class $e \in \mathbb{Z}$. In general, the quotient space obtained by collapsing each circle to a point is a two dimensional *orbifold*. All Seifert fibre spaces have a geometric structure.
- The 10 Euclidean 3-manifolds fit into the general theory of flat manifolds developed by Bieberbach. Bieberbach showed that a compact Euclidean manifold of dimension n is finitely covered by an n -torus. Bieberbach's results also apply to Euclidean orbifolds, producing the 219 types of 3-

dimensional crystallographic groups known to chemists.

- The *mapping cylinder* construction produces an n -manifold M from any automorphism θ of an $(n - 1)$ -manifold F as the quotient $M = F \times [0, 1]/(x, 1) \equiv (\theta(x), 0)$. In the case that F is a 2-torus, the automorphism is determined up to isotopy by an element of the group $GL(2, \mathbb{Z})$. These give 3-manifolds with the *Solv*, *Nil* and Euclidean geometries. When the genus of F is more than 1 there is a (possibly trivial) torus decomposition into geometric pieces.
- A *Haken* manifold, M , is a compact, irreducible 3-manifold which contains a closed embedded surface with infinite fundamental group that injects under the map induced by inclusion into the fundamental group of M . Haken manifolds include many important classes of 3-manifolds, and a great deal is now known about these manifolds through the work of Haken, Waldhausen, Thurston and many others. In particular they have geometric decompositions. However, Hatcher [38] showed that all but finitely many Dehn surgeries on a knot give a non-Haken manifold. More recently, Cooper and Long [20] showed that all but finitely many such fillings give a 3-manifold containing an essential *immersed* surface.

The next step in the classification program is to decompose along *essential* embedded tori. The JSJ decomposition (of Jaco-Shalen and Johannson) gives a canonical splitting of a compact 3-manifold by cutting out a maximal Seifert fibred piece.

Thurston [80] introduced the idea of “hyperbolic Dehn surgery” which is a method of *continuously* changing one 3-manifold into another with a different topology. The intermediate spaces are **cone-manifolds** with a hyperbolic metric everywhere except along a knot or link called the *singular locus*. The set of manifolds form a discrete subset, contained in the larger subset of *orbifolds*. This method of continuously changing topology and geometry only works in dimension three. The computer program SnapPea developed by Jeff Weeks [88] allows one to put this philosophy into practice. Many insights and theorems have developed from this point of view.

Roughly speaking an **orbifold** is the quotient of a manifold by a finite group of diffeomorphisms. Actually an orbifold has the *local structure* of such a space. It is the natural object to consider when one is studying *discrete symmetry groups*. Compact two dimensional orbifolds are classified in a similar way to surfaces, using an orbifold version of Euler characteristic. This classification encompasses the classification of the regular solids (finite subgroups of the orthogonal group $O(3)$), the classification of the 17 wallpa-

per groups, and of periodic tessellations of the hyperbolic plane. There are, however, four families of *bad* or *non-geometric* two-dimensional orbifolds that do not arise globally as the quotient of a manifold by a finite group. However they do arise quite naturally as the base-orbifolds of certain Seifert fibrations. In fact the base orbifold of a Seifert fibration is *bad* if and only if the fibration is not isotopic to one with the fibres geodesic in a geometric structure on the Seifert fibre space.

The *Orbifold Theorem* characterizes when a 3-dimensional orbifold with 1-dimensional singular locus has a *geometric structure*, in other words, when it is the quotient of a homogeneous space by a discrete group of isometries. This theorem has many consequences, for example an irreducible, atoroidal, closed orientable 3-manifold which admits a symmetry with 1-dimensional fixed set is geometric. It follows that all 3-manifolds of Heegaard genus two have a geometric decomposition.

Chapter 1

Geometric Structures

The aim of this memoir is to give an introduction to the statement and main ideas in the proof of the “Orbifold Theorem” announced by Thurston in late 1981 ([83], [81]). The Orbifold Theorem shows the existence of geometric structures on many 3-dimensional orbifolds, and on 3-manifolds with a kind of topological symmetry.

The main result implies a special case of the following Geometrization Conjecture proposed by Thurston in 1976 as a framework for the classification of 3-manifolds. For simplicity, we state the conjecture only for compact, orientable 3-manifolds.

Conjecture 1.1 (Geometrization Conjecture). ([81]) *The interior of every compact 3-manifold has a canonical decomposition into pieces having a geometric structure.*

The kinds of *decomposition* needed are:

1. prime (or connected sum) decomposition, which involves cutting along separating 2-spheres and capping off the pieces by gluing on balls.
2. torus decomposition, which involves cutting along certain incompressible non-boundary parallel tori.


The meaning of *canonical* is that the pieces obtained are unique up to ordering and homeomorphism. The spheres used in the decomposition are not unique up to isotopy, but the tori are unique up to isotopy.

A *geometric structure* on a manifold is a complete Riemannian metric which is locally homogeneous (i.e. any two points have isometric neighbourhoods). A *geometric decomposition* is a decomposition of this type into

pieces whose interior have a geometric structure. There are essentially eight kinds of geometry needed; of these hyperbolic geometry is the most common and the most interesting.

1.1 Geometry of surfaces

We would like to generalize the well-known topological classification of closed 2-manifolds (surfaces). The *orientable* surfaces are just:



The *non-orientable* closed surfaces (those containing Möbius strips) are: the real projective plane $P = \mathbb{R}P^2$, the Klein bottle $K = P\#P, P\#P\#P, \dots$. (See [58] or [4] for details.)

These surfaces are easy to distinguish by their orientability and Euler characteristic given by

$$\chi = \#(\text{vertices}) - \#(\text{edges}) + \#(\text{faces}),$$


for any decomposition of the surface into polygons.

It has been known since the nineteenth century that there is a very close relationship between geometry and topology in dimension two. Each surface can be given a spherical, Euclidean or hyperbolic structure, that is, a Riemannian metric of constant curvature $K = +1, 0$, or -1 . Further, the topology of the surface is determined by the geometry via the Gauss-Bonnet formula:

$$2\pi\chi(M) = \int_M K dA.$$

Exercise 1.2. Prove this formula for constant curvature surfaces, using the fact that the angle sum of a (geodesic) triangle in a space of constant curvature K is $\pi + KA$, where A is the area of the triangle.

Example 1.3. A 2-sphere clearly has a spherical metric — just take the round sphere S^2 in Euclidean 3-space \mathbb{E}^3 . A torus can be given a Euclidean metric: take a square (or any parallelogram) in the Euclidean plane \mathbb{E}^2 and glue together opposite edges. (Note that the corners fit together to form a small Euclidean disk since the angles of the parallelogram add up to 2π .) Similarly, a closed surface of genus $g \geq 2$ can be given a hyperbolic metric by taking a regular $4g$ -gon in the hyperbolic plane \mathbb{H}^2 with angles $2\pi/4g$ and gluing together edges in pairs in the usual combinatorial pattern.



Exercise 1.4. Show that there exists a regular $4g$ -gon in the hyperbolic plane \mathbb{H}^2 with angles $2\pi/4g$ for each $g \geq 2$.

1.2 Geometry of 3-manifolds

It now seems natural to ask whether there is a similar division of 3-manifolds into different geometric types, but this question was not considered until the work of Thurston starting in about 1976.

Question: What kinds of geometries are needed to deal with 3-manifolds?

We would like to find geometric structures (or metrics) on 3-manifolds which are *locally homogeneous*. Roughly, this means the space should look locally the same near every point; more precisely: any two points have isometric neighbourhoods. Our spaces should also be *complete* as metric spaces, i.e. every Cauchy sequence converges. Intuitively, this means you can't fall off the edge of the space after going a finite distance!

First, we obviously have 3-dimensional spaces of constant curvature: Euclidean geometry \mathbb{E}^3 , spherical geometry S^3 and hyperbolic geometry \mathbb{H}^3 . These geometries look the same near every point and in every direc-

tion. Many examples of 3-manifolds with these geometries are discussed in Thurston's book [82].

Not all closed 3-manifolds can be modelled on the constant curvature geometries. For example, the universal cover of a 3-manifold with one of these geometries is topologically either \mathbb{R}^3 or S^3 . So $S^2 \times S^1$, with universal cover $S^2 \times \mathbb{R}$, cannot have such a geometry. Nevertheless, it does have a very nice homogeneous metric: take the natural product metric on $S^2 \times S^1$.

Formally, we say that a manifold M has a *geometric structure* if it admits a complete, locally homogeneous Riemannian metric. This gives a way of measuring the length of smooth curves by integrating an element of arc length ds , and we can talk about geodesics, angles, volume etc. Then the universal cover X of M has a complete *homogeneous* metric, i.e. the isometry group G acts transitively on X . Further, the stabilizer $G_x = \{g \in G : gx = x\}$ of each point $x \in X$ is *compact*, since it is a closed subgroup of $O(n)$. Then the manifold M is isometric to a quotient space X/Γ , where Γ is a discrete subgroup of G . (This can be proved by an “analytic continuation” argument using the “developing map” discussed below, see also [80]).

1.3 Thurston's eight geometries

Following the viewpoint of Klein's Erlangen program (from 1872), we can also regard *geometry* as the study of the properties of a space X which are invariant under a group of transformations G . The geometry (G, X) is *homogeneous* if G acts transitively on X . The geometry is *analytic* if each transformation in G is uniquely determined by its restriction to any non-empty open subset of X . For example, groups of isometries of manifolds are analytic.


The geometries needed for studying 3-manifolds are pairs (G, X) where X is a simply connected space, and G is a group acting transitively on X with compact point stabilizers. To avoid redundancy, we require that G is a *maximal* such group. Finally, we restrict to geometries which can model compact 3-manifolds: G contains a discrete subgroup Γ such that X/Γ is compact.

Thurston showed that there are *exactly eight* such geometries on 3-manifolds. The most familiar 3-dimensional geometries are the constant curvature geometries: Euclidean geometry \mathbb{E}^3 (of constant curvature $K = 0$), spherical geometry S^3 (of constant curvature $K = +1$), and hyperbolic geometry \mathbb{H}^3 (of constant curvature $K = -1$). The other geometries are the

product geometries $S^2 \times \mathbb{E}^1$, $\mathbb{H}^2 \times \mathbb{E}^1$; and three “twisted products” called *Nil*, *PSL* and *Solv* geometries. (See [73], [81] for detailed discussions of these geometries.)

1.4 Developing map and holonomy

A (G, X) *geometric structure* on a manifold M is given by a covering of M by open sets U_i and diffeomorphisms $\phi_i : U_i \rightarrow X$ to open subsets of X , giving coordinate charts on M , such that all the transition maps are restrictions of elements in G . (If G acts by isometries on X , this means that M is locally isometric to X .)




Given an analytic (G, X) structure on M , analytic continuation of coordinate charts gives a “global coordinate chart”, called a *developing map*

$$\text{dev} : \tilde{M} \rightarrow X$$

defined on the universal cover \tilde{M} of M .

This is constructed as follows: Begin with an embedding $\phi_1 : U_1 \subset M \rightarrow X$ giving a coordinate chart on M . If $\phi_2 : U_2 \rightarrow X$ is another coordinate chart with $U_1 \cap U_2$ connected and non-empty, there is a unique $g \in G$ such that $g \circ \phi_2 = \phi_1$ on $U_1 \cap U_2$. So ϕ_1 extends to a map $\phi : U_1 \cup U_2 \rightarrow X$, with $\phi = \phi_1$ on U_1 and $\phi = g \circ \phi_2$ on U_2 . In this way, we can extend ϕ_1 by analytic continuation along paths in M .



Since the result of the analytic continuation only depends on the homotopy class of the path involved, we obtain a well defined map $\text{dev} : \tilde{M} \rightarrow X$. Then dev is a local diffeomorphism satisfying the equivariance condition

$$\text{dev} \circ \gamma = h(\gamma) \circ \text{dev}$$

for each deck transformation γ in $\pi_1(M)$, where

$$h : \pi_1(M) \rightarrow G$$

is a homomorphism called the *holonomy representation* for the geometric structure. (See Thurston [82, Chapter 3] for more details.)

Note that dev and h are not uniquely defined; changing the original coordinate chart ϕ_1 by an element $g \in G$ gives a new developing map $g \circ \text{dev}$ with corresponding holonomy representation $g \circ h \circ g^{-1}$. It can be shown that the pair (dev, h) determines the (G, X) -structure on M .

As a simple example, let M be the Euclidean surface obtained as follows. Let D the subset of the Euclidean plane bounded by two distinct rays starting at a point x and making an angle θ . Identify the two sides of D by an isometry and delete the point x . The resulting surface is not complete. The image of the holonomy is discrete if and only if θ is a rational multiple of π . The developing map has image the complement of x and is not injective.

The developing map is an important tool for analyzing (G, X) structures. For instance, it can be used to prove the following important completeness criterion (see [82]).

Theorem 1.5. *Let M be a manifold with a geometric structure modelled on a geometry (G, X) where G is a group of isometries of X . Then M is complete as a metric space if and only if the developing map $\text{dev} : \tilde{M} \rightarrow X$ is a covering map.*

If X is simply connected, then such a complete manifold M is isometric to X/Γ where the holonomy group $\Gamma = h(\pi_1 M) \cong \pi_1(M)$ is a discrete subgroup of $\text{Isom}(X)$ which acts freely and properly discontinuously on X .

The proof of the Orbifold Theorem involves deformations through *incomplete* structures. This means that discrete group techniques cannot be used; however, the developing map and holonomy again play a key role.

1.5 Evidence for the Geometrization Conjecture

We begin by restating Thurston's Geometrization Conjecture. For simplicity, we will assume that all manifolds are *compact* and *orientable*.


Conjecture 1.6 (Geometrization Conjecture). *Let M be a compact, orientable, prime 3-manifold. Then there is a finite collection of disjoint, embedded incompressible tori in M (given by the Johannson, Jaco-Shalen torus decomposition), so that each component of the complement admits a geometric structure modelled on one of the eight geometries discussed in section 1.3.*

There is a great deal of evidence for this conjecture. Here are some of the main results.


Theorem 1.7 (Thurston). *The Geometrization Conjecture is true if M is a Haken manifold.*

Recall that M is *Haken* if it is irreducible and contains an incompressible surface. For example, this theorem applies if M is irreducible and ∂M contains a surface $\not\cong S^2$. This result is proved by a difficult argument using hierarchies; Thurston developed many wonderful new geometric ideas and techniques to carry this out. The article of Morgan in [64] provides a good overview of the proof.

An important application of this result is to knot complements. Let K be a knot (i.e. an embedded circle) in $S^3 = \mathbb{R}^3 \cup \infty$. Then K is called a *torus knot* if it can be placed on the surface of a standard torus. It is easy to see that $S^3 - K$ is then a Seifert fibre space with ≤ 2 exceptional fibres (see section 2.8).



A knot K' is called a *satellite knot* if it is obtained by taking a non-trivial embedding of a circle in a small solid torus neighbourhood of a knot K . (*Non-trivial* means that the circle is not contained in a 3-ball, and is not isotopic to K). Then $S^3 - K'$ contains an incompressible torus T which is the boundary of the solid torus around K . (This follows since the exterior of every non-trivial knot has incompressible boundary, by the loop theorem.)



Corollary 1.8. *Let K be a knot in S^3 . Then $S^3 - K$ has a geometric structure if and only if K is not a satellite knot. Further, $S^3 - K$ has a hyperbolic structure if and only if K is not a satellite knot or a torus knot.*

Thus, “most” knot complements are hyperbolic. Similarly, “most” link complements are hyperbolic.

The unresolved cases of the Geometrization Conjecture are for closed, orientable, irreducible 3-manifolds M which are non-Haken. (It is known that this is a large collection of 3-manifolds!) Such manifolds fall into 3 categories:

1. Manifolds with $\pi_1(M)$ finite.
2. Manifolds with $\pi_1(M)$ infinite, and containing a $\mathbb{Z} \times \mathbb{Z}$ subgroup.
3. Manifolds with $\pi_1(M)$ infinite, and containing no $\mathbb{Z} \times \mathbb{Z}$ subgroup.

For manifolds of type (1), the Geometrization Conjecture reduces to the following:

Conjecture 1.9 (Orthogonalization Conjecture). *If $\pi_1(M)$ is finite, then M is spherical; thus M is homeomorphic to S^3/Γ where Γ is a finite subgroup of $O(4)$ which acts on S^3 without fixed points.*

(This includes the Poincaré Conjecture as the special case where $\pi_1(M)$ is trivial.)


For manifolds of type (2), the recent work of [30] and [15] shows that the manifolds are either *Seifert fibre spaces* (see 2.8) or Haken; hence the Geometrization Conjecture holds.

For manifolds of type (3), the Geometrization Conjecture reduces to the following:

Conjecture 1.10 (Hyperbolization Conjecture). *If M is irreducible with $\pi_1(M)$ infinite and containing no $\mathbb{Z} \times \mathbb{Z}$ subgroup, then M is hyperbolic.*

Further important evidence for the hyperbolization conjecture is provided by Thurston’s “Hyperbolic Dehn surgery theorem”, to be discussed in more detail in 5.6 below.

Definition 1.11. First we recall the construction of 3-manifolds by *Dehn filling*. Let M be a compact 3-manifold with boundary consisting of one or more tori T_1, \dots, T_k . We can then form *closed* 3-manifolds by attaching solid tori to T_1, \dots, T_k using arbitrary diffeomorphisms between their boundary tori. These 3-manifolds are said to be obtained from M by *Dehn filling*. The resulting manifold depends only on the isotopy classes of the *surgery curves* $\gamma_i \subset T_i$ which bound discs in the added solid tori; we denote it $M(\gamma_1, \dots, \gamma_k)$. If L is any link in S^3 , we can apply this construction to the link exterior M , obtained by removing an open tubular neighbourhood of L from S^3 . Then we say that the resulting 3-manifold is obtained from S^3 by *Dehn surgery* along L .



Theorem 1.12. (Lickorish [56], Wallace [87]). *Every closed, orientable 3-manifold can be obtained by Dehn surgery along some link in S^3 .*

Before stating Thurston's result, we note that an orientable hyperbolic 3-manifold which has *finite volume* but is non-compact is homeomorphic to the interior of a compact 3-manifold \overline{M} which is compact with boundary $\partial\overline{M}$ consisting of tori. We then call M a *cusped* hyperbolic manifold, and we write $M(\gamma_1, \dots, \gamma_k)$ for the manifolds obtained by Dehn filling on \overline{M} .

Theorem 1.13 (Hyperbolic Dehn Surgery Theorem). *If M is a cusped hyperbolic 3-manifold, then “almost all” manifolds obtained from M by Dehn filling are hyperbolic. (More precisely, only a finite number of surgeries must be excluded for each cusp.)*

In particular, if M has one cusp, then the Dehn filled manifolds $M(\gamma)$ are closed hyperbolic manifolds for all but a finite number of isotopy classes of surgery curves γ .

Since every closed 3-manifold can be obtained by Dehn filling from a hyperbolic link complement ([56], [87], [65]), this shows that in some sense “most” closed 3-manifolds are hyperbolic! (However, it is not currently known how to make this into a precise statement.)

In fact the number of non-hyperbolic surgeries is usually very small. The worst known case is the figure knot complement which has 10 non-hyperbolic surgeries (see section 5.7).

The Geometrization Conjecture, and the special cases proved so far, has had a profound effect on 3-manifold topology, including major roles in the solution of several old conjectures (e.g. The Smith Conjecture, see [64].)

The existence of a geometric structure on a given manifold provides a great deal of information about that manifold. For example, its fundamental group is residually finite, so has a solvable word problem. For hyperbolic 3-manifolds, the Mostow rigidity theorem shows that the hyperbolic structure is unique; thus geometric invariants of hyperbolic 3-manifolds are actually topological invariants. This provides very powerful tools for understanding 3-dimensional topology.

1.6 Geometric structures on 3-manifolds with symmetry

In late 1981, Thurston announced that the Geometrization Conjecture holds for 3-manifolds with a kind of topological symmetry (see also Thurston's Theorem A stated on page 153 of the appendix.)

Theorem 1.14 (Symmetry Theorem). (Thurston [83]) *Let M be an orientable, irreducible, closed 3-manifold. Suppose M admits an action by a finite group G of orientation preserving diffeomorphisms such that some non-trivial element has a fixed point set of dimension one. Then M admits a geometric decomposition preserved by the group action.*

Later, we'll state a more general version in terms of *orbifolds* (see section 2.13 and Thurston's Theorem B stated on page 153).

This theorem has many applications to the study of group actions on 3-manifolds and to the existence of geometric structures on 3-manifolds. Here is an important special case.


Theorem 1.15. *Assume M and G are as in the symmetry theorem above. If M contains no incompressible tori, then M has a geometric structure such that this action of the group G is by isometries. In particular, the fixed point set of each group element is totally geodesic.*

Taking $M = S^3$ and the group to be cyclic gives the **Smith Conjecture**: If ϕ is a periodic, non-free, orientation preserving diffeomorphism of S^3 then ϕ is conjugate to a rotation. In particular, the fixed point set of ϕ is an unknotted circle.

1.7 Some 3-manifolds with symmetry


Next we give some applications to 3-manifolds, constructing various classes of manifolds with symmetry to which the main theorem can be applied. Rolfsen's book [69] is an excellent general reference for these constructions.

Every orientable 3-manifold has a *Heegaard decomposition*: a representation as the union of two handlebodies glued together along their boundaries.




Glue boundaries together by a homeomorphism

Proof. Triangulate the manifold and take a regular neighbourhood of the 1-skeleton as one handlebody, and its complement (a neighbourhood of the dual 1-skeleton) as the other handlebody. \square



Example 1.16. The only manifold with a genus 0 Heegaard decomposition is S^3 , while the only manifolds with genus 1 Heegaard decompositions are the lens spaces and $S^2 \times S^1$. The manifolds S^3 and $S^2 \times S^1$ are clearly geometric. Every lens space $L(p, q)$ is the quotient of $S^3 = \{(z_1, z_2) \in \mathbb{C}^2 : |z_1|^2 + |z_2|^2 = 1\}$ by the cyclic group \mathbb{Z}_p of isometries generated by $(z_1, z_2) \mapsto (tz_1, t^q z_2)$, where $t = \exp(2\pi i/p)$. So $L(p, q)$ has a spherical structure.

Example 1.17. Heegaard decompositions of genus 2 have a certain kind of 2-fold symmetry which can be seen as follows. A surface F_2 of genus two has a special “hyper-elliptic” involution τ which takes every simple closed curve to an isotopic curve (possibly reversing orientation).



Using Lickorish’s result that every diffeomorphism is isotopic to a product of Dehn twists around simple closed curves [56], we obtain:

Proposition 1.18. *Every diffeomorphism $\phi : F_2 \rightarrow F_2$ is isotopic to ϕ' such that $\phi'\tau = \tau\phi'$.*

Extending τ over each handlebody and adjusting the gluing map by this proposition, we obtain an involution on any manifold M with a genus 2 Heegaard decomposition. (The fixed point set is 1-dimensional, and covers a 3-bridge knot in the quotient $M/\tau = S^3$.)


Corollary 1.19. *Every 3-manifold of Heegaard genus two admits a geometric decomposition.*

Proof. If the genus 2 manifold is irreducible then theorem 1.14 applies directly. Otherwise, a result of Haken [34] shows that the manifold is a connected sum of two genus one manifolds, which are geometric by example 1.16. \square

Next we consider the construction of 3-manifolds by *Dehn surgery* as described in 1.11. If a link in S^3 has a suitable kind of symmetry, the manifolds obtained by Dehn surgery on the link often exhibit a similar kind of symmetry.

Example 1.20. The *two bridge* (or *rational*) knots and links in S^3 are the links which can be cut by a 2-sphere into two pairs of unknotted, unlinked arcs. Alternatively, these are the links which have a projection with exactly 2 local maxima and 2 local minima. These have been extensively studied and classified by Schubert [72].

Each 2-bridge link has a 2-fold symmetry: it can be arranged in \mathbb{R}^3 so that it is invariant under a 180 degree rotation. Further, this involution always extends over the solid torus added in Dehn surgery.



Corollary 1.21. *Every Dehn surgery on a 2-bridge knot in S^3 gives a manifold admitting a geometric decomposition.*

Proof. If the Dehn surgered manifold M is irreducible this follows from the symmetry theorem. Otherwise, note that the Heegaard genus of M is at most 2 (**Exercise**). So theorem 1.19 applies. \square

A similar argument proves

Corollary 1.22. *Every Dehn filling on a bundle over the circle with once-punctured torus as fibre admits a geometric decomposition.*


Proof. Such a manifold again has a suitable 2-fold symmetry. (**Exercise**.)

1.8 3-manifolds as branched coverings


Finally we consider the construction of 3-manifolds by *branched coverings*.

Let M and N be 3-manifolds. A map $f : M \rightarrow N$ is a *branched covering*, *branched over* $L \subset N$, if

1. the restriction $f : M - f^{-1}(L) \rightarrow N - L$ is a covering, and
2. any point $x \in f^{-1}(L)$ has a neighbourhood homeomorphic to $D \times I$, where D is the unit disc in \mathbb{C} , on which f has the form $f : D \times I \rightarrow D \times I, (z, t) \mapsto (z^n, t)$, for some integer $n \geq 2$.



Example 1.23. A solid torus is the 2-fold branched covering of a 3-ball branched over two unknotted, unlinked arcs as shown below.



Theorem 1.24. (Alexander [2]). *Every closed orientable 3-manifold can be obtained as a branched covering of S^3 branched along a link.*

In fact, one can always take the branching set in S^3 to be the figure eight knot by results of Hilden et. al. [42].

A branched covering f is *regular* if the covering transformations act transitively on the fibres $f^{-1}(x)$; then N is the quotient of M by the group of covering transformations. (The branched coverings in the above theorems are typically irregular.)

Theorem 1.25. *Any regular branched covering space, M , of S^3 branched over a knot or link L has a geometric decomposition.*

Proof. If M is irreducible then this follows from the symmetry theorem. Otherwise it follows from the Orbifold Theorem (see 2.58 below). Let Q be the quotient orbifold with underlying space S^3 and singular locus L ; each component of L is labelled with the branch index on that component. It is an easy exercise to check the hypotheses of the Orbifold Theorem in this setting. Thus Q has a geometric decomposition. Lifting this gives a geometric decomposition of M . \square

For example, any manifold of Heegaard genus two is a (regular) 2-fold covering of S^3 branched over a 3-bridge knot or link (see Example 1.17).

A knot K in S^3 is *prime* if there is no 2-sphere which meets K transversally in exactly two points and separates K into two non-trivial knotted arcs. A *Conway sphere* is an incompressible 4-punctured sphere in the exterior of the knot whose boundary consists of four meridians of the knot. Now let M be the n -fold cyclic branched cover of S^3 , branched along the knot K . It follows from the Orbifold Theorem that M is prime if and only if K is prime. Also if M contains an essential torus then either K is a torus knot or satellite knot, or else there is a Conway sphere and covering degree is 2.


Corollary 1.26. *Let K be any knot in S^3 which is not a torus knot or a satellite knot. Then the n -fold cyclic branched cover of S^3 over K is hyperbolic for all $n \geq 3$, except for the 3-fold cover of the figure eight knot (which has a Euclidean structure).*

Proof. The preceding remarks show that the branched cover is irreducible and atoroidal hence has a geometric structure by the symmetry theorem. Dunbar has classified those orbifolds with non-hyperbolic geometric structures and underlying space the three sphere. This provides the one exception noted in the theorem. (See Dunbar [26], Bonahon-Siebenmann [9],[10].) \square


Remark 1.27. The condition on K holds if and only if $S^3 - K$ has a hyperbolic structure.

In contrast, the 2-fold branched coverings of many knots do not admit hyperbolic structures. For example, each 2-bridge knot has a 2-fold branched cover which is a lens space, so has a spherical structure. (The branched covering is obtained by gluing together two copies of the example in 1.23.)

Definition 1.28. A *Montesinos knot or link* is a link L contained in an unknotted solid torus V in S^3 such that there are disjoint meridional disks D_1, \dots, D_n in V which separate V into components whose closures are balls and with the property that L intersects each such ball in a pair of unknotted arcs. In addition it is required that the boundary of V be incompressible in $V - L$; in other words, every meridional disc of V intersects L . The 2-fold cover of S^3 branched over a Montesinos knot or link is a Seifert fibre space.



For knots with non-trivial Conway decomposition the 2-fold cover contains an essential torus.



Chapter 2

Orbifolds

An *orbifold* is a space locally modelled on the quotients of Euclidean space by finite groups. Further, these local models are glued together by maps compatible with the finite group actions. Natural examples are obtained from the quotient of a manifold by a finite group, but not all orbifolds arise in this way.

Given any idea in 3-dimensional topology, one can try to quotient out (locally) by finite group actions. This leads to orbifold versions of concepts including: covering space, fundamental group, submanifold, incompressible surface, prime decomposition, torus decomposition, bundle, Seifert fibre space, geometric structure. This chapter will review some of these basic concepts for orbifolds, and state the main result in terms of orbifolds. See Thurston [84], Scott [73] and Bonahon-Siebenmann [11] for more details.

2.1 Orbifold definitions

Formally, a (smooth) orbifold consists of local models glued together with orbifold maps. A *local model* is a pair (\tilde{U}, G) where \tilde{U} is an open subset of \mathbb{R}^n and G is a finite group of diffeomorphisms of \tilde{U} . We will abuse this terminology by saying the quotient space $U = \tilde{U}/G$ is the local model. An *orbifold map* between local models is a pair $(\tilde{\psi}, \gamma)$ where $\tilde{\psi} : \tilde{U} \rightarrow \tilde{U}'$ is smooth, $\gamma : G \rightarrow G'$ is a group homomorphism and $\tilde{\psi}$ is *equivariant*, that is $\tilde{\psi}(g\tilde{x}) = \gamma(g)\tilde{\psi}(\tilde{x})$ for all $g \in G$ and $\tilde{x} \in \tilde{U}$. Then $\tilde{\psi}$ induces a map $\psi : \tilde{U}/G \rightarrow \tilde{U}'/G'$ and we will abuse terminology by saying it is an orbifold map. If γ is a monomorphism and $\tilde{\psi}, \psi$ are both injective we say that ψ is an *orbifold local isomorphism*.

An *n-dimensional orbifold* Q consists of a pair (X_Q, \mathcal{U}) where X_Q is the

underlying space which is a Hausdorff, paracompact, topological space and \mathcal{U} is an *orbifold atlas*. Sometimes we will abuse notation by using Q to denote the underlying space. The atlas consists of a collection of *coordinate charts* (U_i, ϕ_i) where the sets U_i are an open cover of Q such that the intersection of any pair of sets in this cover is also in the cover. For each chart there is a local model \tilde{U}_i/G_i and a homeomorphism $\phi_i : U_i \rightarrow \tilde{U}_i/G_i$. These charts must satisfy the *compatibility condition* that whenever $U_i \subset U_j$ the inclusion map is an orbifold local isomorphism.

The orbifold atlas is not an intrinsic part of the structure of an orbifold; two atlases define the *same* orbifold structure if they are *compatible*: if there is an atlas which contains both of them.


The *local group* G_x at a point x in a local model \tilde{U}/G is the isotropy group of any point $\tilde{x} \in \tilde{U}$ projecting to x . This is well defined up to conjugacy. The *singular locus* $\Sigma(Q)$ of Q is $\{x \in X_Q : G_x \neq \{1\}\}$. An orbifold is a *manifold* if all local groups are trivial, i.e. $\Sigma(Q)$ is empty.

An orbifold is *locally orientable* if it has an atlas (U_i, ϕ_i) where each local model is a quotient $U_i = \tilde{U}_i/G_i$ by an orientation preserving group G_i . It is *orientable* if, in addition, the inclusion maps $U_i \subset U_j$ are induced by orientation preserving maps $\tilde{U}_i \rightarrow \tilde{U}_j$.

To describe an orbifold Q pictorially, we show the underlying space X_Q , mark the singular locus Σ , and label each point x of Σ by its local group G_x . Usually we just use a label n to denote a cyclic group of rotations of order n .


Many examples of orbifolds are provided by quotient spaces M/G where G is a finite group of diffeomorphisms of M , or more generally G is a group acting properly discontinuously on M .

Example 2.1. A Euclidean torus has a symmetry which is a rotation of order 2 with 4 fixed points. The quotient is a *pillowcase* P ; the underlying space is a 2-sphere with 4 singular points where the local group is \mathbb{Z}_2 .




Alternatively, tessellate the plane with parallelograms. Define Γ to be the group generated by π -rotations centred at the midpoints of edges. Then

$$P = \mathbb{R}^2/\Gamma.$$



Example 2.2. $\mathbb{R}^n/(\text{reflection})$ gives an orbifold with *mirror* or *reflector* or *silvered* points corresponding to the hyperplane of fixed points.



An *orbifold with boundary* Q is similarly defined by replacing \mathbb{R}^n by the closed half space \mathbb{R}_+^n . The *orbifold boundary* $\partial_{orb}Q$ of Q corresponds to points in the boundary of \mathbb{R}_+^n in the local models; thus a point x is in $\partial_{orb}Q$ if there is a coordinate chart $\phi : U \rightarrow \tilde{U}/G$ with $x \in U$ such that $\phi(x) \in (\tilde{U} \cap \partial\mathbb{R}_+^n)/G$. An orbifold is (*orbifold*) *closed* if it is compact and the orbifold boundary is empty.

Note that the orbifold boundary is generally *not* the same as the manifold boundary of the underlying space. The set of points in the singular locus of an orbifold Q which are locally modelled on the quotient of \mathbb{R}^n by a reflection (as in example 2.2) is called the *mirror singular locus* or *silvered boundary* $\Sigma_{mirror}(Q)$. The boundary of the underlying topological manifold is $\partial_{top}X_Q = \partial_{orb}Q \cup \Sigma_{mirror}(Q)$. A compact orbifold with boundary can be made into a closed orbifold by making the boundary into mirrors.

2.2 Local structure

To avoid local pathologies such as wild fixed point sets (see [7]), we assume that our orbifolds are differentiable: modelled on \mathbb{R}^n modulo finite groups of *diffeomorphisms* rather than homeomorphisms.


Theorem 2.3. *A differentiable n -orbifold is locally modelled on \mathbb{R}^n modulo a finite subgroup G of the orthogonal group $O(n)$. Thus a neighbourhood of a point is a cone on the spherical $(n - 1)$ -orbifold S^{n-1}/G .*

Proof. Given a point p in \mathbb{R}^n and a finite group G of diffeomorphisms of \mathbb{R}^n fixing p we will show there is a G -invariant neighbourhood U of p such that the action of G on U is conjugate to a linear action. There is a Riemannian metric on \mathbb{R}^n invariant under G , obtained by starting with any Riemannian metric and averaging over the finite group G . The exponential map then gives a diffeomorphism between a neighbourhood of zero in $T_p\mathbb{R}^n$ and a neighbourhood U of p in \mathbb{R}^n . This Riemannian metric restricted to $T_p\mathbb{R}^n$ is an inner product. The action of G on $T_p\mathbb{R}^n$ is linear and preserves this inner product, so we may regard G as a subgroup of the group $O(n)$ of isometries of this inner product. The exponential map provides a conjugacy of this action to the action of G on U . \square

This theorem gives us a description of the local structure of orbifolds in low dimensions.

In dimension 1, each point stabilizer G_x is trivial or cyclic of order 2 generated by a reflection; so we have regular points and *mirror* points.

In dimension 2, G_x is a finite subgroup of $O(2)$. Thus either: G_x is a cyclic rotation group \mathbb{Z}_k giving a *cone point* of angle $2\pi/k$, or G_x is a reflection group of order 2 giving *mirror* points, or G_x is a dihedral group D_{2k} of order $2k$ (generated by reflections in two lines meeting at an angle π/k) giving a *corner point*. So the underlying spaces of 2-orbifolds are 2-manifolds with cone points and mirrors along the boundary, meeting at corners.




If F is a surface then $F(n_1, n_2, \dots, n_k)$ will denote the 2-orbifold with underlying space F and with points p_i in the interior of F which are locally modelled on \mathbb{R}^2 modulo a group of rotations of order n_i . A *football* is $S^2(n, n)$, a *teardrop* is $S^2(n)$ with $n > 1$, and a *spindle* is a 2-orbifold $S^2(n, m)$ with $n > m > 1$.


The local structure of orientable 3-orbifolds is determined by the following classical result.

Theorem 2.4. *A finite subgroup G of $SO(3)$ is cyclic, dihedral, or the group of rotational symmetries of a regular solid. The quotient orbifold S^2/G has underlying space S^2 with 2 or 3 cone points and is one of the following:*

- $S^2(n, n)$ if $G =$ cyclic of order n
- $S^2(2, 2, n)$ if $G =$ dihedral of order $2n$
- $S^2(2, 3, 3)$ if $G =$ symmetries of regular tetrahedron
- $S^2(2, 3, 4)$ if $G =$ symmetries of cube or octahedron
- $S^2(2, 3, 5)$ if $G =$ symmetries of icosahedron or dodecahedron.

(A proof is outlined in exercise 2.20 below.)






Since each point in an orientable (or locally orientable) 3-orbifold has a neighbourhood which is a cone on one of these spherical 2-orbifolds, we have

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

This shows that there are many orientable 3-orbifolds; they can be specified by giving a trivalent graph in an orientable 3-manifold, with edges labelled so the above conditions hold at vertices.

Example 2.6. The *double* of a cube gives a Euclidean 3-orbifold whose underlying space is S^3 and singular locus is the 1-skeleton of the cube with all edges labelled 2. Along each edge the local groups are \mathbb{Z}_2 ; at each vertex the local group is $\mathbb{Z}_2 \times \mathbb{Z}_2$ and a neighbourhood of the vertex is a cone on $S^2(2, 2, 2)$.




2.3 Orbifold coverings


An *orbifold covering* $f : \tilde{Q} \rightarrow Q$ is a continuous map $X_{\tilde{Q}} \rightarrow X_Q$ such that each point $x \in X_Q$ has a neighbourhood $U = \tilde{U}/G$ for which each component V_i of $f^{-1}(U)$ is isomorphic to \tilde{U}/G_i , where G_i is a subgroup of G . Further, $f|_{V_i} : V_i \rightarrow U$ corresponds to the natural projection $\tilde{U}/G_i \rightarrow \tilde{U}/G$. A covering is *regular* if the orbifold covering transformations act transitively on each fibre $f^{-1}(x)$.

Example 2.7.

(a) *Branched coverings* give orbifold coverings. For example, a genus 2 handlebody double covers a ball containing 3 unknotted arcs labelled 2. (Compare example 1.23.)



(b) $M/G' \rightarrow M/G$ is a regular orbifold covering, when G' is a subgroup of a properly discontinuous group G . For example $G = \mathbb{Z}_2$ acts on $M = S^1$ by reflection with quotient orbifold an interval with mirrored endpoints $I(2,2)$.



Example 2.8. The last example generalizes as follows. Suppose that Q is an orbifold with mirror singular locus $\Sigma_{\text{mirror}}(Q)$. There is a 2-fold cover \tilde{Q} of Q obtained by taking two copies of Q and identifying these along the mirror singular locus. This is the *local-orientation double cover*.


Theorem 2.9. *Every orbifold Q has a universal covering $\pi : \tilde{Q} \rightarrow Q$ with the property that for any covering $p : R \rightarrow Q$ there is $\pi' : \tilde{Q} \rightarrow R$ such that $\pi = p \circ \pi'$.*

Proof. (For the case where Σ has codimension 2.) The *regular set* is $Y = Q - \Sigma(Q)$. For each codimension-2 cell e of $\Sigma(Q)$ let $n_e \in \mathbb{N}$ be the label and μ_e a *meridian* loop linking e . Let G be the normal subgroup of $\pi_1 Y$ generated by $\{ \mu_e^{n_e} : e \in \Sigma \}$, and let \tilde{Y} be the ordinary covering corresponding to G . Now choose a path metric on Q and lift this to a metric on \tilde{Y} . Then \tilde{Q} is the *metric completion* of \tilde{Y} . \square

The *orbifold fundamental group* $\pi_1^{\text{orb}}(Q)$ of Q is the group of covering transformations of the universal cover. In the above construction we have $\pi_1^{\text{orb}}(Q) = \pi_1(Y)/G$. For example, if the underlying space of Q is a surface and the singular locus has only cone points, then the orbifold fundamental group is the topological fundamental group of the complement of the singular locus with relations that kill certain powers of loops going round the cone points.

Example 2.10. For the pillowcase we have

$$\pi_1^{\text{orb}}(\text{pillowcase}) = \langle a, b, c, d \mid a^2 = b^2 = c^2 = d^2 = abcd = 1 \rangle.$$




An orbifold is (*very*) *good* if it is (finitely) covered by a *manifold*. Otherwise it is *bad*.

Example 2.11. Teardrops, $S^2(n)$, and spindles, $S^2(m, n)$ with $m \neq n$, are bad.

Proof. To get a *manifold* covering of the top half of a teardrop we must take an n -fold covering; but there is no connected n -fold covering of the

bottom half: a disc. Similarly for a spindle, the universal cover is a g -fold covering, where $g = \gcd(n, m)$.



2.4 Orbifold Euler characteristic

If a surface M has a cell decomposition with V vertices, E edges and F faces the Euler characteristic is $\chi(M) = V - E + F$. It follows from this formula that if \tilde{M} is a d -fold covering of M that $\chi(\tilde{M}) = d\chi(M)$ since \tilde{M} has a cell decomposition obtained by lifting cells from M , and each cell of M has d lifts to \tilde{M} . Thus, we can regard the Euler characteristic as a way of *counting* the cells in M which multiplies under taking coverings. We would like to extend this to orbifolds. Consider the universal cover $\tilde{D} \rightarrow D^2(n)$. Topologically this is an n -fold cyclic branched cover of the disc branched over the cone point p in $D^2(n)$. This suggests that p should only “count” as $1/n$ of a vertex instead of as 1 vertex. The group of orbifold covering transformations \mathbb{Z}_n of \tilde{D} fix a point \tilde{p} which projects to p . Thus in some sense p might be thought of as having n “distinct lifts” which each occupy the fraction $1/n$ of the point \tilde{p} . This motivates the following definition.

The *orbifold Euler characteristic* of Q is

$$\chi(Q) = \sum_{\sigma \in Q} (-1)^{\dim(\sigma)} / |\Gamma(\sigma)|,$$

where σ ranges over (open) cells in X_Q and $\Gamma(\sigma)$ is the local group assigned to points in σ . The cell decomposition used is chosen so that two points in the same open cell have isomorphic local groups. For example, an open 1-cell is either contained in a mirror edge or disjoint from all mirror edges.

Exercise 2.12. Prove that the definition of orbifold Euler characteristic does not depend on the cell decomposition used.

Proposition 2.13.

- (a) For any d -fold orbifold covering $\tilde{Q} \rightarrow Q$, $\chi(\tilde{Q}) = d\chi(Q)$.
 (b) For a closed orientable 2-orbifold Q with cone points of orders m_i ,

$$\chi(Q) = \chi(X_Q) - \sum_i (1 - 1/m_i).$$

Exercise 2.14. Prove proposition 2.13.

2.5 Geometric structures on orbifolds

An orbifold is a *combinatorial* gadget: an underlying space together with groups labelling cells. Often we have more structure, for example a space modulo a group of isometries.

Let X be a space and G a group acting transitively on X . Then a (G, X) -orbifold is locally modelled on X modulo finite subgroups of G . In particular an orbifold is *hyperbolic* if $X = \mathbb{H}^n$ is hyperbolic n -space and $G = \text{Isom}(X)$. Euclidean orbifolds ($X = \mathbb{E}^n$) and spherical orbifolds ($X = S^n$) are defined similarly.


Exercise 2.15. If (G, X) is a geometry where $G \subset \text{Isom } X$, then $Q = X/\Gamma$ is a (G, X) -orbifold for any discrete subgroup Γ of G . (Theorem 2.26 below gives a converse to this.)

A closed, orientable geometric 2-orbifold has a Riemannian metric of constant curvature except at singular points, where the metric is cone like.


Example 2.16.

- (a) The pillowcase in example 2.1 is a quotient $P = \mathbb{E}^2/\Gamma$ of the Euclidean plane \mathbb{E}^2 by discrete group of isometries Γ , so is a Euclidean orbifold.

Alternatively, the surface of a tetrahedron in Euclidean space with opposite edges of equal length has an intrinsic Euclidean metric with four cone points with cone angle π . So this is also gives a Euclidean structure on the pillowcase orbifold $S^2(2, 2, 2, 2)$.



(b) A hyperbolic torus with cone point angle π is a hyperbolic 2-orbifold.



The following orbifold version of the Gauss-Bonnet theorem relates the orbifold Euler characteristic and Gaussian curvature K .

Proposition 2.17. For any closed orientable 2-orbifold Q

$$\int_{X_Q} K dA = 2\pi\chi(Q).$$

Remark: The integral is taken over the non-singular points of Q . This can be combined with part (b) of 2.13 as follows:

$$\int_{X_Q} K dA + \sum_i 2\pi(1 - 1/m_i) = 2\pi\chi(X_Q).$$

Here the left hand side represents the “total curvature” of the orbifold; the “correction term” $2\pi(1 - 1/m_i)$ represents the curvature concentrated at a cone point with cone angle $2\pi/m_i$.

Exercise 2.18. Prove proposition 2.17.

Exercise 2.19.

- (a) Enumerate the closed orientable 2-orbifolds with $\chi = 0$.
- (b) Show each is $\mathbb{R}^2/(\text{group of isometries})$.
- (c) Show each *compact* one is torus/(finite group).

Exercise 2.20.

- (a) Enumerate the closed orientable 2-orbifolds with $\chi > 0$.
- (b) Use this to prove theorem 2.4.

Example 2.21. Let P be a polygon in \mathbb{H}^2 with geodesic sides such that for each vertex there is an integer $n > 0$ such that the interior angle at that vertex is π/n . Let G be the group of isometries generated by reflections in the sides of P . Then G is a discrete group and the quotient \mathbb{H}^2/G is a hyperbolic orbifold isometric to P and with mirror edges. A corner at which the angle is π/n has local group D_{2n} . A similar construction works for Euclidean and spherical polygons.

Theorem 2.22. *Let Q be a closed 2-orbifold which is not a teardrop or spindle, or the quotient of one of these by an involution. Then*

Q has a hyperbolic structure if and only if $\chi(Q) < 0$,

Q has a Euclidean structure if and only if $\chi(Q) = 0$,

Q has a spherical structure if and only if $\chi(Q) > 0$.

Sketch of proof. If the orbifold has some mirror singular locus, take the local-orientation double cover as in 2.8. One can construct a geometric structure on the cover that is invariant under the covering reflection. Using the enumeration of orbifolds with non-negative orbifold Euler characteristic 2.19, 2.20 gives the result in the non-hyperbolic case. If $\chi(Q) < 0$ a continuity argument can be used to show there is a hyperbolic polygon and a pairing of the sides by isometries which gives the required orbifold. The ‘only if’ parts follow immediately from the Gauss-Bonnet theorem 2.17. \square

Combining exercise 2.19 and theorem 2.22 gives:

Theorem 2.23. *The closed orientable Euclidean 2-orbifolds are the torus, pillowcase $S^2(2, 2, 2, 2)$, and turnovers $S^2(2, 3, 6)$, $S^2(2, 4, 4)$, and $S^2(3, 3, 3)$ (obtained by doubling Euclidean triangles with all cone angles π/n_i , $n_i \in \mathbb{N}$).*

Note: An *apple turnover* is a triangular pastry with apple inside. To make one: double, along the boundary, a triangle made of pastry. But remember to place the apple inside before doubling!

Next we note that the idea of developing map extends to orbifolds with analytic geometric structures.

Theorem 2.24. *Let Q be a (G, X) -orbifold, where (G, X) is an analytic geometry. Then there is a developing map*

$$\text{dev} : \tilde{Q} \longrightarrow X$$

defined on the universal orbifold cover \tilde{Q} of Q , and a holonomy representation $h : \pi_1^{\text{orb}}(Q) \rightarrow G$ such that

$$\text{dev} \circ \gamma = h(\gamma) \circ \text{dev}$$

for each deck transformation γ in $\pi_1^{\text{orb}}(Q)$.

Proof. The idea is to construct the developing map and universal cover simultaneously, by piecing together the maps $\tilde{U}_i \rightarrow X$ given by coordinate charts $U_i = \tilde{U}_i/G_i$ on Q . We actually consider all germs of the (local) inverse maps $X \rightarrow \tilde{U}_i \rightarrow U_i = \tilde{U}_i/G_i \subset Q$. These form a manifold $G(Q)$ which fibres over $X \times Q$ with the isotropy group G_x as fibre. Further, the local maps fit together by *analytic continuation* to give a foliation of $G(Q)$. Each leaf projects by local homeomorphisms to both X and Q ; this gives the desired universal cover and the developing map. See Thurston [84] for the details. \square

Example 2.25. Given a Euclidean tetrahedron as in 2.16(a), we can just roll it around the plane to see its developing map. More generally, if we draw a pattern on any closed 2-dimensional Euclidean orbifold, the developing map gives a wallpaper pattern in the plane.

The following important result generalizing Theorem 1.5 is a version of Poincaré's polyhedron theorem.

Theorem 2.26. *Let Q be an orbifold with a geometric structure modelled on an analytic geometry (G, X) where G is a group of isometries of X . If Q is complete as a metric space, then the developing map $\text{dev} : \tilde{Q} \rightarrow X$ is a covering map; hence Q is good. If X is simply connected, then the holonomy representation $h : \pi_1^{\text{orb}}(Q) \rightarrow G$ is an isomorphism onto a discrete subgroup Γ of G which acts properly discontinuously on X , and Q is isometric to the quotient X/Γ .*

In the above theorem, if G is a *linear group* it follows from Selberg's Lemma (see for example [68]) that every complete (G, X) orbifold is *very good*, i.e. finitely covered by a manifold. In particular, this applies to the three 2-dimensional constant curvature geometries, and to Thurston's eight 3-dimensional geometries.

Corollary 2.27. *Complete, geometric 2-orbifolds and 3-orbifolds are very good.*

The following results follow immediately from theorems 2.26 and 2.22.

Corollary 2.28. *The only bad closed 2-orbifolds are spindles, $S^2(m, n)$ with $m \neq n$, teardrops, $S^2(n)$, and the quotients of one of these by a reflection.*

Corollary 2.29. *Every good closed 2-orbifold is very good. Equivalently every connected 2-orbifold except those listed in 2.28 is finitely covered by a manifold.*

2.6 Some geometric 3-orbifolds

Here are some examples of branched covering spaces giving interesting geometric 3-orbifolds.

Example 2.30. 2-bridge knots.

A 2-bridge knot, K , has 2-fold branched cover which is a union of two solid tori, so is a *lens space* $L(p, q)$ with a spherical structure (see example 1.16). Further, this branched covering can be realized as a quotient map $S^3/\mathbb{Z}_p \rightarrow S^3/D_{2p}$ where $\mathbb{Z}_p \subset D_{2p}$ are groups of isometries of S^3 . Hence S^3 with K , labelled 2, is a spherical 3-orbifold.

In fact, \mathbb{Z}_p is the cyclic group of isometries of


$$S^3 = \{(z_1, z_2) \in \mathbb{C}^2 : |z_1|^2 + |z_2|^2 = 1\}$$

generated by $g : (z_1, z_2) \mapsto (tz_1, t^q z_2)$ where $t = \exp(2\pi i/p)$, and the dihedral group D_{2p} is generated by g and the involution $\tau : (z_1, z_2) \mapsto (\bar{z}_1, \bar{z}_2)$. Note that $L(p, q) = S^3/g$, and that τ induces an involution on $L(p, q)$ with 1-dimensional fixed point set.

Exercise 2.31. Verify that g and τ both preserve the Heegaard solid tori $V = \{(z_1, z_2) \in S^3 : |z_1| \leq |z_2|\}$ and $V' = \{(z_1, z_2) \in S^3 : |z_1| \geq |z_2|\}$. Further, the quotient map induced by τ on V/g and V'/g is the usual 2-fold branched covering of a solid torus over a 3-ball branched over two unknotted, unlinked arcs as in example 1.23. Hence the quotient orbifold S^3/D_{2p} is topologically S^3 with a 2-bridge knot or link $K(p, q)$ labelled 2 as singular locus.

Example 2.32. The Borromean rings.

The *Borromean rings* labelled 2 is a Euclidean orbifold. Start with a Euclidean cube and fold each face in half. This is the quotient of \mathbb{E}^3 by the group $\Gamma < \text{Isom}(\mathbb{E}^3)$ generated by 180 degree rotations about 3 sets of orthogonal axes.




The Borromean rings labelled 4 is a hyperbolic orbifold. This can be constructed as above, starting with regular hyperbolic dodecahedron with angles $\pi/2$.


Example 2.33. Figure eight knot.

The figure eight knot labelled 2 is a spherical orbifold. This follows from example 2.30 since the figure eight knot is a 2-bridge knot; in fact, the 2-fold branched cover is the lens space $L(5, 3)$.

The figure eight knot labelled 3 is a Euclidean orbifold: Start with Borromean rings labelled 2 as in 2.32 above. Dividing out by the order 3 symmetry gives a symmetric link in S^3 , with two components labelled 2 and 3. Taking the 2-fold branched cover gives the figure 8 knot labelled 3.



The orbifold fundamental group is a Euclidean group generated by 120 degree rotations about disjoint diagonals of two adjacent cubes.



The figure eight knot labelled k is hyperbolic for $k \geq 4$. The hyperbolic structures for $k \geq 5$ are given in Thurston's analysis of hyperbolic Dehn surgery on the figure eight knot (see [80, chapter 4]); a direct geometric or arithmetic construction can be given for $k = 4$ (see [41], [40]).


2.7 Orbifold fibrations

Informally an *orbifold fibre bundle* is locally the quotient of a bundle by a finite group action which preserves the bundle structure. More formally an orbifold bundle with total space E and generic fibre F over a base orbifold B is an orbifold map $p : E \rightarrow B$ between the underlying spaces such that each point $b \in B$ has a neighbourhood $U = \tilde{U}/G$ and an action of G on F so that $p^{-1}(U) = (\tilde{U} \times F)/G$ where G acts diagonally. This local product structure must be compatible with p ; thus the following diagram commutes:

$$\begin{array}{ccc}
 \tilde{U} \times F & \longrightarrow & p^{-1}(U) \\
 \downarrow \text{proj}_1 & & \downarrow p|_U \\
 \tilde{U} & \longrightarrow & U
 \end{array}$$


Consider the case when the base space is the 1-orbifold the unit interval I with mirror endpoints. Then an orbifold bundle $p : E \rightarrow I$ with fibre F is $F \times I$ modulo an involution on $F \times 0$ and another involution on $F \times 1$. Any such bundle has a 2-fold covering by an F -bundle over S^1 . As an example if the fibre is \mathbb{R} and the involution at each end is reflection, then $E = \mathbb{R}^2(2, 2)$ is an *infinite pillowcase*.

Example 2.34. The torus is a circle bundle over a circle. $S^1 \rightarrow S^1 \times S^1 \rightarrow S^1$. There is an involution of the circle with two fixed points. It is covered by an involution of the torus with four fixed points (shown in example 2.1) and has quotient a pillowcase P . Thus the quotient of the original bundle by this \mathbb{Z}_2 action gives a bundle with total space a pillowcase and with generic fibre a circle $S^1 \rightarrow P \rightarrow I$.



The fibre over each interior point of the interval I is a circle, and the fibre over each boundary point is an interval. Topologically P is decomposed into two arcs, called *singular fibres* connecting the singular points together with a foliation of the complementary annulus by circles.

Example 2.35. Consider a torus bundle over a circle $T^2 \rightarrow M \rightarrow S^1$ with monodromy ϕ . Let σ be the involution of T shown represented by a half rotation around the centre of the square. Then σ is central in the mapping class group of the torus and therefore extends to an involution on M preserving the projection onto S^1 . Let τ be the involution of T given by the reflection of the square shown. Assume that $\phi\tau = \tau\phi^{-1}$. Then τ extends to an involution on M which covers a reflection of the base space S^1 . Thus we have an action of $\mathbb{Z}_2 \times \mathbb{Z}_2$ on M . The quotient is an orbifold bundle over the interval (with mirror endpoints) and with generic fibre a pillowcase $P \rightarrow M \rightarrow I$.



Theorem 2.36 (Geometric bundles over 1-orbifolds).

- (a) A 3-dimensional orbifold which fibres over a compact 1-orbifold with spherical fibre has a geometric structure modelled on $S^2 \times \mathbb{E}^1$.
- (b) A 3-dimensional orbifold which fibres over a connected, compact 1-orbifold with Euclidean fibre has a geometric structure modelled on either \mathbb{E}^3 , Nil or Solv geometry.


Exercise 2.37. Deduce this from the corresponding result for 3-manifolds which fibre over S^1 proved in Thurston [82, 3.8.10, 4.7.1].

2.8 Orbifold Seifert fibre spaces

In 1933, H. Seifert introduced and classified the 3-dimensional manifolds now known as *Seifert fibre spaces* in a beautiful paper [74]; an English translation of his paper is contained in the book [75]. See [73] for an excellent treatment of Seifert fibre spaces and their geometries.

A Seifert fibre space (SFS) is a generalization of the concept of a bundle over a surface with fibre the circle. Such a bundle is foliated by circles and each circle has a neighbourhood diffeomorphic to the solid torus $D^2 \times S^1$ foliated by the circles $\{x\} \times S^1, x \in D^2$. A *Seifert fibre space* is a 3-manifold foliated by circles (called *fibres*), where each fibre has a neighbourhood diffeomorphic to a quotient $(D^2 \times S^1)/G$ where G is a finite group preserving both factors and acting *freely*; the fibres are the images of $\{x\} \times S^1$.

In the orientable case, each fibre has a neighbourhood which is a *fibred solid torus* foliated by circles obtained by taking a solid cylinder $D^2 \times I$ foliated by intervals $\{x\} \times I$ and identifying the two ends after a rotation of $2\pi\alpha/\beta$ where α, β are coprime integers. When $\beta > 1$, the core circle of the solid torus is an *exceptional fibre* of order β ; it wraps once around the solid torus before closing up, while all nearby *regular fibres* wrap β times around before closing up. This corresponds to a quotient $(D^2 \times S^1)/G$ where $G \cong \mathbb{Z}_\beta$ is the group of diffeomorphisms of $D^2 \times S^1$ generated by $(z, t) \mapsto (\lambda^\alpha z, \lambda t)$ where $\lambda = \exp(2\pi i/\beta)$ and we regard D^2, S^1 as subsets of \mathbb{C} .



A *Seifert fibred orbifold* or *orbifold Seifert fibre space* (OSFS) is a 3-orbifold E which fibres over a 2-dimensional orbifold B . Each fibre has a neighbourhood modelled on

$$(D^2 \times S^1)/G$$

where G is a finite group preserving both factors. The fibres are the images

of $\{x\} \times S^1$; these are either circles or intervals. By identifying each fibre to a point we obtain the bundle projection $p : E \rightarrow B$, where the *base space* B is a 2-orbifold. For more detailed treatments, see [84], [9], [10].


Example 2.38. Any quotient of a Seifert fibre space by a finite group action which preserves the foliation by circles in an orbifold Seifert fibre space.

Example 2.39. Let $Q = X/\Gamma$ be a geometric 2-orbifold, where Γ is a subgroup of $G = \text{Isom } X$. Then Γ acts naturally on the *unit tangent bundle* $UT(X)$ of X and the quotient orbifold $UT(X)/\Gamma$ is the *unit tangent bundle* $UT(Q)$ of Q . This is an orbifold Seifert fibre space which fibres over Q with circle as generic fibre.


Each fibre in an orientable orbifold Seifert fibre space has a fibred neighbourhood $(D^2 \times S^1)/G$ which is either a (*singular*) *fibred solid torus* or a (*singular*) *fibred folded ball*. These will now be described.

A (*singular*) *fibred solid torus* is fibred the same way as the fibred solid torus described above, however the singular locus of the orbifold structure is the core circle of the solid torus. This is labelled by some integer $n \geq 1$. The fibred solid torus is called *singular* if $n > 1$. In this case $G \cong \mathbb{Z}_n \times \mathbb{Z}_\beta$.

A (*singular*) *folded ball* is the \mathbb{Z}_2 quotient of a (singular) solid torus by the involution fixing two arcs shown in 1.23. The underlying space of the quotient is a 3-ball. The singular locus of a folded ball is two unknotted arcs each labelled 2, while the singular locus of a *singular* folded ball is an **H** graph. The first diagram shows a fibred folded ball and three different fibres. The angle between the lines of singular locus is $\pi/2$. There is one fibre along the *soul* (defined later in 7.5) of the folded ball of length 1. A regular fibre has length 4. The remaining fibres start and end on the same strand of singular locus and have length 2. The base orbifold is a sector of a circle. The two radii are mirrors and the centre is a corner reflector with angle $\pi/2$. The group at the corner is dihedral of order 4.



In general, a regular fibre has length $2n$, a fibre lying over an edge of the base orbifold has length n and the fibre lying over the centre has length 1. In this case the base orbifold contains a corner reflector with angle π/n and local group dihedral of order $2n$. The next diagram (taken from [84]) shows a folded ball where the angle between the lines of singular locus is $\pi/3$. The base orbifold has a corner with this angle.



The singular locus $\Sigma = \Sigma_V \cup \Sigma_H$ of an orbifold Seifert fibre space is a union of the *vertical* singular locus Σ_V , and the *horizontal* singular locus Σ_H . The subset $\Sigma_V \subset \Sigma$ is the union of all arcs in Σ which are contained in a fibre, and Σ_H is the closure of $\Sigma - \Sigma_V$ and consists of all arcs in Σ transverse to the fibres. The projection $p(\Sigma_H)$ to the base orbifold B of the horizontal singular set is the *mirror singular set* $\Sigma_{mirror}(B)$ of B .

We now describe this locally. Given a fibred neighbourhood $(D^2 \times S^1)/G$, the action of G on D^2 may include orientation reversing elements with 1-dimensional fixed set. In this case $\Sigma_{mirror}(D^2/G) \neq \emptyset$. The fibre over a

point $x \in D^2/G$ is either a circle or an interval. Furthermore the points for which the fibre is an interval are precisely the points in $\Sigma_{\text{mirror}}(D^2/G)$.

The fixed set $\text{Fix}(g)$ of a non-trivial element $g \in G$ is a 1-manifold. The components of $\text{Fix}(g)$ are either fibres of the fibration and are called *vertical*, or else transverse to the fibres and called *horizontal*. If some component of $\text{Fix}(g)$ is horizontal then g has order 2. The horizontal fibres are precisely those points which project to $\Sigma_{\text{mirror}}(D^2/G)$.

If Q is an orbifold SFS, then there is a fibred neighbourhood (which is usually *not* a regular neighbourhood), $\mathcal{N}(\Sigma)$, of the singular locus such that the complement is a Seifert fibred *manifold*. The union of all fibres which meet Σ consists of those fibres contained in Σ_V together with a set of fibres that are intervals with both endpoints in Σ_H . There is a regular neighbourhood of this set which is a union of fibres. The components of $\mathcal{N}(\Sigma)$ are of two types. The first type is a singular solid torus; the core curve is contained in Σ and may or may not be an exceptional fibre. The second type is a union of (singular) folded balls with the property that if two intersect then each component of the intersection is $D^2(2, 2)$. The boundary $\partial\mathcal{N}(\Sigma)$ is a Euclidean 2-manifold since it is fibred by circles.

Our proof of the *collapsing theorem* 7.13 (Collapsing theorem) proceeds by constructing the orbifold Seifert fibration in \mathcal{N} . The idea is suggested by the following exercise.

Exercise 2.40. Show that the orbifold (S^3, K) with singular locus a Montesinos knot K labelled 2 is an OSFS with regular fibre a meridian of the solid torus V described in Definition 1.28.


An orbifold Seifert fibre space Q (with no mirror boundary) such that the orbifold singular locus is all vertical has underlying space a manifold which is Seifert fibred by the Seifert fibration on Q . The following result implies that every OSFS is either such an OSFS or is the quotient by an involution of such an OSFS.

Proposition 2.41. *If Q is an orbifold SFS and $\Sigma_H(Q) \neq \emptyset$ then there is an orbifold covering \tilde{Q} of degree 2 such that the singular set of \tilde{Q} is vertical. The base orbifold of \tilde{Q} is the double along $\Sigma_{\text{mirror}}(B)$ of the base orbifold B of Q . Every regular fibre lifts to the covering.*

Proof. Let $N \subset \mathcal{N}(\Sigma)$ be the fibred neighbourhood of the horizontal singular locus. There is an orbifold cover \tilde{N} of degree 2 which has no horizontal singular locus. This covering is constructed locally. A (singular) folded ball has a unique degree 2 orbifold cover by a (singular) solid torus.

These covers are compatible on the intersection of two (singular) folded balls. Now $\partial\tilde{N}$ consists of two disjoint lifts of ∂N . Thus an orbifold cover of Q is obtained by gluing to N two copies of $Q - N$ one onto each lift of ∂N . Note that ∂N may not be connected, so that there are more than two possible lifts of ∂N , but given one such lift, there is a unique second lift which is disjoint from the first. \square

Example 2.42. The figure below illustrates the proposition when the base orbifold is a disc with mirror boundary. The orbifold Q has underlying space S^3 and the singular locus is an unlink of two components each labelled 2. The union of exceptional fibres is an annulus with boundary the singular locus. Each exceptional fibre is a radius of the annulus. The regular fibres are circles which link the annulus. The 2-fold orbifold cover \tilde{Q} is $S^2 \times S^1$ with the product Seifert fibres.



Example 2.43. Start with the product fibration on $S^2 \times S^1$ and make an orbifold structure by labelling one of the circle fibres with a cyclic group. This gives a bad 3-orbifold which is an orbifold SFS. The base space is a teardrop orbifold.

However most orbifold Seifert fibre spaces are very good:

Proposition 2.44. *Suppose that Q is an orbifold SFS with base space a good 2-orbifold B . Then Q is finitely orbifold covered by a circle bundle over a surface. In particular Q is very good.*

Proof. Since B is good it is very good by (2.29), so it has a finite orbifold cover by a manifold \tilde{B} . The orbifold bundle $p : Q \rightarrow B$ induces an epimorphism $\pi_1^{orb} Q \rightarrow \pi_1^{orb} B$ and the subgroup $\pi_1^{orb} \tilde{B}$ determines a finite orbifold cover \tilde{Q} of Q . Then \tilde{Q} is an orbifold SFS with base space \tilde{B} . The singular locus in an orbifold SFS maps to singular locus in the base space. Since \tilde{B} is a manifold it follows that \tilde{Q} is a manifold. \square

If B is not good, then by (2.28) B is a teardrop or spindle, or the quotient of one of these by a reflection. Then the underlying space of Q is homeomorphic to a manifold of Heegaard genus at most 1, i.e. a lens space, S^3 or $S^2 \times S^1$. In these cases it is easy to determine whether or not Q is very good.

Corollary 2.45. *If Q is an orbifold SFS with non-empty orbifold boundary then Q is very good, i.e. finitely covered by a manifold.*

Proof. Suppose that $p : Q \rightarrow B$ is the orbifold SFS projection. Then $p(\partial_{orb}Q)$ is a non-empty subset of $\partial_{orb}B$. By (2.28) a bad 2-orbifold is always a closed orbifold. (The quotient of a spindle or teardrop by a reflection has underlying space a disc with “mirror boundary” but the orbifold boundary is empty.) Hence B is good. The result now follows from (2.44). \square

It follows from the discussion of (oriented) orbifold Seifert fibred spaces that every point has a neighbourhood which is fibred as a fibred (singular) solid torus or a fibred (singular) folded ball. Conversely, if a 3-dimensional orbifold Q is decomposed into circles and intervals such that every point has a neighbourhood which is a fibred (singular) solid torus or (singular) folded ball, then Q is an orbifold Seifert fibre space.

A *meridian* of a (singular) folded ball, V , is a curve in ∂V with pre-image a meridian in the (singular) solid torus that double covers V . The next three lemmas will be used in the proof of the collapsing theorem.

Lemma 2.46. *Suppose that V is a (singular) solid torus or a (singular) folded ball. An orbifold fibration on the boundary of V extends to an orbifold Seifert fibration of V unless a circle fibre on the boundary is orbifold-compressible in V . This happens exactly when such a circle is a meridian of V .*

Proof. This result is obvious for a (singular) solid torus. A (singular) folded ball has a 2-fold orbifold cover which is a (singular) solid torus. The result for a (singular) folded ball now follows from that for a (singular) solid torus. \square

Observe that in the above lemma, it is of no consequence whether or not a folded ball is singular or not.

Lemma 2.47. *The only orientable orbifold Seifert fibre spaces with orbifold compressible boundary are the (singular) solid torus and the (singular) folded ball.*

Proof. Suppose Q is an orbifold SFS with orbifold compressible boundary. Since $\partial_{orb}Q \neq \emptyset$ it follows from (2.45) that Q is very good. Hence Q is finitely orbifold covered by a Seifert fibred manifold with compressible boundary. Any Seifert fibred manifold with compressible boundary has base space a disc with at most one exceptional fibre, and is therefore a solid torus or solid Klein bottle. Hence any orientable orbifold Seifert fibre space with orbifold-compressible boundary is the quotient of a solid torus by a finite group, i.e. a (singular) folded ball or (singular) solid torus. \square

Definition 2.48. Let X be a closed orientable Euclidean 2-orbifold, thus X is either a torus, pillowcase or turnover by theorem 2.23. Then a *thick* X is a product $X \times I$, where $I = \mathbb{R}$ or $[-1, 1]$ is an interval. A *folded thick* X is a quotient $X \times I/\bar{\tau}$ where $\bar{\tau}$ is an orientation preserving isometric involution of $X \times I$ that reverses the ends of I . Thus $\bar{\tau}(x, t) = (\tau x, -t)$ where τ is an orientation reversing isometric involution of X . (These will play an important role in sections 7.5 and 7.6 below.)

Lemma 2.49. *Every folded thick pillowcase is an orbifold SFS.*

Proof. Let T be the 2-fold orbifold cover of a pillowcase P by a torus and σ the orbifold covering transformation of T such that $P = T/\sigma$.

Let $\pi : \mathbb{E}^3 \rightarrow V$ be the universal orbifold cover of a folded thick pillowcase $V = P \times \mathbb{R}/\bar{\tau}$ as above. Let $G \cong \pi_1^{orb}(V)$ be the group of orbifold covering transformations of V . The Euclidean plane $\mathbb{E}^2 \equiv \pi^{-1}(P) \subset \mathbb{E}^3$ is invariant under G . The restriction $\pi| : \mathbb{E}^2 \rightarrow P$ is the universal orbifold cover of the pillowcase P . Let $\tilde{\tau}, \tilde{\sigma} \in \pi_1^{orb}(V)$ be the orbifold covering transformations corresponding to τ, σ and $\tilde{\tau}|, \tilde{\sigma}|$ the restrictions to $\pi^{-1}(P)$.

Now G has a torsion free abelian subgroup of index 4 consisting of translations. The rotational part of the holonomy restricted to \mathbb{E}^2 is a map $\theta : G \rightarrow O(2)$ which has image a group of order 4. Now $\theta(\tilde{\sigma})$ is multiplication by -1 and $\theta(\tilde{\tau})$ has eigenvalues ± 1 . Hence each eigenvector of $\theta(\tilde{\tau})$ is preserved by $\theta(G)$. The foliation of \mathbb{E}^2 by lines parallel to an eigenvector is G -invariant. This foliation by lines extends to a G -invariant foliation on \mathbb{E}^3 . The projection of this foliation to V gives the orbifold Seifert fibration. \square

Finally we note that orbifold Seifert fibre spaces are always geometric, provided they contain no bad 2-dimensional orbifolds.

Theorem 2.50 (Geometric Orbifold Seifert Fibre Spaces).

A 3-dimensional orbifold E which fibres over a 2-dimensional geometric orbifold B has a geometric structure modelled on either S^3 , $S^2 \times \mathbb{E}^1$, \mathbb{E}^3 ,

Nil, PSL or $\mathbb{H}^2 \times \mathbb{E}^1$ geometry. Further, all the fibres are geodesics in this geometry. The kind of geometry involved is determined by the Euler characteristic χ of the base orbifold (obtained by identifying each fibre to a point), and the Euler number e of the fibration as follows.

	$\chi > 0$	$\chi = 0$	$\chi < 0$
$e \neq 0$	S^3	<i>Nil</i>	<i>PSL</i>
$e = 0$	$S^2 \times \mathbb{E}^1$	\mathbb{E}^3	$\mathbb{H}^2 \times \mathbb{E}^1$

Remarks:

- (a) $e = 0$ if and only if the orbifold bundle is finitely covered by a product bundle.
- (b) If the base space B is a bad orbifold then the bundle is a bad orbifold if $e = 0$. If B is bad and $e \neq 0$ the bundle has S^3 geometry. However, the corresponding Seifert fibration is *not geometric* — it is *not* isotopic to a fibration where the fibres are geodesics in the S^3 geometry.
- (c) The proof of 2.50 is similar to the arguments for manifolds. There is a geometric proof (see [84]) which involves constructing a bundle type metric with constant curvature on the base together with a constant curvature connection. For manifolds, there are also more algebraic proofs ([73], [48]) which use explicit presentations for the fundamental groups.

2.9 Suborbifolds

A k -dimensional *suborbifold* of an n -orbifold is locally modelled on the inclusion $\mathbb{R}^k \subset \mathbb{R}^n$ modulo a finite group. An *orbifold disc (or ball)* is the quotient of a 2-disc (or 3-ball) by a finite group.

A 2-dimensional suborbifold P of a 3-orbifold Q is *orbifold-incompressible* or *essential* if

1. $\chi(P) > 0$ and P does not bound an orbifold ball in Q , or
2. $\chi(P) \leq 0$ and any 1-suborbifold on P which bounds an orbifold disc in $Q - P$ also bounds an orbifold disc in P .

Example 2.51. Euclidean and hyperbolic turnovers, $S^2(p, q, r)$ with $1/p + 1/q + 1/r \leq 1$, are always incompressible since every 1-suborbifold in a turnover bounds an orbifold disc.

A 3-orbifold is *orbifold-irreducible* if it contains no bad 2-suborbifold or essential spherical 2-suborbifold. It is *orbifold-atoroidal* if every essential Euclidean 2-orbifold is boundary parallel.

An *orbifold isotopy* is an isotopy $F : P \times [0, 1] \rightarrow Q$ such that for each $t \in [0, 1]$ the map $F_t : P \rightarrow Q$ is an orbifold isomorphism onto a suborbifold of Q .

2.10 Spherical decomposition for orbifolds

We now want to describe a geometric decomposition for a 3-orbifold Q . In this section, we describe an orbifold version of the *prime* or *connected sum* decomposition for 3-manifolds.

Step 1: If Q contains any bad 2-suborbifold S , then Q is bad. (In fact, any manifold covering of Q induces a manifold covering of S .) Thus if Q contains a bad 2-suborbifold, we stop here — we won't try to prove anything for such orbifolds!

Example 2.52. This occurs if Q is an orbifold with underlying space $S^1 \times S^2$ and singular set Σ meeting a sphere $p \times S^2$ in one point.

Step 2: Prime decomposition.

2(a) If there are any essential spheres or *footballs* $= S^2(n, n)$ choose a maximal non-parallel set. Cut along them, and fill in each boundary component with an orbifold ball.

2(b) Look for essential $S^2(p, q, r) \subset Q$. Cut along elliptic ones ($\chi > 0$) and add *cones*.

Remark: The nature of a turnover in a hyperbolic orbifold depends on its Euler characteristic. A spherical turnover must bound a ball neighbourhood of a vertex. A Euclidean turnover must bound a cusp neighbourhood, and one with negative Euler characteristic is isotopic to a totally geodesic suborbifold.

The prime decomposition theorems of Kneser [51] and Milnor [62] for 3-manifolds and Schubert [71] for links generalize to the following result for 3-orbifolds. (See Bonahon-Siebenmann [11], p. 445.)

Theorem 2.53 (Prime decomposition theorem).


Every compact 3-orbifold containing no bad 2-suborbifold can be decomposed into irreducible pieces by cutting along a finite collection S of spherical 2-suborbifolds and capping off the boundaries by orbifold balls. If S is a minimal such collection, then the resulting collection of irreducible 3-orbifolds is canonical (but the collection of spherical 2-suborbifolds is not unique up to isotopy).

For example if this is applied to the orbifold (S^3, K) where the singular locus is a knot K labelled 2 this gives a decomposition of the knot K as a connected sum of knots.

2.11 Euclidean decomposition for orbifolds

A closed 2-orbifold with $\chi = 0$ is a torus modulo a finite group. We want a generalization to orbifolds of the torus splitting theorem of Johannson and of Jaco-Shalen for 3-manifolds, and of the *characteristic submanifold* which is a *maximal* Seifert fibred submanifold (bounded by tori).

Example 2.54. A simple example is two 3-orbifolds each with a pillowcase boundary component glued together. We then cut apart to obtain the Euclidean orbifold decomposition.



If there are any essential suborbifolds with $\chi = 0$, we cut out the *characteristic suborbifold* of Bonahon and Siebenmann [11]. This is an *orbifold Seifert fibre space* and is unique up to orbifold-isotopy.

Theorem 2.55 (Characteristic Suborbifold Theorem). [11]

A compact, orbifold-irreducible 3-orbifold contains a collection \mathcal{T} of disjoint non-parallel incompressible Euclidean suborbifolds with the property that each component of the closure of the complement of a regular neighbourhood of \mathcal{T} is either an orbifold Seifert fibre space or is orbifold-atoroidal. Further, if \mathcal{T} is minimal with respect to these properties then \mathcal{T} is unique up to (orbifold) isotopy.

For example, for an orbifold (S^3, K) with a knot K labelled 2 as singular locus, Conway spheres cutting the knot in 4 points give a Euclidean decomposition of the orbifold along pillowcases $S^2(2, 2, 2, 2)$.

2.12 Graph orbifolds

Recall that an (orientable) *graph manifold* is a 3-manifold which contains a collection of disjoint **incompressible** tori such that the complementary components are Seifert fibre spaces. (Some authors do not require that the tori are incompressible.) A *graph orbifold* is a 3-orbifold which can be cut along a finite set of orbifold incompressible Euclidean 2-suborbifolds into pieces which are orbifold Seifert fibre spaces. From 2.50 it follows that this decomposes the graph orbifold into geometric pieces, and each of these is very good by corollary 2.27. Generalizing earlier work of Hempel [39], McCullough and Miller show in [59] that every 3-orbifold with a geometric *decomposition* is finitely covered by a manifold. Thus a graph orbifold is the quotient of graph manifold by a finite group action which preserves the graph manifold structure. (However if the graph orbifold Q is an OSFS with a bad base orbifold, the OSFS structure on Q obtained from this procedure is different from the original.)

Proposition 2.56. *If Q is a graph orbifold then either Q is a bad orbifold SFS or else Q is very good.*

Sketch of proof. In view of (2.44) and the remarks after it we may assume that the decomposition involves more than one orbifold SFS, and hence that each of the components has non-empty boundary. If some component has Σ_H non-empty then there is a 2-fold orbifold cover of Q which restricts to the cover given by (2.41) on each component with $\Sigma_H \neq \emptyset$. This is because in a component with $\Sigma_H \neq \emptyset$ the cover restricted to the boundary components is uniquely determined: each torus boundary component lifts and each pillowcase has the unique 2-fold orbifold cover by a torus. We may therefore assume that $\Sigma_H = \emptyset$.

The base 2-orbifold for each component has non-empty *orbifold* boundary and by (2.28) is therefore very good. It follows from (2.44) that each component is very good. In fact the covering constructed in (2.41) has the property that the regular fibres lift to the covering. Following [39] one shows that there is an integer $k > 0$ and a finite covering, \tilde{C} , of each component, C , so that every component $T \subset \partial C$ and every pre-image $\tilde{T} \subset \partial \tilde{C}$ the covering $\partial \tilde{T} \rightarrow T$ corresponds to the characteristic subgroup given by the kernel of $\pi_1 T \rightarrow H_1(T; \mathbb{Z}_k)$. Then the coverings of the components can be glued together to give a manifold cover of Q . \square

2.13 The Orbifold Theorem

By means of the prime decomposition theorem 2.53 for orbifolds the classification of compact 3-orbifolds may be reduced to those that are orbifold-irreducible. By cutting along compressing orbifold discs we may also reduce to the case that the boundary is orbifold incompressible. Finally by means of the Characteristic Suborbifold theorem 2.55 we may further reduce the classification to the orbifold-atoroidal case.

A **geometric structure** on an orbifold Q , possibly with non-empty orbifold boundary, is an orbifold isomorphism from the orbifold interior $Q - \partial_{orb} Q$ to an orbifold X/G where X is one of the eight 3-dimensional geometries and G is a discrete subgroup of isometries of X .

Since a manifold may be regarded as an orbifold without singular locus the geometrization conjectures for manifolds and orbifolds can be combined into a single conjecture.

Conjecture 2.57 (Orbifold Geometrization Conjecture). *Every compact orbifold-irreducible, orbifold-atoroidal 3-orbifold admits a geometric structure.*

Thurston announced a proof of this conjecture in the case that the singular locus has dimension at least 1 (see the appendix, page 153). We will now discuss the version of Thurston's Orbifold theorem whose proof we will outline in chapter 7.

Theorem 2.58 (Thurston's Orbifold theorem). *Let Q be a compact, orientable 3-orbifold which is orbifold-irreducible and orbifold-atoroidal. Assume $\Sigma(Q)$ has dimension 1, and $\partial_{orb} Q$ consists of orbifold-incompressible Euclidean 2-orbifolds. Then Q has a geometric structure.*

The case where the orbifold boundary is orbifold-incompressible but not Euclidean can be deduced from this case by a doubling argument. The case

where there is 2-dimensional singular locus can probably be deduced from this case and Thurston's theorem for Haken manifolds. The non-orientable case presents a number of new cases to consider at various stages of an argument which already consists of many cases.

Chapter 3

Cone-Manifolds

To find a geometric structure on a topological 3-orbifold Q , we will typically start with a complete hyperbolic structure on $Q - \Sigma$ (a Kleinian group) and try to *deform* this to a hyperbolic structure on the 3-orbifold Q (another Kleinian group). The intermediate stages will be hyperbolic metrics with cone-type singularities — **3-dimensional hyperbolic cone-manifolds**.

3.1 Definitions

An *n-dimensional cone-manifold* is a manifold, M , which can be triangulated so that the link of each simplex is piecewise linear homeomorphic to a standard sphere and M is equipped with a complete path metric such that the restriction of the metric to each simplex is isometric to a geodesic simplex of constant curvature K . The cone-manifold is *hyperbolic*, *Euclidean* or *spherical* if K is -1 , 0 , or $+1$.

Remark: We could allow more general topology, for example M a rational homology n -manifold. Most arguments still apply in that setting.

The *singular locus* Σ of a cone-manifold M consists of the points with no neighbourhood isometric to a ball in a Riemannian manifold. It follows that

- Σ is a union of totally geodesic closed simplices of dimension $n - 2$.
- At each point of Σ in an open $(n - 2)$ -simplex, there is a *cone angle* which is the sum of dihedral angles of n -simplices containing the point.
- $M - \Sigma$ has a smooth Riemannian metric of constant curvature K , but this metric is *incomplete* if $\Sigma \neq \emptyset$.

We will see that many of the techniques and results from Riemannian geometry also apply to cone-manifolds. Useful references on the geometry of cone-manifolds and more general spaces of piecewise constant curvature include [12] and [47].

Example 3.1.


Let Δ be the lune (or bigon) in S^2 contained between two geodesics making an angle α . The double of Δ gives a cone-manifold structure on S^2 with two *cone points* each with cone angle 2α . This cone-manifold is called a *football*.

Example 3.2.

Let Δ be a triangle with angles $0 < \alpha, \beta, \gamma < \pi$ in

- \mathbb{H}^2 if $\alpha + \beta + \gamma < \pi$,
- \mathbb{E}^2 if $\alpha + \beta + \gamma = \pi$,
- S^2 if $\alpha + \beta + \gamma > \pi$.


Doubling the triangle Δ gives a cone-manifold structure on S^2 with three *cone points* with cone angles $2\alpha, 2\beta, 2\gamma$. These cone-manifolds are called *turnovers* and are hyperbolic, Euclidean or spherical respectively. These cone-manifolds are denoted $S^2(2\alpha, 2\beta, 2\gamma)$.



Remark: In the spherical case, not all angles satisfying $\alpha + \beta + \gamma > \pi$ are possible: $\pi - \alpha, \pi - \beta, \pi - \gamma$ represent the edge lengths of the “dual” or “polar” spherical triangle, so must satisfy three triangle inequalities (see, for example, [5]).

Example 3.3.

A closed, orientable, constant curvature orbifold of dimension 2 or 3 is a cone-manifold with cone angles of the form $2\pi/m$, where $m \geq 2$ is an integer. The singular locus locally looks like:




In 3.6 below, we will see that if M is an orientable 3-dimensional cone-manifold with cone angles at most π then the singular locus is again a trivalent graph, and the link of every point in the singular locus is a football or turnover.

Example 3.4.


Let $T^3 = S^1 \times S^1 \times S^1$ be a 3-torus and let Σ consist of three simple closed curves in orthogonal directions, say $S^1 \times p \times q, q \times S^1 \times p, p \times q \times S^1$ where p, q are distinct points in S^1 . Then $T^3 - \Sigma$ is homeomorphic to the complement of the Borromean rings, as can be seen in 2.32.

We will construct hyperbolic cone-manifold structures on T^3 with arbitrary cone angles $0 \leq \alpha, \beta, \gamma < 2\pi$ along the three components of Σ . (Cone angle “zero” corresponds to a cusp.) Further, these hyperbolic structures on T^3 degenerate as any angle approaches 2π .

We begin with a polyhedron Q in \mathbb{H}^3 as shown below.




Such a polyhedron can be constructed whenever α, β, γ are less than 2π . For example, Q can be obtained from four copies of a hyperbolic cube C with dihedral angles as shown below; such a cube exists by a theorem of Andreev [3] which characterizes the convex polyhedra in \mathbb{H}^3 with all dihedral angles $\leq \pi/2$.




Gluing together the opposite vertical faces of Q gives a hyperbolic structure on $T^2 \times I$ bent along a curve with angle $\alpha/2$ on top and a curve with angle $\beta/2$ on the bottom, and with a vertical axis with cone angle γ . Finally, doubling this gives a hyperbolic cone-manifold structure on T^3 with cone angles α, β, γ along the components of Σ .

Here is a direct construction of Q , which also gives some idea of the variation in the shape of Q as the angles α, β, γ are varied. Begin with a “vertical” geodesic in \mathbb{H}^3 and two pairs of planes: one pair meeting in a ridge line with angle $\alpha/2$, the other meeting along a valley line with angle $\beta/2$. Further, the geodesics defining the ridge and the valley should be orthogonal to the vertical geodesic and to each other (after vertical translation).



There is a one parameter family of planes orthogonal to the ridge meeting one side of the valley; with angles of intersection varying monotonically from


0 to $\pi - \beta/4$. Hence, there is a unique such plane for which the angle is $\pi/2$.




Similarly, we find three other vertical planes; by symmetry these meet at the same angle γ . Finally, it is easy to see that the angle γ varies from $\pi/2$ to 0 as the vertical distance from the valley to ridge varies from 0 to ∞ .

From this description of Q we can see how the hyperbolic cone-manifold structures on T^3 degenerate as any cone angle approaches 2π .


(a) As $\alpha, \beta \rightarrow 2\pi$, with $\gamma < 2\pi$ fixed, the polyhedra flatten out to give a 2-dimensional hyperbolic limit: a hyperbolic torus with one cone point of angle γ .



(b) As $\alpha, \beta, \gamma \rightarrow 2\pi$, the geometric limit (see chapter 6) can be either a point, circle or a line, depending on the exact mode of convergence. In each case there is a limiting Euclidean structure after rescaling.



(c) If $\gamma \rightarrow 2\pi$ with $\alpha, \beta < 2\pi$ fixed, then the hyperbolic structures have diameter going to infinity, while two tori (with cone angle γ) become smaller and smaller, looking more and more like Euclidean tori as $\gamma \rightarrow 2\pi$. The limiting polyhedra give a hyperbolic structure on the manifold obtained from T^3 – (two horizontal curves) by splitting open along two incompressible tori.



(This process is a 3-dimensional analogue of the process of creating a cusp in a hyperbolic surface by pinching a curve to a point, in going to the boundary of Teichmüller space. Compare example 6.9.)

Exercise 3.5.

(a) Show that there are hyperbolic cone-manifold structures on S^3 with the Borromean rings as the singular locus, with arbitrary cone angles satisfying

$0 \leq \alpha, \beta, \gamma < \pi$. [Hint: reassemble 8 copies of a cube C of the type described above.]

(b) Describe how these hyperbolic structures degenerate as any angle approaches π .

(c) Show that $T^3 - \Sigma$ is homeomorphic to the complement of the Borromean rings in S^3 .

Many other examples of 3-dimensional hyperbolic cone-manifolds arise from Thurston's theory of hyperbolic Dehn surgery. A detailed discussion of this will be given in section 5.6 below.

3.2 Local structure

Each point in an n -dimensional cone-manifold has a neighbourhood called a *standard cone neighbourhood* which is an *open cone* $\text{Cone}_K(S; R)$ with constant curvature K and radius R , based on a *spherical cone-manifold* S of dimension $(n - 1)$.

Topologically, $\text{Cone}_K(S; R)$ is $S \times [0, R)$ with $S \times \{0\}$ collapsed to a point. The metric is

$$ds^2 = dr^2 + s_K^2(r) d\theta^2$$

where $d\theta^2$ denotes the metric on S , $r \in [0, R)$ and

$$s_K(r) = \begin{cases} \frac{1}{\sqrt{K}} \sin(\sqrt{K}r) & \text{if } K > 0, \\ r & \text{if } K = 0, \\ \frac{1}{\sqrt{|K|}} \sinh(\sqrt{|K|}r) & \text{if } K < 0. \end{cases}$$

If $K \leq 0$ then $\text{Cone}_K(S; R)$ is defined for $0 < R \leq \infty$. If $K > 0$ then $\text{Cone}_K(S; R)$ is defined for $0 < R \leq \frac{\pi}{\sqrt{K}}$ and we define the *suspension* $\text{Susp}_K(S)$ of S to be the completion of $\text{Cone}_K(S; R)$ where $R = \frac{\pi}{\sqrt{K}}$. This suspension is obtained by gluing together two closed cones of radius $\frac{\pi}{2\sqrt{K}}$; the centres of these cones are called the *suspension points*. These cones and suspensions are analogues of standard balls and spheres for constant curvature cone-manifolds.


3.3 Standard cone neighbourhoods

Given a point p in a n -dimensional cone-manifold, let $S_r(p)$ denote the set of points at distance r from p . Then for r sufficiently small, $S_r(p)$ with the

induced path metric is an $(n - 1)$ -dimensional cone-manifold of constant positive curvature.

Let $\mathbb{H}^n(K)$ denote the complete, simply connected n -manifold of constant curvature K . Then each *non-singular point* of an n -dimensional cone-manifold has a neighbourhood which is a metric ball in $\mathbb{H}^n(K)$.

If $n = 2$, each *singular point* has a cone neighbourhood which is a metric ball as shown below:



In polar coordinates the metric has the form:


$$ds^2 = dr^2 + s_K^2(r)d\theta^2,$$

where $0 \leq \theta \leq \alpha$, $r \geq 0$.

If $n = 3$, then each point has a neighbourhood which is a cone on a spherical 2-orbifold. We will see in 4.2 that every orientable spherical cone 2-manifold with cone angles at most π is either a sphere, a football (as in example 3.1) or a turnover (as in example 3.2).

Proposition 3.6. *Every point in the singular locus of an orientable 3-dimensional cone-manifold with cone angles at most π has neighbourhood which is the cone on a football or turnover. Thus the singular locus is a trivalent graph.*

A cone neighbourhood of *singular point* which is not a vertex looks like:




In cylindrical coordinates the metric has the form:

$$ds^2 = dr^2 + s_K^2(r)d\theta^2 + c_K^2(r)dz^2,$$

where $0 \leq \theta \leq \alpha$, $r \geq 0$, $z \in \mathbb{R}$, and $c_K(r) = s'_K(r)$. (For example, $s_K(r) = \sinh(r)$, $c_K(r) = \cosh(r)$ in the hyperbolic case where $K = -1$.)


A vertex has a neighbourhood which is obtained by taking two isometric tetrahedra in $\mathbb{H}^3(K)$ and identifying along the three faces containing the vertex. A standard cone neighbourhood is a metric ball centred on the vertex inside this neighbourhood.



3.4 Geodesics

A cone-manifold has a *path metric* in the sense of Gromov and Alexandrov: the *distance* between two points is the infimum of lengths of paths joining the points. A *geodesic* in a cone-manifold is a curve which is locally length minimizing. At each point p of a cone-manifold M there is a *tangent cone* T_pM isometric to a Euclidean cone: T_pM is the union of the Euclidean tangent cones to all the n -dimensional simplices containing p . The subset

S_p of unit tangent vectors at p is a spherical cone-manifold — the visual sphere at p .



Next we outline some important results on geodesics in cone-manifolds; see [47] and [12] for more details.

Lemma 3.7 (Hopf-Rinow). *Let M be a complete, connected cone-manifold.*

(i) *Then any two points in M can be joined by a geodesic of length equal to the distance between the points.*


(ii) *Given any vector v in $T_p M$, there is a constant speed geodesic g_v in M with initial tangent vector v , and g_v is uniquely defined in some neighbourhood of p .*

Proof. Part (i): Take a sequence of constant speed curves $\gamma_i : [0, 1] \rightarrow M$ joining p to q , with $\text{length}(\gamma_i) \rightarrow \text{dist}(p, q)$ as $i \rightarrow \infty$. This is bounded and equicontinuous, so has a convergent subsequence by Ascoli-Arzelà. The limit is a geodesic. Part (ii) is easy. \square


Exercise 3.8. Check the details in the above argument.

Lemma 3.9 (Lemma A). *In a cone-manifold with all cone angles $< 2\pi$, if the interior of a geodesic contains a point in the singular locus then the entire geodesic is contained in a single stratum of the singular locus.*

First we will sketch the main idea of the proof in the 2-dimensional case. If a curve meets Σ at a point p and is not entirely contained in Σ , then (the smallest) angle between the incoming and outgoing arcs at p is $< \pi$. Hence, the curve can be shortened by smoothing the corner at p .



Note: This result fails if cone angles larger than 2π are allowed. In this case length minimizing geodesics may pass through cone points. In fact, there is a pencil of extensions of any geodesic at any cone point with angle $> 2\pi$, consisting of all outgoing geodesic arcs making an angle $\geq \pi$ with the incoming arc.



Lemma 3.9 implies that if a vector v is not tangent to Σ then the geodesic g_v can be extended until it meets the singular locus. If v is tangent to Σ then g_v is contained in Σ and can be extended until it meets a different strata of the singular locus.

Lemma 3.10 (Lemma B). *A connected spherical cone-manifold of dimension at least 2 and curvature 1 with all cone angles $< 2\pi$ has diameter $\leq \pi$, with equality if and only if the cone-manifold is a suspension. Further, the only two points at distance π apart are the suspension points.*

Proof. We prove Lemmas 3.10 and 3.9 simultaneously by induction on dimension, following an argument of Thurston (compare [44]). Let $(A_n), (B_n)$ be the statements of Lemmas A and B for n -dimensional cone-manifolds. The statement (A_1) is trivially true. Lemma B does not apply in dimension 1.

Assume (A_{n-1}) and (B_{n-1}) are true for $n \geq 2$. Let g be a length minimizing geodesic in an n -dimensional cone-manifold. Suppose p is a singular point in the interior of g , and let v_- and v_+ be the unit tangent vectors to g at p , directed away from p . Then v_- and v_+ give two points in the unit tangent cone S_p , which is a spherical cone-manifold of dimension $(n-1)$. The angle between these tangent vectors is the distance between v_- and v_+ measured in S_p so is $\leq \pi$, by the hypothesis on cone angles if $n = 2$ and by the induction hypothesis (B_{n-1}) if $n > 2$. If this angle is $< \pi$ then the length of g could be reduced by smoothing the corner at p , so g would not be locally length minimizing. If the angle is equal to π then $n > 2$ and it follows from (B_{n-1}) that g is tangent to Σ at p , hence contained in Σ . This proves (A_n) .

To prove (B_n) , we consider a length minimizing geodesic g in an n -dimensional spherical cone-manifold S of curvature 1. By (A_n) , either the interior of g is disjoint from Σ or the entire geodesic g is contained in a single stratum of Σ . Elementary spherical geometry therefore shows that g has length $\leq \pi$; hence S has diameter $\leq \pi$. Further, if g has length π then its interior has a neighbourhood which is a suspension. Now suppose that S contains two points p, q at distance π apart. Let U be the set of unit tangent vectors $v \in S_p$ such that the geodesic g_v with initial tangent vector v is length minimizing on $[0, \pi]$ and joins p to q . Then U is an open subset of S_p by our previous remark. Suppose that v is a vector in the closure of U . We claim that the geodesic g_v is defined on $[0, \pi]$. If not, then g_v ends at some singular point $p_0 \in \Sigma$. Let g' be a shortest geodesic from p_0 to q . Then $g_v \cup g'$ has length $\leq \pi$ since v lies in \bar{U} . But this curve has a corner at p_0 , so can be shortened to give a path from p to q of length less than π , contradicting the choice of p and q . Since the exponential map is continuous, it now follows that $v \in U$. Hence, $U = S_p$ and M is the suspension of S_p , with p and q as suspension points. \square


3.5 Exponential map

Recall that g_v denotes the constant speed geodesic starting at p with initial tangent vector v . Let $\mathcal{D}_p \subset T_p M$ be the subset consisting of all v in $T_p M$ such that g_v is defined and globally length minimizing up to time 1. We also say that $g_v([0, 1])$ is a *segment*.

Then there is a well-defined, continuous *exponential map*

$$\exp : \mathcal{D}_p \rightarrow M,$$

defined by $\exp(v) = g_v(1)$.



From the previous lemmas we obtain:

Lemma 3.11. *If M is a complete, connected cone-manifold with all cone angles $< 2\pi$ then the exponential map $\exp : \mathcal{D}_p \rightarrow M$ is onto.*

To study $\partial\mathcal{D}_p$ and the *cut locus* $\exp(\partial\mathcal{D}_p)$, we need to understand:
How can a geodesic from p stop minimizing length?


If $f : M \rightarrow N$ is a continuous map between cone-manifolds, then it is possible to consider directional derivatives at each point of M . If all such directional derivatives exist, then one obtains a map $df : T_p M \rightarrow T_{f(p)} N$. We say that q is a *conjugate point* of p in M if $d\exp(v) = 0$ for some non-zero $v \in TM_p$ such that $\exp(v) = q$.

Lemma 3.12 (Key Lemma). *A point $q \in M$ lies in the cut locus $\exp(\partial\mathcal{D}_p)$ if and only if at least one of the following cases occurs:*

- (1) *There are two minimizing geodesics from p to q in M .*
- (2) *q is a conjugate point of p .*
- (3) *q is in the singular locus Σ and there is a length minimizing geodesic from p to q which doesn't extend past the point q .*


Proof. Use the usual arguments from Riemannian geometry plus Lemma 3.9.

Note: (2) is very special in our context: it only occurs when the curvature $K > 0$, and the diameter of the cone-manifold is π/\sqrt{K} .



Further, if (2) occurs then (1) also occurs for our constant curvature spaces.

Picture for case (1):




3.6 Dirichlet domains

Define the (*open*) *Dirichlet domain* $\overset{\circ}{D}_p$ at a point p as the subset of $q \in M$ such that there is a unique minimal geodesic from p to q , and if $q \in \Sigma$ then the entire geodesic is in Σ .

By the previous lemma, this is just the image of the interior of \mathcal{D}_p under the exponential map: $\overset{\circ}{D}_p = \exp(\overset{\circ}{\mathcal{D}}_p)$. Hence $\overset{\circ}{D}_p$ is a union of geodesic segments starting at p , so is *star-shaped* with respect to p .

Example 3.13. A Dirichlet domain for the Euclidean orbifold $S^2(2, 4, 4)$ (the double of a triangle with angles $\pi/2, \pi/4, \pi/4$.)



Let M be a cone-manifold of curvature K , let p be a point in M , and S_p the unit tangent cone (or visual sphere) at p . Define the *global cone* C_p at p to be the infinite cone $\text{Cone}_K(S_p; \infty)$ if $K \leq 0$, or the suspension $\text{Susp}_K(S_p)$ if $K > 0$.

Now assume that p has no conjugate points. Then the exponential map $\exp : \mathcal{D}_p \rightarrow M$ has no critical points, and we can pull back the metric on M to obtain a constant curvature metric on \mathcal{D}_p .

Since \mathcal{D}_p is star-shaped, it also embeds by a local isometry into the global cone C_p . The image D_p of this embedding will be called the *Dirichlet domain* for M based at the point p . Then D_p is isometric to \mathcal{D}_p , and there is a natural map

$$q : D_p \rightarrow M$$

which is *onto* by Lemma 3.11 if all cone angles are at most 2π .

If p is a non-singular point, then the global cone C_p is just the simply connected n -manifold $\mathbb{H}^n(K)$ of constant curvature K , and D_p is a subset of $\mathbb{H}^n(K)$. If M is a complete non-singular manifold, then D can be obtained by taking the usual Dirichlet region based at any preimage of p in the universal cover $\tilde{M} \cong H^n(K)$. Note that this is always a convex polyhedron in $H^n(K)$.

Special Case: If p has a conjugate point, then Lemma 3.10 shows that M is a suspension and p is a suspension point. In this case, we will regard $C_p = \text{Susp}_K(S_p) = M$ as the Dirichlet domain for M at p , and consider this to be a convex polyhedral subset of C_p .

Next we examine the structure of the Dirichlet domain in more detail. The following result, in particular the convexity conclusion, will play a key role at several places in the proof of the orbifold theorem.

Proposition 3.14. *Let D be the Dirichlet domain based at a point p of a cone-manifold M of constant curvature K .*

(1) *If M has cone angles $< 2\pi$, then D is a star-shaped (geodesic) polyhedron in the global cone C_p . (i.e. D is a union of totally geodesic simplices with one vertex at p .) Further, M is obtained from D by identifying pairs of faces by isometries.*


(2) *If M has cone angles $\leq \pi$, then the polyhedron D is convex (i.e. any two points of D can be joined by a minimal geodesic lying in D .)*

Before giving a detailed proof, we will sketch the main ideas in the 2-dimensional case.

Sketch of proof (assuming $K \leq 0$): We look at the local picture near a point q in the cut locus $\exp(\partial\mathcal{D}) \subset M$.


Case 1: q is non-singular. Then there are at least two shortest geodesic in M from p to q , say $\gamma_1, \dots, \gamma_k$.

These come out from q in k different directions, and end at points p_1, \dots, p_k (corresponding to different “lifts” of p .)




Then the shortest geodesics from p to points x near q are perturbations of $\gamma_1, \dots, \gamma_k$ joining p_1, \dots, p_k to x . So the part of $\overset{\circ}{D}$ near q consists of the *Voronoi regions*

$$D_i = \{x: d(x, p_i) \leq d(x, p_j) \text{ for all } j \neq i\}.$$




These are bounded by planes equidistant from two points, so the boundary of D is polyhedral and locally convex near q .

Case 2: q is singular and there is one shortest geodesic from p to q . Then a wedge of angle $2\pi - \text{cone angle}$ is excluded from the Dirichlet domain near Σ .



Again the boundary of D is polyhedral near q . It is locally convex if and only if the cone angle is $\leq \pi$.

Another view:



Case 3: q is singular and there are at least two shortest geodesics from p to q .

Exercise. Prove case 3 by combining the arguments from cases 1 and 2.

Proof of 3.14. We will prove a slightly more general result by induction on $\dim M$. Given a finite set of points $P = \{p_i\}$ in a cone-manifold M then the (open) *Voronoi region at p_i* is the set $V(p_i)$ consisting of points $q \in M$ satisfying:

- (1) $d(q, p_i) \leq d(q, p_j)$, for all $j \neq i$,
- (2) there is a unique shortest geodesic γ from q to P , and
- (3) if $q \in \Sigma$, then the global cone C_q is a suspension, and the unit tangent vector to γ at q gives a suspension point in C_q .

We also define the *cut locus of P* , $\text{Cut}(P)$, to be the complement of $\cup_i V(p_i)$ in M .

If P contains a single point p , then the cut locus is $\text{Cut}(p) = \exp(\partial \mathcal{D}_p)$, and $V(p)$ is just the complement of $\text{Cut}(p)$ in M , i.e. an “open Dirichlet domain” at p .

We inductively prove the following statement.

Convex(n): Let M be a cone-manifold of dimension n with cone angles $\leq 2\pi$. Let $P = \{p_1, \dots, p_k\}$ be a set of points in M . Also assume that if $k = 1$ then the point in P has no conjugate point in M .

(a) Then each Voronoi region $V(p_i)$ is a (geodesic) polyhedron contained in the global cone at p_i .

(b) Further, if M has cone angles $\leq \pi$, then each Voronoi region is convex.

Proof. If $\dim M = 1$, then each $V(p_i)$ is an interval so the result is clear.

To show that Convex($n - 1$) \Rightarrow Convex(n), we study the local structure of the cut locus $\text{Cut}(P)$ near a point q_0 . The condition on P means that q_0 is not conjugate to any point p_i . Hence, there exists $\delta > 0$ such that only finitely many geodesics of length $< d(p, q_0) + \delta$ join q_0 to P . In particular, there are finitely many shortest geodesics from q to P for all q near q_0 and these are obtained by *small perturbations* of the shortest geodesics from q_0 to P . (Compare [8], Lemma 1.3.)

Let γ be a shortest geodesic joining q_0 to a point $p \in P$, with length d_0 . Then if γ' is a geodesic sufficiently close to γ joining a point q to p , then the law of cosines in $H^n(K)$ expresses $\text{length}(\gamma')$ as an explicit function depending only on d_0 , the distance from q_0 to q , and the angle $\theta = \angle pq_0q$ between γ and q_0q . Further, $\text{length}(\gamma')$ is an increasing function of θ for $0 \leq \theta \leq \pi$, if we fix d_0 and $d(q_0, q)$.

Let $v_1, v_2, \dots, v_k \in S_{q_0}$ be initial unit tangent vectors of the shortest geodesics from q_0 to P . Then the observation above gives us an explicit local description of the sets $V(p_i)$ near q_0 as cones on the Voronoi regions of the collection of points $\{v_1, v_2, \dots, v_k\}$ on S_{q_0} . (Note that this is a spherical cone-manifold of dimension $n - 1$.) To finish the argument, we look separately at the geometry of these cones for the cases where $n = 2$ and $n \geq 3$.

If $n = 2$, S_{q_0} is circle of length equal to the cone angle α of M at q_0 . If we have at least two shortest geodesics from q_0 to P , then we have $k \geq 2$ points v_i on a circle of length $\alpha \leq 2\pi$. Hence each Voronoi region $V(v_i)$ has length $\leq \pi$, and the cone on this is polyhedral with angle $\leq \pi$ at the cone point.

If q_0 is a *singular* point of M , it is possible that there is only *one* shortest geodesic from q_0 from P . In this case, the Voronoi region of the corresponding point on S_{q_0} is an interval of length α , and the cone on this is polyhedral with angle α at the cone point. This will be locally convex if and only if $\alpha \leq \pi$.

So for $n = 2$ we see that the Voronoi regions in M are locally polyhedral,

and have corner angles $\leq \pi$, *provided* that all cone angles are $\leq 2\pi$.

If $n \geq 3$, we use the inductive hypothesis directly. By Lemma 3.12 and the following remarks, the points v_1, \dots, v_k satisfy the assumption for Convex($n - 1$) applied to the $(n - 1)$ -dimensional spherical cone-manifold S_{q_0} . By induction each of these Voronoi regions $V(v_i)$ is polyhedral and is convex if the cone angles are $\leq \pi$. But the cone on such a region is always polyhedral and locally convex if $n \geq 3$. This shows that the $V(p_i)$ satisfies the conclusions of Convex(n) locally near each point q_0 in $Cut(P)$ and completes the proof by induction. \square

3.7 Area and volume of cone-manifolds

For 2-dimensional cone-manifolds we have the following following version of the Gauss-Bonnet theorem relating topology and geometry:

Theorem 3.15. *Let M be a closed, 2-dimensional cone-manifold of constant curvature K with n cone points with cone angles $\theta_1, \dots, \theta_n$. Then*

$$\int_M K dA + \sum_i (2\pi - \theta_i) = 2\pi\chi(M). \quad (1)$$

Exercise 3.16. Prove this Gauss-Bonnet formula.

In particular, this gives an explicit formula relating the area A to the cone angles for a 2-dimensional cone-manifold of constant curvature K :

$$KA = 2\pi\chi(M) - \sum_i (2\pi - \theta_i).$$

For higher dimensional cone-manifolds, there is a remarkable formula for the *variation* of volume when constant curvature cone-manifolds are deformed. The basic ingredient needed is the following theorem, originally proved in the spherical case by Schläfli [70] in 1858. For modern proofs see, for example, [86], [63], or [21].

Theorem 3.17 (Schläfli Formula for Polyhedra). *Let P_t be a smooth one-parameter family of polyhedra in the simply-connected n -dimensional space $\mathbb{H}^n(K)$ of constant curvature K . (Thus the faces of X_t are totally geodesic planes in $\mathbb{H}^n(K)$ which vary smoothly with t .) Then the derivative of the volume V_t of P_t satisfies the equation:*

$$(n - 1)K \frac{dV_t}{dt} = \sum_F V_{n-2}(F) \frac{d\theta_F}{dt} \quad (2)$$

where the sum is over all codimension-2 faces of X_t , and V_{n-2} denotes the $(n-2)$ -dimensional volume, and θ_F denotes the dihedral angle at F .

In the 2-dimensional case, the theorem is equivalent to the Gauss-Bonnet theorem for constant curvature polygons. The formula becomes

$$K \frac{dA_t}{dt} = \sum \frac{d\theta_v}{dt}$$

where θ_v is the angle at vertex v and A denotes area. Integrating this gives

$$KA = \sum \theta_v + \text{constant}.$$

Exercise 3.18. Determine the constant of integration by considering the case where the polygon flattens out to a straight line segment.

Remark: As Milnor notes in [63], the theorem also applies to the case of hyperbolic polyhedra with some ideal vertices. When $n \neq 3$, no change in the statement of the theorem is needed. In the 3-dimensional case, some edge lengths $V_{n-2}(F)$ become infinite. However, the theorem remains valid if we remove small horoball neighbourhoods of the ideal vertices before measuring edge lengths. (The right hand side of (2) is easily seen to be independent of the choice of horoballs, using the fact that the sum of dihedral angles at an ideal vertex is constant.)

Exercise 3.19. Use the Schläfli formula for ideal simplices to prove that the regular ideal simplex is the unique simplex of maximal volume in \mathbb{H}^3 .

Following Hodgson [43] we can apply the Schläfli formula to study the volume of constant curvature cone-manifolds. The following theorem shows that the variation in volume for a family of cone-manifold structures is completely determined by the changes in geometry *along the singular locus*.

Theorem 3.20 (Schläfli Formula for Cone-Manifolds).

Let C_t be a smooth family of cone-manifold structures of constant curvature K . Assume that the underlying space and singular locus are of fixed topological type. Then the derivative of volume V_t of C_t satisfies

$$(n-1)K \frac{dV_t}{dt} = \sum_i V_{n-2}(\Sigma_i) \frac{d\theta_i}{dt} \quad (3)$$

where the sum is over all strata Σ_i of the singular locus Σ , and θ_i is the cone angle along Σ_i .

Proof. Divide C_t into geometric simplices, varying smoothly with t , such that the singular locus remains a subcomplex. Applying the Schläfli formula (2) to each simplex and adding shows that the variation of volume is given by

$$(n-1)K \frac{dV_t}{dt} = \sum_F V_{n-2}(F) \frac{d\theta_F}{dt}$$

summed over all codimension 2 faces of the triangulation of C_t , where θ_F denotes the cone angle along F . However, at any non-singular face F the cone angle is 2π for all t so $\frac{d\theta_F}{dt} = 0$. So the right hand side reduces to a sum over faces F in the singular locus Σ . \square

As a corollary we see that the volumes of cone-manifolds satisfy the following remarkable monotonicity property:

Corollary 3.21. *Let C_t be a smooth family of cone-manifolds of constant curvature K . If $K > 0$, then the volume increases strictly monotonically as any cone angle is increased. If $K < 0$, then the volume decreases strictly monotonically as any cone angle is increased.*

Chapter 4

Two-dimensional Cone-Manifolds

In this chapter, we discuss the classification of 2-dimensional geometric cone-manifolds. The main ingredients needed for the classification are the Dirichlet domains and Gauss-Bonnet theorem for cone-manifolds 3.15 discussed in the previous chapter, and the developing map and holonomy representation.

4.1 Developing map and holonomy

Let M be a constant curvature cone-manifold. Although M itself has singularities, if we remove the singular set, Σ , the remainder, $M - \Sigma$, has a smooth (G, X) structure. In particular, can still define the *developing map* on the universal cover $\widetilde{M - \Sigma}$ of $M - \Sigma$:

$$\text{dev} : \widetilde{M - \Sigma} \rightarrow X$$

There also exists a *holonomy* homomorphism


$$h : \pi_1(M - \Sigma) \rightarrow G$$

satisfying $\text{dev} \circ \gamma = h(\gamma) \circ \text{dev}$ for each deck transformation γ in $\pi_1(M - \Sigma)$.

However, the difference between this case and that for a compact *manifold* is that the holonomy group need not be discrete and the developing map need not be a diffeomorphism. In particular this means that neither M nor $M - \Sigma$ can be described as X/Γ , where Γ is the image of the holonomy homomorphism.

The key property that distinguishes the two cases is that of *completeness*. When N is a complete constant curvature manifold, then $\text{dev} : \tilde{N} \rightarrow X$ is a covering map, hence a diffeomorphism. However, if $N = M - \Sigma$, then it will be *incomplete*. Even though dev is a local diffeomorphism, it will not satisfy the unique path lifting property.

Example 4.1. Developing map for a cone with angle α .




The element of the holonomy group corresponding to a loop linking a component of the singular locus is a rotation by angle α , where α is the cone angle of that component.

4.2 Two-dimensional spherical cone-manifolds

Proposition 4.2. *A 2-dimensional orientable spherical cone-manifold M with cone angles $\leq \pi$ is a 2-sphere with 0, 2 or 3 cone points. The metrics are described as follows:*

- (i) 0 cone points: $M = S^2$
- (ii) 2 cone points: A spherical football, with two equal cone angles α . (M is the double of lune of angle $\alpha/2$)
- (iii) 3 cone points: A spherical turnover, with cone angles α, β, γ such that $\alpha + \beta + \gamma > 2\pi$. (M is the double of a spherical triangle with angles $\alpha/2, \beta/2, \gamma/2$.)

Note: The local models for 3-dimensional cone-manifolds with cone angles $\leq \pi$ are cones on these.



Proof. Note that a spherical cone-manifold with all cone angles less than 2π is compact. This follows from the fact that every minimal length geodesic has length no more than $\pi (\frac{\pi}{\sqrt{K}}$ for curvature K). Thus the diameter of the Dirichlet domain is at most 2π and its closure is compact.

The Gauss-Bonnet theorem 3.15 gives


$$\int_M K dA + \sum_i (2\pi - \theta_i) = 2\pi\chi(M),$$

where $\theta_1, \dots, \theta_k$ are the cone angles. Since the curvature $K = 1$ and the cone angles are at most π it follows that $\chi(M) > 0$, thus $M = S^2$ and $\sum_i (2\pi - \theta_i) < 2\pi\chi(S^2) = 4\pi$. Since the cone angles are at most π , there are at most 3 cone points.

The developing map $\text{dev} : (\widetilde{M - \Sigma}) \rightarrow S^2$ induces a holonomy representation $h : \pi_1(M - \Sigma) \rightarrow SO(3)$. The holonomy of a loop around a cone point is a rotation r_θ by an angle equal to the cone angle θ . Using this fact, one sees that

- (i) it is impossible to have a single cone point.
- (ii) if there are two cone points, then the two cone angles are equal.
- (iii) if there are three cone angles α, β, γ , then the cone angles are twice the angles of some spherical triangle.

The figure below proves (i) and (ii). To see (iii), observe that $\pi_1(M - \Sigma)$ has a presentation $\langle A, B, C \mid AB = C \rangle$, where A, B, C are meridians. Given the fixed points of the rotations A and B we can construct the product $C = AB$ as follows. Write $\text{hol}(A) = R_1 \circ R_2$, $\text{hol}(B) = R_2 \circ R_3$, where R_i are reflections with axis(R_2) joining $\text{Fix}(A)$ to $\text{Fix}(B)$; axis(R_1) through $\text{Fix}(A)$ at angle $\alpha/2$ to axis(R_2); and axis(R_3) through $\text{Fix}(B)$ at angle $\beta/2$ to axis(R_2). Then one sees immediately that $\text{hol}(C) = R_1 \circ R_3$ is a rotation through twice the third angle of the spherical triangle formed by the axes.




This shows that the *holonomy* has one of the forms indicated in the proposition, but we don't yet know the cone-manifold structure has the form claimed. To see this we show that M is obtained from a convex spherical polygon and therefore has a "standard" structure.

If there are two cone points choose a geodesic joining them. Cut M along this geodesic to obtain D . Then D is simply connected so the developing map sends D to a region in the sphere bounded by two geodesics. The angle between these geodesics is the cone angle at each cone point. Since this cone angle is at most π the developing map embeds D . Hence D is isometric to a lune and M is a football.

If there are three cone angles α, β, γ , then take minimal geodesics from one of the singular points to each of the other two. These are disjoint; otherwise we can cut, paste and straighten to get a shorter geodesic.

The complement of these two geodesics is simply connected and thus maps into S^2 via the developing map. As before this map is an embedding. Its image in S^2 is bounded by two copies of each geodesic. From the figure we see that $\alpha_1 = \alpha_2$.



This implies that the region consists of two isometric triangles. The boundary of each triangle consists of the images of the two geodesics from x , together with another geodesic connecting the other two singular points. \square

A 2-dimensional spherical cone-manifold is an orbifold if and only if the cone angle at each cone point is of the form $2\pi/n$ for some integer $n \geq 2$. Then it is the quotient of S^2 by a discrete group of isometries by theorem 2.26 (Poincaré’s theorem).

Since a neighbourhood of a point in a 3-dimensional cone-manifold M is a cone on a 2-dimensional spherical cone-manifold it follows that the underlying topology of $(M, \Sigma(M))$ is the same as that of an orbifold when the cone angles in M are at most π . However this is no longer the case when cone angles larger than π are allowed. For example take the cone on the double of any spherical polygon.

Exercise 4.3. Extend proposition 4.2 to the non-orientable case.


4.3 Two-dimensional euclidean cone-manifolds

We can classify 2-dimensional Euclidean cone-manifolds by the techniques we used for spherical cone-manifolds.

Proposition 4.4. *A complete (orientable) 2-dimensional Euclidean cone-manifold C with cone angles $\leq \pi$ is one of the following:*

- (a) compact
 - (i) a Euclidean turnover (the double of an acute angled Euclidean triangle)
 - (ii) a pillowcase (the quotient of a Euclidean torus by an isometric involution with 4 fixed points as in example 2.1)

- (iii) a torus.
 (b) non-compact
 (i) infinite cylinder,
 (ii) infinite pillowcase,
 (iii) infinite cone,
 (iv) plane.




Proof. Assume first that M is compact. If M has no singular points then M is a torus. Otherwise, the Gauss-Bonnet theorem 3.15 gives $2\pi\chi(M) = \sum (2\pi - \text{cone angle})$. Hence $\chi(M) > 0$ so M is a sphere, and there are 3 or 4 cone points. When there are 3 cone points the same argument as in the spherical case shows that it is the double of a triangle.

When there are 4 cone points, they will all have angle π and M is a geometric orbifold with group \mathbb{Z}_2 at each singular point. The orbifold

fundamental group has a presentation


$$\langle a, b, c, d \mid a^2 = b^2 = c^2 = d^2 = abcd = 1 \rangle,$$

as in example 2.10. The kernel of the map to \mathbb{Z}_2 sending each generator to 1 determines a 2-fold regular orbifold covering space which is a torus.




Suppose M is non-compact. If M has no cone points then M is an infinite cylinder or plane. If there is a cone point x , then since M is non-compact there is a (length minimizing) ray γ starting at x and going out to infinity. (Choose a sequence of points in M going out to infinity. Join these points to the base point; then the limiting position of the line segments gives a ray.) We claim that by choosing a suitable basepoint, p , there is a Dirichlet domain P for M with γ contained in ∂P . To see this there is a neighbourhood U of x which is a standard cone. Then $U \cap \gamma$ is a geodesic. Choose p in U close to x and symmetric with respect to γ . Clearly the cut locus for p inside U is $\gamma \cap U$. If γ were not contained in ∂P this would contradict that γ is length minimizing.

We now regard P as a convex subset of the plane \mathbb{E}^2 . Observe that ∂P contains two copies of γ meeting at an angle equal to the cone angle at x .




If the cone angle is $< \pi$, P contains the entire wedge between the rays by convexity; so M is a cone. If the angle is π , then ∂P contains a line L made up of two copies of γ . Hence either P is a half-plane and M is a cone,


or P is an infinite strip lying between two parallel lines.



M is obtained from P by isometric identifications on ∂P . If P is a strip, both lines must be identified to themselves by an order 2 rotation. The order 2 points must both lie on a common perpendicular to the two parallel lines in ∂P , otherwise γ would not be length minimizing. Hence M is an infinite pillowcase. \square



Remark: The infinite pillowcase has a geodesic connecting the cone points. It is the image of an affine subspace invariant under the holonomy group. The orthogonal affine subspaces are also invariant, forming the orbifold normal bundle to the geodesic.



Exercise 4.5. Extend proposition 4.4 to the non-orientable case.


Exercise 4.6. Show that every pillowcase is isometric to some Euclidean tetrahedron which has opposite edges of equal length.

4.4 Euclidean examples with large cone angles


If the cone angles are allowed to be greater than 2π , the topology, even in the compact case, is uncontrolled, i.e. the genus can be anything. Such Euclidean metrics arise naturally from *quadratic differentials* on a Riemann surface of genus $g \geq 2$ (see Strebel [77, chapter 2]). A quadratic differential can be written locally as $\omega = \phi(z) dz^2$ where z is a local complex coordinate on the surface and ϕ is a holomorphic function. Then

$$ds = |\sqrt{\omega}| = |\phi(z)|^{1/2} |dz|$$


gives a well-defined Euclidean metric with cone points at the zeros of ϕ .



If the cone angles are allowed to be $\pi < \alpha < 2\pi$ and the space is compact then the Gauss-Bonnet theorem implies that the underlying space is still a sphere, but the number of singular points is unbounded.



In the non-compact Euclidean case, with $\pi < \alpha < 2\pi$, there are infinitely many examples with infinitely many cone points going out to infinity. These examples are not either a bundle or a cone.



4.5 Spaces of cone-manifold structures

The previous theorems describe the topological structure of certain cone-manifolds. What parameters are there in their geometric structures? Smooth (G, X) -structures on a manifold N are locally parametrized by their holonomy representations $\pi_1(N) \rightarrow G$, up to conjugacy. This will be discussed in more detail in chapter 5.

We can consider a geometric cone-manifold structure on M as an incomplete smooth structure on $M - \Sigma$ whose metric completion is M . The holonomy along a *meridian curve* linking one component of Σ determines the cone angle on that component, as in example 4.1.

Example 4.7. The footballs and spherical turnovers described in Proposition 4.2 are determined by the number of singular points and their cone angles. Similarly for the Euclidean turnovers. The corresponding holonomy representations are either abelian (2 singular points) or of the form

$$\langle A, B, C \mid AB = C \rangle$$

where A, B, C are all rotations of the model space X .

Example 4.8. The *non-compact* Euclidean cone-manifolds of Proposition 4.4 are parametrized by length of closed geodesic in cases (i) and (ii), or by angle in case (iii).

Example 4.9. Euclidean tori and pillowcases are parametrized by lattices in the plane (up to isometry). Representations of $\pi_1(\text{torus}) = \mathbb{Z} \oplus \mathbb{Z}$ are determined by two translation vectors; the pillowcase group is a \mathbb{Z}_2 extension of the torus group.

4.6 Two-dimensional hyperbolic cone-manifolds


The uniformization theorem shows that compact two dimensional hyperbolic *surfaces* are parametrized by their underlying conformal structures.

For a closed surface of genus g , the space of these structures has dimension $6g - 6$.

Two dimensional hyperbolic *cone-manifolds* are parametrized by their underlying conformal structures (including position of singular points) and cone angles. Subject to the restriction on cone angles from the Gauss-Bonnet theorem 3.15, all possible angles and conformal structures occur. (See Troyanov [85], McOwen [60].) For fixed genus and number of cone points, the angle constraint determines a convex set of angles including the origin:


$$\sum_i (2\pi - \alpha_i) > 2\pi\chi.$$

Note that a *cusps* can arise as limit of cone points as cone angle $\alpha \rightarrow 0$.






Finally we mention some of the complications that arise when cone angles $\geq \pi$ are allowed.

Cone points can't collide when $\alpha < \pi$. However, two angle π cone points can become arbitrarily close together; for example, approximating an infinite pillowcase cusp.







When all cone angles are $< \pi$, there is a pair of pants decomposition, obtained by surrounding pairs of singular points by a geodesic. (The resulting “pants” may have some boundary curves replaced by cone points.) When angles are between π and 2π , the shortest curve will pass through the singular points.




If cone angles are only restricted to the interval $(0, 2\pi)$, then two cone points of angles α, β can collide without creating a thin collar. They just become a single cone point of angle $\alpha + \beta - 2\pi$.

Fundamental domains



In dimension 3 the topology of such a collision is much more complicated (a link becomes a graph with vertices). This is a key reason why cone angles at most π are much easier to deal with than cone angles less than 2π .



In three dimensions two circle components of Σ can collide when cone angles are $\geq \pi$

Chapter 5

Deformations of Hyperbolic Structures

5.1 Introduction

Let Q be a compact 3-dimensional orbifold with link singularities. In most cases the manifold $Q - \Sigma$ obtained by removing the singular locus has a finite volume hyperbolic structure. If Q has a hyperbolic structure, it can be viewed as a cone-manifold structure with cone angles of the form $2\pi/m$. In the proof of the Orbifold Theorem we will attempt to connect the complete structure on $Q - \Sigma$, viewed as a cone-manifold with angles 0, with the desired orbifold structure via a family of cone-manifolds.

To study hyperbolic cone-manifold structures on Q with singularities along Σ , we first remove a neighbourhood of Σ from Q to obtain a compact manifold M with boundary consisting of tori. First we investigate when deformations of a hyperbolic structure on M exist. We will show that hyperbolic (or general (G, X)) structures on a compact manifold M are locally in 1–1 correspondence with nearby holonomy representations $\pi_1(M) \rightarrow G$ up to conjugacy. (If M has boundary, we may have to restrict to the complement of a small neighbourhood of the boundary ∂M .)

We then study how the deformed hyperbolic structures behave near the boundary of M . We will see that to find a nearby cone-manifold structure, it suffices to find a nearby holonomy representation for which the holonomy of each meridian is elliptic.

Next we discuss Thurston’s analysis of representation spaces for 3-manifold groups into $PSL(2, \mathbb{C})$ and his theory of hyperbolic Dehn surgery. In particular, this implies that hyperbolic cone-manifold structures on Q with cone

angles, α_i , along the components of Σ exist for all sufficiently small values of α_i .


We finish the chapter with some examples, and some general conjectures on the global structure of hyperbolic Dehn surgery spaces.

5.2 Deformations and degenerations of surfaces

The proof of the orbifold theorem involves deforming a hyperbolic cone metric in an attempt to produce a hyperbolic orbifold. In the process the hyperbolic cone structure can degenerate. By analyzing how this happens one produces some other kind of geometric structure. In this section we will describe some two-dimensional analogues of the degenerations that occur in the proof of the orbifold theorem. The first example shows how Euclidean and spherical structures can arise and the second one suggests how an orbifold Seifert fibre structure can arise.


Example 5.1. Euclidean/Spherical transition

Given $\theta \in [0, 2\pi]$ there is a turnover $M(\theta) = S^2(\theta, \theta, \theta)$ with 3 cone angles θ . This is the double of a triangle of constant curvature with all angles $\theta/2$. It is hyperbolic, Euclidean or spherical depending on whether $\theta/2$ is less than, equal to, or greater than $\pi/3$. When the cone angle $\theta = 0$ this is interpreted as the double of an ideal hyperbolic triangle which gives a hyperbolic three-punctured sphere with three cusps. As $\theta \rightarrow 2\pi/3$ the diameter of $M(\theta)$ approaches 0. We may rescale the metric on $M(\theta)$ by multiplying the metric by a constant $\lambda = \text{diam}(M(\theta))^{-1}$ to obtain $M'(\theta) = \lambda \cdot M(\theta)$ of diameter 1. The curvature of $M'(\theta)$ is $\pm 1/\lambda^2$ and this goes to zero as $\theta \rightarrow 2\pi/3$. In this way one obtains a continuous family of cone metrics $M'(\theta)$ of varying constant curvature for $0 \leq \theta \leq 2\pi$.



Example 5.2. *Degeneration to an Orbifold fibration*

The Gauss-Bonnet theorem implies there is no hyperbolic metric on a torus. It follows that a pillowcase is not a hyperbolic orbifold. In other words there is no hyperbolic cone metric on the sphere with 4 cone points each with cone angle π . Otherwise the 2-fold cover branched over these points would give a hyperbolic metric on the torus. However for each $\theta \in [0, \pi)$ there is a hyperbolic cone-manifold $M(\theta)$ with underlying space the sphere and four cone angles θ obtained by doubling a hyperbolic quadrilateral with all corner angles $\theta/2$.



This quadrilateral may be chosen to have diameter 1. Then the area of $M(\theta)$ approaches 0 as $\theta \rightarrow \pi$. The limit of $M(\theta)$ is an interval of length 1. For θ close to π there is an orbifold foliation of $M(\theta)$ by short circles plus two intervals joining the cone points. There is a map of $M(\theta)$ to the interval which collapses each circle to a point.

The orbifold fundamental group of a pillowcase has an infinite cyclic normal subgroup. A hyperbolic orbifold cannot have such a subgroup. As the cone angles increase to the orbifold angle of π , loops representing this subgroup shrink to points. In some sense this is forced by the holonomy representation which more and more nearly is a representation of the orbifold fundamental group. This produces an orbifold fibration. It is a two-dimensional version of the kind of collapsing that can happen with a three-dimensional orbifold.

5.3 General deformation theory

Of basic importance in the deformation theory of geometric structures on manifolds is the following observation. Given the holonomy representation $\rho : \pi_1(M) \rightarrow G$ for a (G, X) -structure on M , all nearby representations also correspond to geometric structures on M . (Compare [89], [80, chap. 5], [57].)

Let $\mathcal{R} = \text{Hom}(\pi_1(M), G)$ denote the space of representations $\pi_1(M) \rightarrow G$; the group G acts on \mathcal{R} by conjugation. If G is an algebraic (or analytic) group then \mathcal{R} has a natural structure as an algebraic (or analytic) variety; this gives a natural topology on \mathcal{R} . (This will be described in more detail in section 5.5 below).

Theorem 5.3 (Deformations exist).

Suppose that M is the interior of a compact manifold with boundary. Let $\rho : \pi_1(M) \rightarrow G$ be the holonomy representation of a (possibly incomplete) (G, X) -structure on M . Then there is a neighbourhood, U , of ρ in the representation space $\mathcal{R} = \text{Hom}(\pi_1(M), G)$ such that for each $\rho' \in U$ there is a (nearby) (G, X) -structure on M with holonomy ρ' .

Sketch of Proof. (See [31], [43] and [14] for more details.) Suppose that M is a connected (G, X) -manifold without boundary. Let \tilde{M} be the universal cover of M . Then there is a developing map $\text{dev} : \tilde{M} \rightarrow X$ which is a local diffeomorphism. This map is equivariant with respect to the holonomy representation $h : \pi_1(M) \rightarrow G$ in the sense that for all covering transformations g of \tilde{M}

$$\text{dev} \circ g = h(g) \circ \text{dev}.$$

Conversely if M is a smooth manifold (without a geometric structure) then given a homomorphism $h : \pi_1 M \rightarrow G$ and a map $\text{dev} : \tilde{M} \rightarrow X$ which is equivariant in the above sense and which is a local diffeomorphism we may use dev to pull-back the (G, X) -structure to \tilde{M} . The equivariance condition ensures this structure is preserved by covering transformations and therefore covers a (G, X) -structure on M . Hence a (G, X) -structure on M is determined by the pair (dev, h) satisfying the equivariance condition. There is an equivalence relation generated by *isotopy* and *thickening*. Two (G, X) -structures on M are *isotopic* if there is a (G, X) -diffeomorphism between them which is isotopic to the identity. Suppose that N is M minus a collar. We call M a *thickening* of N .

We now outline another approach to describing a geometric structure on a manifold (see Thurston [82, chapter 3]). Given only a holonomy representation $\rho : \pi_1(M) \rightarrow G$ we can construct a bundle $E = E_\rho \rightarrow M$ with fibre X . The developing map gives a section $s : M \rightarrow E$. This will now be described. There is a diagonal action of $\pi_1 M$ on $\tilde{M} \times X$ where the action on the first factor is by covering transformations and the action on the second factor is by the holonomy. The quotient $E = (\tilde{M} \times X)/\pi_1 M$ is a bundle as stated. The graph of the developing map:

$$\text{Graph}(\text{dev}) = \{ (\tilde{x}, \text{dev}(\tilde{x})) : \tilde{x} \in \tilde{M} \} \subset \tilde{M} \times X$$

is preserved by the action of $\pi_1 M$ and therefore defines a section $s : M \rightarrow E$. This section is defined by $s(x) = [(\tilde{x}, \text{dev}(\tilde{x}))]$. The notation $[(\tilde{x}, y)]$ denotes the projection of a point in $\tilde{M} \times X$ to E .

The product structure on $\tilde{M} \times X$ is preserved by the action of $\pi_1 M$ and so there is a *horizontal foliation* of E by leaves which are the projections of $\tilde{M} \times y$ for $y \in X$. This foliation determines a *flat connection* on E . A bundle over M with fibre X has a universal cover which is a bundle over \tilde{M} and the covering transformations act diagonally if and only if there is a foliation of E transverse to the fibres. Furthermore, the action on X is by elements of G if and only if the holonomy of the bundle is in G .

The statement that the developing map is a local diffeomorphism is equivalent to the statement that the graph of dev is transverse to the horizontal foliation. This is equivalent to the statement that the section s is transverse to the horizontal foliation of E .

It follows that a (G, X) -structure on M is determined by a flat X -bundle over M with holonomy in G together with a section s of this bundle which is transverse to the horizontal foliation.

Given a hyperbolic metric on M with holonomy ρ and a representation ρ' close to ρ then there is a bundle $E_{\rho'}$. The section $s : M \rightarrow E_\rho$ given by the (G, X) -structure on M is transverse to the horizontal foliation of E_ρ . The bundle $E_{\rho'}$ is close to E_ρ . One may construct a section $s' : M \rightarrow E_{\rho'}$ close to s . Since transversality is an open condition, if one restricts to a compact subset N of M which contains M minus a collar, then $s'|_N$ will still be transverse to the horizontal foliation, provided ρ' is close enough to ρ . Then the pair (ρ, s') determines a (G, X) -structure on $N \cong M$ with holonomy ρ' . \square


The previous result can be refined as follows. Let $\text{Def}(M)$ denote the *deformation space* consisting of (G, X) -structures on M modulo the equivalence relations of isotopy and thickening described above. This has a natural topology, where two points are close if they correspond to structures with developing maps $\tilde{M} \rightarrow X$ which are close in the C^∞ topology on compact subsets of \tilde{M} . The deformation space $\text{Def}(M)$ is locally parametrized by nearby holonomy representations in \mathcal{R} modulo the action of G by conjugation. (See Goldman [31] for a precise statement.)

5.4 Deforming hyperbolic cone-manifolds

Suppose we have a hyperbolic cone-manifold structure on Q with a link as the singular locus Σ . Theorem 5.3 applies to a (possibly incomplete) hyperbolic structure with no singularities on a manifold without boundary. To apply it to a cone-manifold structure, cut out a singular solid torus neighbourhood of each component of the cone locus. We then deform what remains, and glue in suitably deformed singular solid tori to get a new cone-manifold structure. This will be described in more detail below.

We will first investigate (incomplete) hyperbolic structures on $M = Q - \Sigma$ near the complete structure.


The *ends* of a complete, orientable hyperbolic 3-manifold M with finite volume are *cusps* which are topologically $T^2 \times [0, \infty)$ and are foliated by horospherical tori. They are obtained by dividing out the foliation of \mathbb{H}^3 by horospheres by a $\mathbb{Z} \oplus \mathbb{Z}$ lattice. Each torus has an induced flat metric which decreases exponentially as one moves out in the cusp; all geodesics in this flat metric have geodesic curvature $+1$ in M .




Near a parabolic isometry there are both elliptic and hyperbolic isome-

tries. In dimension 2, the unique fixed point at infinity becomes an interior fixed point or an invariant axis. Annular regions between horocycles develop into regions between equidistant curves.

Notice also that the geometry of equidistant curves varies continuously: the geodesic curvature is $\coth r$ at distance r from a point, $\tanh r$ at distance r from a geodesic axis, and $+1$ on any horocycle.




In dimension 3 both elliptic and hyperbolic isometries have an invariant axis. Commuting elements share the same axis. The *complex length*, \mathcal{L} , of such an element of $PSL(2, \mathbb{C})$ is defined to $\ell + i\theta$, where ℓ is the translation distance along the axis and θ is the angle of rotation around the axis. It satisfies $\text{tr} = 2 \cosh(\mathcal{L}/2)$. Thus, a group element is elliptic if and only if $|\text{tr}| < 2$ and tr is real. Further, the angle of rotation θ is related to the trace by $\text{tr} = 2 \cos(\theta/2)$.




If one deforms the complete structure on the 3-manifold $M = Q - \Sigma$, the region between two parallel horocyclic tori will develop into a region between two equidistant surfaces. The quotient under the holonomy group of the torus will contain a foliation by equidistant tori. If the holonomies of the meridians are elliptic the tori can be filled in to a cone-manifold structure on X_Q .

Again the geometry of the equidistant surfaces varies continuously: the surface at distance r from an axis has an intrinsic flat metric with principal curvatures $\tanh r$ and $\coth r$, while all principal curvatures on a horosphere are $+1$.



Combining this discussion with the general deformation theory of section 5.3 shows that to find a nearby cone-manifold structure, it suffices to find a nearby holonomy representation for which the holonomy of the meridian is elliptic.

To see this, we remove a neighbourhood of the singular locus. The developing image of a neighbourhood of the boundary will lie in a neighbourhood of a nearby axis. We then fill the neighbourhood in to obtain a cone-manifold structure, with the new cone angle.



Exercise 5.4. Extend this discussion to the case where the singular locus contains trivalent vertices.

5.5 Representation spaces

We now want to show that one can always deform the holonomy representation, keeping the meridians elliptic and varying the cone angles independently. First we need to find non-trivial deformations of a hyperbolic structure on M . To do this we estimate the dimension of the representation variety, $\mathcal{R} = \text{Hom}(\pi_1(M), G)$, where $G = \text{PSL}(2, \mathbb{C})$.

Let $\mathcal{R} = \text{Hom}(\Gamma, G)$ denote the set of all representations (homomorphisms) from Γ into G . If Γ is a finitely generated group, and G is a Lie group this has the structure of a real analytic variety (an algebraic variety if G is an algebraic group). If $\gamma_1, \dots, \gamma_g$ is a set of generators for Γ then \mathcal{R} embeds in G^g , via the evaluation map

$$\mathcal{R} \rightarrow G^g, \quad \rho \mapsto (\rho(\gamma_1), \dots, \rho(\gamma_g)).$$

The image is the analytic subset of G^g satisfying the relations of Γ .

We'll be interested in representations $\rho : \Gamma \rightarrow \text{PSL}(2, \mathbb{C})$; in this case we get a *complex* analytic variety. (It becomes a complex algebraic variety, if we lift the representations into $SL(2, \mathbb{C})$ — see [24].)

Proposition 5.5. *If a group π is described in terms of g generators and r relations, the dimension of the variety of representations of π into a complex analytic Lie group G is at least*

$$(g - r) \dim G.$$

Sketch of proof. If π has a presentation

$$\pi = \langle a_1, \dots, a_g \mid w_1 = \dots = w_r = e \rangle,$$

define $\tau : G^g \rightarrow G^r$ by $\{\rho(a_i)\} \mapsto \{\rho(w_j)\}$. Then $\mathcal{R} = \tau^{-1}(e, e, \dots, e)$. Each equation $w_i = e$ is given locally by $\dim G$ complex analytic equations so reduces the complex dimension by at most $\dim G$. \square

For any 3-manifold M with $\partial M \neq \emptyset$, there is a deformation retraction of M to a 2-complex K (obtained by collapsing 3-cells away from free boundary faces, starting at the boundary). This gives a presentation of $\pi_1(M) \cong \pi_1(K)$ with g generators and r relations, where $g - r = 1 - \chi(K) = 1 - \chi(M)$.

For 3-dimensional hyperbolic structures we have $G = PSL(2, \mathbb{C})$ and $\dim_{\mathbb{C}} G = 3$, so we get the estimate

$$\dim_{\mathbb{C}} \mathcal{R} \geq 3(1 - \chi(M)).$$

If the holonomy representation ρ has trivial centralizer

$$Z(\rho) = \{g \in G : g\rho(\gamma)g^{-1} = \rho(\gamma) \text{ for all } \gamma \in \pi_1(M)\}$$

then conjugation determines a 3-complex dimensional subvariety of equivalent structures so we obtain:

$$\dim_{\mathbb{C}} \text{Def}(M) \geq -3(\chi(M)).$$

Exercise 5.6. Show that the centralizer is trivial for any holonomy representation of a finite volume hyperbolic structure.

Exercise 5.7. If M is a compact n -manifold with n odd, then $\chi(M) = \frac{1}{2}\chi(\partial M)$. [Hint: consider the double of M .]

Using this we obtain

$$\dim_{\mathbb{C}} \text{Def}(M) \geq -\frac{3}{2}\chi(\partial M).$$


So if ∂M is a union of tori, this only gives 0 as a lower bound. A subtler argument of Thurston (see [80], [22]) gives:

Theorem 5.8. *Let $\rho : \pi_1(M) \rightarrow PSL(2, \mathbb{C})$ be an irreducible representation (i.e. $\rho(\pi_1(M))$ has no fixed point on S_{∞}^2). Then each irreducible component of $\mathcal{R} = \text{Hom}(\pi_1(M), PSL(2, \mathbb{C}))$ containing ρ has complex dimension $\geq 3 - \frac{3}{2}\chi(\partial M) + t$, where t is the number of torus boundary components.*

Hence,

$$\dim_{\mathbb{C}} \text{Def}(M) \geq -\frac{3}{2}\chi(\partial M) + t.$$

Idea of proof. If ∂M consists of a single torus, drill out a (suitable) properly embedded arc from M giving a new manifold M' with $\partial M'$ of genus 2. Then $\dim_{\mathbb{C}} \text{Def}(M') \geq 3$. Thurston shows that we can kill off the fundamental group of a 2-handle to obtain a representation of $\pi_1 M$ by adding just two complex relations: that *two* carefully chosen elements have trace equal to 2. \square



We also need to know the behaviour of holonomies for meridians as representations are deformed. If ∂M consists of t tori, T_i , and γ_i are meridian curves on T_i , define:

$$Tr : \text{Def}(M) \rightarrow \mathbb{C}^t$$

by

$$Tr(\rho) = (\text{tr}\rho(\gamma_1), \dots, \text{tr}\rho(\gamma_t))$$

By *Mostow-Prasad rigidity* there is a unique complete hyperbolic structure on M . It follows that the holonomy of the complete structure ρ_0 gives an isolated point in $Tr^{-1}(\pm 2, \dots, \pm 2)$. Hence the polynomial functions $\text{tr}\rho(\gamma_i)$ are non-constant near ρ_0 , and it can be shown that Tr gives an open mapping whose image contains a neighbourhood of $(\pm 2, \dots, \pm 2)$. (See [22], [23].) In fact, with some additional work it can be shown that $\dim_{\mathbb{C}} \text{Def}(M) = t$ near ρ_0 .

From the previous section, we conclude that all representations near ρ_0 whose meridians have traces in the open interval $(-2, 2)$ correspond to cone-manifold structures with cone angles α_i given by $\text{tr}\rho(\gamma_i) = 2 \cos(\alpha_i/2)$.

Corollary 5.9. *Suppose that M is a hyperbolic cone-manifold with singular locus a 1-manifold and whose holonomy is on the component of the representation variety containing ρ_0 . If the cone angles of M are α_i , there is $\epsilon > 0$ such that if $|\alpha'_i - \alpha_i| < \epsilon$ for all i then there is a hyperbolic cone-manifold structure on M close to the original structure and with these cone angles. Furthermore the holonomy of this structure is on the same component of the representation variety as ρ_0 .*

The proof of the Orbifold Theorem consists of a study of what can happen at the boundary of the realizable angle set.

Theorem 5.10. (Hodgson-Kerckhoff [45]) *Suppose that N is a closed 3-manifold and L is a closed 1-manifold in N . Finite volume hyperbolic cone-manifold structures on (N, L) (i.e. structures on N with singularities along L) are locally parametrized by the cone angles on the components of L when all cone angles are $\leq 2\pi$.*

The proof of this result uses quite different, analytic techniques: infinitesimal deformations give cohomology classes which can be represented by harmonic forms. These are studied by the use of a Bochner formula and a Fourier series type analysis of asymptotic behaviour of harmonic forms near the singular locus. A survey of this approach is given in [49].

Theorem 5.11. (Kojima [53]) *Suppose that N is a closed 3-manifold and L is a closed 1-manifold in N . Two hyperbolic cone-manifold structures on (N, L) with corresponding cone angles equal are isometric, provided all angles are $\leq \pi$,*

Sketch of proof. The proof uses part of the proof of the Orbifold Theorem, Mostow rigidity for the complete structure, and the local parametrization by cone angle of theorem 5.10. Given two structures, the idea is to decrease the angles to zero, giving complete hyperbolic structures on $N - L$. By Mostow-Prasad rigidity, these structures are equal. By the local parametrization theorem, they were equal throughout the deformation. \square

5.6 Hyperbolic Dehn filling

Given a 3-manifold M with boundary a union of tori T_1, \dots, T_k , and a choice $\gamma = (\gamma_1, \dots, \gamma_k)$ of a non-trivial simple closed curve γ_i on each T_i , one can do γ -Dehn filling on M by attaching a solid torus to each T_i so that γ_i bounds a disk. The result is denoted by $M(\gamma)$ as in section 1.11.

With a choice of generators for each $\pi_1(T_i)$, the γ_i correspond to pairs of relatively prime integers (p_i, q_i) .

In the following discussion, we *assume* there is a single boundary torus $T = T_1$ for simplicity. It is shown in [80], [68] that the complex translation length for elements in the boundary torus can be lifted to \mathbb{C} so that rotation angle is lifted from S^1 to \mathbb{R} , and each parabolic (at the complete structure) has complex length 0. If μ, λ are the complex lengths for a chosen set of generators for $\pi_1(T)$, then the complex length of the (p, q) -curve is $p\mu + q\lambda$. A solution to $p\mu + q\lambda = \alpha i$ near the complete structure gives a cone-manifold structure on $M(p, q)$ with cone angle α ; it is a smooth structure if $\alpha = 2\pi$.


Define

$$DS : \text{Def}(M) \rightarrow (\mathbb{R}^2 \cup \infty) / \pm 1$$

by

$$DS(\rho) = (x, y) \text{ if } x\mu + y\lambda = 2\pi i.$$

Then points along lines of rational slope in $\mathbb{R}^2 \cup \infty$ correspond to hyperbolic cone-manifolds.



Theorem 5.12. (Thurston [80]) *DS maps onto a neighbourhood of ∞ in $\mathbb{R}^2 \cup \infty$. Thus all but finitely many Dehn fillings on M are hyperbolic.*

The number of the *exceptional* (non-hyperbolic) surgeries is not effectively computable from this proof. However, the computer program Snappea developed by Jeff Weeks [88] provides a powerful tool for studying examples, and can estimate the number of exceptional surgeries quickly.


There are universal bounds on the number of Dehn surgeries without *negatively curved* metrics given by the “length 2π ” Theorem of Gromov-Thurston, (see [33], [6], [1]). The recent “length 6” theorem of Agol and Lackenby (see [1], [57]) gives new bounds on the number of surgeries giving manifolds whose fundamental group is not *word hyperbolic*.

Recently, Hodgson-Kerckhoff have obtained the first *universal* bounds on the number of *non-hyperbolic* surgeries (see [46], [49]).

Finally we mention some conjectures on the global structure of hyperbolic Dehn Surgery space.

Conjecture 1. The Dehn surgery coordinate map $DS : \text{Def}(M) \rightarrow (\mathbb{R}^2 \cup \infty) / \pm 1$ should be a (local) diffeomorphism onto its image.


Conjecture 2. Hyperbolic Dehn surgery space, \mathcal{H} , should be star-like with respect to rays from infinity towards the origin. In particular, it should be a connected set.




If both conjectures are true, then this implies global (Mostow-Kojima) rigidity for all hyperbolic cone-manifolds. (The proof of theorem 5.11 sketched above for cone angles $\leq \pi$, would again apply.)

5.7 Dehn surgery on the figure eight knot

Thurston's Princeton University notes [80] and Hodgson's thesis [43] include detailed studies of the hyperbolic Dehn surgery space for the complement $M = S^3 - K$ of the figure eight knot K in S^3 .



In the following discussion the Dehn surgery coordinate (p, q) refers to $p\mu + q\lambda$ where μ is a meridian and λ a standard longitude for the figure eight knot. Thurston shows that the lightly shaded region shown below consists of hyperbolic structures obtained by gluing together positively oriented ideal tetrahedra. Hodgson shows that the “algebraic volume” associated with representations into $PSL(2, \mathbb{C})$ goes to zero along the solid curve shown below. This curve consists of straight line segments corresponding to representations into $\text{Isom}(\mathbb{H}^2)$ and curves corresponding to representations into $SO(3)$. It is conjectured that this represents the true boundary of the hyperbolic Dehn surgery space, but currently hyperbolic structures with Dehn surgery type singularities are only known for some special points within the darkly shaded region.



On the boundary of the hyperbolic region degenerations of the following kinds occur.

- (1) Dehn surgery coordinates $(m, 1)$, $-4 < m < 4$.


Here there are limiting representations $\pi_1(M) \rightarrow PSL(2, \mathbb{R})$, corresponding to foliations with transverse hyperbolic structures. The foliations can be seen directly, since we have two positively oriented simplices flattening out simultaneously. It was shown explicitly in [80] that this is part of the exact boundary of hyperbolic Dehn surgery space. The manifold points in the boundary are as follows. (Note that $M(p, q)$ and $M(-p, q)$ are oppositely oriented copies of the same manifold, since the figure eight knot has an orientation reversing symmetry.)

- (a) The manifold $M(0, 1)$ is a torus bundle over S^1 with Anosov gluing map with matrix $\Phi = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$; this gives rise to a Solv geometry structure.

Here there are hyperbolic cone-manifold structures on $M(0, 1)$ for cone angles $\theta < 2\pi$ which collapse as $\theta \rightarrow 2\pi$ to a circle whose length is $\log \lambda$, where $\lambda > 1$ is the larger eigenvalue of the matrix Φ . (See [43], [41], [78].)

- (b) Each manifold $M(n, 1)$ for $n = \pm 1, \pm 2, \pm 3$ is a Seifert fibre space over a hyperbolic 2-orbifold which is sphere with 3 cone points. There are hyperbolic cone-manifold structures on $M(n, 1)$ for cone angles $\theta < 2\pi$ which collapse as $\theta \rightarrow 2\pi$ to the 2-dimensional hyperbolic structure on the base orbifold. The singular locus is transverse to the fibres of the Seifert fibration and projects to a geodesic in the base orbifold.

For example, $M(\pm 1, 1)$ is the unit tangent bundle of the $(2, 3, 7)$ spherical orbifold (see example 2.39) and the singular locus Σ is the horizontal lift of a geodesic through the order 2 cone point (double covering a geodesic in base):



Similarly, $M(\pm 2, 1)$ is the unit tangent bundle to $S^2(2, 4, 5)$, and $M(\pm 3, 1)$ is the unit tangent bundle to $S^2(3, 3, 4)$.

Each manifold $M(\pm 4, 1)$ is a graph manifold containing an incompressible torus which splits the manifold into the union of a trefoil knot complement and the non-trivial I -bundle over the Klein bottle. In this case, the hyperbolic cone-manifold structures for cone angles $\theta < 2\pi$ split along an essential torus and collapse to give a limiting cusped Seifert fibred structure on the complementary pieces. (Compare example 3.4.)

Remark: These non-hyperbolic manifolds resulting from Dehn surgery on the figure eight knot can be identified using the Kirby calculus (see [43]), or via the “Montesinos trick” which divides out by a 180 degree rotational symmetry and studies the quotient orbifold. (For examples of this technique see [6].)

(2) There are orthogonal representations corresponding to Dehn surgery coordinates on a curve from $(3.618\dots, 0.809\dots)$ passing through $(3, 0)$ to $(3.618\dots, -0.809\dots)$.

Special case: The orbifold $M(3, 0)$ has a Euclidean structure described earlier (see 2.33). In this case there are hyperbolic cone-manifold structures for cone angles $\theta < 2\pi/3$ which collapse to a point as $\theta \rightarrow 2\pi/3$. After rescaling the hyperbolic metrics, these converge to the Euclidean orbifold structure with cone angle $2\pi/3$. Further, there are spherical cone-manifold

structures for $2\pi/3 < \theta \leq \pi$. All of these cone-manifold structures can be constructed directly from suitable polyhedra by identifying pairs of faces by isometries — see [41], [78].

Hodgson's work in [43] implies that near this point, the curve where $volume = 0$ corresponds to Euclidean structures with Dehn surgery singularities, and is *locally* the boundary of hyperbolic Dehn surgery space. (See also [67].) Furthermore, each Euclidean structure can also be approximated by spherical structures with Dehn surgery type singularities. If we consider the larger space $GS(M)$ of all constant curvature geometric structures on M with Dehn surgery type singularities; then we locally obtain a manifold. The Dehn surgery coordinates give a local diffeomorphism from geometric structures in $GS(M)$, up to rescaling of metrics, to a neighbourhood of $(3, 0)$ in \mathbb{R}^2 . The Euclidean structures correspond to the codimension one subspace, where $volume = 0$.

(3) Dehn surgery coordinates on the straight line from $(3.618\dots, .809\dots)$ to $(4, 1)$ and on the straight line from $(-3.618\dots, 0.809\dots)$ to $(-4, 1)$ also correspond to representations into $\text{Isom}(\mathbb{H}^2)$ since simplices are flat. But little is currently known about the geometric meaning of these representations.

Remark: In the proof of the orbifold theorem, we will encounter analogues of all the kinds of degeneration of hyperbolic structures described for the figure eight knot complement. However, the Seifert fibre spaces arising will be orbifolds which have both intervals and circles as fibres.

Chapter 6

Limits of Metrics Spaces

A key ingredient in the proof of the Orbifold Theorem is the analysis of limits of metric spaces. In this chapter we give a short account of Gromov's theory of limits of metric spaces as re-interpreted using ϵ -approximations by Thurston. See Gromov's book [32] for a detailed treatment with many interesting applications.

Roughly, there is an " ϵ -approximation" between two metric spaces if the spaces look the same if we ignore details of size ϵ or smaller. From this we define the *Gromov-Hausdorff distance* between two compact metric spaces, and convergence of sequences of metric spaces.

This generalizes the classical notion of *Hausdorff distance* between two subsets A, B of a metric space X :

$$d_H(A, B) = \inf\{\epsilon > 0 : A \subset N(B, \epsilon; X) \text{ and } B \subset N(A, \epsilon; X)\},$$

where

$$N(A, r; X) = \{ x \in X : \exists a \in A \ d(x, a) < r \}$$

denotes the (open) neighbourhood of radius r around A in X . (See [5] for a detailed discussion of the Hausdorff distance and many geometric applications).


The *Gromov-Hausdorff distance* also generalizes the notion of *Lipschitz distance* between homeomorphic metric spaces. A bijection $f : X \rightarrow Y$ is *K -bilipschitz* if

$$\forall x \neq x' \in X, \quad 1/K \leq d_Y(fx, fx')/d_X(x, x') \leq K.$$

Two metric spaces are close in the Lipschitz sense if there is a $(1 + \epsilon)$ -bilipschitz map between them with ϵ small.

Let X be a metric space. Then $A \subset X$ is an ϵ -net or ϵ -dense if for all $x \in X$ there exists $a \in A$ such that $d(x, a) < \epsilon$.

The basic idea is to approximate a *compact* metric space X by a *finite* ϵ -net A ; we want information about a metric space accurate to within ϵ . If X, Y are compact metric spaces, we regard them as “close” if there are finite ϵ -nets $A \subset X$ and $B \subset Y$ and a $(1 + \epsilon)$ -bilipschitz map $f : A \rightarrow B$.



6.1 ϵ -approximations

Definition: An ϵ -approximation between metric spaces X and Y is a relation $R \subset X \times Y$ such that

- (1) the projections $p_X : R \rightarrow X$ and $p_Y : R \rightarrow Y$ are both onto,
- (2) if xRy and $x'Ry'$ then $|d_X(x, x') - d_Y(y, y')| \leq \epsilon$.

This defines a relation on metric spaces which is symmetric and almost transitive: if X is an ϵ -approximation to Y and Y is an ϵ' -approximation to Z , then X is an $(\epsilon + \epsilon')$ -approximation to Z .

We begin with some examples.


Example 6.1. (a) Let A be an ϵ -net for X . Then we can define a 2ϵ -approximation $R \subset A \times X$ by: $aRx \Leftrightarrow d(a, x) < \epsilon$.
 (b) A 0-approximation is an isometry.

Example 6.2. Suppose that X, Y are subsets of a metric space Z . Define a relation $R \subset X \times Y$ by xRy if $x \in X$, $y \in Y$ and $d_Z(x, y) \leq \epsilon$. This is a 2ϵ -approximation if $Y \subset N(X, \epsilon; Z)$ and $X \subset N(Y, \epsilon; Z)$. This leads to the Hausdorff metric on closed subsets of Z .


We also say that a sequence $\{X_n\}$ of metric spaces *converges* to Y and write $X_n \rightarrow Y$ if there are ϵ_n -approximations between X_n and Y with $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$.

Example 6.3.

(a) $T_n = S^1 \times (1/n)S^1$ converges to the circle S^1 .



(b) $X_n =$ Euclidean solid torus obtained from a cylinder of height $1/n$ by gluing ends with 180° twist converges to a 2-dimensional Euclidean cone with angle π .



Exercise 6.4. What happens if we modify example (b) by varying the twist angles θ ?

Proposition 6.5. Let X, Y be compact metric spaces. Assume for all $\epsilon > 0$ there is an ϵ -approximation between X and Y . Then X is isometric to Y .

Proof. Let $R_n \subset X \times Y$ be an $(1/n)$ -approximation. There exists a countable dense set $A = \{x^k\} \subset X$. Choose $y_n^k \in Y$ with $x^k R_n y_n^k$. Choose a subsequence of y_n^1 converging to y^1 , then a sub-subsequence of y_n^2 converging to y^2 , etc. After repeating this process, we define $f : A \rightarrow Y$ by $f(x^k) = y^k$.

We claim that f is an *isometry* onto $f(A)$ and therefore 1-1. This is

because $x^i R_n y_n^i$ & $x^j R_n y_n^j$ implies

$$|d_X(x^i, x^j) - d_Y(y_n^i, y_n^j)| < 1/n.$$

Taking limits gives $d_X(x^i, x^j) = d_Y(y^i, y^j)$ which proves the claim. Since f is uniformly continuous and A is dense in X and Y is complete, it follows that there exists a unique continuous extension $f : X \rightarrow Y$. Further, $f(A)$ is dense in Y because A is dense in X thus the image of A under R_n is $1/n$ -dense in Y . Since X is compact it follows that $f(X)$ is closed. Hence f is onto. \square

Corollary 6.6. *One can define a metric on the family \mathcal{F} of isometry classes of compact metric spaces, by putting*

$$d(X, Y) = \inf\{\epsilon \mid \text{there is an } \epsilon\text{-approximation between } X \text{ and } Y\}.$$

Further, metric spaces with finitely many points are dense.

(Note that $\text{diam} < \infty$ for compact spaces implies d is always finite.)

Remarks on related definitions:

For subsets of a fixed metric space, our definition of distance is closely related to the Hausdorff distance on closed subsets. Gromov uses this in [32] to define ‘‘Hausdorff convergence’’ of metric spaces. Gromov’s definition of convergence of metric spaces is somewhat stronger than ours. Gromov’s distance $d_G(X, Y)$ between metric spaces X and Y is the infimum of Hausdorff distance between $f(X)$ and $g(Y)$ over all isometric embeddings $f : X \rightarrow Z$, $g : Y \rightarrow Z$ in metric spaces Z . It’s clear that $2d_G(X, Y) \geq d(X, Y)$ by example 6.2 above.

Exercise 6.7. Show that $2d_G(X, Y) = d(X, Y)$, as remarked by Bridson and Swarup in [13]. [Hint: Given an ϵ -approximation $R \subset X \times Y$, construct a suitable metric \tilde{d} on the disjoint union $Z = X \cup Y$ which agrees with the given metrics d_X on X , d_Y on Y .]


6.2 Limits with basepoints

For non-compact space we introduce *basepoints*.

Definition: $(X_n, x_n) \rightarrow (Y, y)$ converges in the *Gromov-Hausdorff* topology if for all $r, \epsilon > 0$ and for all n sufficiently large, there is an ϵ -approximation R_n between $N(x_n, r; X_n)$ and $N(y, r; Y)$ such that $x_n R_n y$.


The idea is that ‘‘big neighbourhoods of the basepoint are almost isometric’’.

Example 6.8. $C_n =$ hyperbolic torus with cone point π/n converges to a complete hyperbolic punctured torus.



Limits will generally depend on choice of basepoint.

Example 6.9. A sequence of hyperbolic genus 3 surfaces with a long thin neck developing can converge to the three limits shown below for different choices of basepoints.



Given an ϵ -approximation $R \subset X \times Y$ if $x_0 R y_0$ we often wish to *restrict* R to an approximation between the r -neighbourhoods $N(x_0, r; X)$

and $N(y_0, r; Y)$. But there is a potential problem here: the projection onto the factors may not be *onto*.

The *smear* of R is the 3ϵ -approximation $R' \subset X \times Y$ given by:
 $xR'y \Leftrightarrow \exists x' \in X \quad y' \in Y$ such that

$$d(x, x') < \epsilon \quad \text{and} \quad d(y, y') < \epsilon \quad \text{and} \quad x'Ry'.$$

If x_0Ry_0 and $r > 0$ then R' restricts to a 3ϵ -approximation between $N(x_0, r; X)$ and $N(y_0, r; Y)$.

With this definition of Gromov-Hausdorff convergence we have: $n^{-1}\mathbb{Z} \rightarrow \mathbb{Q}$ and $n^{-1}\mathbb{Z} \rightarrow \mathbb{R}$. To get *unique* limits, we need to restrict the metric spaces involved. A metric space is called *proper* if *every closed ball is compact*. Note that every proper metric space is complete.

Corollary 6.10. *Let X, Y be proper metric spaces. If for all $r, \epsilon > 0$ there is an ϵ -approximation between $N(x, r; X)$ and $N(y, r; Y)$ then (X, x) is isometric to (Y, y) .*

Corollary 6.11. *Let X_i, Y, Y' be proper metric spaces. If $(X_i, x_i) \rightarrow (Y, y)$ and $(X_i, x_i) \rightarrow (Y', y')$ then (Y, y) is isometric to (Y', y') .*


Remark 6.12.

- (a) One can always assume that limits are complete, since the distance between a metric space and its completion is zero.
- (b) If closed balls in X_i are compact for all i , and $X_i \rightarrow Y$, with Y complete, then all closed balls in Y are also compact. This follows from the fact that a metric space is compact if and only if it is complete and totally bounded (i.e. for all $\epsilon > 0$, there is a finite covering by ϵ -balls).
- (c) The condition that closed balls are compact fails for infinite dimensional spaces. For example, consider \mathbb{R}^∞ with the supremum metric. Then the ball of radius 1 is non-compact; e.g. the sequence $(1, 0, 0, \dots), (0, 1, 0, \dots), \dots$ doesn't converge.

Example 6.13.

- (a) $n^{-1}\mathbb{Z} \rightarrow \mathbb{E}^1$
- (b) $nS^1 \rightarrow \mathbb{E}^1$
- (c) $S^1 \times (n^{-1}S^1) \rightarrow S^1$.
- (d) Let $\mathbb{H}^n(K)$ be “hyperbolic space” of constant curvature $K < 0$, obtained by *rescaling* the metric on \mathbb{H}^n . Then $\mathbb{H}^n(K) \rightarrow \mathbb{E}^n$ as $K \rightarrow 0$.
- (e) Let C_n be the cone on n points with a path metric such that each edge

has length 1. Then the sequence $\{C_n\}$ does not converge. (There is no finite approximation to the ball of radius 1 in the “limit”.)



Exercise 6.14. Describe the possible geometric limits of sequences of 2-dimensional Euclidean tori. (Recall: these correspond to isometry classes of lattices in \mathbb{R}^2 .)

6.3 Gromov's compactness theorem

We now examine the question: When does a sequence of metric spaces have a convergent subsequence?

Definitions: The ϵ -count, $\#(\epsilon, X)$, of a metric space X is the minimum number of balls of radius ϵ needed to cover X .

A collection of based metric spaces (X_n, x_n) is *uniformly totally bounded* if for all $\epsilon, r > 0$ there exists $K > 0$ such that $\#(\epsilon, N(x_n, r; X_n)) < K$.

Theorem 6.15 (Gromov's Compactness Theorem). ([32]) *If (X_i, x_i) is a sequence of proper metric spaces then the following are equivalent:*

- (1) *there is a subsequence $(X_{n_i}, x_{n_i}) \rightarrow (Y, y)$ with Y complete.*
- (2) *there is a subsequence (X_{n_j}, x_{n_j}) which is uniformly totally bounded.*

Proof. We show that (2) implies (1): Assume the sequence satisfies (2). For each ϵ and each i , we can choose K_ϵ ϵ -balls covering $B_{1/\epsilon}(x_i) \subset X_i$. Let $P_{\epsilon,i}$ be the set of centres of these balls together with the base point x_i . Then each $P_{\epsilon,i}$ is a finite set (containing $\leq K_\epsilon + 1$ points), and is an 3ϵ -approximation to $B_{1/\epsilon}(x_i)$. The metric on $P_{\epsilon,i}$ is described by a distance function

$$d_i : \{1, 2, \dots, K_\epsilon + 1\}^2 \rightarrow [0, 2/\epsilon].$$

By compactness there is a convergent subsequence of d_i ; hence there is a subsequence P_{ϵ,i_j} converging to a limiting (finite) metric space L_ϵ (containing $\leq K_\epsilon + 1$ points). (Possibly some points coalesce in the limit, but this won't matter.) Let l_ϵ be the limit of the base points x_{i_j} .

Now choose a collection of $K_{\epsilon/2}$ $\epsilon/2$ -balls covering $B_{2/\epsilon}(x_i)$, and let

$$P_{\epsilon/2,i} = \{\text{centres of these } \epsilon/2 \text{ balls}\} \cup P_{\epsilon,i}.$$

We can choose a further subsequence such that $P_{\epsilon/2,i} \rightarrow L_{\epsilon/2}$; then L_ϵ embeds isometrically in $L_{\epsilon/2}$.

Continuing in this way, we obtain

$$L_\epsilon \subset L_{\epsilon/2} \subset L_{\epsilon/4} \subset \dots$$

Let L be the metric completion of $\bigcup_n L_{\epsilon/2^n}$ and $l \in L$ the limit of the base points. We claim that this is the limit of a (diagonal) subsequence of the X_i .

By our construction, we have a subsequence X_j such that for all $\epsilon > 0$, there is an ϵ -approximation between $B_{1/\epsilon}(x_j)$ and a subset $L_{\epsilon/3}$ of L . Further, each $L_{\epsilon/2^n}$ is an $\epsilon/2^n$ -approximation to $L_{\epsilon/2^{n+1}}$, so is an $\epsilon/2^{n-1}$ -approximation to L . Hence, $X_j \rightarrow L$. \square

Exercise 6.16. Complete the above proof by showing that (1) implies (2).

6.4 Limits of hyperbolic cone-manifolds

We want to know :

- (1) when a sequence of 3-dimensional hyperbolic cone-manifolds has a subsequence which converges to a complete metric space Y .
- (2) when Y is a 3-dimensional hyperbolic cone-manifold.

Let M be a 3-dimensional hyperbolic cone-manifold; then in particular M is complete. Let D be a Dirichlet domain for M , and $x_0 \in D \subset \mathbb{H}^3$. If M has cone angles $< 2\pi$, then the natural quotient map $q : D \rightarrow M$ is onto. But

$$\#(\epsilon, N(q(x_0), r; M)) \leq \#(\epsilon, N(x_0, r; D)) \leq \#(\epsilon/2, N(x_0, r; \mathbb{H}^3)).$$

Therefore, 3-dimensional hyperbolic cone-manifolds with cone angles in $(0, 2\pi]$ are uniformly totally bounded. Then Gromov compactness implies that (1) always holds.

Note: This is false if cone angles $> 2\pi$ are allowed. One may construct a hyperbolic surface with more and more cone points with cone angles larger than 2π in a small region. This can be done so that the area of this region increases without bound.

We have seen that a sequence of n -manifolds can converge to a space of dimension $< n$. We need a way to rule out this kind of behaviour.

If M is a Riemannian n -manifold and $x \in M$, the *injectivity radius at x* is the radius of the “largest embedded ball” in M with centre x .

Proposition 6.17. *Suppose (M_k, x_k) is a sequence of n -manifolds with constant curvature K . Suppose for all points $x \in M_k$ that $\text{inj}(x) > \text{inj}_0$. If $(M_n, x_n) \rightarrow (Y, y_0)$ then Y is an n -manifold of curvature K .*

Proof. Given $y \in Y$ choose $x_n \in M_n$ with $x_n R_n y$. After smearing we can assume that $N(x_n, \text{inj}_0; M_n) \rightarrow N(y, \text{inj}_0; Y)$. But $N(x_n, \text{inj}_0; M_n)$ is isometric to $N(p, r; \mathbb{H}^n(K))$, therefore $N(y, \text{inj}_0; Y)$ is isometric to $N(p, r; \mathbb{H}^n(K))$. \square


We want to adapt this proof to *cone-manifolds*. This raises the question of what is an appropriate notion of *injectivity radius* in a cone-manifold? (With the usual definition $\text{inj}(x) \rightarrow 0$ as $x \rightarrow \Sigma$.) The role of injectivity radius in the above proof is to provide a *standard neighbourhood* of a certain size. The proof uses that the limit of such neighbourhoods is again such a neighbourhood. In a cone-manifold the local geometry is that of a *cone*. We will see that there is a compact family of such neighbourhoods, and this is what is used to extend the proof above.

A *cone* in a n -dimensional cone-manifold M is a subset isometric to a cone on a spherical $(n-1)$ -dimensional cone-manifold, i.e. a “standard cone neighbourhood” as defined in section 3.2.

Definition: $\text{inj}(x)$ is the largest r for which $N(x, r; M)$ is contained in a cone:


$$\text{inj}(x) = \sup\{ r > 0 : \exists x' \in M \exists r' > 0 \text{ s.t. } N(x, r; M) \subset N(x', r'; M) \\ N(x', r'; M) = \text{a cone} \}.$$

Note that we *do not* assume that the standard neighbourhood is *centred* at the point x . This is to avoid difficulties near cone points: a point x near the cone locus has only a small standard ball centred at x ; however there may be much larger standard *cones* centred at cone points which contain x .



In general, the injectivity radius at a point x will be small if there is a short geodesic loop based at x , or if two *different* pieces of the singular locus are close to x .


Example 6.18. The injectivity radius at some points on an infinite Euclidean pillowcase are shown below. Maximal open balls about the points contained in standard cones are also indicated.



With this definition we obtain:

Proposition 6.19. *Let (M_n, x_n) be a sequence of cone-manifolds, with M_n of constant curvature $\kappa_n \in [-1, 0]$ and $\kappa_n \rightarrow \kappa_\infty$. Suppose that there are $\theta_0, \text{inj}_0 > 0$ such that all cone angles are in the range $[\theta_0, \pi]$ and $\text{inj}(x, M_n) > \text{inj}_0$ for all $x \in M_n$ and all n . Then there is a subsequence converging to a cone-manifold (M_∞, x_∞) of curvature κ_∞ .*

The essential reason is that the cones of fixed radius r form a *compact* set of metric spaces, namely the cones on: $S^2(\alpha, \beta, \gamma)$, $S^2(\alpha, \alpha)$ or S^2 with $\alpha, \beta, \gamma \in [\theta_0, \pi]$.



6.5 Bilipschitz convergence

Theorem 6.20. *Suppose compact hyperbolic 3-manifolds M_n converge in the Gromov-Hausdorff topology to a compact hyperbolic 3-manifold M_∞ . Then for all $\epsilon > 0$ and for all sufficiently large n there is a $(1 + \epsilon)$ -bilipschitz map $f : M_\infty \rightarrow M_n$.*

Proof. (Sketch) There exists $inj_0 > 0$ such that the injectivity radius of every point in every M_n and in M_∞ is bigger than inj_0 . (There is a lower bound on injectivity radius in M_∞ because it is compact. Since $M_n \rightarrow M_\infty$, for large n the injectivity radius in M_n is nearly equal to that in the limit.) Let K be a *geodesic triangulation* of M_∞ , such that every simplex $\sigma \in K$ is *small* compared to inj_0 . Let $V = \{v_1, \dots, v_k\}$ be the vertices of K . Let R_n be a $1/n$ -approximation between M_n and M_∞ . Choose $v_i^n \in M_n$ with $v_i^n R_n v_i$. If $v_a, v_b, v_c, v_d \in V$ span a 3-simplex $\sigma \in K$ then $v_a^n, v_b^n, v_c^n, v_d^n \in M_n$ span a *small geodesic simplex* in a standard metric ball. Thus we get a geodesic triangulation K_n of M_n combinatorially the same as K . Also, the corresponding edge lengths are nearly equal. Use ϵ_n -approximations with $\epsilon_n \ll \epsilon \cdot$ (edge lengths of K) to map the vertices into M_n . This extends to a simplicial map $f : K \rightarrow K_n \subset M_n$ which is $(1 + \epsilon)$ -bilipschitz where $\epsilon \rightarrow 0$ as $n \rightarrow \infty$. \square

Extensions:

- (1) The previous result extends easily to cone-manifolds: this just requires some extra care constructing a *thin* geodesic triangulation near Σ . Note that bilipschitz convergence also implies convergence of cone angles, volume, etc.
- (2) To extend the result to the case of a non-compact limit M_∞ requires estimates on the *decay of injectivity radius* described in theorem 7.8. In this case we get bilipschitz maps from any *compact subset* of M_∞ into the M_n .

(3) The result also extends easily to a sequence of cone manifolds M_n , where each has constant curvature K_n lying in the interval $[-1, 0]$.

Theorem 6.21. *Suppose that M_n is a sequence of complete hyperbolic cone 3-manifolds. Suppose that (M_n, x_n) converges in the Gromov-Hausdorff topology to a complete hyperbolic cone 3-manifold (M_∞, x_∞) . Then given $\epsilon > 0$ and $R > 0$, for all sufficiently large n there is a $(1 + \epsilon)$ -bilipschitz map $f : N(x_\infty, R, M_\infty) \rightarrow M_n$ with $d(f(x_\infty), x_n) < \epsilon$. Furthermore f maps singular set to singular set.*

6.6 Convergence of holonomy

If a sequence (M_n, x_n) of 3-dimensional hyperbolic cone-manifolds converges to a hyperbolic cone-manifold (M_∞, x_∞) , we want convergence of the *holonomy* representations

$$h_n : \pi_1(M_n - \Sigma(M_n)) \rightarrow PSL(2, \mathbb{C}).$$

Theorem 6.22. *Assume that for all large n there are $(1 + \epsilon_n)$ -bilipschitz homeomorphisms*

$$\phi_n : (M_\infty, \Sigma(M_\infty)) \rightarrow (M_n, \Sigma(M_n))$$

with $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Let

$$\phi_{n*} : \pi_1(M_\infty - \Sigma(M_\infty)) \rightarrow \pi_1(M_n - \Sigma(M_n))$$

be the induced homomorphism of fundamental groups, and assume that $\pi_1(M_\infty - \Sigma(M_\infty))$ is finitely generated. Then $h_n \circ \phi_{n*} \rightarrow h_\infty$ in the algebraic topology. This means there are $A_n \in PSL(2, \mathbb{C})$ such that $A_n h_n(\phi_{n*}\alpha) A_n^{-1} \rightarrow h_\infty(\alpha)$ for all $\alpha \in \pi_1(M_\infty - \Sigma(M_\infty))$.

Idea of Proof: Write $X_n = M_n - \Sigma_n$. Then the bilipschitz convergence implies that the developing maps $\text{dev}_n : \tilde{X}_n \rightarrow \mathbb{H}^3$ can be adjusted by isometries g_n so that

$$g_n \circ \text{dev}_n \circ \tilde{\phi}_n \rightarrow \text{dev}_\infty : \tilde{X}_\infty \rightarrow \mathbb{H}^3$$

uniformly on compact subsets. Applying this to a large compact subset of \tilde{X}_∞ containing lifts of loops generating $\pi_1(X_\infty)$ gives the result. \square

Chapter 7

Proof of the Orbifold Theorem

In this section we will give a sketch of the proof of the Orbifold Theorem. We have assumed that the reader is familiar with the definitions and general ideas presented so far in this memoir and have tried to highlight the main theorems needed in the proof of this result. For a complete proof the reader may consult [19]. For an alternative proof of a somewhat different version of the Theorem, see [8].

7.1 Topological preliminaries

The Orbifold Theorem states that if a compact, orientable, orbifold-irreducible orbifold has a 1-dimensional singular locus, then it can be cut along a (possibly empty) Euclidean 2-orbifold so that each of the resulting components has a geometric structure. As discussed in Chapter 2, this is the orbifold version of the Geometrization Conjecture 2.57 for orientable 3-orbifolds with the important assumption that the singular locus is non-empty.

Theorem 7.1 (The Orbifold Theorem).

Suppose that \mathcal{O} is a compact, orientable, orbifold-irreducible 3-orbifold with (possibly empty) orbifold-incompressible boundary consisting of Euclidean 2-orbifolds. Suppose that $\Sigma(\mathcal{O})$ is a non-empty graph. Then there is an incompressible Euclidean 2-suborbifold \mathcal{T} (possibly empty) such that each component of $\mathcal{O} - \mathcal{T}$ is a geometric orbifold.

To begin the proof of the Orbifold Theorem, we need an orbifold version of the torus decomposition of a 3-manifold. This will provide the incom-

pressible Euclidean 2-suborbifold \mathcal{T} along which the orbifold is decomposed. The statement of the decomposition theorem given in 2.55 has been specialized to the orbifolds that arise in the version of the Orbifold Theorem given above.

Theorem 7.2 (Euclidean Decomposition Theorem).

Suppose that \mathcal{O} is a compact, orientable, orbifold-irreducible 3-orbifold with (possibly empty) boundary consisting of orbifold-incompressible Euclidean 2-orbifolds. Then there is a (possibly empty) closed, orientable, incompressible Euclidean 2-suborbifold $\mathcal{T} \subset \mathcal{O}$ such that if P is the closure of a component of $\mathcal{O} - \mathcal{T}$, then P is either an orbifold Seifert fibre space or it is orbifold-atoroidal. If P has non-empty boundary, that boundary will be orbifold-incompressible.

For the remainder of this paper we will assume that we have decomposed the orbifold in this manner. Seifert fibred orbifolds that are orbifold-irreducible are easily seen to have geometric structures (see 2.50) so we may assume that the orbifold \mathcal{O} is orbifold-irreducible, orbifold-atoroidal, with (possibly empty) boundary consisting of orbifold-incompressible Euclidean 2-orbifolds. The Orbifold Theorem is equivalent to the statement that such an orbifold \mathcal{O} is geometric as long as the singular set $\Sigma(\mathcal{O})$ is non-empty and 1-dimensional.

Consider the complement of an open regular neighbourhood of the singular locus, $\Sigma(\mathcal{O})$, in \mathcal{O} . This is a compact, orientable manifold with non-trivial boundary. It is easy to check that it is irreducible and atoroidal. Thurston has shown that such manifolds (since they are Haken) have a geometric structure. (See 1.7 for a more general version of this theorem from which this one follows.) In particular, the following holds:

Theorem 7.3 (Thurston's Theorem for Manifolds with Boundary).

Suppose that M is a compact, orientable, irreducible, atoroidal 3-manifold with $\partial M \neq \emptyset$. Either the interior of M admits a complete hyperbolic structure or M is Seifert fibred.

The case when the complement of the singular locus is Seifert fibred can be dealt with because irreducibility of \mathcal{O} allows one to conclude that \mathcal{O} itself is Seifert fibred. This is essentially a matter of showing that the fibration can be extended over a neighbourhood of the singular locus in a manner consistent with the local group action.

Proposition 7.4 (Complement Seifert Fibred).

Suppose that \mathcal{O} is a compact, orientable, orbifold-irreducible 3-orbifold and

that \mathcal{O} , with an open regular neighbourhood of $\Sigma(\mathcal{O})$ removed, is a Seifert fibred 3-manifold. Then \mathcal{O} is a Seifert fibred 3-orbifold.

As noted above, orbifold-irreducible Seifert fibred orbifolds are easily seen to have geometric structures (see 2.50) so we may assume that the complement of the singular locus has a complete hyperbolic structure.

However, if the boundary of \mathcal{O} (which, if non-empty, is assumed to be Euclidean) contains any 2-orbifolds that are *not* tori, there will be components of the singular locus that go out to the boundary of \mathcal{O} . Then the complement of a neighbourhood of the singular locus will have higher genus boundary components that are not tori. Higher genus boundary components will also arise if the singular set has vertices. The complete hyperbolic structure on the complement of the singular locus will not have finite volume nor will it be unique. The deformation theory for such hyperbolic structures is quite different from that of complete hyperbolic structures with finite volume.

To avoid this situation we first remove neighbourhoods of the vertices; this introduces spherical turnover boundary components. Then we double the resulting orbifold along its non-tori boundary components. The singular locus of the double consists of simple closed curves; removing a tubular neighbourhood of this doubled singular locus results in an irreducible manifold with only torus boundary components. It can be given a finite volume hyperbolic structure except when it is Seifert fibred or an I -bundle. These exceptional cases only arise when the original orbifold \mathcal{O} is itself Seifert fibred or when \mathcal{O} is an I bundle over a Euclidean 2-orbifold. We have already dealt with Seifert fibred orbifolds and I -bundles are easily seen to be geometric. Thus we will assume that this doubled orbifold has a complete, finite volume hyperbolic structure on the complement of its singular locus.

We record this process of vertex removal and/or doubling in the following definition and theorem. In the later sections we will want to have the topological conclusions of the theorem even when only some of the vertices are removed. Thus we will allow for this possibility in the construction.

Definition. Suppose that \mathcal{O} is a compact, orbifold-irreducible 3-orbifold with boundary consisting of Euclidean 2-orbifolds. Let $\mathcal{O}_{\{v_i\}}$ denote \mathcal{O} with an open neighbourhood of a subset $\{v_i\}$ of its vertices removed. Let $\partial_{NT}\mathcal{O}_{\{v_i\}}$ denote the subset of $\partial\mathcal{O}_{\{v_i\}}$ consisting of all components that are not tori. Thus each component of $\partial_{NT}\mathcal{O}_{\{v_i\}}$ is a turnover or pillowcase. Let $D\mathcal{O}_{\{v_i\}}$ be the double of $\mathcal{O}_{\{v_i\}}$ along $\partial_{NT}\mathcal{O}_{\{v_i\}}$. (The notation $D\mathcal{O}$ will be used when no vertices are removed.) We will regard $\partial_{NT}\mathcal{O}_{\{v_i\}}$ as a

sub-orbifold of $D\mathcal{O}_{\{v_i\}}$ which separates $D\mathcal{O}_{\{v_i\}}$ into two copies of \mathcal{O} with neighbourhoods of the vertices, $\{v_i\}$, removed.

Theorem 7.5 (Doubling Trick).

Suppose that \mathcal{O} is a compact, orientable, orbifold-irreducible, orbifold-atoroidal 3-orbifold with orbifold-incompressible Euclidean boundary. Furthermore, suppose that \mathcal{O} is not an orbifold Seifert fibre space or an I -bundle over a Euclidean 2-orbifold. As above, let $D\mathcal{O}_{\{v_i\}}$ denote \mathcal{O} , with neighbourhoods of some of its vertices removed, doubled along its non-torus boundary components. Then $D\mathcal{O}_{\{v_i\}}$ has orbifold-incompressible boundary and every orbifold-incompressible pillowcase or turnover in $D\mathcal{O}_{\{v_i\}}$ is orbifold-isotopic to $\partial_{NT}\mathcal{O}_{\{v_i\}}$. If $\{v_i\}$ consists of all the vertices in \mathcal{O} , then $D\mathcal{O}_{\{v_i\}} - \Sigma(D\mathcal{O}_{\{v_i\}})$ admits a finite volume, complete hyperbolic structure.

7.2 Deforming hyperbolic structures

After this preliminary topological preparation, including using vertex removal and/or doubling if necessary, we can assume that we have an orbifold Q_0 that has a non-empty link Σ as its singular locus and that has a complete, finite volume hyperbolic structure on the complement of the singular locus.

Thinking of this complete hyperbolic structure as a cone-manifold structure on Q_0 with all cone angles equal to 0, we begin to deform the structure through a continuous family M_t of hyperbolic cone-manifold structures on Q_0 with increasing cone angles. (By definition, a cone-manifold structure on an orbifold means that there is a cone-manifold structure on the underlying manifold whose singular set is contained in that of the orbifold.) The family is parametrized so that on a component L_i of the singular locus, the cone angle in M_t along L_i is $t\theta_i$, where θ_i corresponds to the orbifold angle. If Q_0 itself has any torus boundary components, these remain cusps (i.e. cone angle 0) for all t .

Such a family, M_t , exists for t in an interval $[0, t_\infty)$ with $t_\infty > 0$ by the results in chapter 5. If the family can be extended to $t = 1$, then Q_0 can be given a hyperbolic structure. The bulk of the proof of the Orbifold Theorem consists of controlling the way the cone structures can degenerate. In the sections that follow, we describe the theorems that are proved in order to study degenerations. Here we first give a preview of the ultimate conclusions of that study.

If Q_0 was obtained by removing vertices and doubling, one type of degeneration that can occur will lead us to replace one or more of the vertices

that were removed. The result will be a hyperbolic cone-manifold structure on a new orbifold Q_1 whose singular locus will be a graph, not just a link. The cone angles on Q_1 will then be increased, with the deformation still parametrized so that at time t the cone angles will be $t\theta_i$, where the θ_i correspond to the orbifold angles. We then need to analyze the possible degenerations of this new family of structures. In order to include such families as well, the theorems below 7.6 (trouble at $t < 1$) and 7.7 (trouble at $t = 1$) are stated for orbifolds that may have vertices and for parametrization that may begin with some value of t bigger than 0.

The process of filling in vertices and the resulting topological changes are discussed in more detail below. Similarly, even if there are no vertices, it is still possible that the orbifold Q is the double, $D\mathcal{O}$, of the original orbifold \mathcal{O} in the Orbifold Theorem that we want to prove is geometric. It is necessary to draw conclusions about \mathcal{O} from the information derived about Q . Again, this is discussed below. However, these issues should be considered secondary, and the reader is encouraged, at the first reading, to assume that the original orbifold, \mathcal{O} , had no vertices and no boundary other than possibly tori. Then Q would equal \mathcal{O} throughout this chapter and all statements and conclusions about Q would apply directly to \mathcal{O} .

Theorem 7.6 (trouble at $t < 1$). *Suppose that Q is a compact, orientable 3-orbifold with singular locus a graph Σ and (possibly empty) boundary consisting of tori. Suppose that $\epsilon < t_\infty < 1$ and that there is a continuous family of hyperbolic cone-manifold structures on Q for $t \in [\epsilon, t_\infty)$. Then one of the following happens:*

- 1 (hyperbolic) *There is a hyperbolic cone-manifold structure on Q with cone angles corresponding to $t = t_\infty$.*
- 2 (Euclidean) *There is a Euclidean cone-manifold structure on Q with cone angles corresponding to $t = t_\infty$.*
- 3 (vertex filling) *Q contains an open subset which has the structure of a finite-volume complete hyperbolic 3-dimensional cone-manifold M' with cone angles corresponding to $t = t_\infty$. The 2-dimensional cross-sections of the ends of M' , when given the orbifold angles of Q , are orbifold-incompressible in Q .*

Case 3 (vertex filling) only occurs for $t_\infty < 1$ when the original orbifold had vertices, neighbourhoods of which were removed; Q was obtained by doubling. Any boundary components of M' that don't come from the boundary of Q itself, arise from the boundary of some of these vertex neighbourhoods. They appear as cusps in the hyperbolic structure on M' . In this case the cone angles can still be increased a small amount in M' and

the boundary components from these vertices can be filled in to give a cone-manifold structure on the new orbifold Q' that includes these vertices. (Topologically Q' is obtained by cutting Q open along some of the turnovers created during the vertex removal and doubling process and then adding cones to the resulting boundary turnovers.) The family of hyperbolic structures on this new cone-manifold can be extended by increasing the cone angles, at least a small amount. Once all the vertices have been filled back in, or if there were none to begin with, this case can no longer occur. The underlying topological structure at this stage will be that of the original orbifold, doubled along any non-torus Euclidean boundary components it might have had.

In case 1 (hyperbolic), there actually is no degeneration. The cone angles can be increased further, and the family can be extended beyond $t = t_\infty$. To see this when the family begins with the complete structure on the complement of the singular locus of Q ($t = 0$), note that, since the family is continuous, the holonomy representations all lie on the same component of the representation variety as that of the complete structure. Using convergence of holonomy (6.22) and the finiteness of the outer automorphism group of $\pi_1(Q - \Sigma)$, we conclude that the holonomy representation of the limiting hyperbolic cone manifold M_{t_∞} is on the same component. By the deformation theory developed in chapter 5, the cone angles can be increased further. When the family doesn't begin at the complete structure, which will occur when Q has vertices (i.e., some of the vertices that were removed have been put back in, via case 3 (vertex filling)), it is still possible to view the holonomy representations as lying on a component of a variety to which the deformation theory in chapter 5 applies. The argument is then the same. However, an explanation of this fact is beyond the scope of this outline; the reader is referred to [19] for details. We note that this situation will only arise when the original orbifold, \mathcal{O} , in the statement of the Orbifold Theorem has vertices.

In case 2 (Euclidean) the results of Hamilton on 3-manifolds with positive Ricci curvature can be applied to conclude that Q has a spherical structure. The argument used to arrive at this conclusion will be explained in Section 7.4. This case cannot occur if the boundary of Q is non-empty or if Q was obtained by doubling.

If an orbifold has a spherical structure, then it will have a finite orbifold fundamental group. Thus, the previous theorem, together with the Ricci curvature argument, implies that if Q has infinite orbifold fundamental group, then no degeneration of the hyperbolic structure is possible for

$t_\infty < 1$, other than that leading to filling in of vertices. However, considerably more can occur at the limiting value $t_\infty = 1$.

Theorem 7.7 (trouble at $t = 1$). *Suppose that Q is a compact, orientable, orbifold-irreducible 3-orbifold with singular locus a graph and (possibly empty) boundary consisting of tori. Suppose that there is a continuous family of hyperbolic cone-manifold structures on Q for $t \in [\epsilon, 1)$, for some $\epsilon < 1$. Then one of the following happens:*

- 1 (hyperbolic) Q contains a finite-volume complete hyperbolic 3-suborbifold whose ends have orbifold-incompressible cross-sections.
- 2 (Euclidean) Q is a compact Euclidean 3-orbifold.
- 3 (graph) Q is a graph orbifold.
- 4 (bundle) Q is an orbifold bundle with generic fibre a pillowcase or a turnover and base a 1-orbifold.

Using this theorem and the discussion after theorem 7.6 (trouble at $t < 1$), we can finish the proof of the Orbifold Theorem as follows:

We begin with the orbifold denoted by Q_0 at the beginning of this section. Its singular locus is a link and the complement of the singular locus has a complete, finite volume hyperbolic structure. If the original orbifold, \mathcal{O} , has no vertices and no pillowcase or turnover boundary components, then $Q_0 = \mathcal{O}$. If it has no vertices but does have either pillowcase or turnover boundary components, $Q_0 = D\mathcal{O}$, which is \mathcal{O} , doubled along its non-torus boundary components. If \mathcal{O} has vertices, $Q_0 = D\mathcal{O}_{\{v_i\}}$, which is obtained from \mathcal{O} by deleting open neighbourhoods of all of its vertices and then doubling along its non-torus boundary components.

By the deformation theory in Chapter 5, there is a continuous family of cone-manifold structures on Q_0 with cone angles $t\theta_i$, where the θ_i are the orbifold angles. The family begins with $t = 0$, the complete structure; if $t = 1$ is reached, then Q_0 has a hyperbolic structure. If $Q_0 = \mathcal{O}$, then \mathcal{O} is hyperbolic, hence geometric as desired. If Q_0 has been obtained from \mathcal{O} by any vertex removal and/or doubling, it will contain incompressible spherical turnovers and/or Euclidean turnovers or pillowcases in its interior. This is not possible for a hyperbolic orbifold, so $Q_0 = \mathcal{O}$ is the only possibility in this case.

Using theorem 7.6 (trouble at $t < 1$) and the discussion after it, we can analyze the types of degeneration that can occur as $t \rightarrow t_\infty < 1$. In case 1 (hyperbolic) of 7.6 there is no degeneration; the family can be extended. So we can assume this case doesn't occur. If case 2 (Euclidean) occurs, Theorem 7.12 (Euclidean/spherical transition) (which depends on the work of Hamilton and whose proof is outlined in Section 7.4) implies that Q_0 is

spherical. If $Q_0 = \mathcal{O}$, then \mathcal{O} is spherical, hence geometric as desired. If Q_0 has been obtained by any vertex removal and/or doubling, it can be shown to have infinite orbifold fundamental group which is not possible for a spherical orbifold. Again, $Q_0 = \mathcal{O}$ is the only possibility in this case.

As discussed after the statement of 7.6 (trouble at $t < 1$), case 3 (vertex filling) of (7.6) can only occur if the original orbifold \mathcal{O} had vertices. Then $Q_0 = D\mathcal{O}_{\{v_i\}}$ where the set $\{v_i\}$ consists of *all* the vertices of \mathcal{O} . Denote by M' the finite volume, complete hyperbolic 3-dimensional cone-manifold structure obtained as the limit at t_∞ . Then some of the cusps M' must have Euclidean turnovers as cross-sections. In Q_0 , these turnover cross-sections, with the orbifold angles, will be spherical since $t_\infty < 1$. By 7.5 (doubling trick) they are orbifold isotopic in Q_0 to some of the spherical turnovers created by removing vertices and doubling. We let C denote the orbifold with boundary obtained by giving the compact core of M' (see Proposition 7.11 (ends at $t < 1$)) the orbifold angles coming from Q_0 . It has boundary consisting of tori and spherical turnovers that are incompressible in Q_0 . It follows, using 7.5 (doubling trick), that C is homeomorphic as an orbifold to Q_0 with open neighbourhoods of some of the spherical turnovers removed. This, in turn, can be viewed as constructed by removing neighbourhoods of *all* of the vertices of the original orbifold, \mathcal{O} , but then not doubling along some of the resulting spherical turnovers, leaving them instead as boundary components.

M' is a cone-manifold structure on the interior of C where the boundary turnovers appear as Euclidean turnover cross-sections of some of the ends. It can be shown that the cone angles can be increased slightly so that the turnover cross-sections become spherical and can be filled in with a cone. The result is a hyperbolic cone-manifold structure on a new orbifold, Q_1 . The orbifold, Q_1 , is obtained from C by attaching cones to the spherical turnover boundary components; in particular, it will have vertices. It can also be viewed as obtained from the original orbifold, \mathcal{O} , by removing open neighbourhoods of only a proper (possibly empty) subset of the vertices of \mathcal{O} and then doubling along the non-torus boundary components.

The cone angles of the cone-manifold structure on Q_1 can now be increased, forming a new continuous family of hyperbolic cone-manifolds, beginning with parameter $t = \epsilon$, where $0 < \epsilon$. We can then apply the same arguments to this family. If Q_1 has also been obtained from \mathcal{O} by removing some vertices and/or doubling, then, again, the only degeneration possible as $t \rightarrow t_\infty < 1$ is case 3 (vertex filling). The only other possibility would be for there to be degeneration as $t \rightarrow t_\infty$, where $t_\infty = 1$. But this is not possi-

ble until all the vertices have been filled in. To see this, note that, if Q_1 has also been obtained by removing some vertices and doubling, it will contain incompressible spherical turnovers. For each such spherical turnover, there will be a value of t strictly less than 1 for which the angles correspond to a Euclidean cone structure on a turnover. One can show that there would actually be a totally geodesic Euclidean turnover in the hyperbolic cone structure on Q_1 ; this is impossible.

Thus we can repeat the same process until all the vertices have been filled in. Case 3 (vertex filling) of Theorem 7.6 (trouble at $t < 1$) can no longer occur. The other possibilities in Theorem 7.6 (trouble at $t < 1$), where $t_\infty < 1$, or in the case when $t = 1$ is attained, have already been shown to give geometric structures on the original orbifold, \mathcal{O} .

We are now reduced to the case when $t_\infty = 1$ and when Q is either the original orbifold, \mathcal{O} or the original orbifold doubled along its non-torus boundary components, (denoted by $D\mathcal{O}$). In particular, Q is orbifold-irreducible. Applying 7.7 (trouble at $t = 1$) and 7.5 (doubling trick) we will now show that either Q is the original orbifold, \mathcal{O} , and is geometric or $Q = D\mathcal{O}$ and we can “undouble” it to find a geometric structure on \mathcal{O} .

We now discuss the cases of 7.7 (trouble at $t = 1$). In case 2 (Euclidean) Q obviously has a geometric structure and, in case 4 (bundle) there is a geometric structure on the bundle by (2.36). In case 3 (graph) either Q is actually Seifert fibred, hence geometric by (2.50), or it has an incompressible, non-peripheral Euclidean 2-suborbifold. Similarly, in case 1 (hyperbolic) either the 3-suborbifold is all of Q and Q is geometric or Q contains an incompressible, non-peripheral Euclidean 2-suborbifold. Since the original orbifold, \mathcal{O} , had no such 2-suborbifolds, then, in both cases (3 graph) and (1 hyperbolic), if Q is not geometric, it must have been obtained by doubling. Thus, in all cases, if Q is the original orbifold, \mathcal{O} , it is geometric.

Now suppose that Q was obtained as a double; then $Q = D\mathcal{O}$, where \mathcal{O} is the original orbifold. This might a priori occur even when Q is geometric. By 7.5 (doubling trick) the incompressible, non-peripheral Euclidean 2-suborbifolds in $D\mathcal{O}$ are precisely those that come from the boundary components along which the doubling occurred. In case 1 (hyperbolic), since no hyperbolic 3-orbifold contains an incompressible, non-peripheral Euclidean 2-suborbifold, these doubling 2-suborbifolds must all be contained in the boundary of the hyperbolic 3-suborbifold. Since there are no other such non-peripheral Euclidean 2-suborbifolds, it must be the case that the hyperbolic 3-suborbifold equals \mathcal{O} and the original orbifold \mathcal{O} is geometric.

Case 4 (bundle) can't occur as a double since all the incompressible, non-

peripheral Euclidean 2-suborbifolds have the property that cutting along them leads to an I -bundle over a Euclidean 2-orbifold, a case that was ruled out since it could be handled directly. It can be shown that the Euclidean 2-suborbifolds created in Q by doubling \mathcal{O} are represented by totally geodesic 2-dimensional sub-cone-manifolds in the approximating hyperbolic cone-manifold structures. If Q is Euclidean (case 2), the doubling suborbifolds will be totally geodesic so \mathcal{O} will also be Euclidean. Similarly, in case 3 (graph), if Q is Seifert fibred, it can be shown that \mathcal{O} is Seifert fibred. In case 3 (graph), if Q is not Seifert fibred, it is a union of Seifert fibred orbifolds glued along incompressible Euclidean 2-suborbifolds. By 7.5 (doubling trick) these must have come from the boundary components of \mathcal{O} since \mathcal{O} was orbifold-atoroidal. Thus the Seifert fibred pieces of the graph orbifold must equal \mathcal{O} , possibly doubled along a subset of its boundary components. By the previous argument, \mathcal{O} itself is Seifert fibred.

This completes the outline of the proof of the Orbifold Theorem, assuming Theorem 7.6 (trouble at $t < 1$) and Theorem 7.7 (trouble at $t = 1$).

7.3 Controlling degenerations

In this section we give an outline of the theorems that are needed to control the family of hyperbolic cone-manifolds that we are studying and explain how they lead to proofs of 7.6 (trouble at $t < 1$) and 7.7 (trouble at $t = 1$).

As above, we begin with an orbifold Q (for example $DO_{\{v_i\}}$) whose singular locus is a non-empty graph and we assume that it has a smooth family of hyperbolic cone-manifold structures, M_t , whose cone angles equal t times the orbifold angles of Q . The set of $t \in [0, 1]$ such that M_t is a hyperbolic cone-manifold is (relatively) open and non-empty. If $t = 1$ is in this set then there is a hyperbolic structure on Q . Otherwise, for some $\epsilon > 0$, there will be a maximal t_∞ in $(\epsilon, 1]$ for which M_t is hyperbolic for all t in $[\epsilon, t_\infty)$. Our goal is to understand the behaviour of M_t as $t \rightarrow t_\infty$.

The primary geometric quantity that controls this behaviour is the injectivity radius. (See Chapter 6.) We begin the analysis by considering the different possibilities for the injectivity radii of points in family of cone-manifolds, M_t .

Case 1 (inj bd everywhere). The injectivity radius is bounded below over **all** points in M_t and all $t \in [\delta, t_\infty)$, for some $\delta < t_\infty$.

Note that we need to stay away from those hyperbolic structures with cusps in order to bound below injectivity radius. There are cusps at $t = 0$

and also just before a vertex is filled in. This is the only role of δ in the statement.

The volumes of the M_t are bounded above; indeed, by the Schläfli formula 3.20, they are decreasing as t increases. Thus, if all the injectivity radii are bounded below, then the M_t are covered by a uniform number of standard balls. So they are all compact, and their diameters are bounded above. As $t \rightarrow t_\infty$, the M_t converge (in the Gromov-Hausdorff topology) to a cone-manifold structure, M_{t_∞} , on a homeomorphic underlying space (M, Σ) .

The only subtlety in this case is that the family of holonomy representations $h_t : \pi_1(M - \Sigma) \rightarrow \text{Isom}(\mathbb{H}^3)$ associated to the structures also converges. By convergence of holonomy (6.22), they will converge up to automorphisms of $\pi_1(M - \Sigma)$. Using the finiteness of the automorphism group in this case, a subsequence will converge.

If $t_\infty = 1$, then Q has a geometric structure and we are done. (This is a simple example of case 1 (hyperbolic) in 7.7 (trouble at $t = 1$).) If $t_\infty < 1$, then the family can be extended. (This is a simple example of case 1 (hyperbolic) in 7.6 (trouble at $t < 1$).)

Note that, if Q has boundary, M_t will have at least one cusp for all t so this case won't occur. However, if Q has no boundary and is orbifold-atoroidal and orbifold-irreducible, this is the "generic" case.

Case 2 (inj bd at base pt). There are points $z_t \in M_t$ at which the injectivity radius is uniformly bounded below for all $t \in [\delta, t_\infty)$, for some $\delta < t_\infty$.

We take the z_t to be our basepoint and consider convergence in the (based) Gromov-Hausdorff topology. The first step in this analysis is to show that, if the injectivity is uniformly bounded below at a sequence of points, then it is uniformly bounded below (with a different bound, of course) in the ball of a fixed radius around those points. Furthermore, the relation between the two bounds can be made independently of the underlying topology. The precise statement is:

Theorem 7.8 (Bounded Decay of Injectivity Radius).

Given $\epsilon, \delta, r > 0$ there is $\eta > 0$ such that if M is any complete 3-dimensional cone-manifold of constant curvature $\kappa \in [-1, 0]$ and with all cone angles in $[\delta, \pi]$ and if $x, z \in M$ with $\text{inj}(z) > \epsilon$ and $d(z, x) < r$ then $\text{inj}(x) > \eta$.

In particular, this theorem implies that no pieces of the singular locus can come together and no sets can collapse within any bounded dis-

tance of the basepoint. It follows that the limit will again be a complete 3-dimensional cone-manifold.

Theorem 7.9 (3d limit). *Let (M_n, z_n) be a sequence of pointed, complete 3-dimensional cone-manifolds with constant curvature $\kappa_n \in [-1, 0]$ and uniformly bounded volume and with all cone angles in $[\delta, \pi]$. Suppose there is an $\epsilon > 0$ such that $\text{inj}(z_n) > \epsilon$ for all n . Then there is a subsequence (M_{n_i}, z_{n_i}) converging in the Gromov-Hausdorff topology to a complete, finite volume 3-dimensional cone-manifold of curvature κ , where $\kappa_{n_i} \rightarrow \kappa$.*

However, if we are not in the previous case 1 (inj bd everywhere) and the injectivity radius goes to 0 at a sequence of points, the diameter will go to infinity and there is no guarantee that the limiting cone-manifold will be homeomorphic to the approximates. One thing that can happen is that a cusp develops. Just before this happens, the approximate cone-manifolds become stretched out so that they contain a submanifold that is topologically the product of a compact 2-dimensional Euclidean cone-manifold with a long interval, where the metrics on the 2-dimensional cross-sections are scaled down exponentially as one moves along the interval. Part of the approximates may then pinch off in the limit.

We show below that this is the only way that the limit can differ topologically from the approximates under the hypothesis that the injectivity radius at the basepoint is bounded below. Furthermore, we show that the creation of new cusps can only occur as a result of pillowcases or turnovers that are either boundary parallel or were created by doubling (possibly after removing vertices).

In order to see that this is the only limiting behaviour that can occur, one notes that the Gromov-Hausdorff topology provides almost isometric maps of larger and larger diameter pieces of the geometric limit into the approximates (6.21). In order to control the limiting behaviour, we need first to understand the ends of the geometric limits and then derive some topological conclusions about the maps and about the topology of the approximates. The following propositions characterize the ends of finite volume, complete hyperbolic cone-manifolds with cone angles at most π . The two cases correspond to the cases $t_\infty = 1$ (in which case the limit is an orbifold, not just a cone-manifold) and $t_\infty < 1$, respectively in our limiting procedure of hyperbolic cone-manifolds M_t as $t \rightarrow t_\infty$.

Proposition 7.10 (ends at $t = 1$). *Suppose that Q is a complete, finite-volume, hyperbolic 3-orbifold. There is a compact (non-convex) core C of Q such that each component, E , of $Q - C$ is isometric to the quotient of a*

torus cusp by a finite group of isometries. Thus the closure of E is orbifold isomorphic to $F \times [0, \infty)$ where F is an orientable, closed, Euclidean 2-orbifold: a turnover, pillowcase or torus.

Proposition 7.11 (ends at $t < 1$). *Suppose that M is a complete, finite-volume, hyperbolic 3-dimensional cone-manifold with cone angles in $(0, \theta_0]$ for some $\theta_0 < \pi$. Then $M = C \cup E_1 \cup \cdots \cup E_n$ where C is compact and each $E_i \cong F_i \times [0, \infty)$ is a cusp. Each F_i is a turnover or torus, and $C \cap E_i = \partial E_i \cong F_i$.*

Proposition 7.10 (ends at $t = 1$) follows easily from the fact that Q is finitely orbifold-covered by a hyperbolic manifold. Proposition 7.11 (ends at $t < 1$) requires knowledge of the possible non-compact 3-dimensional Euclidean cone-manifolds. A discussion of this topic appears in the last section of this chapter.

If such a cusp develops as t approaches t_∞ , then the Gromov-Hausdorff topology implies that there are almost isometric maps of large compact pieces of the geometric limit into the cone-manifolds. If $t_\infty < 1$ and the geometric limit has an end with a non-torus cross-section, then 7.11 (ends at $t < 1$) implies that there are turnovers in the cone-manifolds whose angle sums approach 2π as $t \rightarrow t_\infty$. Since $t_\infty < 1$ the turnovers must be spherical in the orbifold; hence they must be the result of removing vertices in the original orbifold and doubling.

If $t_\infty = 1$ (or if there are only torus cross-sections when $t_\infty < 1$), all the cross-sections of ends in the geometric limit will be Euclidean orbifolds. We claim that, for t sufficiently close to t_∞ , the images of these cross-sections will be **incompressible** in the orbifold Q . Assuming this claim, we can finish case 2 (inj bd at base pt).

Using the Gromov-Hausdorff topology, we obtain embeddings of the compact core C , as described in 7.10 (ends at $t = 1$) and 7.11 (ends at $t < 1$), of the limit cone-manifold or orbifold. If $t_\infty < 1$, the image of the boundary of C consists of tori that are incompressible in Q , and turnovers that, with their angles replaced by the orbifold angles, are spherical in Q . The tori, since they are incompressible, must be boundary parallel and, by 7.5 (doubling trick), the turnovers must be orbifold-isotopic to those created by removing vertices. It follows that Q contains a finite-volume complete hyperbolic 3-dimensional cone-manifold M' , (homeomorphic to the interior of C), with angles corresponding to $t = t_\infty$. This is case 3 (vertex filling) of 7.6 (trouble at $t < 1$). Note that, as discussed in the previous section, if the original orbifold \mathcal{O} had no vertices or if they have all been filled in,

Q will not contain any spherical turnovers. Thus, all the boundary of C is boundary parallel in Q and M' will be homeomorphic to Q . Contrary to the choice of t_∞ , Q has a hyperbolic cone-manifold structure with cone angles corresponding to $t = t_\infty$. No degeneration has occurred and the family can be extended.

Similarly, if $t_\infty = 1$, the image of the compact core, C , will have boundary consisting of incompressible tori (which must be boundary parallel) and incompressible Euclidean turnovers and pillowcases (which are the result of doubling). It follows that Q contains an orbifold-incompressible finite-volume complete 3-dimensional hyperbolic suborbifold. This is case 1 (hyperbolic) in 7.7 (trouble at $t = 1$).

It remains to be seen why the tori, pillowcases, and Euclidean turnovers are *incompressible* in Q . The turnovers are trivially incompressible because every simple closed curve in them bounds a (singular) disk in the turnover itself. We will first consider the case when there are no vertices and when there are only tori in the boundary of C . We then sketch the changes necessary when there are vertices or pillowcases.

The complement of the singular locus in Q is irreducible so, if an embedded torus is compressible in the complement of the singular locus, it is either contained in a ball or bounds a solid torus. Since $Q - \Sigma(Q)$ is also atoroidal, an incompressible torus must be boundary parallel.

The holonomy of elements of the fundamental group of the boundary of the core C are all parabolic in the geometric limit, so, by convergence of holonomy, any given element must become arbitrarily close to parabolic in the approximates. This is not possible if the torus is contained in a ball, in which case the holonomies are all trivial. Thus, if the torus is compressible in Q , it must bound a (singular) solid torus. The meridian curve will be represented by either the trivial element, if the solid torus is non-singular, or an elliptic element with rotation bounded away from 0, if the solid torus is singular. Therefore, by convergence of holonomy, in a sequence of approximates any given curve on a torus boundary of C can be a meridian at most a finite number of times.

The subtlety here is that, a priori, in a sequence of approximations, C could be mapped into Q in topologically distinct ways so that an infinite sequence of distinct curves bound (singular) disks in Q . Let Q_n denote the image of C under the n th approximating map union the (singular) solid tori bounded by the compressible tori. Then Q_n is obtained from C by Dehn filling along the n th meridians, μ_n^i . Viewed as a suborbifold of Q , the boundary of Q_n is incompressible. It is not difficult to show that there are

only a finite number of 3-dimensional sub-orbifolds of Q with incompressible boundary, up to isotopy. After taking a subsequence, we can assume that the Q_n are all diffeomorphic to the same orbifold, Q_∞ .

If $t_\infty < 1$ the cone angles will be less than the orbifold angles in Q and in Q_∞ . To avoid using Mostow-Kojima rigidity for hyperbolic cone-manifolds 5.11 (see [53]) which depends on arguments similar to those in the proof of the Orbifold Theorem (including those in the next few paragraphs) and on local rigidity of cone-manifolds 5.10 (see [45]), we remove a neighbourhood of the singular locus.

Let N_∞ denote Q_∞ with a regular neighbourhood of its singular locus removed and let \hat{C} denote C with a regular neighbourhood of its singular locus removed. It is not hard to see that the complement of the singular locus of a finite volume hyperbolic cone-manifold can be given a complete metric of strictly negative curvature. (See, e.g., [53]; a further argument is required in the case with vertices.) Hence, \hat{C} is irreducible and atoroidal and by 7.3 (Thurston's Haken theorem) it has a finite volume hyperbolic metric.

Suppose that N_∞ is homeomorphic to \hat{C} ; i.e., that the only solid tori added are singular. An infinite number of distinct curves on the boundary of \hat{C} are mapped to the curves on the boundary of N_∞ that bound singular disks in Q . This implies that \hat{C} has an infinite number of self-homeomorphisms which are homotopically distinct. But, since \hat{C} can be given a complete finite volume hyperbolic metric, Mostow rigidity implies that the group of self-homotopy equivalences which are homeomorphisms on the boundary is finite, a contradiction.

Thus N_∞ is obtained from \hat{C} by an infinite sequence of Dehn fillings where, on each filled-in torus, the same curve is a meridian at most a finite number of times. But, by Thurston's hyperbolic Dehn surgery theorem, all but a finite number of these Dehn fillings result in finite volume hyperbolic manifolds. Furthermore, the hyperbolic structures on such a sequence of fillings has arbitrarily short closed geodesics. By Mostow rigidity there is a unique hyperbolic structure on N_∞ ; by discreteness and finite volume, there is a shortest geodesic. This gives a contradiction. Thus, all the images of the boundary tori of \hat{C} , hence of C , are incompressible in Q as claimed.

When C has pillowcase boundary components, the argument is very similar. If the image of such a pillowcase in Q is compressible, it either is contained in a singular ball with a single unknotted arc of singular locus or it bounds a folded ball. In the first case the holonomy of the image pillowcase would be a single elliptic element. As before, this is impossible

by convergence of holonomy for approximating maps sufficiently far out in the sequence.

If the image of a pillowcase bounds a folded ball, a simple closed curve, called the *meridian*, in the pillowcase which does not bound a (singular) disk in the pillowcase does bound a non-singular disk in the folded ball. Convergence of holonomy again implies that a single curve can be a meridian at most a finite number of times since, in C , the holonomy of every element in the orbifold fundamental group of the pillowcase is parabolic and non-trivial. If there are pillowcases in C , then $t_\infty = 1$ and C is an orbifold, not just a cone-manifold. It has a complete, finite volume hyperbolic structure on its interior. We again conclude that the same orbifold, Q_∞ , is obtained by orbifold Dehn fillings on infinitely many distinct meridians on each component. The theory of hyperbolic Dehn filling, as extended to orbifolds by Dunbar-Meyerhoff ([64]), leads, as before, to a contradiction, using Mostow rigidity applied to Q_∞ .

When there are vertices, the argument is essentially the same. However, when the limit, C , is not an orbifold but only a cone manifold, a further argument beyond that contained in [53] is required in order to show that there is a hyperbolic structure on the complement of the singular locus where the holonomies around the edges connecting the vertices are all parabolic. This fact is used to show it is not possible to obtain the same manifold by Dehn filling on infinitely many distinct curves on each torus boundary component of \hat{C} . The argument then proceeds as before. This completes the outline of the proof of Case (2 inj bd at base pt).

Case 3 (inj \rightarrow 0 everywhere). The injectivity radius goes to 0 for all $x \in M_t$ as $t \rightarrow t_\infty$.

In this case, the diameter of M_t may actually go to 0. If so, we rescale so that the diameter is 1. If not, we don't rescale. There are 2 subcases here, depending on whether or not the injectivity radius in the (possibly) rescaled metric goes to 0 at all points.

Case 3a (rescaled inj bd). The injectivity radius does not go to 0 all $x \in M_t$ when the diameter is scaled to equal $\max(1, \text{diam } M_t)$.

We are assuming that we are not in Case (2 inj bd at base pt), so the injectivity radius goes to 0 everywhere in the unscaled metric. The metric must have been scaled to have diameter 1 in this case. By 7.8 (decay of inj), since the diameter is bounded above, the injectivity radius must be

uniformly bounded below at all points in the rescaled metric. By 7.9 (3d limit), the limit as $t \rightarrow t_\infty$ will be a compact Euclidean cone-manifold.

If $t_\infty = 1$, we are at the orbifold angles and Q has a Euclidean structure. This is case 2 (Euclidean) of 7.7 (trouble at $t = 1$).

If $t_\infty < 1$, we will argue, using the work of Hamilton, that Q has a spherical structure. This is case 2 (Euclidean) of 7.6 (trouble at $t < 1$).

We record this step as the following theorem. The argument will be outlined in the next section.

Theorem 7.12 (Euclidean/spherical transition). *Let Q be a compact orbifold with a Euclidean cone structure with some cone angles strictly less than the orbifold angles. Then Q has either a spherical structure or a $S^2 \times \mathbb{R}$ structure. If the Euclidean cone structure arises as a rescaled limit (in the Gromov-Hausdorff topology) of hyperbolic cone structures on Q , then Q has a spherical structure.*

Case 3b (collapsing). The injectivity radius goes to 0 for all $x \in M_t$ when the diameter is scaled to equal $\max(1, \text{diam } M_t)$.

This is the most complicated case in the analysis, which we refer to as the “collapsing case”. In the manifold context there has been considerable analysis (see [17], [18], [66], [29]) of the topology of manifolds that admit a sequence of metrics with curvature bounds where the injectivity radius goes to 0 at every point. Such manifolds are shown to possess a generalized Seifert fibred structure called an “F-structure”. A 3-dimensional manifold with an F-structure is a graph manifold.

The theorems below may be viewed as a generalization to cone-manifolds of these theorems. However, it is not apparent at this time that the techniques in the manifold context generalize directly.

We say that a 3-dimensional orbifold Q has an ϵ -collapse if there is a 3-dimensional hyperbolic cone-manifold M with $\text{inj}(x) < \epsilon \cdot \min(1, \text{diam}(M))$ for all $x \in M$, and a homeomorphism $f : Q - \partial Q \rightarrow M$ such that $f(\Sigma(Q - \partial Q)) = \Sigma(M)$. In addition, for every edge e of $\Sigma(Q)$, the difference between the cone angle on e in Q and the cone angle on $f(e)$ in M is less than ϵ . It is often convenient to use orbifold terminology when referring to M so we will pass back and forth between M and Q .

Theorem 7.13 (Collapsing Theorem for Cone-manifolds). *Suppose that Q is a compact, orientable, orbifold-irreducible 3-orbifold with non-empty 1-dimensional singular locus and with orbifold-incompressible Euclidean boundary. Then there is $\epsilon > 0$ such that if Q has an ϵ -collapse*

then either


- (1) Q is a graph orbifold or
- (2) Q is an orbifold bundle with generic fibre a turnover or pillowcase and base a 1-orbifold.

Furthermore there is an edge of $\Sigma(Q)$ labelled 2.

We discuss the theorems used in the proof of this theorem in the final section of this chapter.

Cases (1) and (2) of the collapsing theorem correspond to cases 3(graph) and 4 (bundle) in 7.7 (trouble at $t = 1$) . Assuming the theorems stated in this section, this concludes the outline of the proofs of 7.6 (trouble at $t < 1$) and 7.7 (trouble at $t = 1$) and, hence, of the Orbifold Theorem.

We summarize the logic in this section in the following flow diagram:



7.4 Euclidean to spherical transition

In this section we outline the proof of 7.12 (Euclidean/spherical transition).

The basic idea behind this theorem is that, if there is a Euclidean cone-manifold structure on Q with angles strictly less than the orbifold angles, then one should be able to spread out some of the concentrated curvature away from the singular locus to obtain a metric with the orbifold angles and some positive curvature on the smooth part of the orbifold. The Ricci flow on Q with this metric should either lead to a spherical orbifold metric or imply that there is an $S^2 \times \mathbb{R}$ structure on Q .

When the singular locus is a link, this process of “smoothing” the metric to the orbifold angles can be done simply and explicitly. When there are vertices, it is less clear how to reach the orbifold angles while maintaining positive curvature on the smooth part of the orbifold so a more ad hoc argument is used. Furthermore, Hamilton’s results are proved only for manifolds so, in all cases, a device for finding an orbifold cover which is a manifold is required. (Hamilton has an unpublished manuscript [36] which generalizes his results to orbifolds, but we won’t use that here.)

The following theorem of Hamilton is in [37].

Theorem 7.14. *A compact 3-manifold, M , with non-negative Ricci curvature which is not everywhere flat is diffeomorphic to a quotient of either S^3 or $S^2 \times \mathbb{R}$ by a group of fixed point free isometries in the standard metrics. Furthermore, if the original metric with non-negative Ricci curvature has a non-trivial group of symmetries, the homogeneous metric on M will possess the same group of symmetries.*

We suppose that Q has a Euclidean cone structure with cone angles strictly less than the orbifold angles. Assume that the singular locus is a link; i.e., assume that there are no vertices. Consider disjoint singular solid tori, each containing a single component of the singular locus. Assume that each torus consists of points a constant distance from the component it contains. The cross-sections perpendicular to the singular loci are Euclidean disks with a single cone point. The metrics on these can be smoothed in a rotationally symmetric way to obtain a smooth metric with non-negative Ricci curvature on the underlying manifold X of Q .

This smoothed metric on X is not flat so 7.14 (Hamilton’s theorem) implies that X is finitely covered by S^3 or $S^2 \times S^1$. Assume that we are in the S^3 case. Viewing Q as X with a link determining the singular locus, we can lift to the topological universal cover of X which is homeomorphic to S^3 . This defines an orbifold cover \tilde{Q} of Q whose underlying space is S^3

and whose singular locus is a link. The homology of the complement of the link is a direct product of infinite cyclic groups, each of which is generated by a meridian around a component of the link. We map onto a product of finite cyclic groups, sending the meridian to a generator and killing the p th power if the singular component has local group \mathbb{Z}_p . The kernel of this homomorphism defines an orbifold covering of \tilde{Q} which is a manifold. We denote this manifold by M . It also is an orbifold cover of Q .

Return to the Euclidean cone-manifold structure on Q . This time we “smooth” the metric so that the cone angles along the singular locus equal the orbifold angles. This is done in the same way as before, in tubular neighbourhoods of each singular component, in a radially symmetric fashion on each transverse disk. The new metric on each disk still has a singular point at the centre of the disk but it has the orbifold angle; the smooth portion of the disk has some positive curvature near the singular point.

This metric lifts to a smooth metric on the manifold cover M with non-negative Ricci curvature. Hamilton’s theorem implies that the Ricci flow converges to a spherical metric possessing any symmetries that the original metric on M had. Thus the spherical metric descends to Q and Q has a spherical structure as desired.

If the underlying manifold X is finitely covered by $S^2 \times S^1$, then Hamilton’s proof shows that the algebraic splitting of the curvature operator that appears in an $S^2 \times \mathbb{R}$ will exist in all the metrics that occur in the Ricci flow for all *positive* times. Since the smoothed metric on X has concentrated positive curvature orthogonal to the original singular locus, this must be compatible with the splitting. From this, it follows that Q is, up to a 2-fold cover, homeomorphic to a bundle over S^1 with fibre S^2 with finitely many singular points. Hamilton’s proof also provides a metric with positive curvature on the fibres for all positive times, so using Gauss-Bonnet and the fact that all the cone angles are less than π , there will be at most 3 singular points in the fibres.

From this description, it is apparent that Q is, in fact, Seifert fibred with the singular locus contained in the fibres. This is easily seen to have a geometric structure. However, it is also not hard to see that such an orbifold cannot have a hyperbolic metric in the complement of its singular locus, so this case does not actually arise in our context.

When there are vertices, it is less clear how to do the smoothing parts of the argument so we resort to a more ad hoc argument which we hope to simplify in the near future.

We will call a vertex whose link is $S^2(2, 2, n)$ a *dihedral vertex*; an edge, labelled n , joining two distinct dihedral vertices is called a *dihedral edge* of order n . Such an edge has a neighbourhood whose 2-fold branched cover over the edges labelled 2 is a tubular neighbourhood of a closed curve labelled n . The process of smoothing such a neighbourhood (both to the orbifold angle and to angle 2π) described above was radially symmetric on each disk cross-section so it is symmetric with respect the order 2 symmetry on the cover.

If Q has a Euclidean cone structure with some of its angles less than the orbifold angles, it can be seen to have finite orbifold fundamental group. Otherwise, the orbifold universal cover is non-compact and would contain an bi-infinite ray. This can be ruled out, using triangle comparison theory. The same argument shows that any orbifold obtained from Q by decreasing the labels on some edges (i.e. increasing the desired orbifold cone angles) will also have finite orbifold fundamental group. Since the only orbifolds with a geometric structure that can have finite orbifold fundamental group are spherical, it suffices to show that such an orbifold is geometric in order to conclude that it is spherical.

To guarantee the existence of a dihedral edge in our orbifold, we change all of the labels to 2's. Denote this orbifold by Q_2 . We claim that it is geometric, hence spherical.

Assuming that Q_2 is spherical, its topological type, including the singular graph, is an OSFS and belongs to one of a few known families (see [26], [28]). The original orbifold Q has the same topology with some of the labels increased. By looking at each family, it is then possible to show that changing the labels leads either to an orbifold with infinite orbifold fundamental group, which is impossible for Q , or to a spherical orbifold.

In order to obtain a spherical structure on Q_2 , we note that there must be a dihedral edge (of order 2) if there are any vertices. We first attempt to find a hyperbolic structure on Q_2 where the holonomy around the dihedral edge remains parabolic. This parabolicity requirement has the effect of removing a neighbourhood of the edge, creating a sphere with 4 cone points on the boundary. The cone angles begin at 0; we attempt to increase them to π . Throughout the deformation the meridian curve will be parabolic. If we reach cone angles π , the boundary will become a pillowcase cusp.

The only way this sequence of hyperbolic structures can degenerate is for it to collapse at time $t = 1$, in which case Q_2 with the neighbourhood removed is a graph manifold. By the arguments in the last section (using the latitude hypothesis) on collapsing, the fibration can be extended over the

folded ball that the pillowcase bounds in Q_2 . Thus Q_2 is a graph manifold; since it has finite orbifold fundamental group it is a spherical OSFS.

If we reach the final angles, we can begin to increase the cone angle around the dihedral edge. (At the angles of the form $\frac{2\pi}{k}$, this amounts to doing hyperbolic Dehn filling along a meridian of the pillowcase.) Since Q_2 can't be hyperbolic, this must degenerate at some stage. Either we again conclude that Q_2 is a graph manifold, hence a spherical OSFS or we obtain a Euclidean structure on Q_2 with cone angles π along all edges except the dihedral edge.

One can then find an orbifold cover which unfolds the angle π edges, leaving one with a link singularity in the cover. This can be smoothed symmetrically with respect to the covering maps as described above. The argument now proceeds as before in the link singularity case. The spherical structure obtained from the Ricci flow will be symmetric and descend to one on Q_2 .

7.5 Analysis of the thin part

We have seen that as long as the injectivity radius is bounded below, controlling the degeneration of the family of hyperbolic cone-manifolds is not difficult. However, when the injectivity radius goes to 0, controlling the topology of the geometric limits becomes more difficult. In the outline, there were two key theorems concerning the topology of the regions where the injectivity radius is small. The first 7.11 (ends at $t < 1$) described the ends of a finite volume hyperbolic cone-manifold. The cross-sections of these ends provided us with 2-dimensional Euclidean sub-cone-manifolds that put strong topological limitations on the orbifolds Q that could degenerate when the injectivity radius was bounded below at the basepoint but went to 0 elsewhere. The second theorem 7.13 (Collapsing theorem) provided a topological classification for those orbifolds Q that could admit a family of metrics where the injectivity radius went to 0 everywhere.

One way to understand the topological and metric structure near a point where the injectivity radius is going to 0 is to rescale the metric so the injectivity radius is 1, using that point as the basepoint. Since the injectivity radius at the basepoint goes to zero, the sequence of scale factors will go to infinity. Applying 7.8 (decay of inj), one sees that the limit of the scaled structures will be a complete Euclidean cone-manifold. If the diameter of the original sequence doesn't go to zero (or goes to zero at a slower rate than the injectivity radius), the diameter of the geometric limit will be infinite

and the cone-manifold will be non-compact.

One of the most important tools in the proof of the Orbifold Theorem is the classification of non-compact 3-dimensional cone-manifolds whose cone angles are at most π . The restriction on the cone angles is crucial to this theorem. If the cone angles are allowed to lie between 0 and 2π , the number of possibilities becomes unbounded.






Theorem 7.15 (Bieberbach-Soul Theorem). *A non-compact, orientable, 3-dimensional Euclidean cone-manifold with cone angles in $(0, \pi]$ is isometric to one of the following.*

- (1) *A cone.*
- (2) *A (possibly singular) solid torus, possibly with a twisted product metric.*
- (3) *The product of a compact orientable 2-dimensional Euclidean cone-manifold with a line:*
 - (i) *torus* $\times\mathbb{R}$,
 - (ii) *pillowcase* $\times\mathbb{R}$, or
 - (iii) *turnover* $\times\mathbb{R}$.
- (4)
 - (i) *A folded ball, or*
 - (ii) *a singular folded ball.*
- (5)
 - (i) $D^2(\pi, \pi) \times S^1$,
 - (ii) *a twisted product $D^2(\pi, \pi) \tilde{\times} S^1$, or*
 - (iii) *a twisted line bundle over a Klein bottle.*
- (6)
 - (i), (ii), (iii) \mathbb{R}^3 *with the singular locus shown, or*
 - (iv) *a twisted line bundle over $\mathbb{R}P^2(\pi, \pi)$.*
- (7)
 - (i), (ii) \mathbb{R}^3 *with the singular locus shown.*

Remark: These are illustrated in the following figure. The reader may wish to refer to Chapter 2 for an explanation of some of the terms in the theorem. In particular, definition 2.48 extends to cone-manifolds in the obvious way.

This theorem is actually a special case of a general theorem about non-compact, orientable, n -dimensional Euclidean cone-manifolds with cone angles at most π . That theorem states that such a Euclidean cone-manifold is, up to a 2-fold branched cover, isometric to a normal bundle of a lower dimensional, compact Euclidean cone-manifold.

This general theorem is analogous to the Bieberbach Theorems for Euclidean manifolds. In particular, it reduces the classification of non-compact,



Euclidean cone-manifolds to that of compact, lower dimensional Euclidean cone-manifolds and involutions on them. For example, the possible non-compact, orientable, 2-dimensional Euclidean cone-manifolds are a cone (normal bundle over a point with some angle; this includes the plane), an infinite cylinder (normal bundle over a circle), and an infinite pillowcase (normal bundle over an interval with angle π attached to its endpoints, which is the circle divided out by an involution). The 3-dimensional theorem above follows from the classification of 0, 1, and 2 dimensional Euclidean cone-manifolds with angles at most π and involutions on them.

The proof of the general theorem starts by following the outline of the proof of the Soul Theorem, due to Cheeger and Gromoll ([16]) which gives a structure theorem for non-compact manifolds with non-negative Ricci curvature. Indeed, in some cases the topology of the underlying space of the Euclidean cone-manifolds can be inferred directly from the Soul Theorem if one can smooth the metric to obtain a positively curved one. (The topology of the soul may change, however.) This argument works for all cone angles at most 2π . However, the more precise isometric description as a normal bundle is only true for cone angles at most π and requires further analysis.

The compact set, C , for which the Euclidean cone-manifold, B , is the normal bundle is called the **soul** of B . It can be described in terms of Busemann functions on B . A Busemann function, b_γ , is determined by an infinite ray γ ; it is defined by $b_\gamma(x) = \lim_{t \rightarrow \infty} d_B(x, \gamma(t)) - t$, where $d_B(\cdot, \cdot)$ denotes distance in B . The soul C is derived from the level set for the maximum value of the function obtained by taking the infimum over all rays emanating from a chosen point, $p \in B$. This construction is at the centre of [16].

The 3-dimensional theorem above gives a list of the possible geometric limits under the scaling process, at least when the limit is non-compact. The Gromov-Hausdorff topology implies that neighbourhoods of points with small injectivity radius in hyperbolic cone-manifolds can be approximated by these. This leads to a structure theorem for the topology of the set with small injectivity radius that generalizes the Margulis Lemma for hyperbolic manifolds. The following theorem, which is a consequence of the generalized Margulis lemma, says that, if a point in a 3-dimensional hyperbolic cone-manifold, M , has a small injectivity radius when M is scaled to have diameter at least 1, it has a neighbourhood that is almost isometric to a neighbourhood of the soul in a non-compact Euclidean cone-manifold. The possible list for such models comes from 7.15 (Bieberbach-Soul theorem), where a few cases have been eliminated using the finite volume hypothesis.

Theorem 7.16 (Local Margulis for Cone-Manifolds). *Let M be a finite volume 3-dimensional hyperbolic cone-manifold with cone angles in the range $(0, \pi]$. Then there is an $\epsilon > 0$ so that if $x \in M$ with $\text{inj}(x) < \epsilon \min(1, \text{diam}(M))$ then x has a compact neighbourhood, containing $N(x, 1000\text{inj}(x))$, which is almost isometric to one of the following:*

- (1) *A (singular) solid torus,*
- (2) *A (singular) folded ball,*
- (3) *A thick torus, pillowcase, or turnover,*
- (4) *A folded thick torus, pillowcase or turnover.*

The main idea in the proofs of both the Collapsing Theorem for cone-manifolds 7.13 (Collapsing theorem) and the structure of the ends of cone-manifolds 7.11 (ends at $t < 1$) is to use the neighbourhoods of points with small injectivity radius whose topology is described by this theorem and analyze how they can be glued together.

The analysis of the ends of hyperbolic cone-manifolds in 7.11 (ends at $t < 1$) is simplified by the fact that, since $t_\infty < 1$ and we have not reached the orbifold angles, all the cone angles are strictly less than π . This reduces the list of possible local models from 7.16 (local Margulis) to a (singular) solid torus, a thick torus, or a thick turnover. The (singular) solid tori can be incorporated into the compact part. The remainder of the proof involves showing that any pair of standard neighbourhoods homeomorphic to a thick turnover (torus) that intersect can be amalgamated into a larger neighbourhood homeomorphic to a thick turnover (torus). The product ends are created in this manner.

7.6 Outline of the Collapsing Theorem

In this section we will outline the proof of the Collapsing Theorem 7.13, which states that, if an orbifold Q has an ϵ -collapse, for sufficiently small ϵ , then it is either a graph orbifold or an orbifold bundle with generic fibre a turnover or a pillowcase.

The proof begins with the local Margulis theorem 7.16 (local Margulis) which provides the local models from which the orbifold is built. There is an analogy to a child's construction kit. The construction kit contains pieces which are Euclidean models that are almost isometric to standard neighbourhoods and we can build orbifolds by fitting together pieces from this kit. The pieces are metric spaces which must be glued by almost isometries along parts of their boundaries. With one exception, which is easily

analyzed separately, the pieces have OSFS structures. The main issue is to show that these structures can be glued together to give the structure of a graph orbifold, except in special cases when Q is an orbifold bundle. The geometry of the collapse provides extra information, called the latitude hypothesis, on the relation between the fibres of the different pieces. This plays a key role in the argument.

With the exception of (folded) thick turnovers, every standard model admits at least one OSFS structure. However, it is easy to see that if there are any (folded) thick turnovers then Q is a bundle.

The idea is that (folded) thick turnovers are the only standard neighbourhoods with “triangular” shaped boundaries. The standard neighbourhoods are almost isometric to Euclidean models, and are glued together by isometries. Thus the corresponding Euclidean model neighbourhoods have almost isometric boundaries. Hence the only standard neighbourhood that can be connected to a (folded) thick turnover is another such. This implies that Q is a union of a finite sequence of (folded) thick turnovers, arranged in either a linear fashion (giving a bundle over an interval with generic fibre a turnover) or a circular fashion (giving a bundle over a circle). This is essentially the same argument that provides the structure of the ends of a cone-manifold whose angles are strictly less than π 7.11 (ends at $t < 1$).

Having dealt with (folded) thick turnovers, we may assume that every standard neighbourhood admits at least one OSFS structure. One of these neighbourhood types, folded thick tori, is easily analyzed. The boundary of a folded thick torus, V , is a torus T which is incompressible in V . If T is incompressible in Q , it must be boundary parallel and Q equals V , which is an OSFS. If T compresses in Q , then $Q - V$ is a (singular) solid torus. Every every folded thick torus admits two OSFS structures, given by the two eigenvectors of the involution that does the folding. By 2.46, a fibration on the boundary of a (singular) solid torus extends over the interior unless the fibre is isotopic to a meridian. This can occur for at most one of the two fibrations of V , so Q is an OSFS.

Thus we are reduced to the case that the only standard neighbourhoods in Q are (folded) thick pillowcases, (singular) folded balls, (singular) solid tori and thick tori. This corresponds to moving down the flowchart at the end of this section past the first two boxes.

The basic strategy is to attempt to fit the OSFS structures on the standard neighbourhoods together to give a single OSFS structure on Q when it is orbifold-atoroidal or, more generally, when Q arises as a double, to give an OSFS structure to each component after Q is cut along incompressible

Euclidean 2-dimensional sub-orbifolds.

The boundaries of these standard neighbourhoods consist of tori or pillowcases. By 2.46 they both admit countably many orbifold Seifert fibrations, parametrized by the slope of a regular fibre. The main tool used to extend an OSFS defined on the boundary over the interior of a (singular) solid torus or (singular) folded ball is the lemma 2.46 which states this can be done unless a regular fibre is a meridian.

In a general topological setting, it is quite possible to have a manifold or orbifold that is the union of two SFS glued along their boundaries which is not a SFS or a graph manifold. For example the exterior, X , of the trefoil knot is a SFS. Glue a solid torus, which also is a SFS, to X along their boundaries so that a meridian curve of the solid torus is glued to a regular fibre in the boundary of X . This is ± 6 Dehn filling depending on whether the trefoil is left- or right-handed. The resulting closed manifold is $L(2, 1) \# L(3, 1)$. It is not difficult to show that it is not a graph manifold.

In our situation there is an additional piece of information coming from the geometry of an almost collapsed orbifold called the **latitude hypothesis**. A *latitude* is the isotopy class of any shortest closed geodesic on a pillowcase or torus. There may be up to three such isotopy classes, though generically there is exactly one. When we attempt to construct an OSFS structure on Q at each stage we will have a finite number of suborbifolds each of which has been given an OSFS structure such that a regular fibre is isotopic to a latitude of the boundary pillowcase (or torus) of some (singular) folded ball (or solid torus) standard neighbourhood. We discuss the pillowcase case here; the torus case is similar. The latitude hypothesis is the statement that if α is a latitude of a standard neighbourhood which is a (singular) folded ball then α is not homotopic in $Q - \Sigma$ to either a point or to a meridian of Σ . Thus all regular fibres appearing in our construction satisfy this condition. This prevents the phenomenon of killing the homotopy class a regular fibre when gluing together two OSFS.

The latitude hypothesis is proved by estimating the holonomy of a latitude. A standard neighbourhood in Q which is a (singular) folded ball is almost isometric to a compact Euclidean cone-manifold. The first step is to show that the Euclidean holonomy of a latitude is the composition of two rotations through π around almost parallel axes. This uses the fact that the diameter of a (singular) folded ball standard neighbourhood is very large compared to the diameter of the soul, i.e. the distance between the two axes of rotation. Roughly speaking, by taking a large finite orbifold cover of the (singular) folded ball, one sees there are two almost parallel rotation axes

in the cover which are not too far apart. Thus the hyperbolic holonomy of a latitude is almost the composition of two rotations through angles almost equal to π around almost parallel distinct axes. The hyperbolic holonomy is non-trivial; hence, the latitude is essential in $Q - \Sigma$. This estimate also shows that the latitude has very small complex translation length. Its holonomy can't be close to a rotation through an angle $2\pi/n$ around some edge of Σ because such an elliptic does not have "very small" rotation angle. Thus a latitude is not homotopic in $Q - \Sigma$ to a meridian of Σ .

The union of standard neighbourhoods meeting Σ is a suborbifold bounded by tori. We denote by \mathcal{N} the union of those components of this orbifold that are not singular solid tori. A foliation argument is used to show that:

Theorem 7.17. *At least one component of the union of the standard neighbourhoods meeting Σ is not a singular solid torus. Hence \mathcal{N} is non-empty.*

This is important because the boundaries of the components of \mathcal{N} are incompressible in \mathcal{N} so they must either be boundary parallel or compressible in Q . A compressible torus must bound a (singular) solid torus since Q is irreducible. It follows that if each component of \mathcal{N} is an OSFS then (using the latitude hypothesis) this OSFS structure can be extended over all of Q . Similarly, if each component of \mathcal{N} is a graph orbifold, then so is Q .


The standard neighbourhoods in \mathcal{N} are of three types: thick pillowcases, (singular) folded balls, and folded thick pillowcases. All three have pillowcases as boundary components. By topological methods, it is possible to combine any neighbourhoods of the same type that intersect. Either it can be arranged that any two neighbourhoods of the same type are disjoint, or a component of \mathcal{N} fibres over a 1-dimensional orbifold with Euclidean fibre, a case that can be easily handled separately. Furthermore, the same argument used previously to deal with folded thick tori shows that either all the folded thick pillowcases are incompressible in Q , or Q is a folded thick pillowcase union a (singular) folded ball. As with the folded thick torus, the latter two cases are readily seen to be OSFS. When a folded thick pillowcase, V , is incompressible, we cut along a pillowcase which is boundary parallel in V . This creates a component which is again homeomorphic to V and is used in the decomposition as a graph orbifold; the remaining component contains a new thick pillowcase with one component on the boundary. We continue to denote the latter piece by \mathcal{N} ; it suffices to show that its components are graph orbifolds.

A component, X , of \mathcal{N} is a union $X = A \cup B$ where A is a disjoint union of thick pillowcases and B is a disjoint union of (singular) folded

balls. Furthermore, it is possible to arrange that each (singular) folded ball component of B intersects A in exactly two components each of which is a $D^2(2, 2)$. The combinatorics of this arrangement may be complicated, so, to simplify the picture, we cut along a (possibly compressible) pillowcase inside each thick pillowcase which does not already have one boundary component on the boundary of X . In the resulting pieces exactly one boundary component of each thick pillowcase is on the boundary; the other boundary component is connected to (singular) folded balls along $D^2(2, 2)$'s.

Specifically, each piece is of the form $A \cup B$ where $B = \bigcup_i B_{2i}$ is a disjoint union of (singular) folded balls and $A = \bigcup_i A_{2i-1}$ is a disjoint union of thick pillowcases. Furthermore B_{2i} intersects only $A_{2i \pm 1}$ and each component of the intersection is isomorphic to $D^2(2, 2)$. Hence each component of X is the union of a finite number of thick pillowcases and (singular) folded balls arranged alternately in a circular way like hollow beads on a string. Each hollow bead is a thick pillowcase. Each piece of string between two beads is a (singular) folded ball. We call this cutting up process "disassembling" \mathcal{N} .

"Disassembly" picture



We will now show that each piece X of the disassembled orbifold admits a OSFS structure. Choose a (singular) folded ball B_1 in X . Then $\partial B_1 \cap \partial X$ is an annulus, as shown in the figure. A core curve of this annulus is a latitude of ∂B_1 . Since this latitude does not compress in B_1 there is a OSFS structure on B_1 with this latitude a regular fibre. Now $C = A_2 \cap B_1 \cong D^2(2, 2)$. The OSFS structure on B_1 may be isotoped to give an orbifold fibration of C with ∂C one of the fibres. This fibration extends productwise over A_2 to an OSFS structure on A_2 . This may be isotoped and then extended over the next (singular) folded ball B_3 attached to A_2 . In this way we can extend


the OSFS structure around the string of beads which makes up X .

We now use these OSFS structures on the pieces to define a graph orbifold structure on Q . The boundary components of each piece has an induced fibration. One boundary component of each piece of the disassembled orbifold is a torus. It is either boundary parallel or bounds a (singular) solid torus in Q . The induced fibration extends over the (singular) solid tori by the latitude hypothesis; fill in these (singular) solid tori with the extended fibration. We now glue back along the pillowcases used to disassemble \mathcal{N} . If such a pillowcase, P , is incompressible, then we do not try to match the OSFS structures on the two sides of P , since P will be part of a set of incompressible pillowcases decomposing Q in OSFS pieces.

If a pillowcase, P , becomes compressible when two pieces are glued together, it must bound a (singular) folded ball in Q (see 2.47). (That it can't be contained in a (singular) ball follows from the Seifert fibres on the pieces.) By the latitude hypothesis, the fibration on P extends over the (singular) folded ball. This may change the fibration that was already on the (singular) folded ball. Continuing in this manner the fibrations are matched up, piece by piece along all of the compressible pillowcases. Thus, every component of \mathcal{N} , cut along incompressible pillowcases, is an OSFS. Hence \mathcal{N} is a graph manifold as desired.

It is useful to note that when Q is the union of two (folded) thick pillowcases glued along their mutual boundary, the process of cutting along incompressible pillowcases decomposes Q into two (folded) thick pillowcases. This is a degenerate version of a graph orbifold where the pieces are each (possibly the quotient by an involution of) a product. In this case Q is a bundle with generic fibre a pillowcase and base a circle or interval. This happens if Q is a Solv orbifold.

The following flowchart summarizes the overall structure of this proof.



Background on the Orbifold Theorem

Many people have obtained partial results and developed related ideas. The following is a selective list. Some of this work is used in our approach, and other parts are used in the approach of Boileau and Porti. We have included some hearsay concerning the events surrounding the Orbifold theorem.

1978 The Smith Conjecture is proved [64].

This was a culmination of the work of many people and used a major part of the theory of 3-manifolds, in particular the work of Bass, Culler & Shalen, Gordon & Litherland, Meeks & Yau, Haken, Waldhausen, and Thurston. It is now a (very special) consequence of the Orbifold theorem.

1981 Thurston announces the Orbifold Theorem [81], [83].

Theorem A. [83] *Let M^3 be a prime, \mathbb{P}^2 -irreducible, compact 3-manifold which admits a diffeomorphism ϕ ($\neq 1$) of finite order whose fixed point set is more than a finite set of points. Then M has a geometric decomposition.*

Theorem B. [83] *Suppose F is a finite group of diffeomorphisms of a compact 3-manifold M , of which some element $\phi \in F$ ($\phi \neq 1$) has more than a finite set of fixed points. Let $O = M/F$ and $\Sigma \subset O$ be the image of the union of fixed point sets of all elements $\phi \neq 1$ of F . Suppose that the 3-manifold $O - \Sigma$ is prime, and that any 2-sided projective plane in $O - \Sigma$ is homotopic to the boundary of a regular neighbourhood of an isolated point of Σ . Then M has a geometric decomposition which is invariant by F . More precisely, there is a collection of disjoint embedded spheres, projective planes, incompressible tori and incompressible Klein bottles, whose union is invariant by F and geometric structures on the pieces obtained by decomposition along the surfaces on which F acts isometrically.*

Interestingly, Thurston's original theorem pre-dated Hamilton's announce-

ment of the results in [35] by a couple of months. Thurston's first version concluded that either an orbifold has a geometric decomposition or else it admits a metric of positive Ricci curvature. Two months later, Thurston heard of Hamilton's result and was thus able to complete his proof of the geometrization theorem for 3-orbifolds with one-dimensional singular locus. Thurston outlined his proof on two occasions in courses at Princeton; in 1982 and again in 1984. On both occasions, due to running out of time, the outline was incomplete in certain aspects at the end of the proof in the collapsing case. In particular the Euclidean/spherical transition in the case of vertices was treated in a few sentences.

1982 Tollefson [79] showed that two involutions of a Haken 3-manifold that are homotopic are in fact conjugate by a diffeomorphism isotopic to the identity provided that the manifold is not a Seifert fibre space and $H_1(M)$ is infinite.

1982 Hamilton [35] classifies 3-manifolds with positive Ricci curvature. At present all known proofs of the orbifold theorem make use of either this result, or else the strengthened 1986 version. It remains an interesting and important question whether there is a proof that does not rely on PDE techniques.

1983,1985 Bonahon & Siebenmann [9],[10] classify orbifold Seifert fibre spaces (OSFS).

1984 Hamilton [37] distributes a preprint giving an orbifold version of his positive Ricci curvature theorem. This version had been suggested by Thurston as a way of completing his proof of the orbifold theorem. Thurston later claimed a proof which avoided appealing to this result. Hamilton's preprint has not yet appeared.

1985 Soma, Ohshika & Kojima [76] give some details of the proof of the Orbifold theorem. In particular they give a somewhat different proof of the classification of non-compact Euclidean 3-dimensional cone-manifolds from the one outlined by Thurston.

1986 Hamilton [36] classifies 3-manifolds with non-negative Ricci curvature.

1986 [43] Hodgson's thesis gives many examples and develops the theory of deformations and change of geometry.

1986 Meeks & Scott [61] show that if a finite group acts on a closed P^2 -

irreducible Seifert fibre space M with infinite fundamental group then there is a homogeneous metric (i.e. geometric structure) preserved by this action. If the action preserves a Seifert fibration up to homotopy then M has an invariant Seifert fibration.

1986 Cheeger & Gromov [17],[18], Fukaya [29] study collapse with bounded curvature and introduce F-structures. The collapsing theorem they prove works in all dimensions. The statement is very similar to that of the collapsing theorem for cone-manifolds used in the proof of the orbifold theorem. However the fact that the local fundamental group of a cone-manifold is not virtually abelian means that different techniques must be used.

1987 Bonahon & Siebenmann [11] develop a JSJ decomposition for orbifolds. This is a characteristic splitting of a 3-orbifold by incompressible Euclidean 2-suborbifolds.

1987 Hodgson [44] gives background and outlines Thurston's proof of the Orbifold theorem. The treatment of the collapsing case was somewhat incomplete, reflecting Thurston's presentation.

1987 McCullough & Miller [59] show that a 3-orbifold with a geometric decomposition has a finite orbifold covering which is a manifold. They deduce that such orbifold fundamental groups are residually finite. They also deduce that an isomorphism between *sufficiently large* 3-orbifold fundamental groups preserving the peripheral structure is induced by an orbifold isomorphism provided the boundary consists of incompressible Euclidean orbifolds.

1988 Dunbar [26], [28] classifies non-hyperbolic geometric 3-orbifolds with underlying space S^3 . In [27] he develops hierarchies for 3-orbifolds.

1990 Zhou's thesis [90] gives some details of the proof of the Orbifold theorem, for the case where the singular locus is a 1-manifold.

1992-1998 Hodgson & Kerckhoff [45], [49], [46] develop a rigidity theory for cone-manifolds, using harmonic deformations.

1995 Kirby [50] lists the Orbifold theorem as a conjecture (Problem 3.46) in his problem list.

1998 Kojima [53] establishes global rigidity for hyperbolic 3-cone-manifolds with cone angles at most π .

1998 Boileau & Porti [8] distribute a preprint with a proof for the case of

an orbifold of the form $Q = M/G$ where M is an irreducible 3-manifold and G a finite group, and the singular locus of Q is a 1-manifold.

They handle the collapsing case by showing that there is a geometric structure on M and using Meeks and Scott to deduce there is a G -equivariant structure, which therefore descends to a geometric structure on Q . They find a curve $\gamma \subset M$ (which, for example, may be a regular fibre of a Seifert fibre space) and show that the Haken manifold $N = M - \mathcal{N}(\gamma)$ is a graph manifold by showing it has Gromov norm zero. This implies there are no hyperbolic pieces. Then by Thurston's theorem for Haken manifolds, N is a graph manifold. The geometry of the collapse is used to give γ and to construct an open cover of N which is used to show the Gromov norm of N (hence of M) is zero.

1998 Cooper, Hodgson & Kerckhoff announce a proof in the case the singular locus is a 1-manifold and the boundary is Euclidean. They outline the proof in a series of 15 lectures at a meeting of the MSJ in Tokyo and distribute a preprint.

The heart of the argument is the **collapsing case** when the injectivity radius goes to zero everywhere. There are two subcases. If one rescales the metrics by multiplying by $\max(1, 1/\text{diameter})$ either the injectivity radius goes to zero somewhere, or else it does not. In the latter case one obtains a Euclidean cone-manifold in the limit. This may be the final orbifold structure. Otherwise one uses Hamilton's theorem to show the final orbifold is spherical. The remaining case is that, even after the rescaling, the injectivity radius goes to zero everywhere. Then every point in one of these almost-collapsed manifolds has a neighbourhood of simple topology. This case is handled by the collapsing theorem 7.13 (Collapsing theorem) where we construct an orbifold Seifert fibration, or orbifold bundle structure on these pieces and fit them all together.

The collapsing case was handled somewhat differently by Thurston, who described the geometry of the collapse in more detail, (in the spirit of the F-structures subsequently developed by Gromov and Cheeger in their proof of a collapsing theorem for manifolds of bounded curvature) and produced the Seifert fibration or Solv structure directly from this geometry. Some of the main ideas are already present in chapter 4 of the original notes of Thurston, [84]. We hope to recover this description from our approach in a future paper.

Bibliography

- [1] I. Agol, *Topology of hyperbolic 3-manifolds*, PhD thesis, UC San Diego, 1998. (Available from <http://math.ucdavis.edu/~iagol>).
- [2] J. W. Alexander, *Note on Riemann spaces*, Bull. Amer. Math. Soc. **26** (1920), 370–372.
- [3] E. M. Andreev, *On convex polyhedra in Lobachevskii space*, Math. USSR, Sbornik, 10 (1970), 413–440.
- [4] M. A. Armstrong, *Basic Topology*, Springer, 1983.
- [5] M. Berger, *Geometry*, Springer, 1987.
- [6] S. Bleiler and C. Hodgson, *Spherical space forms and Dehn filling*, Topology **35**(1996), 809–833.
- [7] R. H. Bing, *Inequivalent families of periodic homeomorphisms of E^3* , Ann. of Math. (2) **80** (1964), 78–93.
- [8] M. Boileau and J. Porti, *Geometrization of 3-orbifolds of cyclic type*, preprint. (Available from <http://xxx.lanl.gov/abs/math/9805073>).
- [9] F. Bonahon and L. Siebenmann, *The classification of Seifert fibred 3-orbifolds*, London Math. Soc. lecture notes **95** (1985), 19–85.
- [10] F. Bonahon and L. Siebenmann, *Seifert fibred orbifolds and their role as natural crystalline parts of arbitrary compact irreducible 3-orbifolds*, preprint, Orsay (1983).
- [11] F. Bonahon and L. Siebenmann, *The characteristic toric splitting of irreducible compact 3-orbifolds*, Math. Ann. **278** (1987), 441–479.
- [12] M. Bridson and A. Haefliger, *Metric Spaces Of Non-positive Curvature*, Springer, 1999.

- [13] M. R. Bridson and G. A. Swarup, *On Hausdorff-Gromov convergence and a theorem of Paulin*, L'Enseign. Math. (2) **40** (1994), no. 3-4, 267–289.
- [14] R. D. Canary, D. B. A. Epstein and P. Green *Notes on notes of Thurston*, Analytical and Geometric Aspects of Hyperbolic Space, London Math. Soc. Lecture Notes **111** (1987), 3–92.
- [15] A. Casson and D. Jungreis, *Convergence groups and Seifert-fibred 3-manifolds*, Inventiones Math. **118** (1994), 441–456.
- [16] J. Cheeger and D. Gromoll, *On the structure of complete manifolds of non-negative curvature*, Ann. of Math **96** (1972), 413–433.
- [17] J. Cheeger and M. Gromov, *Collapsing Riemannian manifolds while keeping their curvature bounded. I*, J. Differential Geom. **23** (1986), 309–346.
- [18] J. Cheeger and M. Gromov, *Collapsing Riemannian manifolds while keeping their curvature bounded. II*, J. Differential Geom. **32** (1990), 269–298.
- [19] D. Cooper, C. Hodgson and S. Kerckhoff, *Geometrization of 3-dimensional orbifolds*, in preparation.
- [20] D. Cooper and D.D. Long, *Some Surface Subgroups Survive Surgery*, preprint, 1997.
- [21] H. S. M. Coxeter, *Non-Euclidean Geometry*, 3rd edition, Toronto, 1957.
- [22] M. Culler and P. B. Shalen, *Varieties of group representations and splittings of 3-manifolds*, Ann. of Math. **117** (1983), 109–146.
- [23] M. Culler and P. B. Shalen, *Bounded, separating, incompressible surfaces in knot manifolds*, Invent. Math **75** (1984), 537–545.
- [24] M. Culler, *Lifting Representations to covering groups*, Adv. Math **59** (1986), 64–70.
- [25] J. Dowty, PhD thesis, Univ. of Melbourne, in preparation.
- [26] W. D. Dunbar, *Geometric orbifolds*, Rev. Mat. Univ. Complut. Madrid **1** (1988), no. 1-3, 67–99.

- [27] W. D. Dunbar, *Hierarchies for 3-orbifolds*, Topology Appl. **29** (1988), 267–283.
- [28] W. D. Dunbar, *Nonfibering Spherical 3-Orbifolds*, Trans Amer. Math. Soc. **341** (1994), 121–142.
- [29] K. Fukaya, *Collapsing Riemannian manifolds to lower dimensional ones*, J. Differential Geom. **25** (1987), no. 1, 139–156.
- [30] D. Gabai, *Convergence groups are Fuchsian groups*, Annals of Math. **136** (1992), 447–510.
- [31] W. M. Goldman, *Geometric Structures on manifolds and varieties of representations*, Contemp. Math. **74**, 169–198, Amer. Math. Soc., Providence, RI, 1988.
- [32] M. Gromov, *Structures métriques pour les variétés riemanniennes*. (Edited by J. Lafontaine and P. Pansu.) Textes Mathématiques 1. CEDIC, Paris, 1981. (Expanded English version: M. Gromov and S. M. Bates, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Birkhäuser, 1999.)
- [33] M. Gromov and W. Thurston, *Pinching constants for hyperbolic manifolds*, Invent. Math., **89** (1987), 1–12.
- [34] W. Haken, *Some results on surfaces in 3-manifolds*, Studies in Modern Topology, pp. 39–98, Math. Assoc. Amer. 1968.
- [35] R. S. Hamilton, *Three-manifolds with positive Ricci curvature*, J. Diff. Geom. **17** (1982), 255–306.
- [36] R. S. Hamilton, *Three-Orbifolds with Positive Ricci Curvature*, preprint, 1984.
- [37] R. S. Hamilton, *Four-Manifolds with Positive Curvature Operator*, J. Diff. Geom. **24** (1986), 153–176.
- [38] A. E. Hatcher, *On the boundary curves of incompressible surfaces*, Pacific J. Math. **99** (1982), no. 2, 373–377.
- [39] J. Hempel, *Residual finiteness for 3-manifolds*, in: S. Gersten and J. Stallings, Eds. Combinatorial Group Theory and Topology, Annals of Math. Studies No. **111** (Princeton University Press, Princeton, 1987), 379–396.

- [40] H. Hilden, M. Lozano, J. Montesinos-Amilibia, *The arithmeticity of the figure eight knot orbifolds*, in: Topology '90 (Columbus, OH, 1990), 169–183, Ohio State Univ. Math. Res. Inst. Publ., 1, de Gruyter, Berlin, 1992.
- [41] H. Hilden, M. Lozano, J. Montesinos-Amilibia, *On a remarkable polyhedron geometrizing the figure eight knot cone manifolds*, J. Math. Sci. Univ. Tokyo **2** (1995), no. 3, 501–561.
- [42] H. Hilden, M. Lozano, J. Montesinos, and W. Whitten, *On universal groups and three-manifolds*, J. Algebra **20** (1972), 250–270.
- [43] C. D. Hodgson, *Degeneration and regeneration of geometric structures on 3-manifolds*, Ph. D. thesis, Princeton Univ., 1986.
- [44] C. D. Hodgson, *Geometric structures on orbifolds: Notes on Thurston's proof*, Preliminary version, 1988.
- [45] C. D. Hodgson and S. P. Kerckhoff, *Rigidity of hyperbolic cone-manifolds and hyperbolic Dehn surgery*, J. Differential Geom. **48**(1998), 1–59.
- [46] C. D. Hodgson and S. P. Kerckhoff, *The shape of hyperbolic Dehn surgery space*, in preparation.
- [47] C.D. Hodgson and J. Tysk, *Eigenvalue estimates and isoperimetric inequalities for cone-manifolds*, Bull. Australian Math. Soc. **47** (1993), 127–144.
- [48] M. Jankins and W. Neumann, *Lectures on Seifert manifolds*, Brandeis Lecture Notes, 1983.
- [49] S. P. Kerckhoff, *Deformations of hyperbolic cone manifolds*, Topology and Teichmüller spaces (Katinkulta, 1995), 101–114, World Sci. Publishing, River Edge, NJ, 1996.
- [50] R. Kirby, *Problems in low-dimensional topology*, pp 35–473, in: AMS/IP Studies in Advanced Mathematics, volume 2, part 2, Amer. Math. Soc., 1997.
- [51] H. Kneser, *Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten*, Jb. Deutsch Math. Ver. **38** (1929), 248–260.

- [52] S. Kojima, *Non-singular parts of hyperbolic 3-cone-manifolds*, preprint.
- [53] S. Kojima, *Deformations of Hyperbolic 3-cone-manifolds*, preprint.
- [54] S. Kojima, *Hyperbolic 3-manifolds singular along knots*, Knot theory and its applications. Chaos Solitons Fractals **9** (1998), no. 4-5, 765–777.
- [55] M. Lackenby, *Word hyperbolic Dehn surgery*, Invent. math. **140** (2000) 2, 243-282.
- [56] W. B. R. Lickorish, *A representation of orientable combinatorial 3-manifolds*, Ann. of Math. **76** (1962), 531–538.
- [57] L. Lok, *Deformations of locally homogeneous spaces and Kleinian groups*, PhD thesis, Columbia Univ., 1984.
- [58] W. S. Massey, *Algebraic Topology: An Introduction*, Springer-Verlag, 1967.
- [59] D. McCullough and A. Miller, *Manifold covers of 3-orbifolds with geometric pieces*, Topology Appl. **31** (1989), no. 2, 169–185.
- [60] R. C. McOwen, *Prescribed curvature and singularities of conformal metrics on Riemann surfaces*, J. Math. Anal. Appl. **177** (1993), no. 1, 287–298.
- [61] W. H. Meeks and G. P. Scott *Finite group actions on 3-manifolds*, Invent. Math. **86** (1986), 287–346.
- [62] J. Milnor, *A unique decomposition theorem for 3-manifolds*, Amer. J. Math. **84** (1962), 1–7.
- [63] J. Milnor, *On the Schäfli differential equality*, in *Collected papers*, volume 1, Publish or Perish, 1994.
- [64] J. Morgan and H. Bass (editors), *The Smith Conjecture*, Academic Press, 1984.
- [65] R. Myers, *Simple knots in compact, orientable 3-manifolds*, Trans. Amer. Math. Soc. **273** (1982), no. 1, 75–91.

- [66] P. Pansu *Effondrement des variétés riemanniennes, d'après J. Cheeger et M. Gromov*. Seminar Bourbaki, Vol. 1983/84. Astérisque No. **121-122** (1985), 63–82.
- [67] J. Porti, *Regenerating hyperbolic and spherical cone structures from Euclidean ones*, *Topology* **37** (1998), no. 2, 365–392.
- [68] J. G. Ratcliffe, *Foundations of Hyperbolic Manifolds*, Springer-Verlag, 1994
- [69] D. Rolfsen, *Knots and Links*, Publish or Perish, 1976.
- [70] L. Schläfli, *On the multiple integral $\int^n dx dy \dots dz$, whose limits are $p_1 = a_1x + b_1y + \dots + h_1z > 0$, $p_2 > 0$, \dots , $p_n > 0$, and $x^2 + y^2 + \dots + z^2 < 1$* , *Quart. J. Math.* **2**(1858), 269–301.
- [71] H. Schubert, *Die eindeutige Zerlegbarkeit eines Knotens in Primknoten*, S.-B. Heidelberger Akad. Wiss. Math.-Nat. Kl. **3** (1949), 57–104.
- [72] H. Schubert, *Knoten mit zwei Brücken*, *Math. Z.* **65** (1956), 133–170.
- [73] P. Scott, *The geometries of 3-manifolds*, *Bull. London Math. Soc.* **15** (1983), 401–487.
- [74] H. Seifert, *Topologie dreidimensionaler gefaserner Räume*, *Acta Math.* **60** (1933) 147–238.
- [75] H. Seifert and W. Threlfall, *A Textbook of Topology* (translation of *Lehrbuch der Topologie*, 1934), Academic Press, 1980.
- [76] T. Soma, K. Ohshika and S. Kojima, *Towards a proof of Thurston's geometrization theorem for orbifolds*. Hyperbolic geometry and 3-manifolds (Japanese) (Kyoto, 1985). Sūrikaisekikenkyūsho Kōkyūroku No. **568** (1985), 1–72.
- [77] K. Strebel, *Quadratic Differentials*, Springer, 1984.
- [78] E. Suárez Peiró, *Dirichlet polyhedra of 3-cone-manifolds and their deformations (Spanish)*, PhD thesis, Universidad Complutense de Madrid, 1998.
- [79] J. L. Tollefson *Involutions of Sufficiently Large 3-Manifolds*, *Topology* **20** (1981), 323–353.

- [80] W. P. Thurston, *The Geometry and Topology of three-manifolds*, Princeton Univ. Math. Dept. (1978). These are available at <http://msri.org/publications/books/gt3m/>.
- [81] W. P. Thurston, *Three dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. **6** (1982), 357–381.
- [82] W. P. Thurston, *Three-dimensional Geometry and Topology, Volume 1*, Princeton Univ. Press, 1997.
- [83] W. P. Thurston, *Three-manifolds with symmetry*, preprint, Princeton University, 1982, 5 pages.
- [84] W. P. Thurston, *Three-dimensional Geometry and Topology, chapter 5*, Princeton Univ. Math. Dept., (1984). (See chapter 13 of [80] for an earlier version.).
- [85] M. Troyanov, *Prescribing curvature on compact surfaces with conical singularities*, Trans. Amer. Math. Soc. **324** (1991), no. 2, 793–821.
- [86] E. B. Vinberg, editor, *Geometry II, Geometry of Spaces of constant curvature*, volume 29 of *Encyclopaedia of Mathematical Sciences*, Springer, 1993.
- [87] A. D. Wallace, *Modifications and cobounding manifolds*, *Canad. J. Math.* **12** (1960) 503–528.
- [88] J. Weeks, *SnapPea, the computer program*, available from <http://thames.northnet.org/weeks/index/SnapPea.html>.
- [89] A. Weil, *On discrete subgroups of Lie groups*, *Ann. of Math.* **72**, (1960) 369–384.
- [90] Q. Zhou, *3-dimensional Geometric Cone Structures*, PhD thesis, UCLA, 1989.

Index

- $DO_{\{v_i\}}$, 121
- $F(n_1, n_2, \dots, n_k)$, 25
- $I(2, 2)$, 27
- $N(A, r; X)$, 107
- $S^2(n, n), S^2(m, n, p)$, 25
- $S^2(\alpha, \beta, \gamma)$, 54
- $\mathbb{H}^n(K)$, 60
- ∂_{orb} , 23
- \mathcal{N} , 42, 148

- algebraic topology, 118
- analytic continuation, 9
- analytic geometry, 8
- apple turnover, 32
- atoroidal, 46

- bad orbifold, 28
- base orbifold, 40
- bi-Lipschitz, 107
- Bieberbach-Soul Theorem, 142
- bilipschitz, 107
- Boileau, 155
- Bonahon, 155
- Borromean rings, 35
- boundary of orbifold, 23
- branched cover, 18
- Busemann function, 144

- characteristic submanifold, 48
- Cheeger, 155
- collapsing case, 135
- Collapsing Theorem, 135
- complete, 7

- complex length, 95
- cone, 24
- cone neighbourhood, 60
- cone-manifold, 53
- cone: $\text{Cone}_K(S; R)$, 59
- conjugate point, 65
- Conway sphere, 19
- corner, 24
- cuspidal, 94
- cut locus, 65

- $\text{Def}(M)$, 94
- deformations count, 98
- Dehn filling, 13, 100
- Dehn Surgery, 13, 17
- developing map, 9, 75
- dihedral edge, 140
- dihedral vertex, 140
- Dirichlet domain, 66
- disassembly, 149
- doubling trick, 122
- Dunbar, 155

- ends, 94
- ends: cone manifold, 131
- ends: orbifold, 130
- ϵ -approximation, 108
- ϵ -collapse, 135
- ϵ -count, 113
- equivalence of (G, X) structures, 92
- equivariant, 21

- essential, 46
- Euclidean cone-manifold, 53
- Euclidean Decomposition Theorem, 120
- Euclidean orbifold, 30
- Euclidean turnover, 79
- Euler number, 46
- exceptional fibre, 39
- exceptional surgery, 101
- exponential map, 64

- figure eight knot, 35
- finite volume hyperbolic, 14
- flat connection, 93
- flowchart of collapsing case, 150
- folded ball, 40
- folded thick X , 45
- folded thick ..., 144
- football, 25, 54
- Fukaya, 155

- G, X orbifold, 30
- G, X structure, 9
- Gauss-Bonnet, 6, 31, 71
- generic fibre, 37
- geodesic, 61
- geometric decomposition, 5
- geometric structure, 5, 8
- geometric structure on an orbifold, 50
- Geometrization Conjecture, 5
- global cone, 66
- good orbifold, 28
- graph manifold, 49
- graph orbifold, 49
- Gromov, 155
- Gromov-Hausdorff distance, 107
- Gromov-Hausdorff topology, 110

- Haken manifold, 11
- Hamilton's Theorem, 138

- Hausdorff distance, 107
- Heegaard splitting, 15, 16
- holonomy, 10, 75
- homogeneous, 7
- Hopf-Rinow, 62
- horizontal singular locus, 42
- hyper-elliptic, 16
- hyperbolic cone-manifold, 53
- Hyperbolic Dehn Surgery Theorem, 14
- hyperbolic orbifold, 30
- hyperbolic structure, 6
- Hyperbolization Conjecture, 13

- incomplete structures, 10
- incompressible, 46
- infinite pillowcase, 82
- injectivity radius, 115
- irreducible, 46
- irreducible representation, 98
- isotopy, 47
- isotropy group, 22

- JSJ decomposition, 48

- Kirby Problem List, 155
- Kojima, 155

- latitude hypothesis, 147
- lens space, 16
- Lipschitz map, 107
- local group, 22
- Local Margulis Theorem, 144
- local model (of orbifold), 21
- local-orientation double cover, 28
- locally orientable, 22

- McCullough, 155
- Meeks, 154
- meridian, 44
- Miller, 155

- mirror points, 23
- Montesinos knot, 20
- non-compact Euclidean, 142
- Ohshika, 154
- open cone, 59
- orbifold atoroidal, 46
- orbifold boundary, 23
- orbifold bundle, 37
- orbifold covering, 27
- orbifold disc, 46
- orbifold Euler characteristic, 29
- orbifold fundamental group, 28
- orbifold geomtric structure, 50
- orbifold incompressible, 46
- orbifold isotopy, 47
- orbifold map, 21
- Orbifold Seifert fibred Space (OSFS), 39
- Orbifold theorem: proof, 119
- Orbifold theorem: statement, 50
- orbifold: definition of, 21
- orientable orbifold, 22
- Orthogonalization Conjecture, 13
- path metric, 61
- pillowcase, 22
- pillowcase, infinite, 82
- Poincaré Conjecture, 13
- Poincaré's polyhedron theorem, 33
- Porti, 155
- prime decomposition, 5
- prime knot, 19
- proper metric space, 112
- rational knot, 17
- reflector points, 23
- regular set, 28
- rescale the metric, 156
- satellite knot, 12
- Schläfli formula: cone-manifolds, 72
- Schläfli formula: polyhedra, 71
- Scott, 154
- segment, 64
- Seifert fibre space (SFS), 39
- Seifert fibred orbifold, 39
- Siebenmann, 155
- silvered points, 23
- singular fibred folded ball, 40
- singular fibred solid torus, 40
- singular folded ball, 40
- singular locus, 22
- smear, 112
- Smith Conjecture, 15
- Soma, 154
- soul, 144
- spherical cone-manifold, 53
- spherical orbifold, 30
- spherical structure, 6
- spindle, 25
- standard cone neighbourhood, 59, 60
- suborbifold, 46
- suspension: $\text{Susp}_K(S)$, 59
- Symmetry Theorem, 15
- tangent cone, 61
- teardrop, 25
- thick X , 45
- thick pillowcase, 144
- thick torus, 144
- thick turnover, 144
- thin part, 141
- Thurston's Theorems A and B, 153
- Tollefson, 154
- torus decomposition, 5
- torus knot, 11
- turnover, 32, 54

two bridge, 17

underlying space, 22

uniformly totally bounded, 113

universal orbifold cover, 28

vertical singular locus, 42

very good, 28

Voronoi region, 68

Zhou, 155