Annals of Mathematics, **162** (2005), 195–290

Geometrization of 3-dimensional orbifolds

By Michel Boileau, Bernhard Leeb, and Joan Porti

Abstract

This paper is devoted to the proof of the orbifold theorem: If $\mathcal{O}$ is a compact connected orientable irreducible and topologically atoroidal 3-orbifold with nonempty ramification locus, then $\mathcal{O}$ is geometric (i.e. has a metric of constant curvature or is Seifert fibred). As a corollary, any smooth orientation-preserving nonfree finite group action on S^3 is conjugate to an orthogonal action.

Contents

1. Introduction

A 3-dimensional orbifold is a metrizable space equipped with an atlas of compatible local models given by quotients of $\mathbb{R}^3$ by finite subgroups of $O(3)$. For example, the quotient of a 3-manifold by a properly discontinuous smooth group action naturally inherits a structure of a 3-orbifold. When the group action is finite, such an orbifold is called very good. We will consider in this paper only orientable orbifolds. The ramification locus, i.e. the set of points with nontrivial local isotropy group, is then a trivalent graph.

In 1982, Thurston [Thu2, 6] announced the geometrization theorem for 3-orbifolds with nonempty ramification locus and lectured about it. Several partial results have been obtained in the meantime; see [BoP]. The purpose of this article is to give a complete proof of the orbifold theorem; compare [BLP0] for an outline. A different proof was announced in [CHK].

The main result of this article is the following uniformization theorem which implies the orbifold theorem for compact orientable 3-orbifolds. A

3-orbifold $\mathcal{O}$ is said to be *geometric* if either its interior has one of Thurston's eight geometries or $\mathcal{O}$ is the quotient of a ball by a finite orthogonal action.

MAIN THEOREM (Uniformization of small 3-orbifolds). *Let $\mathcal{O}$ be a compact connected orientable small 3-orbifold with nonempty ramification locus. Then $\mathcal{O}$ is geometric.*

An orientable compact 3-orbifold $\mathcal{O}$ is *small* if it is irreducible, its boundary $\partial\mathcal{O}$ is a (perhaps empty) collection of turnovers (i.e. 2-spheres with three branching points), and it does not contain any other closed incompressible orientable 2-suborbifold.

An application of the main theorem concerns nonfree finite group actions on the 3-sphere S^3; see Section 2.3. It recovers all the previously known partial results (cf. [DaM], [Fei], [MB], [Mor]), as well as the results about finite group actions on the 3-ball (cf. [MY2], [KS]).

COROLLARY 1.1. *An orientation-preserving smooth nonfree finite group action on S^3 is smoothly conjugate to an orthogonal action.*

Every compact orientable irreducible and atoroidal 3-orbifold can be canonically split along a maximal (perhaps empty) collection of disjoint and pairwise nonparallel hyperbolic turnovers. The resulting pieces are either *Haken* or *small* 3-orbifolds (cf. Section 2). Using an extension of Thurston's hyperbolization theorem to the case of Haken orbifolds (cf. [BoP, Ch. 8]), we show that the main theorem implies the orientable case of the orbifold theorem:

COROLLARY 1.2 (Orbifold Theorem). *Let $\mathcal{O}$ be a compact connected orientable irreducible 3-orbifold with nonempty ramification locus. If $\mathcal{O}$ is topologically atoroidal, then $\mathcal{O}$ is geometric.*

Any compact connected orientable 3-orbifold, that does not contain any bad 2-suborbifold (i.e. a 2-sphere with one branching point or with two branching points having different branching indices), can be split along a finite collection of disjoint embedded spherical and toric 2-suborbifolds ([BMP, Ch. 3]) into irreducible and atoroidal 3-orbifolds, which are geometric if the branching locus is nonempty, by Corollary 1.2. Such an orbifold is the connected sum of an orbifold having a geometric decomposition with a manifold. The fact that 3-orbifolds with a geometric decomposition are finitely covered by a manifold [McCMi] implies:

COROLLARY 1.3. *Every compact connected orientable 3-orbifold which does not contain any bad 2-suborbifolds is the quotient of a compact orientable 3-manifold by a finite group action.*

The paper is organized as follows. In Section 2 we recall some basic terminology about orbifolds. Then we deduce the orbifold theorem from our main theorem.

The proof of the main theorem is based on some geometric properties of cone manifolds, which are presented in Sections 3–5. This geometric approach is one of the main differences with [BoP].

In Section 3, we define cone manifolds and develop some basic geometric concepts. Motivating examples are geometric orbifolds which arise as quotients of model spaces by properly discontinuous group actions. These have cone angles $\leq \pi$, and only cone manifolds with cone angles $\leq \pi$ will be relevant for the approach to geometrizing orbifolds pursued in this paper. The main result of Section 3 is a compactness result for spaces of cone manifolds with cone angles $\leq \pi$ which are thick in a certain sense.

In Section 4 we classify noncompact Euclidean cone 3-manifolds with cone angles $\leq \pi$. This classification is needed for the proof of the fibration theorem in Section 10. It also motivates our results in Section 5 where we study the local geometry of cone 3-manifolds with cone angles $\leq \pi$; there, a lower diameter bound plays the role of the noncompactness condition in the flat case. Our main result, cf. Section 5.2, is a geometric description of the thin part in the case when cone angles are bounded away from π and 0 (Theorem 5.3). As consequences, we obtain thickness (Theorem 5.4) and, when the volume is finite, the existence of a geometric compact core (Theorem 5.5). The other results relevant for the proof of the main theorem are the geometric fibration theorem for thin cone manifolds with totally geodesic boundary (Corollary 5.37) and the thick vertex lemma (Lemma 5.10) which is a simple result useful in the case of platonic vertices.

We give the proof of the main theorem in Section 6. Firstly we reduce to the case when the smooth part of the orbifold is hyperbolic. We view the (complete) hyperbolic structure on the smooth part as a hyperbolic cone structure on the orbifold with cone angles zero. The goal is to increase the cone angles of this hyperbolic cone structure as much as possible. In Section 6.2 we prove first that there exist such deformations which change the cone angles (*openness theorem*).

Next we consider a sequence of hyperbolic cone structures on the orbifold whose cone angles converge to the supremum of the cone angles in the deformation space. We have the following dichotomy: either the sequence collapses (i.e. the supremum of the injectivity radius for each cone structure goes to zero) or not (i.e. each cone structure contains a point with injectivity radius uniformly bounded away from zero).

In the *noncollapsing* case we show in Section 6.3 that the orbifold angles can be reached in the deformation space of hyperbolic cone structures, and therefore the orbifold is hyperbolic. This step uses a *stability theorem* which shows that a noncollapsing sequence of hyperbolic cone structures on the orbifold has a subsequence converging to a hyperbolic cone structure on the orbifold. We prove this theorem in Section 7.

Then we analyze the case where the sequence of cone structures collapses. If the diameters of the collapsing cone structures are bounded away from zero, then we conclude that the orbifold is Seifert fibred, using the *fibration theorem* which is proved in Section 10. Otherwise the diameter of the sequence of cone structures converges to zero. Then we show that the orbifold is geometric, unless the following situation occurs: the orbifold is closed and admits a Euclidean cone structure with cone angles strictly less than its orbifold angles.

We deal with this last case in Sections 8 and 9 proving that then the orbifold is spherical (*spherical uniformization* theorem). For orbifolds with cyclic or dihedral stabilizer, the proof relies on Hamilton's theorem [Ha1] about the Ricci flow on 3-manifolds. In the general case the proof is by induction on the number of platonic vertices and involves deformations of spherical cone structures.

Acknowledgements. We wish to thank J. Alze, D. Cooper and H. Weiß for useful conversations and remarks. We thank the RiP-program at the Mathematisches Forschungsinstitut Oberwolfach, as well as DAAD, MCYT (Grants HA2000-0053 and BFM2000-0007) and DURSI (ACI2000-17) for financial support.

2. 3-dimensional orbifolds

2.1. *Basic definitions.* For a general background about orbifolds we refer to [BMP], [BS1, 2], [CHK], [DaM], [Kap, Ch. 7], [Sco], and [Thu1, Ch. 13]. We begin by recalling some terminology from these references.

A compact 2-orbifold F^2 is said to be *spherical*, *discal*, *toric* or *annular* if it is the quotient by a finite smooth group action of respectively the 2-sphere S^2, the 2-disc D^2, the 2-torus T^2 or the annulus $S^1 \times [0, 1]$.

A compact 2-orbifold is *bad* if it is not good (i.e. it is not covered by a surface). Such a 2-orbifold is the union of two nonisomorphic discal 2-orbifolds along their boundaries.

A compact 3-orbifold $\mathcal{O}$ is *irreducible* if it does not contain any bad 2-suborbifold and if every orientable spherical 2-suborbifold bounds in $\mathcal{O}$ a discal 3-suborbifold, where a *discal* 3-orbifold is a finite quotient of the 3-ball by an orthogonal action.

A connected 2-suborbifold F^2 in an orientable 3-orbifold $\mathcal{O}$ is *compressible* if either F^2 bounds a discal 3-suborbifold in $\mathcal{O}$ or there is a discal 2-suborbifold Δ^2 which intersects transversally F^2 in $\partial\Delta^2 = \Delta^2 \cap F^2$ and is such that $\partial\Delta^2$ does not bound a discal 2-suborbifold in F^2.

A 2-suborbifold F^2 in an orientable 3-orbifold $\mathcal{O}$ is *incompressible* if no connected component of F^2 is compressible in $\mathcal{O}$.

A properly embedded 2-suborbifold F^2 is *∂-parallel* if it co-bounds a product with a suborbifold of the boundary (i.e. an embedded product $\overline{F}\times[0,1]\subset\mathcal{O}$ with $\overline{F}\times 0 = F^2$ and $\overline{F}\times 1\subset\partial\mathcal{O}$), so that $\partial\overline{F}\times[0,1]\subset\partial\mathcal{O}$.

A properly embedded 2-suborbifold $(F,\partial F)\hookrightarrow(\mathcal{O},\partial\mathcal{O})$ is *∂-compressible* if:

- either $(F,\partial F)$ is a discal 2-suborbifold $(D^2,\partial D^2)$ which is ∂-parallel,
- or there is a discal 2-suborbifold $\Delta\subset\mathcal{O}$ such that $\partial\Delta\cap F$ is a simple arc α which does not cobound a discal suborbifold of F with an arc in ∂F, and $\Delta\cap\partial\mathcal{O}$ is a simple arc β with $\partial\Delta=\alpha\cup\beta$ and $\alpha\cap\beta=\partial\alpha=\partial\beta$.

A properly embedded 2-suborbifold F^2 is *essential* in a compact orientable irreducible 3-orbifold, if it is incompressible, ∂-incompressible and not ∂-parallel.

A compact 3-orbifold is *topologically atoroidal* if it does not contain an embedded essential orientable toric 2-suborbifold.

A *turnover* is a 2-orbifold with underlying space a 2-sphere and ramification locus three points. In an irreducible orientable 3-orbifold, an embedded turnover either bounds a discal 3-suborbifold or is incompressible and of non-positive Euler characteristic.

An orientable compact 3-orbifold $\mathcal{O}$ is *Haken* if it is irreducible, if every embedded turnover is either compressible or ∂-parallel, and if it contains an embedded orientable incompressible 2-suborbifold which is not a turnover.

Remark 2.1. The word Haken may lead to confusion, since it is not true that a compact orientable irreducible 3-orbifold containing an orientable incompressible properly embedded 2-suborbifold is Haken in our meaning (cf. [BMP, Ch. 4], [Dun1], [BoP, Ch. 8]).

An orientable compact 3-orbifold $\mathcal{O}$ is *small* if it is irreducible, its boundary $\partial\mathcal{O}$ is a (perhaps empty) collection of turnovers, and $\mathcal{O}$ does not contain any essential orientable 2-suborbifold. It follows from Dunbar's theorem [Dun1] that the hypothesis about the boundary is automatically satisfied once $\mathcal{O}$ does not contain any essential 2-suborbifold.

Remark 2.2. By irreducibility, if a small orbifold $\mathcal{O}$ has nonempty boundary, then either $\mathcal{O}$ is a discal 3-orbifold, or $\partial\mathcal{O}$ is a collection of Euclidean and hyperbolic turnovers.

A 3-orbifold $\mathcal{O}$ is *geometric* if either it is the quotient of a ball by an orthogonal action, or its interior has one of the eight Thurston geometries. We quickly review those geometries.

A compact orientable 3-orbifold $\mathcal{O}$ is *hyperbolic* if its interior is orbifold-diffeomorphic to the quotient of the hyperbolic space $\mathbb{H}^3$ by a nonelementary discrete group of isometries. In particular I-bundles over hyperbolic 2-orbifolds are hyperbolic, since their interiors are quotients of $\mathbb{H}^3$ by nonelementary Fuchsian groups.

A compact orientable 3-orbifold is *Euclidean* if its interior has a complete Euclidean structure. Thus, if a compact orientable and ∂-incompressible 3-orbifold $\mathcal{O}$ is Euclidean, then either $\mathcal{O}$ is an I-bundle over a 2-dimensional Euclidean closed orbifold or $\mathcal{O}$ is closed.

A compact orientable 3-orbifold is *spherical* when it is the quotient of the standard sphere $\mathbb{S}^3$ or the round ball B^3 by the orthogonal action of a finite group.

A *Seifert fibration* on a 3-orbifold $\mathcal{O}$ is a partition of $\mathcal{O}$ into closed 1-suborbifolds (circles or intervals with silvered boundary) called fibers, such that each fiber has a saturated neighborhood diffeomorphic to $S^1 \times D^2/G$, where G is a finite group which acts smoothly, preserves both factors, and acts orthogonally on each factor and effectively on D^2; moreover the fibers of the saturated neighborhood correspond to the quotients of the circles $S^1 \times \{*\}$. On the boundary $\partial\mathcal{O}$, the local model of the Seifert fibration is $S^1 \times D^2_+/G$, where D^2_+ is a half-disc.

A 3-orbifold that admits a Seifert fibration is called Seifert fibred. A Seifert fibred 3-orbifold which does not contain a bad 2-suborbifold is geometric (cf. [BMP, Ch. 1, 2], [Sco], [Thu7]).

Besides the constant curvature geometries $\mathbb{E}^3$ and $\mathbb{S}^3$, there are four other possible 3-dimensional homogeneous geometries for a Seifert fibred 3-orbifold: $\mathbb{H}^2 \times \mathbb{R}$, $\mathbb{S}^2 \times \mathbb{R}$, $\widetilde{\mathrm{SL}_2(\mathbb{R})}$ and *Nil*.

The geometric but non-Seifert fibred 3-orbifolds require either a constant curvature geometry or *Sol*. Compact 3-orbifolds with *Sol* geometry are fibred over a closed 1-dimensional orbifold with toric fiber and thus they are not topologically atoroidal (cf. [Dun2]).

2.2. *Spherical and toric decompositions.* Thurston's geometrization conjecture asserts that any compact, orientable, 3-orbifold, which does not contain any bad 2-suborbifold, can be decomposed along a finite collection of disjoint, nonparallel, essential, embedded spherical and toric 2-suborbifolds into geometric suborbifolds.

The topological background for Thurston's geometrization conjecture is given by the spherical and toric decompositions.

Given a compact orientable 3-orbifold without bad 2-suborbifolds, the first stage of the splitting is called *spherical* or *prime* decomposition, and it expresses the 3-orbifold as the connected sum of 3-orbifolds which are either homeomorphic to a finite quotient of $S^1 \times S^2$ or irreducible. We refer to [BMP, Ch. 3], [TY1] for details.

The second stage (*toric* splitting) is a more subtle decomposition of each irreducible factor along a finite (maybe empty) collection of disjoint and nonparallel essential, toric 2-suborbifolds. This collection of essential toric 2-suborbifolds is unique up to isotopy. It cuts the irreducible orbifold into topologically atoroidal or Seifert fibred 3-suborbifolds; see [BS1], [BMP, Ch. 3].

By these spherical and toric decompositions, Thurston's geometrization conjecture reduces to the case of a compact, orientable 3-orbifold which is irreducible and topologically atoroidal.

Our proof requires a further decomposition along turnovers due to Dunbar ([BMP, Ch. 3], [Dun1, Th. 12]). A compact irreducible and topologically atoroidal 3-orbifold has a maximal family of nonparallel essential turnovers, which may be empty. This family is unique up to isotopy and cuts the orbifold into pieces without essential turnovers.

2.3. *Finite group actions on spheres with fixed points.*

Proof of Corollary 1.1 *from the main theorem.* Consider a nonfree action of a finite group Γ on S^3 by orientation-preserving diffeomorphisms. Let $\mathcal{O} = \Gamma\backslash S^3$ be the quotient orbifold.

If $\mathcal{O}$ is irreducible then the equivariant Dehn lemma implies that any 2-suborbifold with infinite fundamental group has a compression disc. Hence $\mathcal{O}$ is small and we apply the main theorem.

Suppose that $\mathcal{O}$ is reducible. Since $\mathcal{O}$ does not contain a bad 2-suborbifold, there is a prime decomposition along a family of spherical 2-suborbifolds; see Section 2.2. These lift to a family of 2-spheres in S^3. Consider an innermost 2-sphere; it bounds a ball $B \subset S^3$. The quotient Q of B by its stabilizer Γ' in Γ has one boundary component which is a spherical 2-orbifold. We close it by attaching a discal 3-orbifold. The resulting closed 3-orbifold $\mathcal{O}'$ is a prime factor of $\mathcal{O}$. The orbifold $\mathcal{O}'$ is irreducible, and hence spherical. The action of Γ' on $\widetilde{\mathcal{O}'} \cong S^3$ is standard and preserves the sphere ∂B. Thus the action is a suspension and Q is discal. This contradicts the minimality of the prime decomposition. $\square$

2.4. *Proof of the orbifold theorem from the main theorem.* This step of the proof is based on the following extension of Thurston's hyperbolization theorem to Haken orbifolds (cf. [BoP, Ch. 8]):

THEOREM 2.3 (Hyperbolization theorem of Haken orbifolds). *Let $\mathcal{O}$ be a compact orientable connected Haken* 3-*orbifold. If $\mathcal{O}$ is topologically atoroidal and not Seifert fibred, nor Euclidean, then $\mathcal{O}$ is hyperbolic.*

Remark 2.4. The proof of this theorem follows exactly the scheme of the proof for Haken manifolds [Thu2, 3, 5], [McM1], [Kap], [Ot1, 2] (cf. [BoP, Ch. 8] for a precise overview).

Proof of Corollary 1.2 (the orbifold theorem). Let $\mathcal{O}$ be a compact orientable connected irreducible topologically atoroidal 3-orbifold. By [BMP, Ch. 3], [Dun1, Th. 12] there exists in $\mathcal{O}$ a (possibly empty) maximal collection $\mathcal{T}$ of disjoint embedded pairwise nonparallel essential turnovers. Since $\mathcal{O}$ is irreducible and topologically atoroidal, any turnover in $\mathcal{T}$ is hyperbolic (i.e. has negative Euler characteristic).

When $\mathcal{T}$ is empty, the orbifold theorem reduces either to the main theorem or to Theorem 2.3 according to whether $\mathcal{O}$ is small or Haken.

When $\mathcal{T}$ is not empty, we first cut open the orbifold $\mathcal{O}$ along the turnovers of the family $\mathcal{T}$. By maximality of the family $\mathcal{T}$, the closure of each component of $\mathcal{O} - \mathcal{T}$ is a compact orientable irreducible topologically atoroidal 3-orbifold that does not contain any essential embedded turnover. Let $\mathcal{O}'$ be one of these connected components. By the previous case $\mathcal{O}'$ is either hyperbolic, Euclidean or Seifert fibred. Since, by construction, $\partial\mathcal{O}'$ contains at least one hyperbolic turnover T, $\mathcal{O}'$ must be hyperbolic. Moreover any such hyperbolic turnover T in $\partial\mathcal{O}'$ is a Fuchsian 2-suborbifold, because there is a unique conjugacy class of faithful representations of the fundamental group of a turnover in $\mathrm{PSL}_2(\mathbb{C})$.

We assume first that all the connected components of $\mathcal{O} - \mathcal{T}$ have 3-dimensional convex cores. In this case the totally geodesic hyperbolic turnovers are the boundary components of the convex cores. Hence the hyperbolic structures on the components of $\mathcal{O} - \mathcal{T}$ can be glued together along the hyperbolic turnovers of the family $\mathcal{T}$ to give a hyperbolic structure on the 3-orbifold $\mathcal{O}$.

If the convex core of $\mathcal{O}'$ is 2-dimensional, then $\mathcal{O}'$ is either a product $T \times [0,1]$, where T is a hyperbolic turnover, or a quotient of $T \times [0,1]$ by an involution. When $\mathcal{O}' = T \times [0,1]$, then the 3-orbifold $\mathcal{O}$ is Seifert fibred, because the mapping class group of a turnover is finite. When $\mathcal{O}'$ is the quotient of $T \times [0,1]$, then it has only one boundary component and it is glued either to another quotient of $T \times [0,1]$ or to a component with 3-dimensional convex core. When we glue two quotients of $T \times [0,1]$ by an involution, we obtain a Seifert fibred orbifold. Finally, gluing $\mathcal{O}'$ to a hyperbolic orbifold with totally geodesic boundary is equivalent to giving this boundary a quotient by an isometric involution. □

3. 3-dimensional cone manifolds

3.1. *Basic definitions.* We start by recalling the construction of metric cones.

Let k and $r > 0$ be real numbers; if $k > 0$ we assume in addition that $r \leq \frac{\pi}{\sqrt{k}}$. Suppose that Y is a metric space with $\mathrm{diam}(Y) \leq \pi$. On the set $Y \times [0,r]$ we define a pseudo-metric as follows. Given $(y_1,t_1),(y_2,t_2) \in Y \times [0,r]$, let $p_0p_1p_2$ be a triangle in the 2-dimensional model space $\mathbb{M}^2_k$ of constant curvature

k with $d(p_0, p_1) = t_1$, $d(p_0, p_2) = t_2$ and $\angle_{p_0} = d_Y(y_1, y_2)$. We put

$$d_{Y\times[0,r]}\big((y_1, t_1), (y_2, t_2)\big) := d_{\mathbb{M}^2_k}(p_1, p_2).$$

The metric space $C_{k,r}(Y)$ obtained from collapsing the subset $Y \times \{0\}$ to a point is called the *metric cone of curvature* k or k-*cone* of radius r over Y. In the special case when $k > 0$ and $r = \frac{\pi}{\sqrt{k}}$, one also has to collapse the subset $Y \times \{\frac{\pi}{\sqrt{k}}\}$ to a point. The point in $C_{k,r}(Y)$ corresponding to $Y \times \{0\}$ is called the *tip* or *apex* of the cone. The *complete* k-cone or simply k-*cone* $C_k(Y)$ over Y is defined as $C_{k,\infty}(Y) := \cup_{r>0} C_{k,r}(Y)$ if $k \le 0$ and as $C_{k,\frac{\pi}{\sqrt{k}}}(Y)$ if $k > 0$. The complete 1-cone over a space is also called its *metric suspension.*

We define cone manifolds as certain metric spaces locally isometric to iterated cones. To make this precise, we proceed by induction over the dimension. We first make the convention that the connected 1-dimensional cone manifolds of curvature 1 are circles of length $\le 2\pi$ or compact intervals of length $\le \pi$.

Definition 3.1 (Cone manifolds). An n-dimensional *conifold* of curvature k, $n \ge 2$, is a complete geodesic metric space locally isometric to the k-cone over a connected $(n-1)$-dimensional conifold of curvature 1.

A *cone manifold* is a conifold which is topologically a manifold.

Conifolds of curvature $k = +1$, $k = 0$ or $k = -1$ are called *spherical*, *Euclidean* or *hyperbolic*, respectively.

Spelled out in more detail, the definition requires that for every point x in a n-conifold X there exists a radius $\varepsilon > 0$ and an isometry from the closed ball $\overline{B}_\varepsilon(x)$ to the k-cone $C_{k,\varepsilon}(\Lambda_x X)$ over a metric space $\Lambda_x X$ carrying x to the tip of the cone. Moreover, $\Lambda_x X$ must be itself an $(n-1)$-conifold of curvature 1.

The metric space $\Lambda_x X$ is called the *space of directions* or *link* of X at x.[1] It can be defined intrinsically as the space of germs of geodesic segments in X emanating from x equipped with the angular metric. It is implicit in the definition that the links $\Lambda_x X$ are complete metric spaces. Since they have curvature 1, it follows that they are compact with diameters $\le \pi$; see the discussion at the end of this section.

We note that all conifolds of dimension ≤ 2 are manifolds. The links in 3-dimensional conifolds are, according to the Gauß-Bonnet Theorem (extended to singular surfaces), topologically 2-spheres, 2-discs or projective planes. If none of the links is a projective plane, then the conifold is a manifold. The wider concept of conifold will play no role in this paper; later on, we will only consider cone *manifolds* of dimensions ≤ 3.

[1]The standard geometric notation would be $\Sigma_x X$, but we already make extensive use of the letter Σ, namely for the singular locus of an orbifold.

Example 3.2 (Geometric orbifolds). A *geometric orbifold* of dimension n and curvature k is a complete geodesic metric space which is locally isometric to the quotient of the model space $\mathbb{M}^n_k$ by a finite group of isometries.

Unlike topological orbifolds, geometric orbifolds are always *global* quotients, i.e. they are (even finite) quotients of manifolds of constant curvature by discrete group actions.

We define the *boundary* of a conifold by induction over the dimension. The boundary points of a 1-conifold are the endpoints of its interval components. The boundary points of a n-conifold, $n \geq 2$, are the points whose links have boundary.

A point x in a conifold X is called a *smooth interior* point if X is locally at x isometric to the model space $\mathbb{M}^n_k$ of the same curvature and dimension as X, or equivalently, if the link $\Lambda_x X$ is a unit sphere. If $\Lambda_x X$ is a unit hemisphere, the point x is a *smooth boundary* point. All other points are called *singular*. We denote by X^{smooth} the subset of smooth points, and by Σ_X its complement, the *singular locus*.

Let us go through this in low dimensions. One-dimensional cone manifolds contain only regular points. If S is a *cone surface*, i.e. a cone 2-manifold, then Σ_S is a discrete subset. A singular point is either a *corner* of the boundary, if its link is an interval of length $< \pi$, or a *cone point* in the interior, if its link is a circle of length $< 2\pi$. In the latter case, the length of the circle is called the *cone angle*.

Consider now a 3-dimensional cone manifold X. In this case, the singular set Σ_X is one-dimensional, namely a geodesic graph. We define $\Sigma_X^{(1)} \subseteq \Sigma_X$ as the subset of singular points x whose link $\Lambda_x X$ is the metric suspension of (complete 1-cone over) a circle. The length of the circle is called the cone angle at x. We call the closure of a component of $\Sigma_X^{(1)}$ a *singular edge*. The cone angle is constant along edges, and we can thus speak of the *cone angle* of an edge. The complement $\Sigma_X^{(0)} := \Sigma_X - \Sigma_X^{(1)}$ is discrete and its points are called *singular vertices*.

Notice that a cone surface or a cone 3-manifold without boundary is a geometric orbifold if and only if all cone angles are divisors of 2π. In particular the cone angles of a geometric 3-orbifold are $\leq \pi$, and due to this fact we will be mostly interested in cone manifolds *with cone angles* $\leq \pi$.

Proposition 3.3. *Conifolds of curvature k are metric spaces with curvature $\geq k$ in the sense of Alexandrov.*

This can be readily seen by induction over the dimension using the following facts: Since conifolds are metrically complete by assumption, a local curvature bound implies a global curvature bound (Toponogov's theorem); the

k-cones over compact intervals of length $\leq \pi$ and circles of length $\leq 2\pi$ are spaces with curvature $\geq k$; the k-cone over a space with curvature ≥ 1 is a space with curvature $\geq k$. Note also that spaces with curvature ≥ 1 have diameter $\leq \pi$, due to the singular version of the Bonnet-Myers theorem; cf. [BGP, Th. 3.6].

All our geometric considerations will take place within the framework of metric spaces with curvature bounded below. For this theory, we refer the reader to the fundamental paper [BGP] and the introductory text [BBI, Ch. 10].

3.2. *Exponential map*, *cut locus*, (*cone*) *injectivity radius.* Consider a connected conifold X of curvature k and dimension ≥ 2.

For a point $p \in X$, according to our requirement on the local geometry of conifolds, there exists $\varepsilon > 0$ such that the cone $C_{k,\varepsilon}(\Lambda_p X)$ canonically embeds into X, its tip O being mapped to p. This embedding extends naturally to a map from a larger domain inside the complete cone $C_k(\Lambda_p X)$ as follows: Let $\mathcal{E}(p) \subseteq C_k(\Lambda_p X)$ be the union of all geodesic segments $\overline{Oy}$, such that there exists a geodesic segment $\overline{px_y}$ in X with the same length and the same initial direction modulo the natural identification $\Lambda_O(C_k(\Lambda_p X)) \cong \Lambda_p X$. The subset $\mathcal{E}(p)$ is star-shaped with respect to O, and we define the *exponential map* in p

$$\exp_p : \mathcal{E}(p) \longrightarrow X$$

as the map sending each point y to the respective point x_y.

The *conjugate radius* is defined, purely in terms of the curvature, as $r_{\text{conj}} := \frac{\pi}{\sqrt{k}}$ if $k > 0$ and $r_{\text{conj}} := \infty$ if $k \leq 0$, i.e. $r_{\text{conj}} = \operatorname{diam}(C_k(\Lambda_p X))$. The *geodesic radius* in a point p, $0 < r_{\text{geod}}(p) \leq r_{\text{conj}}$, is the radius of the largest ball in $C_k(\Lambda_p X)$ around O on which $\exp_p$ is defined.

Let x be an interior point of a geodesic segment $\sigma = \overline{pq}$. Then $\Lambda_x X$ has extremal diameter π and, by the Diameter Rigidity Theorem, is a metric suspension with the directions of σ in x as poles. The equator of the suspension consists of the directions at x perpendicular to σ.

For any $0 < d < \min(d(p,q), r_{\text{conj}})$ there exists a sufficiently small $\delta > 0$ such that the "thin" cone $C_{k,d}(B_\delta(\Lambda_p\sigma))$ is contained in $\mathcal{E}(p)$ and embeds via $\exp_p$ locally isometrically into X. Here $\Lambda_p\sigma \in \Lambda_p X$ denotes the direction of σ at its endpoint p.

If σ has length $< r_{\text{conj}}$, and if $\sigma' = \overline{pq'}$ is sufficiently Hausdorff close to σ, then there exists an isometrically immersed (2-dimensional) triangle of constant curvature k with σ and σ' as two of its sides. It follows that there do not exist other geodesic segments with the same endpoints as σ and arbitrarily Hausdorff close to σ.

We now focus our attention on *minimizing* geodesic segments. Let p and q be points with $d(p,q) < r_{\text{conj}}$. Our discussion implies that there are at most finitely many minimizing geodesic segments $\sigma_1, \ldots, \sigma_m$ connecting them.

If x is a point sufficiently close to q, then for every i there exists a locally isometrically embedded triangle Δ_i with x as vertex and σ_i as opposite side. Moreover, any minimizing segment $\tau = \overline{px}$ is Hausdorff close to one of the segments σ_i and coincides with the side $\overline{px}$ of the corresponding triangle Δ_i. So, there exists a minimizing segment $\overline{px}$ Hausdorff close to σ_j if and only if $\angle_q(\sigma_j, x) = \min_i \angle_q(\sigma_i, x)$.

Let $\mathcal{D}(p) \subseteq \mathcal{E}(p)$ be the union of all geodesic segments $\overline{Oy}$ in $\Lambda_p X$ whose images $\overline{px_y}$ under $\exp_p$ are *minimizing* segments. Let $\dot{\mathcal{D}}(p) \subseteq \mathcal{D}(p)$ be the subset consisting of O and all interior points of such segments $\overline{Oy}$. Note that $\dot{\mathcal{D}}(p)$ is open and its closure equals $\mathcal{D}(p)$. We have $\mathcal{D}(p) - \dot{\mathcal{D}}(p) = \partial\mathcal{D}(p)$ except in the special case when $k > 0$ and X is a metric suspension with tip p.

Definition 3.4 (Cut locus). The subset $\mathrm{Cut}_X(p) = \mathrm{Cut}(p) := \exp_p(\mathcal{D}(p) - \dot{\mathcal{D}}(p)) \subset X$ is called the *cut locus* with respect to the point p.

In other words, $\mathrm{Cut}(p)$ is the complement of the union of p and all minimizing half-open segments $\gamma : [0, l) \to X$ with initial point $\gamma(0) = p$. More generally, one can define in this way the cut locus $\mathrm{Cut}(F)$ with respect to a *finite set* $F \subset X$. Our discussion above implies:

PROPOSITION 3.5 (Local conicality of cut locus). *For any point $q \in \mathrm{Cut}(p)$ with $d(p, q) < r_{\mathrm{conj}}$ there exists $\varepsilon > 0$ such that*

$$\mathrm{Cut}_X(p) \cap B_\varepsilon(q) = C_{k,\varepsilon}(\mathrm{Cut}_{\Lambda_q X}(F))$$

where $F \subset \Lambda_q X$ is the finite set of directions of minimizing segments between p and q.

If $k > 0$ and X is a metric suspension with tip p, then $\mathrm{Cut}(p)$ consists of just one point, namely the antipode of p.

In all other cases, induction over the dimension, by Proposition 3.5, yields that $\mathrm{Cut}(p)$ is a possibly empty, locally finite, piecewise totally geodesic polyhedral complex of codimension one, and $\mathcal{D}(p)$ is a locally finite polyhedron in $C_k(\Lambda_p X)$ with geodesic faces. The conifold X arises from $\mathcal{D}(p)$ by identifications on the boundary, namely by isometric face pairings.

Definition 3.6 (Dirichlet polyhedron). $\mathcal{D}(p) \subseteq C_k(\Lambda_p X)$ is called the *Dirichlet polyhedron* with respect to p.

In dimension 2, the Dirichlet polyhedra are polygons. If X is a cone surface, then the vertices of $\mathcal{D}(p)$ correspond to either smooth interior points of X with ≥ 3 minimizing segments towards p, to boundary points or to cone points. In the latter cases there may exist only one minimizing segment to p. If this happens for a cone point, then the angle at the corresponding vertex of $\mathcal{D}(p)$ equals the cone angle. This is the only way, in which *concave* vertices of the Dirichlet polygon can occur: Every vertex of $\mathcal{D}(p)$ with angle $> \pi$ corresponds to a cone point which is connected to p by exactly one minimizing segment.

The discussion in dimension 3 is analogous. In particular, if X is a 3-conifold then edges of Dirichlet polyhedra with dihedral angles $> \pi$ project via the exponential map to (parts of) singular edges with cone angles $> \pi$. Therefore we have the following strong restriction on the geometry of Dirichlet polyhedra for cone angles $\leq \pi$:

Proposition 3.7 (Convexity). *In the case of cone angles $\leq \pi$, the Dirichlet polyhedra are* convex*.*

The exponential map is a local isometry near the tip O of $C_k(\Lambda_p X)$.

Definition 3.8 (Injectivity radius). The *injectivity radius* in p, $0 < r_{\text{inj}}(p) \leq r_{\text{geod}}(p)$, is the radius of the largest open ball in $C_k(\Lambda_p X)$ around O on which $\exp_p$ is an embedding; i.e., it is maximal with the property that all geodesic segments of length $< r_{\text{inj}}(p)$ starting in p are minimizing.

Since the cut locus $\text{Cut}(p)$ is closed, there exist cut points q at minimal distance $r_{\text{inj}}(p)$ from p. The minimizing segments $\overline{pq}$ must have angles $\geq \frac{\pi}{2}$ with the cut locus. Since $\text{diam}(\Lambda_q X) \leq \pi$, there can be at most two minimizing segments $\overline{pq}$. If there are two, they meet with maximal angle π at q and form together a geodesic loop with base point p and midpoint q. If there is a unique minimizing segment $\overline{pq}$ and if q does not belong to the boundary, then q must lie on a (closed) singular edge with cone angle $\geq \pi$. Note that this alternative cannot occur for cone angles $< \pi$.

The injectivity radius varies continuously with p on the smooth part and along singular edges. However it converges to zero, e.g. along sequences of smooth points approaching the singular locus. In the singular setting, the injectivity radius is not the right measure for the simplicity of the local geometry. In order to measure up to which scale the local geometry is given by certain simple models, the following modification turns out to be useful, at least as long as the cone angles are $\leq \pi$.

Definition 3.9 (Cone injectivity radius). The *cone injectivity radius* $r_{\text{cone-inj}}(p)$ in p is the supremum of all $r > 0$ such that the ball $B_r(p)$ is contained in a standard ball, i.e. such that there exist $q \in X$ and $R > 0$ with the following property: $B_r(p) \subseteq B_R(q)$ and $\overline{B}_R(q) \cong C_{k,R}(\Lambda_q X)$.

3.3. *Spherical cone surfaces with cone angles $\leq \pi$.* In this section we will discuss closed cone surfaces Λ with curvature 1 and cone angles $\leq \pi$, whose underlying topological surface is a 2-sphere. They occur as links of 3-dimensional cone manifolds with cone angles $\leq \pi$, the class of cone manifolds mostly relevant for us in this paper.

Proposition 3.10 (Classification). *Let Λ be a spherical cone surface with cone angles $\leq \pi$ which is homeomorphic to the 2-sphere. Then Λ is isometric to either*

- *the unit 2-sphere S^2,*
- *the metric suspension $S^2(\alpha, \alpha)$ of a circle of length $\alpha \leq \pi$, or to*
- *$S^2(\alpha, \beta, \gamma)$, the double along the boundary of a spherical triangle with angles $\frac{\alpha}{2}, \frac{\beta}{2}, \frac{\gamma}{2} \leq \frac{\pi}{2}$.*

Proof. The assertion is clear in the smooth case and we therefore assume that Λ has cone points. Due to Gauß-Bonnet, there can be at most three cone points.

If Λ has only one cone point c, then $\Lambda - \{c\}$ is simply connected and hence can be developed (isometrically immersed) into S^2. A circle of small radius centered at c cannot close up under the developing map and we obtain a contradiction. Thus Λ must have two or three cone points.

If Λ has two cone points, we connect them by a minimizing segment σ. By cutting Λ open along σ we obtain a spherical surface which is topologically a disc and whose boundary consists of two edges of equal length. It can be developed into S^2 as well, and it follows that the surface is a spherical bigon, i.e. the metric suspension of an arc. We obtain the second alternative of our assertion.

If Λ has three cone points, we connect any two of them by a minimizing geodesic segment. The segments do not intersect and they divide Λ into two spherical triangles. The triangles are isometric because they have the same side lengths, and we obtain the third alternative. □

A consequence of the classification is the following description for the local geometry of a cone 3-manifold with cone angles $\leq \pi$.

Corollary 3.11. *If p is an interior point in a cone 3-manifold with cone angles $\leq \pi$, then a sufficiently small ball $B_\varepsilon(p)$ is isometric to one of the following (see Figure 1):*

- *a ball of radius ε in a smooth model space $\mathbb{M}^3_k$,*
- *a singular ball $C_{k,\varepsilon}(S^2(\alpha, \alpha))$ with a singular axis of cone angle α, or*
- *a singular ball $C_{k,\varepsilon}(S^2(\alpha, \beta, \gamma))$ with three singular edges emanating from a singular vertex in the center.*

In particular, the singular locus Σ_X is a trivalent *graph; i.e., its vertices have valency at most three.*

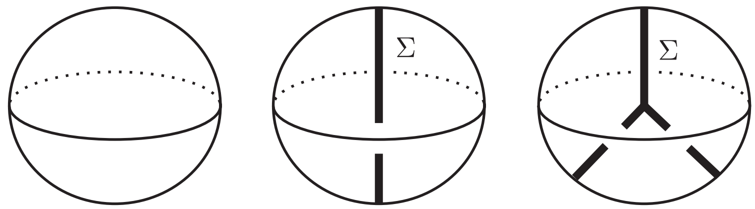

Figure 1

In the remainder of this section, we collect some properties of spherical cone surfaces used later.

LEMMA 3.12. *Let* Λ *be as in Proposition* 3.10. *Then* Λ *does not contain three points with pairwise distances* $> \frac{2\pi}{3}$.

Proof. This is a direct implication of the lower curvature bound 1 because the circumference of geodesic triangles has length $\leq 2\pi$. □

Definition 3.13 (Turnover). A *turnover* is a cone surface which is homeomorphic to the 2-sphere and which has three cone points, all with cone angle $\leq \pi$.

Geometrically, a turnover is the double along the boundary of a triangle in a 2-dimensional model space $\mathbb{M}^2_k$ with angles $\leq \frac{\pi}{2}$.

LEMMA 3.14. (i) *A spherical turnover* Λ *has diameter* $\leq \frac{\pi}{2}$.

(ii) *If* Λ *is a spherical turnover with* $\operatorname{diam}(\Lambda) = \frac{\pi}{2}$, *then at least two of the three cone angles equal* π. *If two points* $\xi, \eta \in \Lambda$ *have maximal distance* $\frac{\pi}{2}$ *then at least one of them, say* ξ, *is a cone point, and* η *lies on the minimizing segment joining the other two cone points, and these must have cone angles* $= \pi$.

Proof. (i) Let $\xi, \eta \in \Lambda$ and suppose that ζ is a cone point $\neq \xi, \eta$. Any geodesic triangle $\Delta(\xi, \eta, \zeta)$ has angle $\leq \frac{\pi}{2}$ at ζ. We denote $\operatorname{rad}(\Lambda, \zeta) := \max d(\zeta, \cdot)$. Since $\operatorname{rad}(\Lambda, \zeta) \leq \frac{\pi}{2}$, hinge comparison implies that $d(\xi, \eta) \leq \frac{\pi}{2}$.

(ii) In the case of equality it follows that the cone angle at ζ equals π and that one of the points ξ or η, say ξ, has distance $\frac{\pi}{2}$ from ζ. If ξ were not a cone point, then it would lie on the segment connecting the two cone points $\neq \zeta$ and only ζ would have distance $\frac{\pi}{2}$ from ξ, contradicting $d(\xi, \eta) = \pi/2$. Hence ξ must be a cone point, and it follows that η lies on the segment joining ζ and the cone point $\neq \zeta, \xi$. □

LEMMA 3.15. *For $\alpha < \pi$ there exists $D = D(\alpha) < \frac{\pi}{2}$ such that: If Λ is a spherical turnover with at least two cone angles $\leq \alpha$ then* $\operatorname{diam}(\Lambda) \leq D(\alpha)$.

Proof. Λ is the double of a spherical triangle Δ with two angles $\leq \alpha/2$ and third angle $\leq \frac{\pi}{2}$. Since the angle sum of a spherical triangle is $> \pi$, all angles of Δ are $> \frac{\pi-\alpha}{2}$. Such triangles can (Gromov-Hausdorff) converge to a point, but not to a segment. Hence the Gromov-Hausdorff closure of the space of turnovers as in the lemma is compact and contains as the only additional space the point. It follows that the diameter assumes a maximum $D(\alpha)$ on this space of turnovers. By part (ii) of Lemma 3.14, we have $D(\alpha) < \frac{\pi}{2}$. ☐

LEMMA 3.16. *For $\alpha < \pi$ and $0 < d \leq \frac{\pi}{2}$ there exists $r = r(\alpha, d) > 0$ such that: If Λ is a spherical turnover with diameter $\geq d$ and cone angles $\leq \alpha$, then it contains an embedded smooth round disc with radius r.*

Proof. The turnover Λ is the double of a spherical triangle Δ with acute angles $\leq \alpha/2$ and a lower diameter bound. Since the angle sum of spherical triangles is $> \pi$, we also have the positive lower bound $\pi - \alpha$ for the angles of Δ. Such triangles have a lower bound on their inradius, whence the claim. ☐

3.4. *Compactness for spaces of thick cone manifolds.* The space of pointed cone 3-manifolds with bounded curvature is *pre*compact in the Gromov-Hausdorff topology by Gromov's compactness theorem; cf. [GLP], because the volume growth is at most as strong as in the model space. The limit spaces in the Gromov-Hausdorff closure are spaces with curvature bounded below. We will show that, under appropriate assumptions, limits of cone 3-manifolds are still cone 3-manifolds.

Definition 3.17 (Thick). For $\rho > 0$, a cone manifold X is said to be *ρ-thick* (at a point x) if it contains an embedded smooth standard ball of radius ρ (centered at x). Otherwise X is called *ρ-thin*.

For $\kappa, i, a > 0$ we denote by $\mathcal{C}_{\kappa,i,a}$ the space of pointed cone 3-manifolds (X, p) with constant curvature $k \in [-\kappa, \kappa]$, cone angles $\leq \pi$ and base point p which satisfies $r_{\text{inj}}(p) \geq i$ and $\operatorname{area}(\Lambda_p X) \geq a$. Let $\mathcal{C}_{\kappa,i} := \mathcal{C}_{\kappa,i,4\pi}$ be the subspace of cone manifolds with smooth base point; they are i-thick at their base points.

THEOREM 3.18 (Compactness for thick cone manifolds with cone angles $\leq \pi$). *The spaces $\mathcal{C}_{\kappa,i}$ and $\mathcal{C}_{\kappa,i,a}$ are compact in the Gromov-Hausdorff topology.*

The main step in the proof of the theorem is the following result.

Proposition 3.19 (Controlled decay of the injectivity radius). *For* $\kappa \geq 0$, $R \geq i > 0$ *and* $a > 0$ *there exist* $r'(\kappa, i, a, R) \geq i'(\kappa, i, a, R) > 0$ *such that the following holds*:

Let X *be a closed cone* 3*-manifold with curvature* $k \in [-\kappa, \kappa]$ *and cone angles* $\leq \pi$. *Let* $p \in X$ *be a point with* $r_{\text{inj}}(p) \geq i$ *and* $\text{area}(\Lambda_p X) \geq a$. *Then for every point* $x \in B_R(p)$ *the ball* $B_{i'}(x)$ *is contained in a standard ball with radius* $\leq r'$. *In particular*, $r_{\text{cone-inj}} \geq i'$ *on* $B_R(p)$.

By a *standard ball* we mean the k-cone over a spherical cone surface homeomorphic to the 2-sphere; cf. Definition 3.9.

Proof. Step 0. It follows from the classification of links, cf. Proposition 3.10, that $\Lambda_p X$ contains a smooth standard disc with radius bounded below in terms of a, and hence the ball $B_i(p)$ contains an embedded *smooth* standard ball with a lower bound on its radius in terms of κ, i and a. We may therefore assume without loss of generality that p is a smooth point.

Step 1. We have a lower bound $\text{vol}(B_R(x) - B_{i/2}(x)) \geq v(\kappa, i) > 0$ because $B_R(x) - B_{i/2}(x)$ contains a smooth standard ball of radius $\geq i/4$. Let $A_x \subseteq \Lambda_x X$ denote the subset of initial directions of minimizing geodesic segments with length $\geq i/2$. The lower bound for the volume of the annulus $B_R(x) - B_{i/2}(x)$ implies a lower bound $\text{area}(A_x) \geq a_1(\kappa, i, R) > 0$.

Step 2. By triangle comparison, there exists for $\varepsilon > 0$ a number $l = l(\kappa, i, \varepsilon) > 0$ such that: Any geodesic loop of length $\leq 2l$ based in x has angle $\geq \frac{\pi}{2} - \varepsilon$ with all directions in A_x. The same holds for the angles of A_x with segments of length $\leq l$ starting in x and perpendicular to the singular locus Σ_X. Thus, if $r_{\text{inj}}(x) \leq l$, then minimizing segments from x to the closest cut points must have angles $\geq \frac{\pi}{2} - \varepsilon$ with all directions in A_x; cf. our discussion of the cut locus in Section 3.2. We use this observation to obtain lower bounds for the injectivity radius.

Lemma 3.20. *For* $a' > 0$ *there exists* $\varepsilon = \varepsilon(a') > 0$. *Let* Λ *be a spherical cone surface homeomorphic to the* 2*-sphere and with cone angles* $\leq \pi$. *Let* $A \subset \Lambda$ *be a subset with* $\text{area}(A) \geq a'$. *Then* $\Lambda = N_{\frac{\pi}{2}-\varepsilon}(A)$ *if* Λ *is a turnover. If* Λ *has* 0 *or* 2 *cone points, then there exists a point* η *such that* $\Lambda - N_{\frac{\pi}{2}-\varepsilon}(A) \subset B_{\frac{\pi}{2}-\varepsilon}(\eta)$. *In the case that* Λ *has two cone points,* η *can be chosen as a cone point.*

Proof. When Λ is a turnover, the description in Lemma 3.14 of segments of maximal length $\frac{\pi}{2}$ implies: Points in Λ with radius (Hausdorff distance from Λ) close to $\frac{\pi}{2}$ must be close to one of the three minimizing segments connecting cone points, i.e., must lie in a region of small area. Hence A contains points with radius $< \frac{\pi}{2} - \varepsilon$ for sufficiently small $\varepsilon > 0$ depending on $\text{area}(A)$.

If Λ has 0 or 2 cone points then it is isometric to the unit sphere S^2 or the metric suspension of a circle with length $\leq \pi$; cf. the classification in Proposition 3.10. In both cases the assertion is easily verified. $\square$

We choose $\varepsilon := \varepsilon(a_1)$ with $a_1 = a_1(\kappa, i, R)$ as in Step 1, and accordingly $l = l(\kappa, i, \varepsilon) = l(\kappa, i, R)$.

Step 3. For a singular vertex x Lemma 3.20 implies that $r_{\text{inj}}(x) \geq i_1 = i_1(\kappa, i, R) := l(\kappa, i, R) > 0$.

Step 4. Assume that x is a singular point with $r_{\text{inj}}(x) \leq i_1 = l$ at distance $\geq i_1/4$ from all singular vertices, and choose the singular direction $\eta_x \in \Lambda_x X$ according to Lemma 3.20. By the assumption on the injectivity radius, there exists a geodesic loop λ of length $\leq 2l$ based at x or a segment $\overline{xy}$ of length $\leq l$ meeting Σ_X orthogonally at a point y. Either of them has angles $\geq \frac{\pi}{2} - \varepsilon$ with the directions in A_x and therefore angles $\leq \frac{\pi}{2} - \varepsilon$ with the direction η_x.

In the case of a loop, consider the geodesic variation of λ moving its base point with unit speed in the direction η_x. Since both ends of the loop have angle $\leq \frac{\pi}{2} - \varepsilon$ with η_x, the first variation formula implies that the length of λ decreases at a rate $\leq -2\sin\varepsilon$. Similarly, in the case of a segment, r_{inj} decreases at a rate $\leq -\sin\varepsilon$. It follows that $r_{\text{inj}}(x) \geq \frac{i_1}{4} \cdot \sin\varepsilon =: i_2 = i_2(\kappa, i, R)$.

Step 5. Suppose now that x is a smooth point with $r_{\text{inj}}(x) \leq i_2$ at distance $\geq i_2/4$ from Σ_X. We choose the direction $\eta_x \in \Lambda_x X$ according to Lemma 3.20. As in Step 4, we see that r_{inj} decays in the direction η_x with rate $\leq -\sin\varepsilon$. It follows that $r_{\text{inj}}(x) \geq \frac{i_2}{4} \cdot \sin\varepsilon =: i_3 = i_3(\kappa, i, R)$.

Conclusion. The assertion holds for $r' := i_1$ and $i' := i_3$. $\square$

Proof of Theorem 3.18. Let (Y, q) be an Alexandrov space in the Gromov-Hausdorff closure of $\mathcal{C}_{\kappa,i,a}$. It is the Gromov-Hausdorff limit of a sequence of pointed cone manifolds $(X_n, p_n) \in \mathcal{C}_{\kappa,i,a}$. For a point $y \in Y$, we pick points $x_n \in X_n$ converging to y. The metric ball $B_\rho(y) \subset Y$ is then the Gromov-Hausdorff limit of the balls $B_\rho(x_n)$ in the approximating cone manifolds X_n.

Proposition 3.19 yields numbers $r' \geq i' > 0$ such that each ball $B_{i'}(x_n)$ is contained in a standard ball $B_{r'_n}(x'_n)$ with radius bounded above by $r'_n \leq r'$. Moreover, the lower bound on the volumes of the balls $B_i(p_n)$ yields a uniform estimate $\text{area}(\Lambda_{x'_n} X_n) \geq a'(\kappa, i, a, d(q, y)) > 0$.

It is clear from the classification of links in Proposition 3.10 that the space $\mathcal{C}^2_{a'}$ of spherical cone surfaces homeomorphic to the 2-sphere with cone angles $\leq \pi$ and area $\geq a'$ is Gromov-Hausdorff compact. Thus, after passing to a subsequence, we have that the links $\Lambda_{x'_n} X_n$ converge to a cone surface $\Lambda \in \mathcal{C}^2_{a'}$. Moreover, $r'_n \to r'_\infty \leq r'$ and $k_n \to k_\infty$ where k_n denotes the curvature of X_n.

It follows that $B_{r'_n}(x'_n) \cong C_{k_n,r'_n}(\Lambda_{x'_n} X_n) \to C_{k_\infty,r'_\infty}(\Lambda)$. This means that Y is a cone manifold locally at y. It is then clear that $Y \in \mathcal{C}_{\kappa,i,a}$. □

In our context, Gromov-Hausdorff convergence implies a stronger type of convergence, namely a version of bilipschitz convergence for cone manifolds. Recall that, for $\varepsilon > 0$, a map $f : X \to Y$ between metric spaces is called a $(1+\varepsilon)$-bilipschitz embedding if

$$(1+\varepsilon)^{-1} \cdot d(x_1, x_2) < d(f(x_1), f(x_2)) < (1+\varepsilon) \cdot d(x_1, x_2)$$

holds for all points $x_1, x_2 \in X$.

Definition 3.21 (Geometric convergence). A sequence of pointed cone 3-manifolds (X_n, x_n) *converges geometrically* to a pointed cone 3-manifold (X_∞, x_∞) if for every $R > 0$ and $\varepsilon > 0$ there exists $n(R, \varepsilon) \in \mathbb{N}$ such that for all $n \geq n(R, \varepsilon)$ there is a $(1+\varepsilon)$-bilipschitz embedding $f_n : B_R(x_\infty) \to X_n$ satisfying:

(i) $d(f_n(x_\infty), x_n) < \varepsilon$,

(ii) $B_{(1-\varepsilon)\cdot R}(x_n) \subset f_n(B_R(x_\infty))$, and

(iii) $f_n(B_R(x_\infty) \cap \Sigma_\infty) = f_n(B_R(x_\infty)) \cap \Sigma_n$.

Note that the definition also implies the inclusion

$$f_n(B_R(x_\infty)) \subset B_{R(1+\varepsilon)+\varepsilon}(x_n).$$

A standard argument (cf. [BoP, Ch. 3.3]) using the strong local structure of cone 3-manifolds and the controlled decay of injectivity radius (Proposition 3.19) shows that within the spaces $\mathcal{C}_{\kappa,i}$ and $\mathcal{C}_{\kappa,i,a}$ the Gromov-Hausdorff topology and the pointed bilipschitz topology are equivalent. We therefore deduce from Theorem 3.18:

Corollary 3.22. *Let (X_n) be a sequence of cone 3-manifolds with curvatures $k_n \in [-\kappa, \kappa]$, cone angles $\leq \pi$, and possibly with totally geodesic boundary. Suppose that, for some $\rho > 0$, each X_n is ρ-thick at a point $x_n \in X_n$.*

Then, after passing to a subsequence, the pointed cone 3-manifolds (X_n, x_n) converge geometrically to a pointed cone 3-manifold (X_∞, x_∞), with curvature $k_\infty = \lim_{n\to\infty} k_n$.

Note that the case with totally geodesic boundary follows from the closed case by doubling along the boundary.

4. Noncompact Euclidean cone 3-manifolds

A heuristic guideline to describe the geometry of the thin part of cone 3-manifolds, (i.e. the possibilities for the local geometry on a *uniform* small scale) is that global results for noncompact Euclidean cone manifolds correspond to local results for cone manifolds of bounded curvature. For instance, in the smooth case, the fact that there is a short list of noncompact Euclidean manifolds reflects the Margulis lemma for complete Riemannian manifolds of bounded curvature.

We show in this section that there is still a short list of noncompact Euclidean cone 3-manifolds with cone angles $\leq \pi$. The corresponding local results for cone manifolds with bounded curvature will be discussed in Section 5.

THEOREM 4.1 (Classification). *Every noncompact Euclidean cone* 3-*manifold E with cone angles $\leq \pi$ belongs to the following list*:

- *smooth flat* 3-*manifolds*, *i.e. line bundles over the* 2-*torus or the Klein bottle*, *and plane bundles over the circle*;
- *complete Euclidean cones $C_0(\Lambda)$* (*over spherical cone surfaces Λ with cone angles $\leq \pi$*) *which are homeomorphic to S^2*;
- *bundles over a circle or a compact interval with fiber a smooth Euclidean plane or a singular plane $\mathbb{M}_0^2(\theta)$ with $\theta \leq \pi$*;
- $\mathbb{R}$ *times a closed flat cone surface with cone angles $\leq \pi$*; *bundles over a ray with fiber a closed flat cone surface with cone angles $\leq \pi$.*

By *bundles* we mean metrically locally trivial bundles. Line bundles refer to bundles with fiber $\cong \mathbb{R}$. In the case of bundles over a ray or a compact interval, the fibers over the endpoints are singular with index two.

We give a short direct proof of the classification without using general results for nonnegatively curved manifolds such as the Soul Theorem or the Splitting Theorem, although the ideas are of course related. The existence of a soul in our special situation is actually a direct consequence of the list given in Theorem 4.1. Recall that a *soul* is a totally convex compact submanifold of dimension < 3 with boundary either empty or consisting of singular edges with cone angle π.

COROLLARY 4.2. *Every noncompact Euclidean cone* 3-*manifold with cone angles $\leq \pi$ is a metrically locally trivial bundle over a soul with fiber a complete cone*, *or a quotient of such a bundle by an isometric involution.*

In particular, *the soul is a point if and only if E is a cone.*

Before giving the proof of Theorem 4.1 we establish some preliminary lemmas. Since E is noncompact, there are globally minimizing rays emanating from every point $x \in E$. We denote by $R_x \subseteq \Lambda_x E$ the closed set of initial directions of rays starting in x.

Lemma 4.3. (i) *R_x is convex, i.e. with any two directions ξ and η, possibly coinciding, R_x contains all arcs $\overline{\xi\eta}$ of length $< \pi$.*

(ii) *If $x \in \Sigma_E$, every cone point of $\Lambda_x E$ at distance $< \frac{\pi}{2}$ from R_x belongs to R_x.*

Proof. (i) The convexity of the Dirichlet polyhedron $\mathcal{D}(x) \subseteq C_0(\Lambda_x E)$, cf. Section 3.2, implies that R_x is convex.

(ii) Suppose that $\xi \in \Lambda_x E$ is a cone point and η is a point in R_x with $d(\xi, \eta) < \frac{\pi}{2}$.

We consider first the case when the cone angle at ξ is $< \pi$. If $\Lambda_x E$ is the metric suspension of a circle then there exists a loop of length $< \pi$ based at η and surrounding ξ. It follows that ξ is contained in the convex hull of η and hence $\xi \in R_x$. If $\Lambda_x E$ is a spherical turnover, we cut $\Lambda_x E$ open along $\mathrm{Cut}(\xi)$ and obtain a convex spherical polygon with ξ as cone point. Inside the polygon we find a loop as before.

We are left with the case that the cone angle at ξ equals π. Let $\rho_\xi \subset C_0(\Lambda_x E)$ be the singular ray in direction ξ. Observe that, if z is a point on ρ_ξ different from its initial point x, $\overline{yz}$ is a segment perpendicular to ρ_ξ and B is a (small) ball around y, then the convex hull of B in $C_0(\Lambda_x E)$ contains z. Now the ray ρ_η is contained in $\mathcal{D}(x)$. Since $\mathcal{D}(x)$ is convex and has nonempty interior, arbitrarily close to every point of ρ_η we find interior points of $\mathcal{D}(x)$. Our observation therefore implies that $\rho_\xi \subset \mathcal{D}(x)$ and $\xi \in R_x$. $\square$

Let x be a point with $r_{\mathrm{inj}}(x) < \infty$, i.e. $\mathrm{Cut}(x) \neq \emptyset$ and R_x is a proper subset of $\Lambda_x E$. We then have as further restriction on R_x that there exists a direction of angle $\geq \frac{\pi}{2}$ with R_x. This follows from the next result by examination of the shortest segments to the cut locus:

Lemma 4.4. *Suppose that $\zeta \in \Lambda_x E$ is the initial direction of a geodesic loop based at x or of a segment $\overline{xy}$ perpendicular to Σ_E at y. Then $\angle_x(\zeta, R_x) \geq \frac{\pi}{2}$.*

Proof. Let $r : [0, \infty) \to E$ be a ray starting in x. In the case of a loop λ, the assertion follows by applying angle comparison to the isosceles geodesic triangle with λ as one of its side and twice the segment $r|_{[0,t]}$ as the other two sides, and by letting $t \to \infty$. Comparison is applied to the angles adjacent to the nonminimizing side λ.

In the second case, the argument is similar. We consider instead the geodesic triangle with sides $\overline{xy}$, $r|_{[0,t]}$ and a minimizing segment $\overline{yr(t)}$ as third side, and observe that every direction at the singular point y has angle $\leq \frac{\pi}{2}$ with $\overline{xy}$. $\square$

Either of the Lemmas 4.3 or 4.4 implies:

LEMMA 4.5. *If v is a singular vertex with* $\operatorname{diam}(\Lambda_v E) < \frac{\pi}{2}$ *then* $R_v = \Lambda_v E$ *and* $\exp_v$ *is a global isometry; i.e.* $E \cong C_0(\Lambda_v E)$.

With respect to nonvertex singular points, Lemma 4.3 implies:

LEMMA 4.6. *Let $x \in \Sigma_E^{(1)}$. Then either there is a singular ray initiating in x, or all rays emanating in x are perpendicular to σ, where σ is the singular edge of $\Sigma_E^{(1)}$ containing x. In the latter case, if the cone angle at σ is $< \pi$, then every direction in x perpendicular to σ is the initial direction of a ray.*

Proof of Theorem 4.1. The smooth case is well-known, and we assume that the singular locus Σ_E is nonempty.

Part 1: *The case when cone angles are* $< \pi$. If E contains a singular vertex, then E is a cone by Lemmas 3.15 and 4.5. If E contains a closed singular geodesic, then Lemma 4.6 implies that the exponential map is an isometry from the normal bundle of σ onto E, i.e. E is a metrically locally trivial bundle over σ with fiber a plane with cone point. We are left with the case that Σ_E consists of lines, i.e. of complete noncompact geodesics. We assume that E is not a cone; i.e. $r_{\text{inj}} < \infty$ everywhere.

Let σ be a singular edge with cone angle θ. Assume that there exists a ray in E perpendicular to σ in a point x. The singular model space $C_0(\Lambda_x E)$ is isometric to the product $\mathbb{M}_0^2(\theta) \times \mathbb{R}$. Note that $\mathbb{M}_0^2(\theta)$ contains no unbounded proper convex subset because $\theta < \pi$. It follows that $\mathcal{D}(x)$ splits metrically as the product of $\mathbb{M}_0^2(\theta)$ with a closed connected subset I of $\mathbb{R}$. Since E is not a cone, I is a proper subset of $\mathbb{R}$ and $\partial\mathcal{D}(x)$ consists of one or two singular planes $\cong \mathbb{M}_0^2(\theta)$. Under our assumption that cone angles are $< \pi$, the points in $\partial\mathcal{D}(x)$ away from the singular axis project to *smooth* cut points. It follows that σ closes up, contradiction.

Hence there are no rays in E perpendicular to σ. Lemma 4.6 leaves the possibility that from each point $x \in \sigma$ emanates at least one *singular* ray. Let us denote by $A, B \subseteq \sigma$ the sets of initial points of singular rays directed to the respective ends of σ. Both subsets A and B are closed, connected and unbounded. So either they have nonempty intersection or one of them, say A, is empty and $B = \sigma$. In the latter case, σ would be globally minimizing and we obtain a contradiction with $A = \emptyset$. Only the first case is possible; i.e., there exists a point x on σ which divides σ into two rays.

Then $\mathcal{D}(x)$ contains the entire singular axis of $\mathbb{M}_0^3(\theta)$ and, by convexity, it splits as $\mathcal{D}(x) \cong \mathbb{R} \times C_x$ where $C_x \subset M_0^2(\theta)$ denotes the cross section through x. Since E is not a cone, C_x is a *proper* convex subset. It follows that C_x is compact and hence a finite-sided polygon with one cone point. Accordingly, $\partial\mathcal{D}(x)$ consists of finitely many strips of finite width.

Away from the edges the identifications on $\partial\mathcal{D}(x)$ are given by an involutive isometry ι, and on the edges by its continuous extension. It must preserve the direction parallel to the singular axis of $M^3_0(\theta)$. Moreover, ι preserves distance from x. It follows that ι maps ∂C_x onto itself and C_x projects to an embedded totally geodesic closed surface $S \subset E$ with at least one cone point. Due to Gauß-Bonnet, S must be a turnover and is in particular two-sided. It follows that $E \cong \mathbb{R} \times S$.

Part 2: *The general case of cone angles* $\leq \pi$. We expand the above analysis and assume again that E is not a cone; i.e. $r_{\text{inj}} < \infty$ everywhere.

For $x \in E$, let us denote by $\dot{\partial}\mathcal{D}(x)$ the smooth part of the boundary of the Dirichlet polyhedron, i.e. the complement of the edges. The identifications on $\partial\mathcal{D}$ are the continuous extension of an involutive self-isometry ι of $\dot{\partial}\mathcal{D}(x)$. Unlike the case of cone angles $< \pi$, ι may now have fixed points; the fixed point set $\mathrm{Fix}(\iota)$ is a union of segments and projects to the interior points on singular edges with cone angle π which are connected to x by exactly one minimizing segment.

Step 1. Let x be an interior point of a singular edge σ with cone angle $\theta \leq \pi$ and suppose that x is not the initial point of a *singular* ray. Then, starting at x, σ remains in both directions minimizing only for finite time; i.e., $\mathcal{D}(x)$ intersects the singular axis of $C_0(\Lambda_x E) \cong \mathbb{M}^3_0(\theta)$ in a compact subsegment I. By convexity, we have $\mathcal{D}(x) \subseteq I \times \mathbb{M}^2_0(\theta)$; compare the proof of part (ii) of Lemma 4.3. The cross section $C_x \subseteq \mathbb{M}^2_0(\theta)$ of $\mathcal{D}(x)$ perpendicular to I through x is an unbounded convex subset.

Step 1a. If $C_x = \mathbb{M}^2_0(\theta)$, then $\mathcal{D}(x) \cong I \times \mathbb{M}^2_0(\theta)$ and $\partial\mathcal{D}(x)$ consists of two copies of $\mathbb{M}^2_0(\theta)$. The involution ι on $\dot{\partial}\mathcal{D}(x)$ either exchanges the boundary planes or it is a reflection on each of them. By a reflection on the singular plane $\mathbb{M}^2_0(\theta)$ we mean an involutive isometry whose fixed point set is the union of two rays emanating from the cone point into "opposite" directions with angle $\frac{\theta}{2}$. Thus E is a bundle with fiber $\cong \mathbb{M}^2_0(\theta)$ over a circle or a compact interval; in the latter case the fibers over the endpoints of the interval are singular with index two, meaning that they are index-two branched subcovers of the generic fiber.

Step 1b. If C_x is a proper subset of $\mathbb{M}^2_0(\theta)$, then $\theta = \pi$ because C_x is unbounded. There is a unique ray $r \subset \mathcal{D}(x)$ with initial point x. Let H be the half-plane in $\mathbb{M}^3_0(\pi)$ bounded by the singular axis and containing r. Cutting $\mathcal{D}(x)$ open along H yields a convex polyhedron $\mathcal{D}'$ which splits as $\mathcal{D}' \cong \mathbb{R} \times P$ where P denotes the cross section containing I. The cross section P is a compact convex polygon with I as one of its sides and angles $\leq \frac{\pi}{2}$ at both endpoints of I.

The cone manifold E arises from $\mathcal{D}'$ by identifications on the boundary. As before, away from the edges they are given by an isometric involution ι' with one-dimensional fixed point set. The involution ι' carries lines to lines and, since it preserves distance from x, it also preserves ∂P. The fixed point set of ι' consists of midlines of strips in $\mathcal{D}'$ and of edges of P, in our situation including I. After performing the identifications, P becomes a compact totally geodesic cone surface $S \subset E$. The boundary ∂S is a union of singular edges with cone angle π. Every corner of ∂S is the initial point of a singular ray perpendicular to S, and the angle at the corner equals half the cone angle of the ray and hence is $\leq \frac{\pi}{2}$. We obtain that E is a line bundle over S with singular fibers (rays) over the boundary.

The cone manifold E can be described as a bundle in a different way. Let us denote by P_t the cross section $\{t\} \times P$ of $\mathcal{D}'$ where we identify P with P_0. Then, for $t > 0$, the union of the two cross sections $P_{\pm t}$ projects to a totally geodesic closed cone surface $S_t \subset E$. All the surfaces S_t are canonically isometric, say to a surface $\hat{S}$. We see that E fibers over $[0, \infty)$ with fiber $\hat{S}$; the singular fiber over 0 is isometric to S and obtained from $\hat{S}$ by dividing out a reflection.

Step 2. In the following we can assume that each singular point initiates a *singular* ray. As a consequence, all singular edges emanating from singular vertices are rays. If there exists a singular vertex v, Lemma 4.3 implies that $R_v = \Lambda_v E$ and E is a cone, contrary to our assumption. Hence E contains no singular vertices and Σ_E is a union of lines.

As in Part 1 it follows that each singular line σ contains a point x dividing it into two rays, and $\mathcal{D}(x) \cong \mathbb{R} \times C_x$ where C_x is the cross section of $\mathcal{D}(x)$ through x. The section C_x is a proper convex subset of $\mathbb{M}^2_0(\theta)$ where $\theta \leq \pi$ is the cone angle at σ. It is proper because E is not a cone.

Step 2a. If C_x is bounded, then $\partial\mathcal{D}(x)$ is a finite union of strips of finite width. We argue as in Step 1b and obtain that E splits off an $\mathbb{R}$-factor or fibers over a ray.

Step 2b. If C_x is unbounded, then $\theta = \pi$ because C_x is a proper subset of $\mathbb{M}^2_0(\theta)$. Moreover, C_x is a Euclidean surface with one cone point of angle π and one boundary line; it can be constructed from a flat strip by identifying one boundary line to itself by a reflection. Hence $\partial\mathcal{D}(x)$ is a smooth Euclidean plane. The involution ι preserves $d(x, \cdot)$ and therefore fixes the unique point on $\partial\mathcal{D}(x)$ closest to x. It follows that ι is a reflection at a line through x, and E is a plane bundle over a compact interval. Over each endpoint of the interval there is a singular fiber isometric to a half-plane and bounded by a singular line with cone angle π.

The proof of Theorem 4.1 is now complete. $\square$

5. The local geometry of cone 3-manifolds with lower diameter bound

5.1. *Umbilic tubes.* We start by describing certain simple cone manifolds which serve as local models and building blocks for the thin part of arbitrary cone 3-manifolds.

The smooth 3-dimensional model space $\mathbb{M}^3_k$ of constant curvature k can be viewed as the complete k-cone over the unit 2-sphere. More generally, we define for a spherical cone surface Λ the *singular* model space $\mathbb{M}^3_k(\Lambda)$ as the complete k-cone $C_k(\Lambda)$. For the metric suspension $\Lambda(\alpha) := C_1(S^1(\alpha))$ of the circle $S^1(\alpha)$ with length $\alpha < 2\pi$ we obtain the model space $\mathbb{M}^3_k(\alpha) := \mathbb{M}^3_k(\Lambda(\alpha))$ of curvature k with a singular axis of cone angle α. The singular model spaces $\mathbb{M}^3_k(\Lambda)$ serve as local models for cone 3-manifolds; cf. Corollary 3.11.

Recall that an embedded connected surface S in a model space $\mathbb{M}^3_k$ is called *umbilic* if in each point both principal curvatures are equal. It follows that the principal curvatures in all points have the same value which we denote by $\mathrm{pc}(S)$. The local extrinsic geometry of the surface is determined by its principal curvature. Its intrinsic Gauß curvature is given by $k_S = k + \mathrm{pc}(S)^2$. We call S *spherical*, *horospherical*, respectively *hyperspherical*, depending on whether $k_S > 0$, $k_S = 0$ or $k_S < 0$.

The model spaces $\mathbb{M}^3_k$ admit the following *umbilic foliations*, i.e. foliations by umbilic surfaces:

- For all k the *spherical* foliation by distance spheres around a fixed point;
- for $k \leq 0$ the *parabolic foliation* by parallel planes if $k = 0$, respectively by horospheres centered at a fixed point at infinity if $k < 0$;
- for $k < 0$ the *hyperbolic foliation* by equidistant surfaces from a fixed totally geodesic plane.

The leaves of these foliations are spherical, horospherical, respectively hyperspherical.

We proceed to construct certain singular spaces with umbilic foliations. Fix a cone surface S with curvature $k_S \geq k$, and let L be a leaf of an umbilic foliation of $\mathbb{M}^3_k$ with curvature k_S. The type of the foliation depends on the sign of k_S. We can develop the universal cover $\widetilde{S^{\mathrm{smooth}}}$ along L; i.e., there exists an isometric immersion $\mathrm{dev} : \widetilde{S^{\mathrm{smooth}}} \to L$. Let N be a unit normal vector field along L and consider the metric obtained from pulling back the Riemannian metric of model space via

$$\begin{array}{rcl} \widetilde{S^{\mathrm{smooth}}} \times \mathbb{R} & \longrightarrow & \mathbb{M}^3_k \\ (x,t) & \mapsto & \exp(tN(\mathrm{dev}(x))). \end{array}$$

We choose the maximal open interval I containing 0 such that the Riemannian metric on $\widetilde{S^{\text{smooth}}} \times I$ is *nonsingular*. This metric has constant curvature k and descends to $S^{\text{smooth}} \times I$.

Definition 5.1 (Complete tubes). We call the cone 3-manifold resulting from metric completion of $S^{\text{smooth}} \times I$ the *complete k-tube over S* and denote it by $\text{Tube}_k(S)$. We refer to the surfaces in $\text{Tube}_k(S)$ arising as the closures of $S^{\text{smooth}} \times \{t\}$ as *cross sections*.

The tubes have natural foliations by umbilic surfaces equidistant from S; the leaves are homothetic to S. To each cone point of S corresponds a singular edge of $\text{Tube}_k(S)$. If $k_S > 0$, then $\text{Tube}_k(S)$ is just the complete k-cone over $k_S^{-1/2} \cdot S$, i.e. the surface S rescaled by the factor $k_S^{-1/2}$. If $k_S \leq 0$ (and hence $k \leq 0$) then $I = \mathbb{R}$.

Definition 5.2 (Complete cusps, necks and cylinders). We call $\text{Tube}_k(S)$ the *complete k-cusp* over S if $k < k_S = 0$, the *complete k-neck* if $k \leq k_S < 0$, and the *complete* (*Euclidean*) *cylinder* if $k = k_S = 0$.

By an *umbilic tube* we mean a closed connected subset of a complete tube which is a union of leaves of the natural umbilic foliation. We will use the following terminology for different types of umbilic tubes: A *standard ball* is a truncated cone over a spherical cone surface which is homeomorphic to the 2-sphere. A *cusp* is a convex umbilic tube inside a complete cusp which is bounded by one umbilic leaf. A *neck* is a convex umbilic tube inside a complete neck bounded by two umbilic leaves; a neck has a totally geodesic *central leaf.* A *cylinder* is an umbilic tube inside a complete Euclidean cylinder bounded by at most two totally geodesic leaves.

5.2. *Statement of the main geometric results.* The main result of this chapter is the following description of the thin part of cone 3-manifolds with lower diameter bound and cone angles bounded away from π. To simplify the exposition, we will also assume a lower bound on cone angles.

Theorem 5.3 (Thin part). *For $\kappa, D_0 > 0$ and $0 < \beta < \alpha < \pi$ there exist constants $i = i(\kappa, \alpha, D_0, \beta) > 0$, $P = P(\kappa, \alpha, D_0) > 0$ and $\rho = \rho(\kappa, \alpha, D_0, \beta) > 0$ such that*:

Let X be an orientable cone 3*-manifold without boundary which has curvature $k \in [-\kappa, 0)$, cone angles $\in [\beta, \alpha]$ and* $\text{diam}(X) \geq D_0 > 0$. *Then X contains a possibly empty, disjoint union X^{thin} of submanifolds which belong to the following list*:

- *smooth Margulis tubes*: *tubular neighborhoods of closed geodesics and smooth cusps of rank one or two*,

- *tubular neighborhoods of closed singular geodesics*,
- *umbilic tubes* (*i.e. standard balls*, *cusps and necks*) *with turnover cross sections and with strictly convex boundary.*

Furthermore, *the boundary of each component of* X^{thin} *is nonempty*, *strictly convex with principal curvatures* $\leq P$, *and each of its* (*at most two*) *components is thick in the sense that it contains a smooth point with injectivity radius* $\geq \rho$ (*measured in* X); *each component of* X^{thin} *contains an embedded smooth standard ball of radius* ρ; *all singular vertices are contained in* X^{thin}, *and on* $X - X^{\text{thin}}$, $r_{\text{cone-inj}} \geq i$.

The proof will be given in Section 5.7.

We call X^{thin} the *thin part* of X and its components *thin submanifolds* or *Margulis tubes.* Notice that some components of X^{thin} may be balls around singular vertices with thick links; one may argue whether such components should be called thin as well.

We deduce two important consequences of Theorem 5.3 which we will use in the proof of the main theorem.

Corollary 5.4 (Thickness). *There exists* $r = r(\kappa, \alpha, D_0, \beta) > 0$ *such that*: *If* X *is as in Theorem* 5.3 *then* X *is* r*-thick*, *i.e. contains an embedded smooth standard ball of radius* r.

Proof. If $X^{\text{thin}} \neq \emptyset$, we find a thick smooth point on ∂X^{thin}. If $X^{\text{thin}} = \emptyset$, there are no singular vertices and the lower bounds on $r_{\text{cone-inj}}$ and the cone angles imply thickness as well. □

Corollary 5.5 (Finiteness). *Let* X *be as in Theorem* 5.3 *and suppose in addition that* $\text{vol}(X) < \infty$. *Then* X *has finitely many ends and all of them are* (*smooth or singular*) *cusps with compact cross sections. In other words*, X *has a compact core with horospherical boundary.*

Proof. According to Theorem 5.3 each thin submanifold contributes a definite quantum to the volume of X. Thus X^{thin} can have only finitely many components. Finiteness of volume implies moreover that thin submanifolds are compact or cusps with compact cross sections.

Consider a (globally minimizing) ray $r : [0, \infty) \to X$. There is a uniform lower bound on the volume of balls with radius i and centers outside X^{thin}, where i is the constant in Theorem 5.3. Hence, by volume reasons, r enters X^{thin} after finite time. A thin submanifold containing a ray is noncompact and must therefore be a cusp. We conclude that the complement of all cusp components of X^{thin} is compact, because otherwise it would contain a ray which would end up in yet another cusp, a contradiction. □

5.3. *A local Margulis lemma for incomplete manifolds.* The results in this section will be applied to the smooth part of cone manifolds.

Let M be an *incomplete* 3-manifold of constant negative sectional curvature $k \in [-\kappa, 0)$. Our discussion could be carried out for arbitrary curvature sign. However, we restrict to negative curvature for simplicity and because this is the only case needed later.

We recall that the *developing map* is a local Riemannian isometry $\mathrm{dev} : \tilde{M} \to \mathbb{M}^3_k$. It is unique up to postcomposition with an isometry and induces the *holonomy* homomorphism $\mathrm{hol} : \mathrm{Isom}(\tilde{M}) \to \mathrm{Isom}(\mathbb{M}^3_k)$. The action $\Gamma := \pi_1(M) \curvearrowright \tilde{M}$ of the fundamental group by deck transformations on the universal cover transfers, via composition with hol, to *holonomy action* $\Gamma \curvearrowright \mathbb{M}^3_k$. Whereas the deck action is properly discontinuous and free, the holonomy action is in general nondiscrete.

Even though $\tilde{M}$ may have complicated geometry, the next result shows that complete distance balls in $\tilde{M}$ are standard; recall the definitions of various radii from Section 3.2.

LEMMA 5.6. *Let $\tilde{x} \in \tilde{M}$ be a lift of $x \in M$. Then $r_{\mathrm{inj}}(\tilde{x}) = r_{\mathrm{geod}}(x)$.*

Proof. We have $r_{\mathrm{inj}}(\tilde{x}) \le r_{\mathrm{geod}}(\tilde{x}) = r_{\mathrm{geod}}(x)$. The immersion $B_{r_{\mathrm{geod}}(x)}(\tilde{x}) \looparrowright \mathbb{M}^3_k$ into model space given by the developing map must be an isometry onto a round ball. Therefore also $r_{\mathrm{inj}}(\tilde{x}) \ge r_{\mathrm{geod}}(x)$. $\square$

We will use these standard balls in $\tilde{M}$ to localize the usual arguments in the Margulis lemma for complete manifolds of bounded curvature.

For $\delta' > 0$ and for a point $y \in \tilde{M}$ with $r_{\mathrm{geod}}(y) > \delta'$, we define $\Gamma_y(\delta') \subset \Gamma$ as the subgroup generated by all elements γ with $d(\gamma y, y) < \delta'$. It is nontrivial if the corresponding point in M has small injectivity radius. For $r, \delta > 0$ and points $y \in \tilde{M}$ with $r_{\mathrm{geod}}(y) > 2r + \delta$ let us moreover define $A_y(r, \delta) \subset \Gamma$ as the subgroup generated by all elements which have displacement $< \delta$ everywhere on the closed ball $\bar{B}_r(y)$. The definition is made so that, if δ is small compared to r, then the generators of $A_y(r, \delta)$ have small rotational part. The groups Γ_y and A_y are locally semi-constant: for z sufficiently close to y, $\Gamma_z \supseteq \Gamma_y$ holds and $A_z \supseteq A_y$. A pigeonhole argument shows that for sufficiently small $\delta' = \delta'(\kappa, r, \delta) > 0$ we have: $A_y(r, \delta)$ is nontrivial if $\Gamma_y(\delta')$ is.

The standard commutator estimate yields:

PROPOSITION 5.7. *For $R > 0$ there exist constants $r \gg \delta > 0$ also depending on κ such that for every point $y \in \tilde{M}$ with $r_{\mathrm{geod}}(y) > R$ the group $A_y(r, \delta)$ is abelian.*

Remark 5.8. In the more general situation of variable curvature one obtains that the groups are nilpotent. We are using that nilpotent subgroups of $\mathrm{Isom}(\mathbb{M}^3_k)$ are abelian.

We fix $R, r, \delta > 0$ so that 5.7 holds. We define the *thin part* $\tilde{M}^{\text{thin}}$ of $\tilde{M}$ as the open subset of points y with $r_{\text{geod}}(y) > R$ and nontrivial $A_y(r, \delta)$, and the thin part M^{thin} as its projection to M.

There is a natural codimension-one locally homogeneous Riemannian *foliation* on the thin part. This can be seen as follows.

Consider a point $y \in M^{\text{thin}}$, i.e. $r_{\text{geod}}(y) > R$ and $A_y = A_y(r, \delta)$ is nontrivial. Let $A'_y \subset \text{Isom}(\mathbb{M}^3_k)$ be the image of A_y under the holonomy homomorphism hol. Observe that by Lemma 5.6, hol is injective on isometries ϕ with $d(\phi\, y, y) < r_{\text{geod}}(y)$, and hence A'_y is still nontrivial. The group A_y is generated by small elements which in particular preserve orientation. The classification of isometries of $\mathbb{M}^3_k$ implies that A'_y either preserves a unique geodesic (axis) or the horospheres centered at a unique point at infinity. In both cases there is a natural choice of a connected abelian subgroup $H_y \subset \text{Isom}(\mathbb{M}^3_k)$ containing A'_y, namely the identity component of the stabilizer of the axis, respectively, the group of translations along the horospheres. Moreover, there is a corresponding A'_y-invariant locally homogeneous Riemannian foliation $\mathcal{F}_y$ of $\mathbb{M}^3_k$, namely by H_y-orbits. The leaves are equidistant surfaces of the axis or they are horospheres. $\mathcal{F}_y$ pulls back by the developing map to a foliation of $\tilde{M}^{\text{thin}}$ near y. The local semi-constancy of A_y implies that these locally defined foliations fit together to form a natural Γ-invariant foliation $\tilde{\mathcal{F}}$ of $\tilde{M}^{\text{thin}}$. There may be one-dimensional singular leaves, namely geodesic segments in $\tilde{M}$ fixed by small deck transformations; for instance, complete A_y-invariant geodesics project to short, closed geodesics in M. $\tilde{\mathcal{F}}$ descends to a foliation $\mathcal{F}$ of M^{thin}. Note that the regular (two-dimensional) leaves of $\mathcal{F}$ are intrinsically flat and extrinsically strictly convex.

5.4. *Near singular vertices and short closed singular geodesics.* In this section, we make the following general assumption:

Assumption 5.9. We assume that X is a cone 3-manifold of constant curvature $k \in [-\kappa, \kappa]$ with cone angles $\leq \pi$ and $\text{diam}(X) \geq D_0 > 0$.

The following result parallels Lemma 4.5:

LEMMA 5.10 (Thick vertex). *For* $0 < d < \frac{\pi}{2}$ *there exists* $i = i(\kappa, D_0, d) > 0$ *such that*: *If* v *is a singular vertex with* $\text{diam}(\Lambda_v X) \leq d$, *then* $r_{\text{inj}}(v) \geq i$.

Proof. Since $\text{diam}(X) \geq D_0$, there exists a point y with $d(y, v) \geq D_0/2$. Let x be a point in $\text{Cut}(v)$ closest to v. Either x is the midpoint of a geodesic loop l of length $2r_{\text{inj}}(v)$ based at v, or x belongs to a singular edge with cone angle π and there is a (unique) minimizing geodesic segment $s = \overline{vx}$ of length $r_{\text{inj}}(v)$ which is perpendicular to the singular locus at x, cf. our discussion of the cut locus in Section 3.2. In both cases, we have a geodesic triangle $\Delta(v, y, x)$ with $\angle_x(v, y) \leq \frac{\pi}{2}$. By our assumption on the diameter of $\Lambda_v X$

moreover, $\angle_v(y,x) \leq d$ holds. Triangle comparison yields a positive lower bound $i(\kappa, D_0, d)$ for $r_{\text{inj}}(v) = d(v,x)$. □

Remark 5.11. Lemma 5.10 allows us to apply the compactness results of Section 3.4 in many situations, for instance to cone manifolds X with a singular vertex v where the cone angles of at least two adjacent singular edges are $\leq \pi - \varepsilon < \pi$, since in this case $\operatorname{diam}(\Lambda_v X) \leq d(\varepsilon) < \frac{\pi}{2}$; cf. Lemma 3.15.

Definition 5.12. The *normal injectivity radius* of a closed (smooth or singular) geodesic γ is the maximal radius $r_{\text{inj}}(\gamma) \in (0, \infty]$ up to which the exponential map on the normal bundle of γ is defined and is an *embedding*, i.e. for every direction ξ perpendicular to γ and for every $0 < l < r_{\text{inj}}(\gamma)$ there exists a geodesic segment of length l with initial direction ξ which minimizes distance from γ.

Parallel to Lemma 4.6 we have:

LEMMA 5.13 (Normal injectivity radius at short singular circles). *For* $0 < \alpha < \pi$ *there exist* $l = l(\kappa, D_0, \alpha) > 0$ *and* $n = n(\kappa, D_0, \alpha) > 0$ *such that*:

A singular closed geodesic σ *with length* $\leq l$ *and cone angle* $\leq \alpha < \pi$ *has normal injectivity radius* $\geq n$.

Proof. We will choose l smaller than $\frac{D_0}{3}$ and hence can pick a point y at distance $d(y,\sigma) \geq \frac{D_0}{3}$ from σ.

Consider a minimizing segment $\tau = \overline{wy}$ from a point $w \in \sigma$ to y. We apply comparison to the geodesic triangle with sides τ, σ, τ. This can be done although the side σ is of course not minimizing. We obtain (for both angles between σ and τ at w):

$$\angle_w(\sigma,\tau) \geq \frac{\pi}{2} - \varepsilon(\kappa, D_0, l) \tag{1}$$

with $\varepsilon = \varepsilon(\kappa, D_0, l) > 0$ and $\varepsilon \to 0$ as $l \to 0$.

We proceed as in the proof of Lemma 5.10. Let x be a point in $\operatorname{Cut}(\sigma)$ closest to σ. Either x is the midpoint of a segment of length $2r_{\text{inj}}(\sigma)$ which is perpendicular to σ at both endpoints, or x belongs to a singular edge with cone angle π and there is a minimizing segment $\overline{wx}$ of length $r_{\text{inj}}(\sigma)$ perpendicular to σ at w and to Σ_X at x. In both cases there exists a point $w \in \sigma$ and a geodesic triangle $\Delta = \Delta(w,y,x)$ with the properties: (i) $d(w,y) \geq D_0/3$; (ii) $d(w,x) = r_{\text{inj}}(\sigma)$; (iii) $\angle_x(w,y) \leq \frac{\pi}{2}$; and (iv) the side $\overline{wx}$ is perpendicular to σ.

We use property (iv) to bound the angle of Δ at w from above: The link $\Lambda_w X$ at w is the metric suspension of a circle of length $\leq \alpha$, and hence (1) implies $\angle_w(y,x) \leq \alpha/2 + \varepsilon$. By choosing $l = l(\kappa, D_0, \alpha) > 0$ sufficiently small, we can assure, for instance, that (v) $\angle_w(y,x) \leq (\alpha+\pi)/4 < \frac{\pi}{2}$. Triangle comparison using the properties (i)-(v) yields a positive lower bound $n(\kappa, D_0, \alpha)$ for $r_{\text{inj}}(\sigma)$. $\square$

5.5. *Near embedded umbilic surfaces.* In this section, X denotes an orientable cone 3-manifold without boundary which has curvature $k \in [-\kappa, \kappa]$ and cone angles $\leq \pi$. We do not need to assume a lower diameter bound.

Definition 5.14 (Umbilic surface). Suppose that $S \subset X$ is an embedded compact connected surface such that $(S - \partial S) \cap \Sigma_X$ is discrete and ∂S is a union of singular edges with cone angle π. We call the surface S *umbilic* if $S^{\text{smooth}} := S - \partial S - \Sigma_X$ is umbilic.

If S is umbilic it follows that $S - \partial S$ meets the singular locus orthogonally in nonvertex singular points. Moreover ∂S can be nonempty only in the totally geodesic case.

Nearby equidistant surfaces of umbilic surfaces are also umbilic. We say that two compact connected embedded umbilic surfaces in a cone manifold are *parallel* if their union bounds an embedded umbilic tube.

In the first part of our discussion, we make the following assumption. Results in the general case will be deduced afterwards.

Assumption 5.15. Suppose that S is *separating* and *not totally geodesic.*

Since S is not totally geodesic, it is two-sided. It has a *convex* and a *concave side* defined as follows: We say that a locally defined unit normal vector field N along S points to the convex side if the principal curvature of S with respect to N is positive, i.e. if the shape operator DN, defined on tangent spaces to S at smooth points, is a positive multiple of the identity. We call the other side of S concave.

Analogously to the cut locus with respect to a point, comparing our discussion in Section 3.2, one can define the *cut locus* $\text{Cut}(S)$ *with respect to the umbilic surface* S. Let $U(S)$ be the union of S and all half-open geodesic segments $\gamma\colon [0, l) \to X$ emanating from S in orthogonal direction, $\gamma(0) \in S$, and minimizing the distance to S. It is an open subset of X. We call the metric completion $\mathcal{D}(S)$ of $U(S)$ the *Dirichlet domain* relative to S. It canonically embeds into $\text{Tube}_k(S)$ and there is a natural quotient map

$$\phi : \mathcal{D}(S) \longrightarrow X. \tag{2}$$

The *cut locus* $\text{Cut}(S)$ is defined as the complement $X - U(S)$. Since S separates X, each connected component of the cut locus is either a locally finite totally geodesic 2-complex or a point corresponding to a tip of $\text{Tube}_k(S)$ contained in $\mathcal{D}(S)$; with every tip, $\mathcal{D}(S)$ contains the entire component of $\text{Tube}_k(S) - S$.

The upper bound π on cone angles implies that $\mathcal{D}(S)$ is *convex.*

We will denote by $X_{\mathrm{conv}}(S)$, $\mathrm{Cut}_{\mathrm{conv}}(S)$, $\mathcal{D}_{\mathrm{conv}}(S)$ and $\partial_{\mathrm{conv}}\mathcal{D}(S)$ the portions of X, the cut locus, Dirichlet domain and its boundary on the convex side of S, and similarly by $X_{\mathrm{conc}}(S)$, $\mathrm{Cut}_{\mathrm{conc}}(S)$, $\mathcal{D}_{\mathrm{conc}}(S)$ and $\partial_{\mathrm{conc}}\mathcal{D}(S)$ the portions on the concave side.

The next two lemmas concern the component $X_{\mathrm{conc}}(S)$ on the concave side.

Lemma 5.16. *If S is spherical or horospherical, then it bounds a standard ball, respectively a cusp embedded in X.*

Proof. The Dirichlet domain $\mathcal{D}(S)$ is convex and therefore contains the convex hull of S in $\mathrm{Tube}_k(S)$. Since S is not hyperspherical, the convex hull fills out the whole component of $\mathrm{Tube}_k(S)$ on the concave side of S. This is a standard ball or cusp, according to whether S is spherical or horospherical, and it embeds into X via the map (2). □

The umbilic surface S can be hyperspherical only if $k < 0$. In this case we define $\rho = \rho(k, \mathrm{pc}(S))$ as the distance from S to the totally geodesic central leaf L_{central} in $\mathrm{Tube}_k(S)$. We denote by T the umbilic tube between S and L_{central}.

Lemma 5.17. (i) *If S is hyperspherical then $d(S, \mathrm{Cut}_{\mathrm{conc}}(S)) \geq \rho$ and the map* (2) *is an embedding on $T - L_{\mathrm{central}}$. It is an embedding on T if $d(S, \mathrm{Cut}_{\mathrm{conc}}(S)) > \rho$.*

(ii) *Rigidity*: *If $d(S, \mathrm{Cut}_{\mathrm{conc}}(S)) = \rho$, then $\partial_{\mathrm{conc}}\mathcal{D}(S) = L_{\mathrm{central}}$. The map* (2) *restricts on L_{central} to a* 2*-fold ramified covering over* $\mathrm{Cut}_{\mathrm{conc}}(S)$. *The corresponding identifications on L_{central} are given by an orientation-reversing isometric involution τ; its fixed point set is either empty or a piecewise geodesic one-manifold and maps homeomorphically onto the boundary of* $\mathrm{Cut}_{\mathrm{conc}}(S)$ *which is a union of singular edges with cone angle π.*

Proof. (i) T is the closed convex hull of S in $\mathrm{Tube}_k(S)$ and therefore belongs to $\mathcal{D}(S)$. This implies the first part of the assertion.

(ii) Note that as soon as $\mathcal{D}(S)$ contains a neighborhood of a point of L_{central}, then it contains a neighborhood of the entire leaf L_{central} and thus $d(S, \mathrm{Cut}_{\mathrm{conc}}(S)) > \rho$. We are using here that S is connected. Therefore, if $d(S, \mathrm{Cut}_{\mathrm{conc}}(S)) = \rho$, then $L_{\mathrm{central}} = \partial_{\mathrm{conc}}\mathcal{D}(S)$. Thus $X_{\mathrm{conc}}(S)$ arises from T by boundary identifications on L_{central}, and L_{central} maps via (2) onto $\mathrm{Cut}_{\mathrm{conc}}(S)$. It is clear that the identifications on L_{central} arise from an isometric involution τ. It must be orientation-reversing because X is orientable by assumption. □

Now we investigate the cut locus on the convex side of S. Let $S_{k,P}$ be a complete umbilic surface with principal curvature $\mathrm{pc}(S_{k,P}) = P > 0$ in the smooth model space $\mathbb{M}^3_k$, and let y be a point on the convex side at distance $h > 0$ from $S_{k,P}$. Consider the convex hull C of $S_{k,P}$ and y. It is rotationally symmetric, and we define $\psi = \psi(k, P, h) \in (0, \frac{\pi}{2}] \cup \{\pi\}$ as its opening angle, i.e. we set $\psi := \pi$ if y is an interior point of C – which can only happen if $k > 0$ – and define ψ as the radius of the disc $\Lambda_y C$ otherwise.

We are interested in lower bounds for ψ. Since the function $\psi(k, P, h)$ is not monotonic in all variables, $\hat{\psi}(\kappa, P, h){:=}\inf_{-\kappa\leq k\leq\kappa, 0<P'\leq P, 0<h'\leq h} \psi(k, P', h')$. Then for all $\kappa, P > 0$:

$$\lim_{h\to 0} \hat{\psi}(\kappa, P, h) = \frac{\pi}{2}. \tag{3}$$

LEMMA 5.18. *If* $\mathrm{pc}(S) \leq P$ *and if* $x \in \mathrm{Cut}(S)$ *with* $d(x, S) \leq h$, *then the angle at* x *between* $\mathrm{Cut}(S)$ *and any shortest segment from* x *to* S *is* $\geq \hat{\psi} = \hat{\psi}(\kappa, P, h)$. *In particular, the angle at* x *between any two shortest segments to* S *is* $\geq 2\hat{\psi}$.

Proof. A shortest segment from x to S corresponds to a point $\bar{x} \in \partial\mathcal{D}(S)$. Let ξ denote the direction at $\bar{x}$ of the perpendicular to S. The Dirichlet domain $\mathcal{D}(S)$ contains the convex hull of S and $\bar{x}$ which in turn contains, locally at $\bar{x}$, the cone over the disc of radius $\hat{\psi}$ around ξ in $\Lambda_{\bar{x}}\mathrm{Tube}_k(S)$. This shows the first assertion, and the second is a direct consequence. □

We use Lemma 5.18 to bound the number of shortest segments from a point x to S and to rule out branching of the cut locus sufficiently close to S. We obtain the following description of the geometry of $\mathrm{Cut}(S)$ near S:

LEMMA 5.19. *There exists* $h = h(\kappa, P) > 0$ *with the following property: If* $\mathrm{pc}(S) \leq P$ *and if* $x \in \mathrm{Cut}(S)$ *with* $d(x, S) < h$, *then there are at most two shortest segments from* x *to* S.

If there are exactly two shortest segments τ_1 *and* τ_2, *then* $\mathrm{Cut}(S)$ *is totally geodesic near* x. *If in addition* x *is singular, then* $\tau_1 \cup \tau_2$ *forms a singular segment orthogonal to* S *at both endpoints and with* x *as midpoint.*

If there is only one shortest segment τ, *then either* τ *is smooth and* x *is an interior point of a singular edge* σ *with cone angle* $\geq 2\hat{\psi}$, *or* x *is a singular vertex,* τ *a singular segment, and the other two singular segments* σ_1 *and* σ_2 *emanating from* x *have cone angles* $\geq 2\hat{\psi}$. *In the first case,* $\mathrm{Cut}(S)$ *is near* x *a totally geodesic half-disc bounded by* σ; *in the second case it is a sector, that is, the* k*-cone over an arc of length* $\leq \frac{\alpha(\tau)}{2} \leq \frac{\pi}{2}$ *bounded by* σ_1 *and* σ_2, *where* $\alpha(\tau)$ *denotes the cone angle at* τ.

Proof. Using (3), we choose $h > 0$ sufficiently small so that $\hat{\psi}(\kappa, P, h) > \frac{\pi}{3}$. According to Lemma 5.18 any two shortest segments from x to S have angle $> \frac{2\pi}{3}$, and hence by Lemma 3.12 there can be at most two of them.

Regarding the second part, the assertion is clear for smooth points x. Suppose therefore that x is singular and that there are two shortest segments τ_1 and τ_2 from x to S. Since $\mathrm{diam}(\Lambda_x X) > \frac{\pi}{2}$, x cannot be a singular vertex; cf. Lemma 3.14. Hence x lies on a singular edge σ and divides it into singular segments σ_1 and σ_2.

Note that if the metric suspension of a circle of length $\leq \pi$ contains two points with distance $> \frac{2\pi}{3}$, then each pole of the suspension lies within distance $< \frac{\pi}{3}$ of one of the points. Thus, after reindexing if necessary, we have $\angle_x(\sigma_i, \tau_i) < \frac{\pi}{3}$. By Lemma 5.18, the σ_i cannot belong to $\mathrm{Cut}(S)$ near x. Hence $\tau_i \subset \sigma_i$.

Suppose now that there is just one shortest segment τ from x to S. If x is an interior point of a singular edge σ with cone angle β then, near x, the cut locus is a totally geodesic half-disc bounded by σ. The angle between τ and $\mathrm{Cut}(S)$ at x is hence $\leq \frac{\beta}{2}$, and Lemma 5.18 implies $\beta \geq 2\hat{\psi}$.

We are left with the case that x is a singular vertex. By Lemma 5.18, the link $\Lambda_x X$ has injectivity radius $> \frac{\pi}{3}$ at the direction tangent to τ, and an area estimate implies that τ must be singular. (A spherical turnover with cone angles $\leq \pi$ has area $\leq \frac{1}{4}\mathrm{area}(\mathbb{S}^2)$, which equals the area of a smooth spherical disc with radius $\frac{\pi}{3}$. Hence the direction of τ at x cannot be a smooth point of $\Lambda_x X$.) Our previous argument shows that the cone angles at singular points near x and not on τ are $\geq 2\hat{\psi}$. The rest follows. $\square$

COROLLARY 5.20. *There exists* $h = h(\kappa, P) > 0$ *such that*: *If* $\mathrm{pc}(S) \leq P$ *then*, *up to distance* h *from* S, $\mathrm{Cut}(S)$ *is a totally geodesic surface*, *possibly with boundary.*

Next, we observe that, due to the convexity of $\mathcal{D}(S)$, $\mathrm{Cut}(S)$ cannot bend away from S too fast. If S has small diameter, or bounded diameter and small principal curvature, this will force the cut locus to close up as soon as it approaches S sufficiently.

LEMMA 5.21. *For* $h > 0$ *there exist* $d_1 = d_1(\kappa, P, h) > 0$ *and* $\tilde{h} = \tilde{h}(\kappa, P, h) > 0$ *such that*:

If $0 < \mathrm{pc}(S) \leq P$, $\mathrm{diam}(S) \leq d_1$ *and* $d(S, \mathrm{Cut}_{\mathrm{conv}}(S)) \leq \tilde{h}$, *then every segment emanating from* S *in the perpendicular direction to the convex side hits the cut locus within distance* $< h$. *Moreover*, $\mathrm{Cut}_{\mathrm{conv}}(S)$ *is a compact totally geodesic surface*, *possibly with boundary*, *which is entirely contained in the tubular neighborhood* $N_h(S)$ *of radius* h *around* S.

Proof. Suppose that there exists a unit speed segment $\tau : [0, h] \to \mathcal{D}(S)$ of length h emanating in the perpendicular direction to the convex side of S. Moreover, consider another such segment $\tilde{\tau} : [0, l] \to \mathcal{D}(S)$ of length l which connects S to the nearest point of $\partial_{\text{conv}}\mathcal{D}(S)$. We then have $d(\tau(0), \tilde{\tau}(0)) \leq \text{diam}(S) \leq d_1$, and, due to the convexity of $\mathcal{D}(S)$, $\angle_{\tilde{\tau}(l)}(\tilde{\tau}(0), \tau(h)) \leq \frac{\pi}{2}$. The segments τ and $\tilde{\tau}$ are opposite sides of a (two-dimensional) quadrangle Q of constant curvature k embedded in $\text{Tube}_k(S)$; the side connecting $\tau(0)$ and $\tilde{\tau}(0)$ is concave with curvature $\text{pc}(S) \leq P$. Elementary geometry in the models $\mathbb{M}^2_k$ implies: If $d_1 = d_1(\kappa, P, h)$ is chosen small enough, then l can be bounded below by a positive constant $\tilde{h}(\kappa, P, h)$.

The second part of the assertion follows from Corollary 5.20. Namely, we replace h by $h' := \min(h, h(\kappa, P))$, where $h(\kappa, P)$ is taken from Corollary 5.20, and adjust d_1 and $\tilde{h}$ accordingly. □

We need the following variant of Lemma 5.21 for umbilic surfaces with small principal curvatures instead of small diameters:

LEMMA 5.22. *For $d_0 > 0$ and $h > 0$ there exist $P_0 = P_0(\kappa, d_0, h) > 0$ and $\tilde{h} = \tilde{h}(\kappa, d_0, h) > 0$ such that*:

If $0 < \text{pc}(S) \leq P_0$, $\text{diam}(S) \leq d_0$ and $d(S, \text{Cut}_{\text{conv}}(S)) \leq \tilde{h}$, then every segment emanating from S in the perpendicular direction to the convex side hits the cut locus within distance $< h$. Moreover, $\text{Cut}_{\text{conv}}(S)$ is a compact totally geodesic surface, possibly with boundary, which is entirely contained in the tubular neighborhood $N_h(S)$.

Proof. The first part of the assertion is proven as for Lemma 5.21. Note that P_0 and $\tilde{h}$ may be chosen monotonically decreasing as h decreases. Thus, to obtain the second part, we may decrease h, if necessary, below the value $h(\kappa, P_0)$ from Corollary 5.20, and then decrease P_0 and $\tilde{h}$ accordingly. □

We now suppose in addition that X has cone angles $\leq \alpha < \pi$. The results discussed above then simplify.

If S is hyperspherical, and hence $k < 0$, we obtain on the concave side:

ADDENDUM 5.23 (to Lemma 5.17). (i) *In the rigidity case of* 5.17, $\text{Cut}_{\text{conc}}(S)$ *is a closed nonorientable totally geodesic surface, and the natural map $S \cong L_{\text{central}} \to \text{Cut}_{\text{conc}}(S)$ is a regular two-fold covering.*

(ii) *There exists $d_0 = d_0(\kappa, \alpha) > 0$ such that*: *If* $\text{diam}(S) \leq d_0$, *then the rigidity case in* 5.17 *cannot occur*, *i.e. the tube T embeds.*

Proof. (i) The orientation-reversing involution τ on L_{central} cannot have fixed points because there are no singular edges with cone angle π.

(ii) Let us assume the rigidity case. We have $\text{diam}(\text{Cut}_{\text{conc}}(S)) < \text{diam}(S)$. We apply the Gauß-Bonnet Theorem to $\text{Cut}_{\text{conc}}(S)$ and note that, for d_0 sufficiently small, the contribution of its smooth part to the curvature integral is a small negative number, say $\in (\alpha - \pi, 0)$, and the contribution of each cone point belongs to the interval $[2\pi - \alpha, 2\pi)$. Since $\mathrm{Cut}_{\mathrm{conc}}(S)$ is nonorientable, it must be a projective plane and have curvature integral 2π. But with one cone point, the curvature integral would amount to $< 2\pi$, and with at least two cone points to $> (\alpha - \pi) + 2(2\pi - \alpha) > 2\pi$. We get a contradiction. $\square$

On the convex side of S, Lemma 5.19 implies for the cut locus near S:

Lemma 5.24. *There exists $h = h(\kappa, P, \alpha) > 0$ such that: If $\mathrm{pc}(S) \leq P$ then, up to distance h from S, $\mathrm{Cut}(S)$ is totally geodesic without boundary.*

Proof. We choose h sufficiently small so that $\hat{\psi}(\kappa, P, h) > \alpha/2$; compare (3). This rules out in 5.19 the possibility of cut points near S with a unique minimizing segment to S. $\square$

From Lemma 5.21 on the closing up of the cut locus near umbilic surfaces of small diameter we deduce:

Lemma 5.25. *There exist $d_2 = d_2(\kappa, P, \alpha) > 0$ and $\tilde{h} = \tilde{h}(\kappa, P, \alpha) > 0$ such that: If $\mathrm{pc}(S) \leq P$ and $\mathrm{diam}(S) \leq d_2$, then $d(S, \mathrm{Cut}_{\mathrm{conv}}(S)) > \tilde{h}$.*

Proof. We use the constant $h = h(\kappa, P, \alpha)$ from Lemma 5.24 and accordingly the constants $d_1 = d_1(\kappa, P, h) = d_1(\kappa, P, \alpha)$ and $\tilde{h} = \tilde{h}(\kappa, P, h) = \tilde{h}(\kappa, P, \alpha)$ from Lemma 5.21.

Suppose that $\mathrm{diam}(S) \leq d_2$ and $d(S, \mathrm{Cut}_{\mathrm{conv}}(S)) \leq \tilde{h}$. Lemmas 5.24 and 5.21 imply for $d_2 \leq d_1$ that $\mathrm{Cut}_{\mathrm{conv}}(S)$ is a closed totally geodesic surface contained in $N_h(S)$. Then $\partial_{\mathrm{conv}}\mathcal{D}(S)$ is a closed totally geodesic surface as well, and it follows that $k > 0$.

Since $\mathrm{Cut}_{\mathrm{conv}}(S)$ is nonorientable, the Gauß-Bonnet theorem and the upper cone angle bound π imply that it is a projective plane with at most one cone point. Hence it is an index two subcover of the complete k-cone of a circle of length $\leq \alpha$ and has diameter $\frac{\pi}{2\sqrt{k}}$. On the other hand, $\mathrm{diam}(\mathrm{Cut}_{\mathrm{conv}}(S)) < \mathrm{diam}(S) + 2h \leq d_2 + 2h$. This yields a contradiction if d_2 is chosen small enough. $\square$

We drop now our assumption 5.15 that S separates X and is not totally geodesic. On the other hand, we restrict to the case of negative curvature and impose an upper cone angle bound $< \pi$. We are interested in the situation when S has small diameter and controlled principal curvature. Our discussion above leads to the following description of the geometry near such surfaces, which is the main result of this section:

Proposition 5.26 (Neighborhoods of umbilic surfaces with small diameter). *For $\kappa, P > 0$ and $\alpha < \pi$ there exists $d = d(\kappa, P, \alpha) > 0$ such that:*

Let X be an orientable cone 3-manifold without boundary which has curvature $k \in [-\kappa, 0)$ and cone angles $\leq \alpha$. Suppose that $S \subset X$ is a (not necessarily separating) umbilic surface with $0 \leq \mathrm{pc}(S) \leq P$ and $\mathrm{diam}(S) < d$. Then S is an umbilic leaf in an embedded umbilic tube $T \subset X$ with convex boundary and the property that each of its at most two boundary components has diameter d.

Remark 5.27. Note that for d sufficiently small, S is either horospherical or a turnover. This follows by applying Gauß-Bonnet as in the proof of Addendum 5.23. In particular, S is always two-sided.

Proof. Step 1. Suppose that S *separates* and is *not totally geodesic*. We choose d smaller than the constant $d_0(\kappa, \alpha)$ in 5.23. By combination of 5.16, 5.17 and 5.23, there exists an embedded umbilic tube $T_0 \subset X$ with S as boundary component and the following properties: T_0 is a ball if S is spherical, a cusp if S is horospherical, and a neck if S is hyperspherical. T_0 has strictly convex boundary with at most two components. Their principal curvatures are $< \sqrt{\kappa}$ if S is hyperspherical and $\leq \mathrm{pc}(S)$ otherwise; hence in all cases they are $\leq P' = P'(\kappa, P) := \max(P, \sqrt{\kappa})$. We may also assume that their diameters are $< d$.

We decrease d below the constant $d_2(\kappa, P', \alpha)$ from 5.25. By applying Lemma 5.25 to the boundary components of T_0 and repeating this procedure finitely many times, we obtain that T_0 can be enlarged to an embedded tube T whose boundary components have diameter $\geq d$ and principal curvature $< P'$.

Step 2. Suppose now that S *does not separate* but still is *not totally geodesic*. Consider the cyclic covering $p : \hat{X} \to X$ associated to the homomorphism $\pi_1(X) \to \mathbb{Z}$ given by the oriented intersection number with S. Any connected component $\hat{S}$ of $p^{-1}(S)$ is an umbilic surface isometric to S which separates $\hat{X}$, and our previous discussion applies. First of all, neither component of $\hat{X}$ obtained by cutting along $\hat{S}$ is a ball or cusp, and thus S is hyperspherical. Furthermore, by Step 1, $\hat{X}$ contains an embedded neck $\hat{T}$ with $\hat{S}$ as an umbilic leaf and boundary components of diameter d.

SUBLEMMA 5.28. *There exists $d_3 = d_3(\kappa, P', \alpha) > 0$ such that: Any two separating umbilic surfaces $S_1, S_2 \subset \hat{X}$ with $\mathrm{pc}(S_i) \leq P'$ and $\mathrm{diam}(S_i) \leq d_3$ are disjoint or coincide.*

Proof. We choose $d_3 \leq \min(d_2, \tilde{h})$ with the constants from 5.25. Suppose that S_1 and S_2 are not disjoint. Then S_1 is contained in $N_{d_3}(S_2) =: Z$ which, by Step 1, is an umbilic tube.

The tube Z, or more precisely the universal cover of Z^{smooth}, develops into a layer of width $2d_3$ in model space $\mathbb{M}^3_k$ bounded by two leaves L_1 and L_2 of an umbilic foliation $\mathcal{F}_{\mathrm{model}}$. The universal cover of S_1^{smooth} develops along a complete umbilic surface U and leaves out at most a discrete set. It follows that U is contained in the layer. If d_3 is sufficiently small, U cannot bound a ball contained in the layer, because U has principal curvature $\leq P'$. Thus U separates the L_i. In the case when the foliation $\mathcal{F}_{\text{model}}$ is not spherical, this already means that U must be one of its leaves, i.e. is parallel to the L_i.

If $\mathcal{F}_{\text{model}}$ is spherical, then L_1, L_2 and also U are round spheres, and we need one more observation to see that U is concentric with the L_i. We consider the function $f = d(L_1, \cdot)$ on model space. Since the development of the universal cover of S_1^{smooth} into U is equivariant with respect to its deck group, the restriction of f to U must have a minimum and maximum point within distance $\leq d_3$. This forces U to be concentric with the L_i if d_3 is small enough. It then follows that S_1 and S_2 are parallel and thus coincide. $\square$

We decrease d further so that $d \leq d_3$. All umbilic leaves of $\hat{T}$ have diameter $\leq d$ and principal curvature $< \sqrt{\kappa} \leq P'$. Sublemma 5.28 therefore implies, that any two translates of $\hat{T}$ by a nontrivial deck transformation of $\hat{X} \to X$ are disjoint. It follows that $\hat{T}$ projects to an embedded neck in X around S, and we are done in this case, too.

Step 3. Finally assume that S is *totally geodesic.* If d is sufficiently small, our assumptions that S is two-sided apply; cf. Remark 5.27. We deduce the claim by applying the above discussion to nearby equidistant surfaces of S.

This completes the proof of Proposition 5.26. $\square$

5.6. *Finding umbilic turnovers.* As in Section 5.5, let X denote an orientable cone 3-manifold without boundary which has curvature $k \in [-\kappa, \kappa]$ and cone angles $\leq \pi$.

We are interested in conditions which imply the existence of umbilic turnovers with small diameter and controlled principal curvature. We will find them as cross sections to minimizing singular segments with cone angle bounded away from π in regions of small injectivity radius; cf. our main result Proposition 5.33.

We start with some observations about the geometry near the middle of minimizing segments in X which express aspects of an almost product structure.

LEMMA 5.29. *For $d, \varepsilon > 0$ there exists $l(\kappa, d, \varepsilon) > 0$ such that:*

Let λ be a (not necessarily shortest) geodesic loop of length $\leq l$ based at x, and let τ be a minimizing segment of length $\geq d$ with x as initial point. Then $\angle_x(\tau, \lambda) \geq \frac{\pi}{2} - \varepsilon$.

This means that the angle of τ with *both* initial directions of λ is $\geq \frac{\pi}{2} - \varepsilon$.

Proof. The assertion follows by application of angle comparison to the triangle with sides τ, λ, τ. This triangle has two minimizing sides, namely twice τ, and the nonminimizing side λ. We apply comparison at the two angles adjacent to λ. $\square$

Lemma 5.30. *For $0 < L < \frac{\pi}{\sqrt{\kappa}}$ and $\varepsilon > 0$ there exists $e = e(\kappa, L, \varepsilon) > 0$ such that*:

If $q_\pm$, $x \in X$ are points satisfying $d(q_-, q_+) = L$ and $d(x, q_\pm) \geq \frac{L}{4}$ and if $\alpha_\pm = \overline{xq_\pm}$ are not necessarily minimizing *geodesic segments such that*

$$\mathrm{length}(\alpha_+) + \mathrm{length}(\alpha_-) \leq d(q_-, q_+) + e, \tag{4}$$

then $\angle_x(\alpha_+, \alpha_-) > \pi - \varepsilon$.

Restricting to nonpositive curvature, we do not need an upper bound for L.

Proof. Step 1. Suppose that one of the segments $\alpha_\pm$ is minimizing. We then can apply angle comparison to the geodesic triangle $\Delta(q_-, x, q_+)$ which has $\alpha_\pm$ as two of its sides and as third side a minimizing segment $\overline{q_-q_+}$. Inequality (4) implies that the comparison angle at x is $\geq \phi(\kappa, L, e)$ with $\lim_{e \to 0} \phi(\kappa, L, e) = \pi$. By choosing e sufficiently small, we obtain that $\angle_x(\alpha_+, \alpha_-) \geq \pi - \frac{\varepsilon}{3}$.

Step 2. The general case can be deduced by considering minimizing segments $\bar{\alpha}_\pm$ from x to $q_\pm$. Then we have $\angle_x(\bar{\alpha}_-, \bar{\alpha}_+) \geq \pi - \frac{\varepsilon}{3}$ and $\angle_x(\bar{\alpha}_\pm, \alpha_\mp) \geq \pi - \frac{\varepsilon}{3}$. It follows that $\angle_x(\bar{\alpha}_\pm, \alpha_\pm) \leq \frac{2\varepsilon}{3}$ and $\angle_x(\alpha_+, \alpha_-) \geq \pi - \varepsilon$, as desired. $\square$

We now focus our attention on *singular* minimizing segments and investigate the cut locus with respect to their midpoints.

Lemma 5.31. *For $L, \varepsilon > 0$ there exists $h = h(\kappa, L, \varepsilon) > 0$ such that*:

Let p be the midpoint of a minimizing singular segment $\sigma = \overline{q_-q_+}$ of length $\geq L$. Let x be a point with $d(p, x) \leq h$. Suppose that there are at least three minimizing segments $\overline{px}$, or that x is singular and there are at least two minimizing segments $\overline{px}$.

Then the cone angle at σ is $\geq \pi - \varepsilon$.

Proof. We may assume that $L < \frac{\pi}{\sqrt{\kappa}}$, say $L \leq \frac{\pi}{2\sqrt{\kappa}}$. This is relevant only in the positive curvature case.

Step 1. We denote the cone angle at σ by θ. The Dirichlet polyhedron $\mathcal{D}(p)$ associated to p can be regarded as a convex polyhedron in the model space $\mathbb{M}^3_k(\theta)$ with singular axis of cone angle θ; cf. Section 3.2.

The minimizing segments σ_i from p to x correspond to points $\bar{x}_i$ in $\partial\mathcal{D}(p)$. Each of the segments σ_i determines a so-called Voronoi cell V_i in the link $\Lambda_x X$.

By definition, V_i consists of those directions at x whose angle with σ_i is strictly less than the angle with all other minimizing segments from x to p. V_i is an open convex spherical polygon. The link $\Lambda_{\bar{x}_i}\mathcal{D}(p)$ is canonically identified with the closed spherical polygon $\hat{V}_i$ arising as the metric completion of V_i.

The Dirichlet polyhedron $\mathcal{D}(p)$ contains at least the subsegment $\overline{q_+q_-}$ of the singular axis, and maybe more. Inside $\mathcal{D}(p)$ there are unique geodesic segments $\overline{\bar{x}_i q_\pm}$. Note that the corresponding segments in X need not be minimizing. However, they are almost minimizing and their initial directions $\eta_{\pm,i} \in \overline{V}_i$ at x are almost gradient directions for the distance functions $-d(q_\pm, \cdot)$. Namely, Lemma 5.30 implies that for any $\varepsilon_1 > 0$ (a constant to be fixed later) we have

$$\angle(\eta_{+,i}, \eta_{-,i}) > \pi - \varepsilon_1 \tag{5}$$

if $h = h(\kappa, L, \varepsilon_1) > 0$ is chosen sufficiently small.

For each pair of points $\bar{x}_i$ and $\bar{x}_j$, $i \neq j$, there are several geodesic segments in $\mathcal{D}(p)$ connecting them. They correspond to loops in X with base point x. At least two of the segments $\overline{\bar{x}_i\bar{x}_j}$ have length $< 2h$. Combining Lemmas 5.29, 5.30 and the proof of the latter one, we obtain, after further decreasing h if necessary, that such segments $\overline{\bar{x}_i\bar{x}_j}$ have angle $\in (\frac{\pi}{2} - 3\varepsilon_1, \frac{\pi}{2} + 3\varepsilon_1)$ with both segments $\overline{\bar{x}_i q_\pm}$. In this sense these segments $\overline{\bar{x}_i\bar{x}_j}$ are *almost horizontal*, where we regard the singular axis as vertical.

Step 2. Let us consider the case that x is smooth and there exist (at least) three minimizing segments $\sigma_1, \sigma_2, \sigma_3$ between p and x.

We next construct an almost horizontal geodesic triangle Δ in $\mathcal{D}(p)$, with vertices $\bar{x}_i$, which winds once around the singular axis. Each point $\bar{x}_i$ lies on a half-plane, or hemisphere if $k > 0$, H_i in $\mathbb{M}^3_k(\theta)$ is bounded by the singular axis. Notice that the H_i are pairwise different because of the horizontality of the segments $\overline{\bar{x}_i\bar{x}_j}$ and the convexity of the Dirichlet polyhedron. Between $\bar{x}_i$ and $\bar{x}_j$ we choose the segment which does not intersect the third half-plane (hemisphere) H_{6-i-j}. The resulting triangle Δ is contained in $\mathcal{D}(p)$ by convexity. Notice that its sides are in general not minimizing. Each side $\overline{\bar{x}_i\bar{x}_j}$ of Δ determines directions $\zeta_{ij} \in \overline{V}_i$, respectively, $\hat{\zeta}_{ij} \in \hat{V}_i$. We saw in Step 1 that

$$|\angle_x(\eta_{\pm,i}, \hat{\zeta}_{ij}) - \frac{\pi}{2}| < 3\varepsilon_1. \tag{6}$$

Let ϕ_i denote the angle of the triangle Δ at the vertex $\bar{x}_i$ measured in $\mathcal{D}(p)$, i.e. the angle between the two directions $\hat{\zeta}_{ij} \in \hat{V}_i$, $j \neq i$. In view of (5) and (6) we conclude that the convex spherical polygon V_i almost contains a bigon with angle ϕ_i. Hence, for ε_1 (and accordingly h) sufficiently small, $\mathrm{area}(V_i) \geq 2\phi_i - \varepsilon/3$. Since $\sum_i \mathrm{area}(V_i) \leq \mathrm{area}(\Lambda_x X) = 4\pi$, we obtain

$$\text{angle sum}(\Delta) \leq 2\pi + \varepsilon/2.$$

On the other hand, Δ is almost horizontal in $\mathcal{D}(p)$. We embed the 2-dimensional singular model $\mathbb{M}^2_k(\theta)$ as a cross section of $\mathbb{M}^3_k(\theta)$ so that it contains, say, $\bar{x}_1$. The nearest point projection of Δ to $\mathbb{M}^2_k(\theta)$ almost preserves angles. Due to Gauß-Bonnet, horizontal triangles with small diameter have angle sum $\simeq \pi + (2\pi - \theta)$, and therefore

$$\text{angle sum}(\Delta) \geq 3\pi - \theta - \varepsilon/2$$

if h is sufficiently small. It follows that $\theta \geq \pi - \varepsilon$ as claimed.

Step 3. The argument is analogous in the case when x is singular and there are two minimizing segments between p and x: Δ becomes an almost horizontal bigon winding once around the singular axis; such bigons have angle sum $\simeq 2\pi - \theta$, i.e. $\geq 2\pi - \theta - \frac{\varepsilon}{2}$ for h sufficiently small; on the other hand, since $\text{area}(\Lambda_x X) \leq 2\pi$, the angle sum must be $\leq \pi + \frac{\varepsilon}{2}$; therefore $\theta \geq \pi - \varepsilon$, as claimed. □

Corollary 5.32. *For $L > 0$ and $0 < \alpha < \pi$ there exists $h = h(\kappa, L, \alpha) > 0$ such that*:

If p is the midpoint of a minimizing singular segment of length $\geq L$ and cone angle $\leq \alpha$, then the ball $B_h(p)$ contains no point x with at least three minimizing segments between x and p, and no singular point x with at least two minimizing segments between x and p. The intersection $\text{Cut}(p) \cap B_h(p)$ is a totally geodesic surface whose boundary (if nonempty) is geodesic and consists of singular segments.

Notice that there are no singular vertices close to p, because the links at vertices have diameter $\leq \frac{\pi}{2}$ (Lemma 3.14) whereas the links at points near p have almost diameter π (Lemma 5.30).

We come to the main result of this section.

Proposition 5.33 (Umbilic cross sections). *For $L, d > 0$ and $0 < \alpha < \pi$ there is $i = i(\kappa, \alpha, L, d) > 0$ such that*:

Let p be the midpoint of a minimizing singular segment σ of length $\geq L$ and cone angle $\leq \alpha$, and assume that $r_{\text{inj}}(p) < i$. Then there exists an umbilic turnover S through p with $\text{diam}(S) \leq d$.

Remark 5.34. Lemma 5.16 implies that the principal curvature of the cross section S is bounded in terms of κ and L. Namely, if $\text{pc}(S)$ were too large, then S would bound a singular ball of radius $< L/2$.

Proof. Let $q_\pm$ denote the endpoints of σ and θ its cone angle. We study the Dirichlet polyhedron $\mathcal{D}(p)$ which we regard as a convex polyhedron in $\mathbb{M}^3_k(\theta)$.

Step 1. *Near p the edges of $\mathcal{D}(p)$ are almost vertical.* By 5.32, near p the interior points on faces of $\mathcal{D}(p)$ correspond to smooth cut points, and the points on boundary edges correspond to singular cut points. (We must also allow degenerate edges with dihedral angle π.) Let $\bar{x}$ be a point on a boundary edge $\bar{\gamma}$ of $\mathcal{D}(p)$ with $d(p, \bar{x}) < h$, h as in 5.32, and let x be the corresponding point in X which lies on a singular edge γ. The segments $\overline{xq_\pm}$ in X corresponding to the segments $\overline{\bar{x}q_\pm}$ have angle $\simeq \pi$ by 5.30. This forces their directions at x to be close to the singular poles of $\Lambda_x X$. This means that $\bar{\gamma}$ is almost vertical.

Step 2. *The cross sections of* $\mathcal{D}(p)$ *are small.* By assumption there exists a cut point $y \in \mathrm{Cut}(p)$ with $d(p, y) < i$. Let $\bar{y}$ be a corresponding point in $\partial\mathcal{D}(p)$. The cross section $C_{\bar{y}}$ of $\mathcal{D}(p)$ through $\bar{y}$ and perpendicular to σ is a convex polygon with cone point of angle θ. Observe that, since $\theta < \pi$, the circumradius and the inradius of the polygon control each other. Hence, for $i > 0$ sufficiently small, the polygon $C_{\bar{y}}$ has small diameter $<< h$. Notice that $C_{\bar{y}}$ has at least one vertex.

Step 3. *The cross sections are bigons.* Our discussion implies that, near p, cross sections of $\mathcal{D}(p)$ are compact convex polygons and $\partial\mathcal{D}(p)$ is a union of finitely many geodesic strips which are almost vertical. According to 5.32, the boundary identifications on $\partial\mathcal{D}(p)$ inside $B_h(p)$ are given by an involutive isometry $\iota : \partial\mathcal{D}(p) \cap B_h(p) \to \partial\mathcal{D}(p) \cap B_h(p)$ which fixes the boundary edges. It follows that there are exactly two strips which are exchanged by ι.

Step 4. $\mathcal{D}(p)$ *is rigid.* Let v_1 and v_2 be the vertices of the cross section $C_{\bar{y}}$, τ and τ' its sides, and let σ_i be the edge of $\partial\mathcal{D}(p)$ through v_i. We denote by H_i the half-plane (respectively, hemisphere if $k > 0$) in $\mathbb{M}^3_k(\theta)$ bounded by the singular axis and containing v_i. Since τ and τ' are exchanged by the isometry ι, they must have the same length. It follows that H_1 and H_2 are opposite to each other in the sense that they meet at σ with angle $\frac{\theta}{2}$. Furthermore, each edge σ_i has equal angles with τ and τ' and thus $\sigma_i \subset H_i$. It follows that ι extends to an isometry of the whole model space $\mathbb{M}^3_k(\theta)$, namely to the reflection at $H_1 \cup H_2$.

Step 5. *Conclusion.* Consider the unique umbilic foliation $\mathcal{F}$ of $\mathbb{M}^3_k(\theta)$ orthogonal to σ and σ_1. It is also orthogonal to the boundary strips of $\mathcal{D}(p)$ and hence to σ_2. Let L denote the leaf of $\mathcal{F}$ through p. Then $L \cap \mathcal{D}(p)$ projects to an umbilic turnover S in X. If i is chosen sufficiently small, we have $\mathrm{diam}(S) < d$. □

5.7. *Proof of Theorem* 5.3: *Analysis of the thin part.* In this section we combine the previous results to analyze the thin part of X. The proof of Theorem 5.3 is organized in five steps:

Step 1. *Around singular vertices.* Let $v \in X$ be a singular vertex. The space of directions $\Lambda_v X$ then has diameter $\leq D(\alpha) < \frac{\pi}{2}$; cf. Lemma 3.15. By the Thick Vertex Lemma 5.10, v is the center of an embedded (closed) standard ball with radius $r_1(\kappa, \alpha, D_0) > 0$. To make the balls around the various vertices disjoint, we define B_v as the closed ball of radius $\frac{r_1}{2}$ centered at v. The umbilic boundary spheres ∂B_v are convex and have principal curvature $\leq P_1(\kappa, r_1) = P_1(\kappa, \alpha, D_0) > \sqrt{\kappa}$.

Step 2. *Organizing small umbilic turnovers.* Let us now consider the umbilic turnovers $S \subset X$ with $\mathrm{pc}(S) < P_1$. If $d_2 = d_2(\kappa, \alpha, P_1) = d_2(\kappa, \alpha, D_0) > 0$ is chosen small enough, then according to Proposition 5.26 any such turnover S with diameter $< d_2$ is a leaf in the natural foliation of an embedded umbilic tube $T_S \subset X$. Moreover, T_S has one or two boundary components which are strictly convex with diameter d_2.

The argument used to prove 5.28 shows that, after decreasing d_2 sufficiently, any two turnovers S in consideration are either disjoint or coincide. The same holds then for the tubes T_S. It shows as well that the T_S are disjoint from the balls B_v with $\mathrm{diam}(\partial B_v) > d_2$. On the other hand, the singular balls B_v with $\mathrm{diam}(\partial B_v) \leq d_2$ are contained in a tube T_S. In the following, we forget about the balls B_v contained in tubes T_S. Denoting by V_1 the union of the remaining balls B_v and the tubes T_S, we see that they are pairwise disjoint.

Notice that for umbilic turnovers as considered here, i.e. with cone angles $\leq \alpha$, controlled principal curvature $\leq P_1$ and small upper diameter bound d_2, diameter and thickness control each other. This is seen as in the proof of Lemma 3.16; the lower bound on the cone angles follows from Gauß-Bonnet and the fact that the Gauß curvature is bounded. As a consequence, there is a lower bound for the thickness of the components of ∂V_1. We also get a lower bound for the thickness of the components of V_1 by leaving out "short" umbilic necks, i.e. necks with central leaves of diameter, say, $> \frac{d_2}{2}$.

Step 3. *Around short closed singular geodesics.* We choose $l_1 = l_1(\kappa, \alpha, D_0) > 0$ small enough so that 5.13 implies that the normal injectivity radius of closed singular geodesics γ with period $\leq 2l_1$ is $> 2n_1(\kappa, \alpha, D_0) > 0$. The closed tubular n_1-neighborhoods $\bar{N}_{n_1}(\gamma)$ around these geodesics are then pairwise disjoint, and we denote their union by V_2. Note that the injectivity radii of their boundaries are everywhere $\leq i_1(\kappa, l_1, n_1)$ with $\lim_{l_1 \to 0} i_1(\kappa, l_1, n_1) = 0$.

By choosing l_1 sufficiently small, we may achieve that

$$V_1 \cap V_2 = \emptyset.$$

This can be seen as follows: The singular closed geodesics of period $\leq 2l_1$ lie outside V_1. If $V_1 \cap V_2$ were nonempty, then an embedded smooth 2-torus T (equidistant to a short singular closed geodesic) with controlled principal curvatures and small area would intersect (touch) an umbilic turnover S' with controlled principal curvature and lower diameter bound. Moreover, the distance between T and Σ_X would be bounded away from zero so that, from the thickness of S', we would get a lower bound on the injectivity radius at the touching point. This contradicts the thinness of the torus T.

Step 4. *Bounding the injectivity radius on the rest of the singular locus.* First we show that, in the spirit of Lemma 4.6, a singular edge either closes up with short period or minimizes up to a certain length.

SUBLEMMA 5.35. *There exists* $l = l(\kappa, \alpha, D_0) > 0$ *such that for every* $l' \leq l$: *If the singular edge* σ *does not close up with period* $\leq 2l'$ *then, for every point* $x \in \sigma$, *there is a minimizing subsegment of length* $> l'$ *with* x *as initial point.*

Proof. Let $\sigma_1 = \overline{xy_1}$ and $\sigma_2 = \overline{xy_2}$ be the maximal minimizing subsegments of σ emanating from x in the two antipodal singular directions. Suppose that both have length $\leq l'$. Due to our diameter assumption, there exists a segment $\tau = \overline{xz}$ of length $\geq D_0/2$ starting in x.

We regard σ_1, σ_2, τ also as segments in the Dirichlet polyhedron $\mathcal{D}(x)$ and denote their respective endpoints on $\partial\mathcal{D}(x)$ by $\bar{y}_1, \bar{y}_2, \bar{z}$.

Since $\mathcal{D}(x)$ has a singular axis with cone angle $\leq \alpha$, the convex hull $C(\bar{z})$ of $\bar{z}$ in $\mathcal{D}(x)$ is a totally geodesic disc which intersects σ orthogonally in a cone point c and which has geodesic boundary with a corner at $\bar{z}$. The convexity of the Dirichlet polyhedron implies that $\angle_{\bar{y}_i}(x, \bar{z}) \leq \frac{\pi}{2}$. By consideration of the geodesic triangle $\Delta(x, \bar{y}_i, \bar{z})$ in $\mathcal{D}(x)$ it follows that $\angle_x(\tau, \sigma_i) \geq \phi(\kappa, D_0, l)$ with $\phi \to \frac{\pi}{2}$ as $l \to 0$. We obtain furthermore that $C(\bar{z})$ contains a disc of radius $r(\kappa, \alpha, D_0, l) > 0$ centered at c.

We investigate $\partial\mathcal{D}(x)$ near the singular axis. For a point $\bar{w} \in \partial\mathcal{D}(x)$ near $\bar{y}_i$, say with $d(x, \bar{w}) \leq 2l$, we consider the convex hull of $\bar{w}$ and $\bar{z}$. It is contained in $\mathcal{D}(x)$, and we obtain that the space of directions $\Lambda_{\bar{w}}\mathcal{D}(x)$ contains a standard disc of radius $\rho(\kappa, \alpha, D_0, l) > 0$. Moreover, $\lim_{l\to 0} \rho(\kappa, \alpha, D_0, l) = \frac{\pi}{2}$. Note hereby that $r(\kappa, \alpha, D_0, l)$ increases as $l \to 0$ and in particular remains uniformly bounded below by a positive constant. We choose $l = l(\kappa, \alpha, D_0)$ small enough so that $\rho(\kappa, \alpha, D_0, l) > \alpha/2$. Then the metric suspensions of circles of length $\leq \alpha$, and hence the links of singular points in X cannot contain a *smooth* ρ-disc. It follows that $\bar{w}$ cannot project to a singular point in $\mathrm{Cut}(x)$ unless it is singular itself, i.e. coincides with a point $\bar{y}_i$. On the other hand, the points y_i are singular, and thus the boundary identifications on $\partial\mathcal{D}(x)$ can identify the points $\bar{y}_i$ at most with each other and with no other points in $\partial\mathcal{D}(x)$. Since all cut points near the y_i are smooth it follows that the y_i cannot be singular vertices. Hence $\bar{y}_1$ and $\bar{y}_2$ have to be identified, and σ closes up with period $\leq 2l'$. Note that the argument also shows that $\mathrm{Cut}(x)$ is totally geodesic near $y_1 = y_2$. □

We further decrease l_1 until $l_1 \leq l$ with the constant l of Sublemma 5.35. This amounts to removing from V_2 some of its components.

We then have that every singular point x outside $V_1 \cup V_2$ is the initial point of a minimizing singular segment σ of length $\geq l_1$. Let m be the midpoint of this segment. If there were an umbilic turnover S through m and $\perp \sigma$, then $\mathrm{pc}(S) < P_1$ and $\mathrm{diam}(S) \geq d_2$ because $m \notin V_1$. Proposition 5.33 implies a lower bound $i_2 = i_2(\kappa, \alpha, l_1, d_2) = i_2(\kappa, \alpha, D_0) > 0$ for $r_{\mathrm{inj}}(m)$.

We now use the *lower* bound β on cone angles to control $r_{\mathrm{inj}}(x)$ in terms of $r_{\mathrm{inj}}(m)$. There is a *smooth* standard ball of radius $\geq r_2(\kappa, \alpha, D_0, \beta) > 0$ embedded in the singular standard ball $B_{i_2}(m)$. Proposition 3.19 implies a lower bound for $r_{\text{cone-inj}}(x)$. Since $x \notin V_1$, this yields a lower bound $i_3(\kappa, \alpha, D_0, \beta)$ > 0 for $r_{\mathrm{inj}}(x)$.

After feeding the baby, if necessary, we choose $r_3(\kappa, \alpha, D_0, \beta)$ with $0 < r_3 < \min(\frac{i_3}{3}, n_1)$ and define V_3 as the closure of the union of all balls of radius r_3 centered at singular points outside $V_1 \cup V_2$. Since $r_3 < n_1$, we have $V_2 \cap V_3 = \emptyset$ because $N_{n_1}(\partial V_2)$ contains no singular points, and so

$$(V_1 \cup V_3) \cap V_2 = \emptyset.$$

Our construction yields that (i) $V_1 \cup V_2 \cup V_3$ contains the tubular $\frac{r_3}{2}$-neighborhood of Σ_X, and (ii) there is a lower bound for r_{inj} on $\partial(V_1 \cup V_3)$. On the other hand, ∂V_2 can become arbitrarily thin.

Step 5. *Foliating the thin part away from the singular locus.* We have $r_{\mathrm{geod}} \geq \frac{r_3}{2}$ on the complement Y of $V_1 \cup V_2 \cup V_3$. The local Margulis lemma, cf. the discussion in Section 5.3, implies that there exist constants $i_4 = i_4(\kappa, r_3) = i_4(\kappa, \alpha, D_0, \beta) > 0$, $l_2 = l_2(\kappa, r_3) = l_2(\kappa, \alpha, D_0, \beta) > 0$, $l_2 << i_4$, and an open subset $Y^{\mathrm{thin}} \subseteq Y$ carrying a natural locally homogeneous codimension-one foliation $\mathcal{F}$, possibly with singular one-dimensional leaves, which enjoy the following properties: (i) $\{r_{\mathrm{inj}} < i_4\} \cap Y \subset Y^{\mathrm{thin}}$; (ii) $\mathcal{F}$ is locally equivalent to a foliation of model space $\mathbb{M}^3_k$ by equidistant surfaces of a geodesic or a horosphere; (iii) intrinsically, the regular leaves of $\mathcal{F}$ are flat, and they admit foliations by parallel geodesics of length $< l_2$ (without singular leaves since X is orientable). Note that these one-dimensional foliations need not be unique on compact leaves. Since r_{inj} is bounded below on $\partial(V_1 \cup V_3)$, we can arrange by choosing the constants i_4, l_2 sufficiently small that

$$Y^{\mathrm{thin}} \cap (V_1 \cup V_3) = \emptyset.$$

It can happen that Y^{thin} intersects a component C of V_2. Then $\partial C \subset \overline{Y^{\mathrm{thin}}}$ and the natural foliation of C by equidistant surfaces around the singular core geodesic extends $\mathcal{F}$. We denote by V_2' the union of all these components C of V_2 and conclude that $\mathcal{F}$ extends to a foliation $\hat{\mathcal{F}}$ of $Y^{\mathrm{thin}} \cup \bar{V}_2' =: \hat{Y}^{\mathrm{thin}}$. Notice that the completeness of X implies that the regular leaves of $\hat{\mathcal{F}}$ are *complete*, since they stay away from $\partial(V_1 \cup V_3)$.

Our discussion implies that the connected components of $\hat{Y}^{\text{thin}}$ are smooth cusps of rank one or two (i.e. quotients of horoballs) or tubular neighborhoods of closed (smooth or singular) geodesics. The singular leaves of $\mathcal{F}$ are short closed smooth geodesics. The boundary of $\hat{Y}^{\text{thin}}$ is a union of complete leaves and we have $r_{\text{inj}} \geq i_4$ on $\partial\hat{Y}^{\text{thin}}$. Each component of $\hat{Y}^{\text{thin}}$ has nonempty boundary, because the leaves of $\hat{\mathcal{F}}$ are strictly convex and r_{inj} increases towards the convex side.

We define X^{thin} as $V_1 \cup V_2 \cup \hat{Y}^{\text{thin}}$ with those components omitted which are tubular neighborhoods of closed smooth geodesics with length, say, $> i_4$. We already removed short umbilic necks earlier. All other components are uniformly thick.

This concludes the proof of Theorem 5.3. $\square$

5.8. *Totally geodesic boundary.* In this section we allow a totally geodesic boundary for X and apply the discussion in Section 5.5 to investigate the geometry near boundary components.

Proposition 5.36 (I-bundle). *For $d_0 > 0$ there exists $\tilde{h} = \tilde{h}(\kappa, d_0) > 0$ such that*:

Let X be a cone 3-manifold of curvature $k \in [-\kappa, 0]$ with totally geodesic boundary and cone angles $\leq \pi$. Suppose that ∂X contains a connected component S with $\operatorname{diam}(S) \leq d_0$ and $d(S, \operatorname{Cut}(S)) \leq \tilde{h}$. Then $\partial X = S$, $\operatorname{Cut}(S)$ is a compact totally geodesic surface, possibly with boundary, $\partial\operatorname{Cut}(S)$ is a union of singular edges, and X carries a natural structure as a singular bundle $X \to \operatorname{Cut}(S)$ with fiber a compact interval.

Proof. Consider the cut locus $\operatorname{Cut}(\partial X)$ with respect to the entire boundary. Note that $d(S, \operatorname{Cut}(\partial X)) \leq d(S, \operatorname{Cut}(S))$. The discussion of Section 5.5 applies since we can replace boundary components of X by nearby equidistant umbilic surfaces. We choose $h > 0$ arbitrarily and then $\tilde{h}$ as the constant provided by Lemma 5.22. Combining Lemma 5.22 and the description of the cut locus in Lemma 5.19 we obtain the assertion. The bundle structure is given as follows: The fiber over $x \in \operatorname{Cut}(S)$ is the union of the (one or two) shortest segments from x to S. $\square$

Note that the boundary identifications on $\mathcal{D}(S)$ correspond, via the nearest point projection $\partial\mathcal{D}(S) \to S$, to an isometric involution ι on S.

If $k < 0$, cut points at maximal distance from S must be corners of $\partial\operatorname{Cut}(S)$, and in particular $\partial\operatorname{Cut}(S) \neq \emptyset$ is a nonempty union of singular edges.

In the case that $k \leq 0$, X is orientable and S is a turnover, ι reverses orientation and two possibilities can occur: Either ι fixes all three cone points and $\mathrm{Cut}(S)$ is a triangle, or ι fixes one cone point and exchanges the other two. In the latter case, $\mathrm{Cut}(S)$ is a disc with one cone point and one corner.

In subsection 6.3 we will need the following version of Proposition 5.36.

Corollary 5.37. *Suppose that $k < 0$. Then for $0 < \beta < \pi$ there exists $\rho = \rho(k, \beta) > 0$ such that: Let X be a cone manifold of curvature k with totally geodesic boundary and cone angles $\in [\beta, \pi]$. Suppose that ∂X contains a turnover S. If X is ρ-thin, i.e. contains no embedded smooth standard ball of radius ρ, then the conclusion of Proposition* 5.36 *holds.*

Moreover, given $\varepsilon > 0$, there exists $\rho_1 = \rho_1(k, \beta, \varepsilon) > 0$ such that if X is ρ_1-thin then the cone angles at the singular edges in the boundary of $\mathrm{Cut}(S)$ *are $\geq \pi - \varepsilon$.*

Proof. The lower bound β on cone angles yields an upper bound $d_0 = d_0(k, \beta)$ for the diameter of the turnover S. We are done by Proposition 5.36 if $d(S, \mathrm{Cut}(S)) \leq \tilde{h}(-k, d_0)$. We suppose therefore that $d(S, \mathrm{Cut}(S)) > \tilde{h}(-k, d_0)$.

We certainly have $\mathrm{Cut}(S) \neq \emptyset$ because X is ρ-thin. Denoting $n := d(S, \mathrm{Cut}(S))$, we consider the family of embedded umbilic surfaces S_r, $0 < r < n$, equidistant to S with distance r from S. They have uniformly bounded principal curvatures $\mathrm{pc}(S_r) < \sqrt{-k}$, and the ratio $\mathrm{diam}(S_r)/\mathrm{diam}(S)$ is a function of k and r, monotonically increasing in r.

We choose h smaller than the constant $h(-k, \sqrt{-k})$ given in Corollary 5.20, and let $d_1 = d_1(-k, \sqrt{-k}, h)$ and $\tilde{h} = \tilde{h}(-k, \sqrt{-k}, h)$ be the constants from Lemma 5.21. By choice of ρ sufficiently small we assure that the surfaces S_r are uniformly thin and, in view of the lower bound on cone angles, have diameter $< d_1$. Lemma 5.21, applied to S_r for $r \to n$, then implies the conclusion of Proposition 5.36. The additional assertion regarding the cone angles at $\partial\mathrm{Cut}(S)$ follows from Lemma 5.19 by choice of the constant h sufficiently small so that $\hat{\psi}(-k, \sqrt{-k}, h) > \frac{\pi}{2} - \frac{\varepsilon}{2}$; cf. (3). $\square$

6. Proof of the main theorem

Let $\mathcal{O}$ be a compact, connected, orientable, small 3-orbifold with nonempty ramification locus Σ. The singular locus Σ is a trivalent graph properly embedded in $|\mathcal{O}|$. Let $\Sigma^{(0)}$ denote the set of vertices of Σ and $\Sigma^{(1)} = \Sigma - \Sigma^{(0)}$ the union of open edges. We regard circle components of the singular locus as edges which close up.

We consider the manifold $M = |\mathcal{O}| - \Sigma^{(1)} - \mathcal{N}(\Sigma^{(0)})$, i.e. we remove the edges of Σ and an open ball neighborhood of each vertex. The manifold M is noncompact, with boundary $\partial M = (\partial\mathcal{N}(\Sigma^{(0)}) \cup \partial\mathcal{O}) - \Sigma^{(1)}$ a finite collection of thrice punctured spheres.

6.1. *Reduction to the case when the smooth part is hyperbolic.* The following proposition allows us to reduce the proof of the main theorem to the case where M admits a complete hyperbolic structure with finite volume and totally geodesic boundary.

PROPOSITION 6.1. *Either the manifold M has a complete hyperbolic structure with finite volume and with totally geodesic boundary, or $\mathcal{O}$ admits a Seifert fibration, an I-bundle structure, or a spherical structure (i.e. a quotient of $\mathbb{S}^3$ or B^3 by an orthogonal action).*

Proof. Let $\overline{M} = \mathcal{O} - \mathcal{N}(\Sigma)$ be a compact core of M. The boundary $\partial\overline{M}$ is the union of compact pairs of pants (which are a compact core of ∂M) together with a collection $P \subset \partial\overline{M}$ of tori and annuli, corresponding to the boundary of a neighborhood of edges in Σ.

LEMMA 6.2. *Either $\overline{M}$ is Seifert fibred or $(\overline{M}, P)$ is an atoroidal pared manifold.*

We recall that an *atoroidal pared manifold* is a pair $(\overline{M}, P)$ such that:

– $\overline{M}$ is a compact orientable irreducible 3-manifold.

– $P \subset \partial\overline{M}$ is a disjoint union of incompressible tori and annuli such that no two components of P are isotopic in $\overline{M}$.

– $\overline{M}$ is homotopically atoroidal and P contains all torus components of $\partial\overline{M}$.

– There is no essential annulus $(A, \partial A) \subset (\overline{M}, P)$.

We remark that with this definition, an atoroidal pared manifold is never Seifert fibred.

Proof. The manifold $\overline{M}$ is irreducible and topologically atoroidal because so is $\mathcal{O}$. With the assumption that $\overline{M}$ is not Seifert fibred, $\overline{M}$ is homotopically atoroidal, and we prove that $(\overline{M}, P)$ is an atoroidal pared manifold. First we show that P is incompressible in $\overline{M}$. A compressible annulus in P would give a teardrop in $\mathcal{O}$, contradicting irreducibility of $\mathcal{O}$. If a torus component of P was compressible, then the irreducibility of $\overline{M}$ would imply that $\overline{M}$ is a solid torus, hence Seifert fibred. It only remains to check that the pair $(\overline{M}, P)$ is anannular. Let $(A, \partial A) \subset (\overline{M}, P)$ be an essential annulus; we distinguish three cases according to whether ∂A is contained in a) torus components of P, b) annulus components of P, or c) a torus and an annulus of P. In the first case, a classical argument using the atoroidality of $\overline{M}$ implies that $\overline{M}$ is Seifert fibred [BS1, Lemma 7]. In case b), adding two meridian discal orbifolds to A along ∂A would give a bad or an essential spherical 2-suborbifold, contradicting the irreducibility of $\mathcal{O}$. Case c) reduces to case b), by consideration of the essential annulus obtained by gluing two parallel copies of A with the annulus $P_0 - \mathcal{N}(\partial A)$, where P_0 is the torus component of P that meets ∂A. □

End of the proof of Proposition 6.1. We consider both possibilities of Lemma 6.2. When $\overline{M}$ is Seifert fibred, then the fibration of $\overline{M}$ extends to a fibration of the orbifold $\mathcal{O}$ by adding the components of Σ as fibers, because $\mathcal{O}$ is irreducible.

When $(\overline{M}, P)$ is an atoroidal pared 3-manifold, since $\overline{M}$ is Haken, by Thurston's hyperbolization theorem for atoroidal Haken pared 3-manifolds (cf. [Thu2, 3, 4, 5], [McM1], [Kap], [MB], [Ot1, 2]), the interior of M admits a complete hyperbolic structure with parabolic locus P. The convex core of this metric may have dimension two or three. If it has dimension three, then this gives a hyperbolic metric on M with totally geodesic boundary and cusp ends, because the boundary is a union of three times punctured spheres, and therefore the Teichmüller space of ∂M is a point.

If the convex core has dimension two, then M is an I-bundle. Since ∂M is a union of three times punctured spheres, and such punctured spheres do not have free involutions, it follows that M is a product of the interval with a three times punctured sphere. Hence there are three possibilities. In the first case $\partial\mathcal{O} = \emptyset$ and $\mathcal{O}$ is the suspension of a turnover. This turnover must be spherical and it is clear that $\mathcal{O}$ is spherical itself. If $\partial\mathcal{O}$ has precisely one component, then $\mathcal{O}$ is a standard quotient of a ball (hence spherical). Finally, the last case happens when $\partial\mathcal{O}$ has two components. In this case $\mathcal{O}$ is an I-bundle over a turnover, the turnover is Euclidean or hyperbolic, and $\mathcal{O}$ is also Euclidean or hyperbolic. □

6.2. *Deformations of hyperbolic cone structures.* From now on we assume that the manifold M admits a complete hyperbolic structure of finite volume with totally geodesic boundary.

Starting with the hyperbolic metric on M, we define in this section a deformation space of hyperbolic cone structures on $\mathcal{O}$ and prove an openness property.

Definition 6.3. A *hyperbolic cone structure* on $\mathcal{O}$ is a hyperbolic cone 3-manifold X with totally geodesic boundary together with an embedding $i\colon (X, \Sigma_X) \hookrightarrow (|\mathcal{O}|, \Sigma)$ such that $|\mathcal{O}| - X$ is a (possibly empty) collection of vertices, open ball neighborhoods of vertices and Euclidean 2-orbifolds in $\partial\mathcal{O}$.

If i is a homeomorphism, we call the cone structure a *cone metric* on $\mathcal{O}$.

The choice of a hyperbolic cone structure on $\mathcal{O}$ assigns cone angles to the edges of Σ. Usually we will consider situations when the cone angles are less than or equal to the orbifold angles. In order to be able to work with small cone angles, we allow in our definition of cone structure small deviations between the topologies of $\mathcal{O}$ and X; i.e., we do not require i to be a homeomorphism. That is, we observe that in a cone manifold the sum of the cone angles of the singular edges adjacent to a vertex of Σ_X is $> 2\pi$, because the link of a vertex is a spherical turnover. Thus, if v is a vertex component of $|\mathcal{O}| - X$, we require the hyperbolic structure on the punctured neighborhood of v to be a cusp; the cone angle sum for the singular edges adjacent to v then equals 2π. If $|\mathcal{O}| - X$ contains the open ball neighborhood B_v of a vertex v, we require the boundary turnover ∂B_v to be totally geodesic; the cone angle sum is $< 2\pi$ in this case. If S is a Euclidean 2-orbifold in $\partial\mathcal{O} - X$, then we request that the hyperbolic structure near S is also a cusp.

Notice that all boundary components of X are turnovers because $\mathcal{O}$ is small.

There are analogous definitions for *Euclidean* and *spherical* cone structures on $\mathcal{O}$; however in these cases it will be enough for us to consider cone metrics.

We regard the complete hyperbolic structure of finite volume and geodesic boundary on M as a hyperbolic cone structure on $\mathcal{O}$ with all cone angles equal to zero.

Let $m_1, \dots, m_q$ be the ramification indices of the edges of Σ (with respect to a fixed numbering). Throughout the proof of the orbifold theorem we will consider the following set of hyperbolic cone structures with fixed ratios for the cone angles. Define:

$$\mathcal{J}(\mathcal{O}) = \left\{ t \in [0,1] \;\middle|\; \begin{array}{l} \text{there exists a hyperbolic cone structure on } \mathcal{O} \\ \text{with cone angles } \left(\frac{2\pi t}{m_1}, \dots, \frac{2\pi t}{m_q}\right) \end{array} \right\}.$$

A hyperbolic cone structure on $\mathcal{O}$ induces a noncomplete hyperbolic structure on M. In particular it has a holonomy representation $\pi_1(M) \to \mathrm{PSL}_2(\mathbb{C})$. The variety of representations $\mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))$ is an affine algebraic set, possibly reducible. The group $\mathrm{PSL}_2(\mathbb{C})$ acts on $\mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))$ by conjugation, and we are interested in the quotient. The topological quotient is not Hausdorff, and one therefore considers the algebraic quotient

$$\mathcal{X}(M) = \mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))//\mathrm{PSL}_2(\mathbb{C})$$

which is again an affine algebraic set. Note that the irreducible representations form a Zariski open subset of $\mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))$. Namely, $\mathrm{Hom}^{\mathrm{irr}}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))$ is the inverse image of a Zariski open subset $\mathcal{X}^{\mathrm{irr}}(M) \subseteq \mathcal{X}(M)$, and $\mathcal{X}^{\mathrm{irr}}(M)$ is the topological quotient of $\mathrm{Hom}^{\mathrm{irr}}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))$. Notice that the holonomy representation ρ_0 of the (metrically) complete hyperbolic structure on M is irreducible.

The polynomial functions on $\mathcal{X}(M)$, one-to-one, correspond to the polynomial functions on $\mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C}))$ invariant under the $\mathrm{PSL}_2(\mathbb{C})$-action.

Given $\gamma \in \pi_1(M)$, we define the trace-like function $\tau_\gamma : \mathcal{X}(M) \to \mathbb{C}$ as the function induced by

$$\begin{aligned} \mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C})) &\to \mathbb{C} \\ \rho &\mapsto \mathrm{trace}(\rho(\gamma)^2). \end{aligned}$$

Let $\mu_1, \dots, \mu_q$ be a family of meridian curves, one for each component of $\Sigma^{(1)}$.

THEOREM 6.4 (Local parametrization). *The map*

$$\tau_\mu = (\tau_{\mu_1}, \dots, \tau_{\mu_q})\colon \mathcal{X}(M) \to \mathbb{C}^q$$

is locally bianalytic at $[\rho_0]$.

This result is the main step in the proof of Thurston's hyperbolic Dehn filling theorem (see [BoP, App. B] or [Kap] for the proof). It implies in particular the following special case of Thurston's Generalized Hyperbolic Dehn Filling Theorem.

COROLLARY 6.5. *The set* $\mathcal{J}(\mathcal{O})$ *contains a neighborhood of* 0.

Proof. We have $\tau([\rho_0]) = (2, \dots, 2)$. Consider the path

$$\begin{array}{rcl} \gamma\colon [0,\varepsilon) & \to & \mathbb{C}^q \\ t & \mapsto & (2\cos\frac{2\pi t}{m_1}, \dots, 2\cos\frac{2\pi t}{m_k}) \end{array} \tag{7}$$

where $\varepsilon > 0$ is sufficiently small. The composition $\tau_\mu^{-1} \circ \gamma$ gives a path of conjugacy classes of representations. It can be lifted to a path of representations $t \mapsto \rho_t$, because there are slices to the action of $\mathrm{PSL}_2(\mathbb{C})$ on the representation variety. (Existence of slices follows from Luna's theorem, as $\mathrm{PSL}_2(\mathbb{C}) \cong \mathrm{SO}(3,\mathbb{C})$ is an affine reductive group.) The representations ρ_t are the holonomies of incomplete hyperbolic structures on M. By construction, the holonomies of the meridians are rotations with angles $2\pi t/m_1, \dots, 2\pi t/m_q$. By a standard result, the deformation of holonomies is, locally near $t = 0$, induced by a deformation of hyperbolic cone structures on $\mathcal{O}$ with cone angles $2\pi t/m_j$. $\square$

LEMMA 6.6. *There exists a unique irreducible curve* $\mathcal{D} \subset \mathbb{C}^q$ *such that* $\gamma([0,1]) \subset \mathcal{D}$.

Proof. For $n \in \mathbb{N}$, we consider the Chebyshev-like polynomial

$$p_n(x) = 2\cos(n \arccos(x/2)).$$

It has the following property:

$$\mathrm{trace}(A^n) = p_n(\mathrm{trace}(A)), \qquad \forall A \in \mathrm{SL}_2(\mathbb{C}),\ \forall n \in \mathbb{N}.$$

An easy computation shows that $p'_n(2) = n$, and therefore

$$\{z \in \mathbb{C}^q \mid p_{m_1}(z_1) = \cdots = p_{m_q}(z_q)\}$$

is an algebraic curve with $(2, \dots, 2)$ as a smooth point. We take $\mathcal{D}$ to be the unique irreducible component containing $(2, \dots, 2)$. Then $\gamma([0, \varepsilon)) \subset \mathcal{D}$ for small $\varepsilon > 0$. Since γ is an analytic curve, it remains in $\mathcal{D}$. □

We define the algebraic curve $\mathcal{C} \subset \mathcal{X}(M)$ to be the irreducible component of $\tau_\mu^{-1}(\mathcal{D})$ that contains $[\rho_0]$. By construction, $[\rho_t] \in \mathcal{C}$ for small $t \geq 0$.

For technical reasons, we define the following variant of $\mathcal{J}(\mathcal{O})$. Here v_0 denotes the volume of the complete hyperbolic structure on M.

(8)

$$\mathcal{J}_0(\mathcal{O}) = \left\{ t \in [0,1] \;\middle|\; \begin{array}{l} \text{there exists a hyperbolic cone structure on } \mathcal{O} \\ \text{with cone angles } \left(\frac{2\pi t}{m_1}, \dots, \frac{2\pi t}{m_q}\right), \text{ holonomy in } \mathcal{C} \\ \text{and volume } \leq v_0 \end{array} \right\}.$$

The condition that the holonomy is in the curve $\mathcal{C}$ will be used in Theorem 6.7, because holomorphic maps on curves are open (openness does not hold for maps on higher dimensional varieties).

Note that $[0, \varepsilon) \subset \mathcal{J}_0(\mathcal{O})$ for small $\varepsilon > 0$ because, according to Schläfli's formula, the volume of the continuous family of cone structures with holonomies ρ_t decreases.

THEOREM 6.7 (Openness). *The set $\mathcal{J}_0(\mathcal{O})$ is open to the right.*

Proof. As remarked above, openness of $\mathcal{J}_0(\mathcal{O})$ at $t = 0$ is a consequence of Thurston's hyperbolic Dehn filling, and we only prove openness at $t > 0$.

Consider the path

$$\begin{aligned} \gamma : [t, t+\varepsilon) &\to \mathcal{D} \subset \mathbb{C}^q \\ s &\mapsto (2\cos(s2\pi/m_1), \dots, 2\cos(s2\pi/m_q)) \end{aligned}$$

defined for some $\varepsilon > 0$. By construction, the image of γ is contained in the curve $\mathcal{D} \subset \mathbb{C}^q$ of Lemma 6.6. Since $\tau_\mu : \mathcal{C} \to \mathcal{D}$ is nonconstant, it is open, and therefore γ can be lifted to $\mathcal{C}$. We can lift it further to a path

$$\begin{aligned} \tilde{\gamma} : [t, t+\varepsilon) &\to \mathrm{Hom}(\pi_1(M), \mathrm{PSL}_2(\mathbb{C})) \\ s &\mapsto \rho_s. \end{aligned}$$

To justify this second lift, notice that the holonomy ρ_t is irreducible (because the corresponding cone structure has finite volume) and therefore the $\mathrm{PSL}_2(\mathbb{C})$-action is locally free. By construction, $\rho_s(\mu_i)$ is a rotation of angle $\frac{2\pi s}{m_i}$. Therefore the cone structure on $\mathcal{O}$ with holonomy ρ_t can be deformed to a continuous family of cone structures on $\mathcal{O}$ with holonomies ρ_s. By Schläfli's formula volume decreases and thus $[t, t+\varepsilon) \subset \mathcal{J}_0(\mathcal{O})$ for $\varepsilon > 0$ sufficiently small. □

A straightforward consequence of Theorem 6.7 is:

Corollary 6.8. *If* $\sup \mathcal{J}_0(\mathcal{O}) \in \mathcal{J}_0(\mathcal{O})$ *then* $1 \in \mathcal{J}_0(\mathcal{O})$ *and* $\mathcal{O}$ *is hyperbolic.*

The next step in the proof is the analysis of degenerating sequences of cone structures on $\mathcal{O}$, namely sequences (t_n) in $\mathcal{J}_0(\mathcal{O})$ that converge to $t_\infty \notin \mathcal{J}_0(\mathcal{O})$. This analysis is carried out in the next section, using the results of Sections 7, 8 and 10.

6.3. *Degeneration of hyperbolic cone structures.* We continue the discussion of deforming hyperbolic cone structures on $\mathcal{O}$ while keeping the ratios of the cone angles fixed.

Let (t_n) be a sequence in $\mathcal{J}_0(\mathcal{O})$. Let $X(t_n)$ be a hyperbolic cone structure on $\mathcal{O}$ corresponding to $t_n \in \mathcal{J}_0(\mathcal{O})$, with the properties as in (8).

Definition 6.9. We say that a sequence of cone 3-manifolds X_n *collapses* if, for every $\rho > 0$, only finitely many X_n are ρ-thick; cf. Definition 3.17.

Proposition 6.10 (Degeneration implies collapse). *If* $t_n \to t_\infty \notin \mathcal{J}_0(\mathcal{O})$, *then the sequence* $(X(t_n))$ *collapses.*

Proof. Assume that $(X(t_n))$ does not collapse. Then, up to passing to a subsequence, the cone manifolds $X(t_n)$ are ρ-thick for some $\rho > 0$. According to Corollary 3.22, $(X(t_n))$ subconverges to a hyperbolic cone 3-manifold X_∞. Since in the definition of $\mathcal{J}_0(\mathcal{O})$ we impose an upper volume bound on the cone structures, we have that $\mathrm{vol}(X_\infty) < \infty$.

Moreover, if $t_\infty < 1$, the cone angles of $X(t_n)$ are all bounded away from π, and if $t_\infty = 1$, then they converge to the orbifold angles of $\mathcal{O}$. This allows us to apply Theorem 7.1 in Section 7 below. According to this theorem, X_∞ yields a hyperbolic cone structure on $\mathcal{O}$ and $t_\infty \in \mathcal{J}_0(\mathcal{O})$, which contradicts the hypothesis. □

We analyse now the situation when the sequence $(X(t_n))$ collapses and treat the cases with and without boundary separately.

The case with boundary is handled by the following geometric fibration result.

Proposition 6.11 (I-bundle). *If* $\partial X(t_n) \neq \emptyset$ *for all* n, *then* $t_\infty = 1$ *and* $\mathcal{O}$ *is a twisted* I*-bundle over the quotient of a turnover by an orientation-reversing involution.*

Proof. Since the sequence $(X(t_n))$ collapses and $\partial X(t_n)$ is a collection of turnovers, for n sufficiently large Corollary 5.37 applies to $X(t_n)$. It shows that the boundary $\partial X(t_n)$ consists of a single turnover and that $X(t_n)$ is a singular

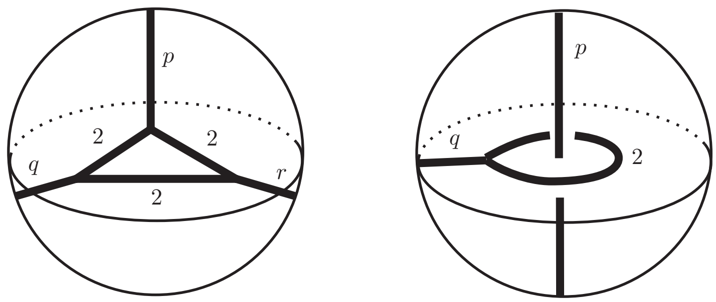

Figure 2. The two I-bundles over the quotient of a turnover

interval bundle over the cut locus $\mathrm{Cut}(\partial X(t_n))$ with respect to $\partial X(t_n)$. The cut locus $\mathrm{Cut}(\partial X(t_n))$ is naturally homeomorphic to the quotient of the turnover $\partial X(t_n)$ by an orientation-reversing isometric involution.

Moreover, since the $X(t_n)$ are negatively curved, $\partial\mathrm{Cut}(\partial X(t_n))$ is a non-empty collection of singular edges; cf. the discussion after Proposition 5.36. By Corollary 5.37, their cone angles converge to π as $n \to \infty$. Thus $t_n \to 1$. It follows that the boundary turnovers of the $X(t_n)$ correspond to a hyperbolic boundary turnover of $\mathcal{O}$. Hence the $X(t_n)$ provide not only cone structures but cone *metrics* on $\mathcal{O}$, and $\mathcal{O}$ is an I-bundle over a quotient of a turnover, as claimed. $\square$

Remark 6.12. A turnover always has an isometric involution which fixes all three cone points and reverses orientation. The quotient is a triangular 2-orbifold.

The turnover is the double along the boundary of a geodesic triangle, and if the triangle has a reflection symmetry, then the turnover has a corresponding orientation-reversing involution which fixes one cone point and exchanges the other two. In this case the quotient is a disc with a corner and one cone point. Figure 2 illustrates the orientable I-bundles over such 2-orbifolds.

Now we discuss the case without boundary, i.e. we assume that $\partial X(t_n) = \emptyset$ for all n. Note that all cusps of the $X(t_n)$ are singular because $\mathcal{O}$ is small. Since, without loss of generality, the sequence (t_n) is strictly increasing, only finitely many $X(t_n)$ can have cusps and the cusps correspond to singular vertices of $\mathcal{O}$. After passing to a subsequence, we may assume that the $X(t_n)$ have no cusps at all; i.e., they are closed cone manifolds, $\mathcal{O}$ is a closed orbifold and the $X(t_n)$ provide cone *metrics* on $\mathcal{O}$. In particular, the pairs $(|\mathcal{O}|, \Sigma)$ and $(X(t_n), \Sigma_{X(t_n)})$ are homeomorphic for all n.

Definition 6.13. We say that the cone 3-manifold X has *ω-thick links*, $\omega > 0$, if in every point $x \in X$ the link $\Lambda_x X$ is ω-thick; cf. Definition 3.17.

Theorem 6.14 (Fibration). *For $\omega, D_0 > 0$ there exists $\delta = \delta(\omega, D_0) > 0$ such that*:

Suppose that $\mathcal{O}$ is closed and admits a cone metric (X, i) of constant curvature $k \in [-1, 0)$ with cone angles less than or equal to the orbifold angles. If the cone 3-manifold X has ω-thick links, $\operatorname{diam}(X) \geq D_0$ and if X is δ-thin, then $\mathcal{O}$ is Seifert fibred.

We postpone the proof of this theorem to Section 10. Assuming it, we continue our argument.

Proposition 6.15. *If $\partial X(t_n) = \emptyset$ for all n, then one of the following holds*:

– *either there exists a Euclidean cone metric on $\mathcal{O}$ with cone angles strictly less than the orbifold angles*,

– *or $\mathcal{O}$ has a Euclidean or a Seifert fibred structure.*

Proof. Case 1. *Collapse with lower diameter bound.* Assume that

$$\operatorname{diam}(X(t_n)) \geq D_0 > 0$$

for all n. In order to apply Theorem 6.14, we check that $X(t_n)$ has ω-thick links for some uniform $\omega > 0$. Since $t_n < t_\infty$ and the cone angles for the cone metrics X_n are proportional to t_n, we have a uniform lower bound $2\pi + \varepsilon$, $\varepsilon > 0$, for the cone angle sum of the singular edges adjacent to any vertex with respect to all cone metrics $X(t_n)$. By Gauß-Bonnet, we obtain a lower bound for the area of links at singular vertices of the $X(t_n)$, and this converts into a lower bound for the thickness of these links. Since the cone angles of the $X(t_n)$ are bounded away from zero, the cone manifolds $X(t_n)$ have uniformly thick links. Hence Theorem 6.14 applies and we obtain that $\mathcal{O}$ is Seifert fibred.

Case 2. *Collapse to a point.* Assume that $\operatorname{diam}(X(t_n)) \to 0$. Then we consider the sequence

$$\overline{X}(t_n) = \frac{1}{\operatorname{diam}(X(t_n))} X(t_n)$$

of rescaled cone 3-manifolds with constant curvature $k_n = -\operatorname{diam}(X(t_n))^2 \in [-1, 0)$ and diameter equal to 1. If the sequence $(\overline{X}(t_n))$ collapses, then for n sufficiently large Theorem 6.14 applies as above to show that $\mathcal{O}$ is Seifert fibred.

If the rescaled sequence $\overline{X}(t_n)$ does not collapse, then by the compactness result 3.22, a subsequence converges geometrically to a closed Euclidean cone 3-manifold X_∞ with diameter 1, and X_∞ yields a Euclidean cone metric on $\mathcal{O}$. Either the cone angles of X_∞ are strictly less than the orbifold angles of $\mathcal{O}$ ($t_\infty < 1$), or X_∞ corresponds to a Euclidean *orbifold* structure on $\mathcal{O}$ with the same branching indices as $\mathcal{O}$ ($t_\infty = 1$), and so $\mathcal{O}$ is Euclidean. □

The last step of the proof of the main theorem is given by:

THEOREM 6.16 (Spherical uniformization). *Let $\mathcal{O}$ be a closed, orientable small 3-orbifold. If there exists a Euclidean cone metric on $\mathcal{O}$ with cone angles strictly less than the orbifold angles, then $\mathcal{O}$ is spherical.*

This theorem is proved in Sections 8 and 9. The flowchart in Figure 3 represents the logic of the proof of the main theorem given in this section.

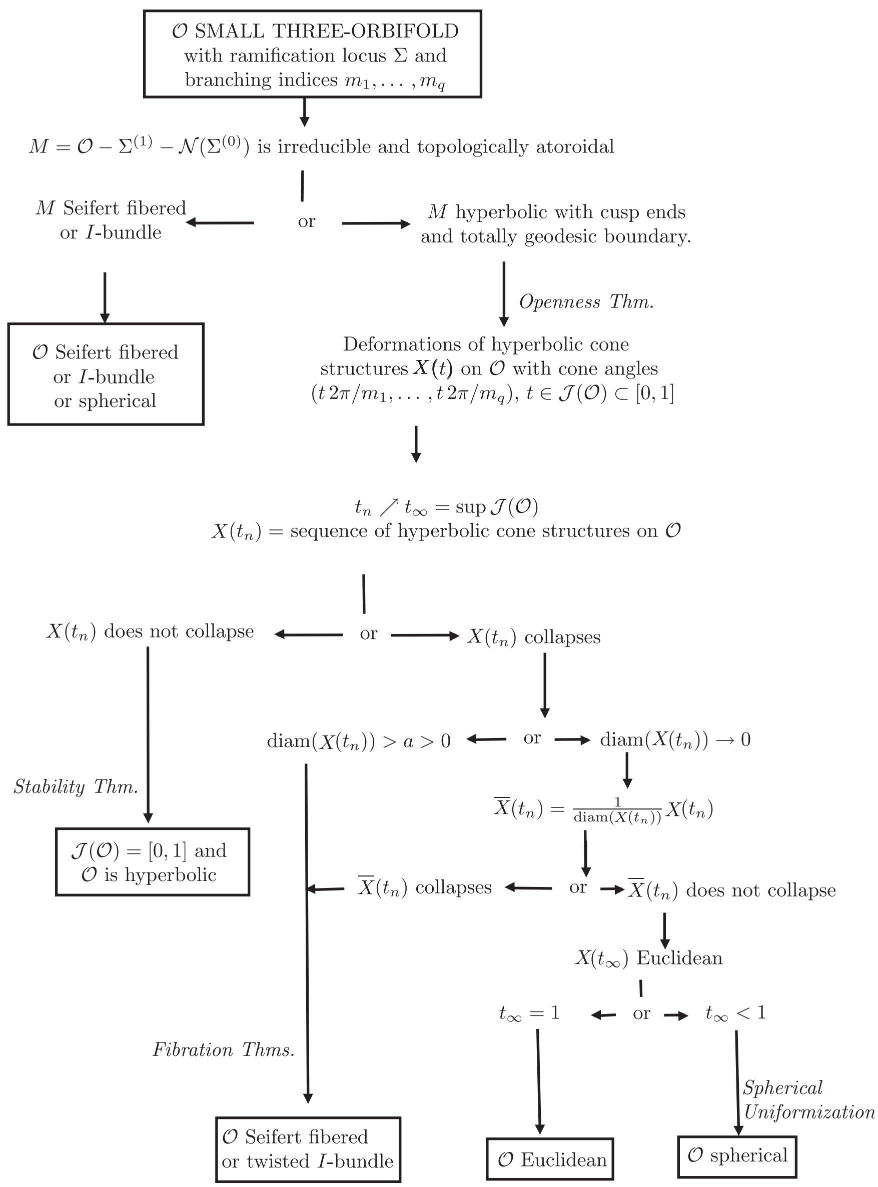

Figure 3. Flowchart of the proof of the main theorem

7. Topological stability of geometric limits

In this section we discuss the change of the topological type of cone manifolds under geometric limits. More specifically, we consider a sequence of compact hyperbolic cone 3-manifolds X_n with cone angles $\leq \pi$. We suppose furthermore that the sequence (X_n) is *noncollapsing*, i.e. that for a uniform radius $\rho > 0$ the X_n contain embedded smooth standard balls $B_\rho(x_n)$.

Due to the compactness theorem (Corollary 3.22), after passing to a subsequence, the pointed cone manifolds (X_n, x_n) geometrically converge

$$(X_n, x_n) \longrightarrow (X_\infty, x_\infty). \tag{9}$$

Their limit X_∞ is again a hyperbolic cone 3-manifold with cone angles $\leq \pi$. We know furthermore that the singular sets converge, $\Sigma_n \to \Sigma_\infty$, and the cone angles converge as well.

If the limit X_∞ is compact, then topological stability holds; that is, (X_n, Σ_n) is homeomorphic to $(X_\infty, \Sigma_\infty)$ for sufficiently large n. In the following, we will study the situation when X_∞ is *noncompact*. In order to obtain topological stability, we need to impose further assumptions. The main result of this section is:

THEOREM 7.1 (Stability in the noncompact limit case). *Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of compact oriented hyperbolic cone* 3*-manifolds with totally geodesic boundary and with cone angles $\in [\omega, \pi]$, $\omega > 0$, which geometrically converges, as in* (9), *to a* noncompact *cone* 3*-manifold . Assume that*:

(1) *The cone manifolds X_n yield hyperbolic cone structures on the same compact* small *orbifold $\mathcal{O}$.*

(2) *There is a uniform upper volume bound* $\mathrm{vol}(X_n) \leq v$.

(3) *Either the cone angles of the X_n are $\leq \alpha < \pi$, or they converge to the orbifold angles of $\mathcal{O}$.*

Then, X_∞ yields a hyperbolic cone structure on $\mathcal{O}$ as well.

Remark 7.2. One can show that the second assumption is implied by the first one using a straightening argument for triangulations under the developing map; cf. [Koj] for the cyclic case. However, this is irrelevant in our later applications, because there we will have a uniform volume bound by construction.

The proof of Theorem 7.1 occupies the entire section and will be divided into several steps.

7.1. *The case of cone angles* $\leq \alpha < \pi$. We consider now the case that the X_n have cone angles $\leq \alpha < \pi$. Here the geometric results of Subsection 5.2 will come to bear.

Geometric finiteness of the limit. The uniform upper volume bound for X_n implies that the limit cone manifold has finite volume:

$$\mathrm{vol}(X_\infty) \leq v.$$

Since $\mathcal{O}$ is small, the boundary components of X_n are turnovers. The lower bound on cone angles yields an upper bound on their diameters, and it follows that the boundary components of X_∞ are also turnovers. Since the cone angles of X_∞ are bounded away from π, finite volume implies geometric finiteness, i.e. X_∞ has finitely many ends and all ends are cusps. Here, we apply the finiteness corollary (Cor. 5.5).

In each cusp $\mathcal{C}_i \subset X_\infty$ we fix a horospherical cross section H_i far out in the thin part, along which we truncate $\mathcal{C}_i$. Denote by N_∞ the resulting compact core of X_∞. The H_i are Euclidean tori or turnovers. By geometric convergence, for sufficiently large n, there are $(1+\varepsilon_n)$-bilipschitz embeddings

$$f_n : N_\infty \hookrightarrow X_n,$$

$\varepsilon_n \to 0$, and we can arrange that $H_{i,n} := f_n(H_i)$ is a leaf of the canonical, locally-homogeneous foliation of the thin part of X_n (namely, either a horosphere in a cusp or a torus equidistant to a short geodesic, possibly singular; cf. Theorem 5.3).

Hyperbolizing the smooth part. We denote by Y the manifold obtained from $X_\infty - \mathcal{N}(\Sigma_\infty^{(0)}) - \Sigma_\infty^{(1)}$ by truncating the singular cusps along horospherical turnovers. Here, we remove *open* singular cusps. The boundary of Y is a union of thrice-punctured spheres because, due to our assumption on cone angles, the cross sections of singular cusps are turnovers.

Proposition 7.3. *The manifold* Y *admits a metrically complete hyperbolic structure with totally geodesic boundary and finite volume.*

Proof. We first deal with the (easy) case that Y has empty boundary. This happens if and only if X_∞ has no singular vertices, no singular cusps and empty boundary. There is nothing to show if X_∞ has no singular locus. If there is a singular locus, one can perturb the hyperbolic metric on the smooth part to a complete Riemannian metric with upper negative curvature bound and finite volume. Since Y is Haken, it follows from Thurston's uniformization theorem that Y admits a complete hyperbolic metric of finite volume.

We assume from now on that $\partial Y \neq \emptyset$. Let $\overline{Y}$ be a compact core of Y. The boundary $\partial\overline{Y}$ is the union of compact pairs of pants (which are a compact core of ∂Y) together with a collection $P \subset \partial\overline{Y}$ of tori and annuli, corresponding to the boundary of a neighborhood of edges and circles in Σ_∞ and to cross sections of smooth cusps. We shall prove that $(\overline{Y}, P)$ is an atoroidal pared manifold (see Lemma 6.2 for the definition). Then Proposition 7.3 follows from Thurston's hyperbolization theorem for atoroidal Haken pared 3-manifolds.

The argument will involve several steps. Part of the information on the topology of Y is obtained from a weak geometric structure, part of it is deduced from the fact that X_∞ arises as a limit of cone manifold structures on small orbifolds.

Step 1. *$\overline{Y}$ is homotopically atoroidal.* We shall show in Proposition 7.10 (in Subsection 7.3) that $\overline{Y}$ admits a nonpositively curved metric, possibly singular, which is negatively curved away from the boundary tori. Consider a π_1-injective map $h : T \hookrightarrow \overline{Y}$ of a 2-torus. The group $\pi_1(T) \cong \mathbb{Z}^2$ acts discretely on $\widetilde{\overline{Y}}$ and hence it preserves a 2-flat (cf. the preliminaries of [LS]). This 2-flat covers a boundary torus and it follows that h can be homotoped into a boundary torus.

Step 2. *$\overline{Y}$ is irreducible.* Let $S \subset \overline{Y}$ be an embedded 2-sphere. We may assume that we have the inclusion of compact cores $\overline{Y} \subset N_\infty$. The sphere $f_n(S)$ bounds a smooth ball B_n in X_n (for n large). We are done, if $B_n \subset f_n(N_\infty)$ for some n. Otherwise, each B_n contains a smooth cross section $H_{i,n}$, which is a leaf of the canonical foliation of the thin part (see above). After passing to a subsequence, we may assume that B_n contains $H_{i_0,n}$ for a fixed i_0. This is absurd because H_{i_0} has nontrivial holonomy, and so has $H_{i_0,n}$ for large n.

Step 3. *P is incompressible.* Each component $P_i \subset P$ corresponds to either cross sections of smooth cusps or tubular neighborhoods of singular edges and circles of the cone manifold X_∞. Hence the holonomy of primitive elements of $\pi_1(P_i)$ is nontrivial.

Step 4. *The pair $(\overline{Y}, P)$ is anannular.* Let $(A, \partial A) \subset (\overline{Y}, P)$ be an essential annulus; we distinguish three cases according to whether ∂A is contained in a) torus components of P, b) annulus components of P, or c) a torus and an annulus of P. In the first case a classical argument using the atoroidality of $\overline{Y}$ implies that $\overline{Y}$ would be Seifert fibred [BS1, Lemma 7], contradicting Step 1. In case b), the annulus A gives rise to an embedded 2-sphere with two cone points in X_∞ and hence (via geometric convergence (9)) in $\mathcal{O}$. Such a 2-sphere bounds an embedded ball with a singular axis because $\mathcal{O}$ is small. It follows that A is parallel to a component of P. Case c) reduces to case b), as in the proof of Lemma 6.2. □

Controlling the geometry of the approaching cone manifolds globally. The Gromov-Hausdorff convergence (9) gives us uniform control on the geometry of all X_n *a priori* only on the portions $f_n(N_\infty)$. Taking into account the structure of the thin part of cone manifolds and using the smallness of $\mathcal{O}$, we will be able to describe also the geometry of the complements $X_n - f_n(N_\infty)$ and see that it is very restricted.

LEMMA 7.4. *Each component of $X_n - f_n(N_\infty)$ is contained in the thin part of X_n, and it is either*

- *a singular ball,*
- *a singular neck containing a turnover $\subseteq \partial X_n$,*
- *or a (singular or smooth) solid torus.*

Remark 7.5. A singular cusp cannot occur in the conclusion of Lemma 7.4 because we assume the X_n to be compact.

Proof. If H_i is a turnover, then $H_{i,n}$ is an umbilic turnover. We go through the three possible cases: If $H_{i,n}$ is spherical, then it bounds a singular ball in X_n. The turnover $H_{i,n}$ cannot be horospherical because X_n is compact. If $H_{i,n}$ is hyperspherical, then there is an umbilic tube bounded by $H_{i,n}$ and a totally geodesic turnover in ∂X_n. Here we use that $\mathcal{O}$ is small.

If H_i is a torus, then $H_{i,n} \subset X_n$ is an almost horospherical torus. It cannot be precisely horospherical, again because X_n is compact. Hence $H_{i,n}$ bounds a (smooth or singular) solid torus $V_{i,n} \subset X_n$. □

If H_i is a torus, denote by $\lambda_{i,n} \subset H_{i,n}$ a geodesic which is a meridian of (i.e. compresses in) the solid torus $V_{i,n}$ bounded by $H_{i,n}$. Furthermore, denote by $\tilde\lambda_{i,n} \subset H_i$ a geodesic such that $f_n(\tilde\lambda_{i,n})$ is homotopic to $\lambda_{i,n}$. The lengths of $\tilde\lambda_{i,n}$ and $\lambda_{i,n}$ are comparable in terms of the bilipschitz constant of f_n.

LEMMA 7.6. *For all i, $\lim_{n\to\infty} \mathrm{length}(\lambda_{i,n}) = \infty$.*

Proof. The radii of the solid tori $V_{i,n}$ tend to ∞ as $n \to \infty$. Using the lower bound on cone angles if $V_{i,n}$ is singular, we obtain the assertion. □

Comparing the topology with the limit. Using these geometric observations, we now describe the change of topology during the transition $(X_\infty, \Sigma_\infty) \rightsquigarrow (X_n, \Sigma_n) = (|\mathcal{O}|, \Sigma)$. We remove from the cone manifolds X_n and the truncated cone manifold N_∞ disjoint small open standard balls around the singular vertices, and denote by Σ'_n and Σ'_∞ the singular loci which remain in the resulting manifolds $X_n - \mathcal{N}_\varepsilon(\Sigma_n^{(0)})$ and $N_\infty - \mathcal{N}_\varepsilon(\Sigma_\infty^{(0)})$. Lemma 7.4 implies that the transition

$$\big(N_\infty - \mathcal{N}_\varepsilon(\Sigma_\infty^{(0)}), \Sigma'_\infty\big) \rightsquigarrow \big(X_n - \mathcal{N}_\varepsilon(\Sigma_n^{(0)}), \Sigma'_n\big) \tag{10}$$

is accomplished by gluing (smooth or singular) solid tori to the boundary tori of N_∞ (i.e. to the smooth H_i).

Correspondingly, the transition

$$\underbrace{N_\infty - \mathcal{N}_\varepsilon(\Sigma_\infty^{(0)}) - \Sigma_\infty}_{=:Y_\infty} \rightsquigarrow X_n - \mathcal{N}_\varepsilon(\Sigma_n^{(0)}) - \Sigma_n \tag{11}$$

between the smooth parts is made by gluing (smooth) solid tori to some of the boundary tori and removing the remaining boundary tori.

Every smooth cusp of X_∞ corresponds to a smooth boundary torus of N_∞. According to Lemma 7.6, we may pass to a subsequence such that the Dehn fillings at every such torus (by smooth or singular solid tori) are pairwise different.

We distinguish three cases:

Case 1. *Infinitely many Dehn fillings by smooth solid tori.* We look at the smooth parts. Then by Proposition 7.3 we have the situation that infinitely many Dehn filings at the same finite volume hyperbolic manifold produce homeomorphic manifolds. This contradicts Thurston's hyperbolic Dehn filling theorem, which implies that those manifolds have different hyperbolic volumes.

Case 2. *X_∞ has smooth cusps and all Dehn fillings use singular solid tori.* Let $Y_\infty := N_\infty - \mathcal{N}_\varepsilon(\Sigma_\infty^{(0)}) - \Sigma_\infty$ and let $Y_n := f_n(Y_\infty)$ be its image under the bilipschitz embedding f_n. By our description of (11), $f_n|_{Y_\infty}$ may be isotoped to an embedding

$$f'_n : Y_\infty \hookrightarrow X_n - \mathcal{N}_\varepsilon(\Sigma_n^{(0)}) - \Sigma_n$$

whose image is the complement of a union of open tubular neighborhoods of singular circles. We use our assumption that the cone manifolds X_n have the same topological type as the orbifold $\mathcal{O}$, i.e. we have an embedding $(X_n, \Sigma_{X_n}) \hookrightarrow (|\mathcal{O}|, \Sigma_{\mathcal{O}})$, that we compose with f'_n to obtain a new embedding

$$\iota_n : Y_\infty \hookrightarrow M := |\mathcal{O}| - \mathcal{N}(\Sigma^{(0)}) - \Sigma^{(1)} \tag{12}$$

onto the complement of a disjoint union of singular solid tori.

The ι_n are homotopy equivalences. Note that Y_∞ is homotopy equivalent to a complete hyperbolic manifold with finite volume and totally geodesic boundary; this follows, without using Proposition 7.3, directly from the original assumption that M is hyperbolic. We consider the homotopy equivalence

$$\iota_n^{-1} \circ \iota_1 : Y_\infty \longrightarrow Y_\infty$$

where ι_n^{-1} denotes a homotopy inverse for ι_n. After passing to a subsequence, and possibly replacing the first embedding ι_1, we may assume that the $\iota_n^{-1} \circ \iota_1$ preserve the toral boundary components. Lemma 7.6 implies that the induced self homotopy equivalences of each boundary torus are pairwise distinct. On the other hand, by Mostow Rigidity, there are, up to homotopy, only finitely many self homotopy equivalences of Y_∞, and we obtain a contradiction.

Case 3. *X_∞ has no smooth cusp.* In this situation, (10) says that X_∞ and X_n have the same topological type. This finishes the proof of Theorem 7.1 when cone angles are bounded above away from π. □

7.2. *The case when cone angles approach the orbifold angles.* In this case X_∞ is a hyperbolic orbifold that has a thin-thick decomposition, by the hypothesis about the cone angles and the bound on the volume. Let N_∞ be a compact core of X_∞ obtained by truncating its cusps along horospherical cross sections H_i. The H_i are now smooth tori, pillows or turnovers. By geometric convergence, for n large enough, we have $(1+\varepsilon_n)$-bilipschitz embeddings

$$f_n\colon (N_\infty, \Sigma_\infty \cap N_\infty) \to (X_n, \Sigma_n),$$

with $\varepsilon_n \to 0$. Since X_n is a cone structure on $\mathcal{O}$, we view the image $f_n(N_\infty) \subset X_n$ as a suborbifold of $\mathcal{O}$, which we denote by $N_n \subset \mathcal{O}$. As a 3-orbifold, N_n is homeomorphic to N_∞.

Lemma 7.7. *For n sufficiently large, each component of $\mathcal{O} - \mathrm{int}(N_n)$ is irreducible.*

Proof. We assume that $\mathcal{O} - \mathrm{int}(N_n)$ contains a spherical 2-suborbifold F^2 which is essential. Since $\mathcal{O}$ is irreducible, F^2 bounds a discal 3-orbifold Δ^3, which contains N_n, for n sufficiently large. Let ρ_n denote the holonomy representation of X_n. We have that $\rho_n \circ f_{n*}$ fixes a point of $\mathbb{H}^3$, because $f_n(N_\infty)$ is contained in a discal 3-orbifold. This is impossible, because $\rho_n \circ f_{n*}$ converges to the holonomy representation of N_∞, which cannot fix a point of $\overline{\mathbb{H}^3}$. □

From the smallness of $\mathcal{O}$ and the previous lemma we obtain:

Corollary 7.8. *Every component of $\mathcal{O} - \mathrm{int}(N_n)$ is either a finite quotient of a solid torus* (*i.e. a solid torus or a solid pillow, possibly singular*) *or a singular neck containing a Euclidean turnover in $\partial\mathcal{O}$.*

According to this corollary, if all the horospherical sections H_i are turnovers then $N_\infty \cong \mathcal{O}$ and X_∞ is a cone structure on $\mathcal{O}$.

Now we assume that some of the horospherical sections $H_1, \ldots, H_p$ are tori or pillows and we look for a contradiction. For $i = 1, \ldots, p$, let $\lambda_{i,n}$ be an essential curve on H_i so that $f_n(\lambda_{i,n})$ bounds a finite quotient of a solid toral component of $\mathcal{O} - \mathrm{int}(N_n)$. First we prove the analogue of Lemma 7.6.

Lemma 7.9. *For each i, $\lim_{n\to\infty} \lambda_{i,n} = \infty$.*

Proof. Suppose that the lemma is false. Then, by passing to a subsequence and changing the indices of the H_i, we can assume that the curves $\lambda_{1,n}$ represent a fixed class $\lambda_1 \in \pi_1(H_1)$ independent of n. Let ρ_n and ρ_∞ denote the holonomy representation of X_n and X_∞, respectively. Since the curves $f_n(\lambda_1)$ are compressible in $\mathcal{O}$, their holonomies $\rho_n(f_n(\lambda_1))$ are either trivial or elliptic with an angle that does not converge to zero. The holonomy $\rho_\infty(\lambda_1)$ is nontrivial and parabolic. Thus we obtain a contradiction because $\rho_\infty(\lambda_1)$ is the limit of $\rho_n(f_n(\lambda_1))$. $\square$

According to this lemma, Dehn fillings on the hyperbolic orbifold N_∞ along infinitely many different meridian curves produce always the same orbifold $\mathcal{O}$. This contradicts the hyperbolic Dehn filling theorem for orbifolds [DuM] (cf. [BoP, App. B]) because the results of surgery along those curves can be distinguished either by an estimate of the volume (obtained from Schläfli's formula) or the length of the shortest geodesics. $\square$

7.3. *Putting a* CAT(-1)*-structure on the smooth part of a cone manifold.* Let X be a hyperbolic cone 3-manifold with cone angles $\leq \pi$ and totally geodesic boundary. To simplify our discussion, we will assume that X is *geometrically finite* in the sense that it has finitely many ends and all ends are cusps. The next proposition is used in Subsection 7.1 to show that the smooth part of X is homotopically atoroidal.

Proposition 7.10. *The compact core of X^{smooth} admits a metric of nonpositive curvature with piecewise totally geodesic boundary. Moreover, the metric is strictly negatively curved away from the boundary tori corresponding to the smooth cusps and singular closed geodesics of X.*

Step 1. *Removing neighborhoods of the singular vertices and truncating singular cusps.* Consider first a vertex $v \in \Sigma_X^{(0)}$. We choose a small positive number $\rho_v < \frac{1}{2} r_{\text{inj}}(v)$ and denote by $w_1, w_2, w_3 \in \Sigma^{(1)}$ the three singular points at distance ρ_v from v. We take the convex hull of $\{w_1, w_2, w_3\}$ inside the closed ball $\overline{B_{\rho_v}(v)}$, and denote by U_v the interior of the convex hull. Its closure $\bar{U}_v$ can be obtained by doubling a hyperbolic simplex along three of its faces; ∂U_v is the union of two geodesic triangles glued along their boundaries.

Consider now a singular cusp $\mathcal{C} \subset X$ with horospherical cross section H. Since X is orientable, H is a Euclidean cone sphere with three or four cone points. As before we form the convex hull of $H \cap \Sigma_X$ inside $\mathcal{C}$ and denote its interior by $U_{\mathcal{C}}$. Then $\partial U_{\mathcal{C}}$ is piecewise geodesic with vertices in $H \cap \Sigma_X$.

By taking out the neighborhoods U_v around all singular vertices v and truncating all singular cusps $\mathcal{C}$, we obtain a hyperbolic cone manifold

$$X' := X - \bigcup_v U_v - \bigcup_{\mathcal{C}} U_{\mathcal{C}}$$

with piecewise totally geodesic concave boundary.

Step 2. *Removing neighborhoods around the singular edges.* The cone manifold X' has no singular vertices any more. Its singular locus $\Sigma_{X'} := \Sigma_X \cap X'$ consists of closed singular geodesics and of singular segments with endpoints in the boundary. We now treat the latter ones.

Consider a singular edge $\sigma = \overline{ww'}$ in $X' \cap \Sigma_X$ with endpoints in $w, w' \in \partial X'$. We will work inside a tubular ρ_σ-neighborhood T of σ in X' with small radius ρ_σ. Choose an interior point m of σ and a little totally geodesic disc Δ orthogonal to σ and centered at m. Consider one of the endpoints of σ, say w. The boundary of X' is concave at w. (This includes the possibility of its being totally geodesic, which we regard as weak concavity.) The link Λ_w of $w \in X'$ is a spherical polygon (in most cases a bigon) with one cone point and concave boundary. We denote its vertices by $\xi_1, \xi_2, \dots, \xi_k$. (In the totally geodesic case, $\partial\Lambda_w$ is a circle and the ξ_i are not well-defined; we then choose two opposite points ξ_1 and ξ_2 on the circle.) For $0 < \delta_i << \rho_\sigma$, let $y_i \in T \cap \partial X'$ be the points with $d(w, y_i) = \delta_i$ and $\overrightarrow{wy_i} = \xi_i$. If the δ_i are sufficiently small, then there exist boundary points $z_i \in \partial X' \cap T$ near w' and geodesic segments $c_i = \overline{y_i z_i}$ inside T intersecting Δ orthogonally. Exchanging the roles of w and w', we construct analogously geodesic segments $c'_1, \dots, c'_l$, with $l \geq 2$. We coordinate the δ_i and δ'_i so that the $k + l$ segments $c_1, \dots, c_k, c'_1, \dots, c'_l$ intersect Δ at the same distance from m. Now we form the convex hull of $c_1, \dots, c_k, c'_1, \dots, c'_l$ inside T and denote its interior by U_σ. Since the $k + l$ segments intersect Δ orthogonally, the boundary of U_σ is the union of $k + l$ totally geodesic quadrilaterals.

We perform this construction for all singular edges σ so that the closed neighborhoods $\bar{U}_\sigma$ are disjoint. Removing the neighborhoods U_σ for all singular edges with endpoints in $\partial X'$ yields a compact cone manifold

$$X^{\text{carved}} := X' - \bigcup_\sigma U_\sigma$$

with piecewise totally geodesic boundary.

LEMMA 7.11. *The cone manifold X^{carved} is locally* CAT(-1) *near the boundary.*

Proof. Since X^{carved} is everywhere locally conical, we have to check that the links at all boundary points are CAT(1). This is nontrivial only at the boundary vertices. In our notation above, these are the points y_i and z_i. The link of both y_i and z_i in X' is a concave bigon when they lie on an edge of $\partial X'$, and a half-sphere when they lie on a totally geodesic piece of $\partial X'$. In the first case, the link in X^{carved} is the intersection of two concave bigons, such that each bigon contains precisely one vertex of the other one. This link is a quadrilateral, with two opposite vertices of angle $> \pi$ and two other vertices of angle $< \pi$. Since each bigon contains precisely one vertex of the other one, the secant between the concave vertices divides the quadrilateral into two convex triangles. Hence the link contains no nonconstant closed geodesics of length $< 2\pi$ — actually no nonconstant closed geodesics at all — and therefore it is CAT(1) due to a criterion by Gromov; cf. [Gr2, §4.2 A and B]. In the other case, when the vertex lies in a totally geodesic piece of $\partial X'$, its link in X^{carved} is the intersection of a half-sphere with a concave bigon (with precisely one vertex in the half-sphere), and the argument is similar. □

Step 3. *Modification near closed singular geodesics at smooth cusps.* We truncate the smooth cusps along horospherical torus cross sections, respectively, removing open tubular neighborhoods of small radii around the singular closed geodesics. Then, by a standard double warped product construction, we perturb the metric locally near the new boundary components to a nonpositively curved Riemannian metric with totally geodesic flat boundary. We obtain a compact nonpositively curved Riemannian manifold X^{core} with piecewise totally geodesic boundary. The boundary components are either totally geodesic (flat or hyperbolic), or piecewise totally geodesic hyperbolic. Lemma 7.11 implies that X^{core} is CAT(0). Topologically, X^{core} is a compact core for X^{smooth}. This concludes the proof of Proposition 7.10. □

8. Spherical uniformization

This section and the following one are devoted to proof of the spherical uniformization theorem:

Spherical uniformization (Theorem 6.16). *Let $\mathcal{O}$ be a closed, orientable connected small* 3-*orbifold. If there exists a Euclidean cone structure on $\mathcal{O}$ with cone angles strictly less than the orbifold angles of $\mathcal{O}$, then $\mathcal{O}$ is spherical.*

Note first that a Euclidean cone structure X on $\mathcal{O}$ could have a boundary. If ∂X is nonempty, then it consists of totally geodesic turnovers. Due to our assumptions, X has cone angles $< \pi$, and the classification of noncompact Euclidean cone manifolds (cf. Section 4 above) implies that X is the product of a Euclidean turnover with an interval. In this case $\mathcal{O}$ is a suspension of a spherical turnover and therefore obviously spherical.

From now on we assume that the Euclidean cone structure X on $\mathcal{O}$ is closed. Then X and $\mathcal{O}$ have the same homeomorphism type (as pairs of topological space and singular locus), and we can consider X as a Euclidean cone *metric* on $\mathcal{O}$.

For the proof of Theorem 6.16 we distinguish three cases which are listed in Definition 8.2.

Definition 8.1. A singular vertex of a (locally) orientable 3-orbifold is called

- *dihedral* if its local isotropy group is a dihedral group, and
- *platonic* otherwise. (Its local isotropy group then is the group of orientation-preserving isometries of a platonic solid.)

Definition 8.2. A (locally) orientable 3-orbifold is called

- of *cyclic type* if its singular locus is not empty and has no vertex,
- of *dihedral type* if it has singular vertices and all vertices are dihedral,
- of *platonic type* if it has a platonic singular vertex.

The cyclic case relies on Hamilton's theorem, as in [BoP]. We reduce the dihedral case to the cyclic one, by using a finite covering argument. The platonic case relies on a deformation argument of spherical cone structures on $\mathcal{O}$ given in Section 9.

Previous to all these cases, in the next subsection we prove that $\pi_1(\mathcal{O})$ is finite.

8.1. *Nonnegative curvature and the fundamental group.* In this section, we relax the condition on cone angles and consider Euclidean cone 3-manifolds with cone angles $\leq \alpha < 2\pi$. In particular, these are Alexandrov spaces of nonnegative curvature, and we will prove the following result in the spirit of the Bonnet-Myers bounded diameter theorem for Riemannian manifolds with lower positive curvature bound.

Proposition 8.3. *Let X be a connected Euclidean cone 3-manifold with cone angles $\leq \alpha < 2\pi$ and nonempty singular set Σ_X. Suppose that Γ is an infinite discrete group acting properly discontinuously on X by isometries. Then Γ is virtually cyclic.*

If Γ acts moreover cocompactly, then X splits isometrically as the product $Y \times \mathbb{R}$ of the real line with a closed Euclidean cone surface Y.

Proof. Step 1. *There are no singular vertices.* Suppose that X contains singular vertices and consider a diverging sequence $x_1, x_2, \ldots$ of distinct vertices in the same Γ-orbit. We fix a base point $p \in X$ and denote by $v_i \in \Lambda_p$ the direction of the segment $\overline{px_i}$. We may assume without loss of generality that the v_i converge,

$$v_i \longrightarrow v_\infty, \tag{13}$$

and that $\frac{d_{i+1}}{d_i} \to \infty$. Applying triangle comparison (in the Euclidean plane) to the triangles $\Delta(p, x_i, x_{i+1})$ for large i, we obtain that

$$\liminf_{i\to\infty}\big(\angle_p(x_i, x_{i+1}) + \angle_{x_i}(p, x_{i+1})\big) \geq \pi.$$

In view of (13), this means that $\angle_{x_i}(p, x_{i+1}) \to \pi$. On the other hand, by the Diameter Rigidity Theorem for CAT(1)-spaces, we have that $\mathrm{diam}(\Lambda_{x_i}) = d < \pi$, a contradiction. This shows that $\Sigma_X^{(0)}$ is empty.

Step 2. *Σ_X has finitely many connected components.* The argument is similar. Assume that there are infinitely many connected components σ_i of Σ_X and consider shortest segments $\overline{px_i}$ from p to σ_i. Let w_i be the direction of $\overline{px_i}$ at x_i. After passing to a subsequence, we conclude as before that $\angle_{x_i}(p, x_{i+1}) \to \pi$. This is absurd because, due to our upper bound on cone angles, we have $\mathrm{rad}(\Lambda_{x_i}, w_i) \leq \max(\frac{\alpha}{2}, \frac{\pi}{2})$ where the radius $\mathrm{rad}(\Lambda_{x_i}, w_i)$ of Λ_{x_i} at w_i is defined as the Hausdorff distance from Λ_{x_i} to the one point subset $\{w_i\}$. Thus Σ_X has finitely many connected components.

Step 3. It follows that a finite index subgroup of Γ preserves one (each) singular component. The discontinuity of the action implies that the components of Σ_X are complete geodesics and that Γ is virtually cyclic.

If the action of Γ is in addition cocompact then X is quasi-isometric to $\mathbb{R}$ and has therefore two ends. We apply the Splitting Theorem, cf. [BBI, Th. 10.5.1], to conclude that X splits as a metric product of $\mathbb{R}$ and an Alexandrov space Y of nonnegative curvature. Then Y must be a closed Euclidean cone surface. □

The first step in the proof of Theorem 6.16 is the following lemma:

LEMMA 8.4. *Let $\mathcal{O}$ be a closed orientable* irreducible 3-*orbifold. If there exists a closed Euclidean cone structure on $\mathcal{O}$ with cone angles strictly less than the orbifold angles of $\mathcal{O}$, then $\pi_1(\mathcal{O})$ is finite.*

Proof. The Euclidean cone structure X on $\mathcal{O}$ (with cone angles strictly less than the orbifold angles of $\mathcal{O}$) lifts to a Euclidean cone structure $\tilde{X}$ on the universal cover $\widetilde{\mathcal{O}}$. The Euclidean cone manifold $\tilde{X}$ has cone angles $\leq \omega < 2\pi$, for some constant $0 < \omega < 2\pi$. In addition, the fundamental group $\pi_1(\mathcal{O})$ acts isometrically on $\tilde{X}$.

If $\pi_1(\mathcal{O})$ is infinite, then Proposition 8.3 shows that $\pi_1(\mathcal{O})$ is virtually cyclic and that $\tilde{X}$ splits as a metric product $\mathbb{R} \times \tilde{Y}^2$, where $\tilde{Y}^2$ is a closed Euclidean cone 2-manifold. Since the action of $\pi_1(\mathcal{O})$ preserves this product, $\tilde{Y}^2$ covers a totally geodesic surface Y in X. The cone surface Y is a Euclidean cone structure on a spherical turnover. This turnover is essential in $\mathcal{O}$, contradicting the irreducibility. □

8.2. *The cyclic case.* Suppose that $\mathcal{O}$ is of cyclic type. The following lemma due to M. Feighn allows us to apply geometrization results for manifolds.

LEMMA 8.5 ([Fei]). *If a closed orientable* 3-*orbifold of cyclic or dihedral type has finite fundamental group, then it is very good.*

Remark 8.6. A closed, orientable, irreducible very good 3-orbifold with finite fundamental group is small. This is a consequence of the equivariant Dehn Lemma (cf. [JR], [MY1]).

We lift the Euclidean cone metric to the universal covering of $\mathcal{O}$, which is a compact manifold denoted by $\widetilde{\mathcal{O}}$. Thus we have a $\pi_1(\mathcal{O})$-invariant Euclidean cone metric on $\widetilde{\mathcal{O}}$ with cone angles $< 2\pi$. This metric can easily be desingularized to a $\pi_1(\mathcal{O})$-invariant smooth Riemannian metric of nonnegative sectional curvature, because the singular components are circles. More precisely, the singular locus is locally a product of a singular disc with $\mathbb{R}$. This disc is isometric to the neighborhood of the tip of a cone in Euclidean space described by the equation $z = k\sqrt{x^2 + y^2}, z \geq 0$. It suffices to round the tip of the cone with a metric of nonnegative curvature invariant by rotations, that only depends on the distance to the singular locus.

By Hamilton's theorem [Ha1, 2], it follows that $\widetilde{\mathcal{O}}$ admits a $\pi_1(\mathcal{O})$-invariant smooth metric locally modelled on S^3, $S^2 \times \mathbb{R}$ or $\mathbb{R}^3$. Only the spherical case is possible because $\pi_1(\mathcal{O})$ is finite. Thus $\mathcal{O}$ is spherical in the cyclic case.

8.3. *The dihedral case.* Suppose that $\mathcal{O}$ is of dihedral type.

There exists a singular edge $e_0 \subset \Sigma$ with the following properties:

– e_0 has two different endpoints, and

– the branching index of every other edge of Σ adjacent to e_0 is 2.

To prove the existence of e_0, notice that any edge with label > 2 satisfies the properties. In addition, if all labels are $= 2$, there is always an edge with different endpoints.

The covering provided in the following result will be useful for the reduction to the cyclic case.

PROPOSITION 8.7. *There exists a finite regular covering*

$$p\colon \hat{\mathcal{O}} \to \mathcal{O}$$

of orbifolds such that $\hat{\mathcal{O}}$ *is small of cyclic type and its branching locus is* $\hat{\Sigma} = p^{-1}(e_0)$. *In addition* p *preserves the ramification index of* e_0.

Proof. Let $\mathcal{O}'$ be the orbifold obtained from $\mathcal{O}$ by removing the open edge e_0 from the branching locus. (This change of the orbifold structure amounts to putting the label 1 on e_0.) Since all edges of $\mathcal{O}$ adjacent to an endpoint of e_0 have label 2, $\mathcal{O}'$ is still an orbifold.

LEMMA 8.8. *The fundamental group of $\mathcal{O}'$ is finite.*

Proof. We use the same Euclidean cone metric on $\mathcal{O}$ as above and consider it as a singular metric of nonnegative curvature on $\mathcal{O}'$. We argue by contradiction as in the proof of Lemma 8.4. If $\pi_1(\mathcal{O}')$ were infinite, $\widetilde{\mathcal{O}'}$ would split metrically as a product $\mathbb{R} \times Y'$. This is absurd because the metric on $\widetilde{\mathcal{O}'}$ has singular vertices. □

Since $\mathcal{O}'$ is of dihedral type and $\pi_1(\mathcal{O}')$ is finite, $\mathcal{O}'$ is very good by Lemma 8.5. The universal covering of $\mathcal{O}'$ induces a finite regular covering $p\colon\hat{\mathcal{O}} \to \mathcal{O}$, where $\hat{\mathcal{O}}$ is a closed orientable 3-orbifold with underlying space the universal covering of $\mathcal{O}'$, and branching locus $\hat{\Sigma} = p^{-1}(e_0)$. Notice that $\hat{\Sigma}$ is a finite collection of disjoint embedded circles, by the choice of e_0. Therefore $\hat{\mathcal{O}}$ is of cyclic type.

LEMMA 8.9. *$\hat{\mathcal{O}}$ is small.*

Proof. Since $\mathcal{O}$ is irreducible and very good, by the equivariant sphere theorem (cf. [DD], [JR], [MY3]) the universal cover $\widetilde{\mathcal{O}}$ of $\mathcal{O}$ (which is also the universal cover of $\hat{\mathcal{O}}$) is an irreducible 3-manifold.

Suppose that $F \subset \hat{\mathcal{O}}$ is a spherical 2-suborbifold. Since $\widetilde{\mathcal{O}}$ is irreducible, F bounds a ball quotient Q which is a cyclic 3-suborbifold. The Smith conjecture [MB] implies that Q is a discal suborbifold. (It also follows from the orbifold theorem in the cyclic case whose proof we have already completed, respectively from Corollary 1.1 in the introduction.) Hence $\hat{\mathcal{O}}$ is irreducible.

Remark 8.10. Irreducibility of $\hat{\mathcal{O}}$ also follows from the irreducibility of $\mathcal{O}$ by an equivariant spherical 2-orbifold theorem ([Mai], [TY2]), whose proof relies on PL least area techniques for 2-orbifolds that generalize the notion of PL least area surfaces introduced in [JR]. However, for completeness, we use here the fact that $\mathcal{O}$ is very good with finite fundamental group to give a direct argument.

To see that $\hat{\mathcal{O}}$ is small, suppose that $F \subset \hat{\mathcal{O}}$ is an essential 2-suborbifold. By irreducibility, F cannot be spherical or bad and therefore has infinite fundamental group. Since $\pi_1(\hat{\mathcal{O}})$ is finite, F lifts to a compressible surface in the universal covering. The equivariant Dehn lemma implies that F has a compressing disc. □

This finishes the proof of Proposition 8.7. □

We consider the compact 3-orbifold $\mathcal{O}_0 = \mathcal{O} - \mathcal{N}(e_0)$ obtained by removing a regular neighborhood of the singular edge e_0.

LEMMA 8.11. *The orbifold $\mathcal{O}_0 = \mathcal{O} - \mathcal{N}(e_0)$ is Haken and topologically atoroidal.*

Proof. We first prove the irreducibility of $\mathcal{O}_0$. Let $S \subset \mathcal{O}_0$ be a spherical 2-suborbifold. It bounds a discal 3-suborbifold in $\mathcal{O}$, since $\mathcal{O}$ is irreducible. If it does not bound a discal 3-orbifold in $\mathcal{O}_0$, then a neighborhood $\mathcal{N}(e_0)$ is contained in the interior of a discal 3-suborbifold of $\mathcal{O}$. This is impossible, since e_0 is a singular edge with two distinct vertices in Σ.

It is clear that $\mathcal{O}_0$ does not contain any Euclidean or hyperbolic turnover, because $\mathcal{O}$ cannot contain such a turnover by smallness.

To see that $\mathcal{O}_0$ is topologically atoroidal, suppose that $T \subset \mathcal{O}_0$ is a nonsingular torus or a pillow which is incompressible in $\mathcal{O}_0$. Since $\mathcal{O}$ is small, T must be compressible in $\mathcal{O}$, and the compression discal 2-suborbifold must meet e_0. Hence, by irreducibility of $\mathcal{O}$, T bounds a solid pillow containing e_0 and thus is parallel to $\partial\mathcal{O}_0$. The case that T is a smooth torus cannot occur. $\square$

Now we can apply Thurston's hyperbolization for Haken orbifolds; cf. [BoP, Ch. 8]. It follows that $\mathcal{O}_0$ is Seifert or hyperbolic.

The Seifert case is quickly treated: $\mathcal{O}$ is obtained from $\mathcal{O}_0$ by gluing a solid pillow P to its boundary. If the meridian of the pillow $\partial\mathcal{O}_0$ is homotopic to the fiber of $\mathcal{O}_0$ then the irreducibility of $\mathcal{O}$ implies that $\mathcal{O}_0$ contains no essential annulus. It follows that $\mathcal{O}_0$ is a solid pillow itself. Solid pillows admit many Seifert fibrations, and we can modify the Seifert fibration so that the fiber is not a meridian of the solid pillow P. Hence the Seifert fibration extends to $\mathcal{O}$. Since a Seifert fibred 3-orbifold with finite fundamental group is spherical, $\mathcal{O}$ is spherical, i.e. the orbifold theorem holds in this case. Hence from now on we make the following:

Assumption 8.12. The orbifold $\mathcal{O}_0$ admits on its interior a complete hyperbolic structure of finite volume.

We proceed now with the proof as in Section 6.2 by starting to increase the cone angle along the singular edge e_0 and by analyzing the degenerations.

We fix some notation. Let m_0 be the ramification index of the edge e_0. For $t \in [0,1]$, let $X(t)$ denote a hyperbolic cone structure on $\mathcal{O}$, having the same prescribed cone angles as the orbifold $\mathcal{O}$ along the edges and circles of $\Sigma - e_0$ and cone angle $\frac{2\pi}{m_0}t$ along the edge e_0. $X(0)$ denotes the complete hyperbolic structure of finite volume on the interior of $\mathcal{O}_0$.

In order to study the deformations of the hyperbolic cone structure $X(t)$ while increasing t, we consider the variety of representations $\mathrm{Hom}(\pi_1(\mathcal{O}_0), \mathrm{PSL}_2(\mathbb{C}))$ and the variety of characters $\mathcal{X}(\mathcal{O}_0)$. As in Theorem 6.4, the (square of the) trace of the meridian around e_0 is a local parameter for $\mathcal{X}(\mathcal{O}_0)$ near the complete structure; cf. [BoP, Th. B.2.7]. Hence the irreducible component $\mathcal{C}$ of $\mathcal{X}(\mathcal{O}_0)$ that contains the holonomy of $\mathcal{O}_0$ is a curve.

As in Section 6.2, we define

$$\mathcal{I}(\mathcal{O}) := \left\{ t \in [0,1] \;\middle|\; \begin{array}{l} \text{there exists a hyperbolic cone structure on } \mathcal{O} \\ \text{with cone angle } \frac{2\pi t}{m_0} \text{ along } e_0 \text{ and cone angles equal} \\ \text{to the orbifold angles at all other edges,} \\ \text{with holonomy in } \mathcal{C} \text{ and volume } \leq v_0 \end{array} \right\}$$

where v_0 denotes the volume of the complete structure.

By hypothesis, $0 \in \mathcal{I}(\mathcal{O})$, hence $\mathcal{I}(\mathcal{O}) \neq \emptyset$. Exactly as in Theorem 6.7 one proves that $\mathcal{I}(\mathcal{O})$ is open to the right. Since $\pi_1(\mathcal{O})$ is finite, $\mathcal{O}$ is not hyperbolic and $1 \notin \mathcal{I}(\mathcal{O})$. Let $t_\infty := \sup \mathcal{I}(\mathcal{O})$. We have that $t_\infty \notin \mathcal{I}(\mathcal{O})$ by (right) openness.

LEMMA 8.13. *For any sequence (t_n) in $\mathcal{I}(\mathcal{O})$ with $t_n \to t_\infty$, the sequence of cone manifolds $(X(t_n))_{n\in\mathbb{N}}$ collapses.*

Proof. Assume that $X(t_n)$ does not collapse. After choosing base points x_n in the thick parts of $X(t_n)$ and passing to a subsequence, $(X(t_n), x_n)$ converges geometrically to a pointed hyperbolic cone 3-manifold (X_∞, x_∞) with finite volume. The manifold X_∞ cannot be compact because $t_\infty \notin \mathcal{I}(\mathcal{O})$.

We use the finite cover $p\colon \hat{\mathcal{O}} \to \mathcal{O}$ of Proposition 8.7 and denote by $\hat{X}(t) = p^{-1}(X(t))$ the lifted hyperbolic cone structure on $\hat{\mathcal{O}}$. The $\hat{X}(t_n)$ converge to a finite regular covering $\hat{X}_\infty$ of X_∞. Since all cone angles of each $\hat{X}(t_n)$ are equal, the stability theorem (Thm. 7.1) applies and $\hat{X}_\infty$ is a hyperbolic cone structure on $\hat{\mathcal{O}}$. Now $\hat{X}_\infty$ is not compact, and its ends are singular cusps which correspond to singular vertices of $\hat{\mathcal{O}}$. But $\hat{\mathcal{O}}$ is of cyclic type, a contradiction. □

Let (t_n) be a sequence as above. We distinguish two cases according to whether the sequence $\operatorname{diam}(X(t_n))$ is bounded below away from zero or not.

If $\operatorname{diam}(X(t_n)) \geq D > 0$ for some uniform D, then the fibration theorem implies that $\mathcal{O}$ is a Seifert fibred 3-orbifold. Since $\pi_1(\mathcal{O})$ is finite, it follows that $\mathcal{O}$ is spherical.

Otherwise, up to taking a subsequence, we can assume that $\operatorname{diam}(X(t_n)) \to 0$. Then we consider the rescaled sequence $\frac{1}{\operatorname{diam}(X(t_n))} X(t_n)$. If this rescaled sequence collapses, then the fibration theorem still implies that $\mathcal{O}$ is Seifert fibred, and hence spherical.

If the rescaled sequence does not collapse, then a subsequence converges to a closed Euclidean cone manifold $X(t_\infty)$ homeomorphic to $\mathcal{O}$. We have $t_\infty < 1$, because $\pi_1(\mathcal{O})$ is finite. Thus $X(t_\infty)$ lifts to a $\pi_1(\mathcal{O})$-invariant Euclidean cone metric on the universal covering $\tilde{\mathcal{O}}$ with singular locus a link and cone angle $t_\infty 2\pi < 2\pi$. We conclude as in the cyclic case that $\mathcal{O}$ is spherical. □

8.4. *The platonic case.* Suppose that $\mathcal{O}$ is of platonic type. The proof of Theorem 6.16 in this case is by induction on the number of platonic vertices in the branching locus.

At each platonic vertex we have one singular edge with label 2 and two edges with label $3, 4$ or 5. We fix a singular edge e of $\mathcal{O}$ with label $n_e > 2$ such that at least one of its adjacent vertices is platonic.

Let $\mathcal{O}'$ be the orbifold obtained from $\mathcal{O}$ by replacing the branching index n_e of e by 2. We want to apply the induction hypothesis to $\mathcal{O}'$, because it has fewer platonic vertices than $\mathcal{O}$. To do it, we need the following lemma:

LEMMA 8.14. *The orbifold $\mathcal{O}'$ is small.*

Proof. By assumption, $\mathcal{O}$ and hence also $\mathcal{O}'$ are closed. The lemma follows from the fact that, for closed orbifolds, smallness is a property independent of the labels of the branching locus. □

Notice that $\mathcal{O}'$ has orbifold angles greater than or equal to those of $\mathcal{O}$ because one label has decreased. Thus the Euclidean cone structure on $\mathcal{O}$ given in the hypothesis of Theorem 6.16, when viewed on $\mathcal{O}'$ still has cone angles strictly less than the orbifold angles. It follows from the induction hypothesis that $\mathcal{O}'$ is spherical.

The induction step, and hence the conclusion of the proof of Theorem 6.16, is due to the following result which we will prove in Section 9.3.

PROPOSITION 8.15. *The spherical structure on $\mathcal{O}'$ can be deformed, through a continuous family of spherical cone metrics, to a spherical structure on $\mathcal{O}$.*

9. Deformations of spherical cone structures

This section is devoted to proving Proposition 8.15. Some preliminaries on varieties of representations are required.

9.1. *The variety of representations into* SU(2). Let Γ be a finitely generated group. The variety of representations $\mathrm{Hom}(\Gamma, \mathrm{SU}(2))$ is compact. The group SU(2) acts on $\mathrm{Hom}(\Gamma, \mathrm{SU}(2))$ by conjugation, and the quotient

$$\mathcal{X}(\Gamma, \mathrm{SU}(2)) = \mathrm{Hom}(\Gamma, \mathrm{SU}(2))/\mathrm{SU}(2)$$

is also compact, but in general not algebraic.

The variety of characters. The action by conjugation of $\mathrm{SL}(2, \mathbb{C})$ on the variety of representations $\mathrm{Hom}(\Gamma, \mathrm{SL}(2, \mathbb{C}))$ is algebraic. The quotient

$$\mathcal{X}(\Gamma, \mathrm{SL}(2, \mathbb{C})) = \mathrm{Hom}(\Gamma, \mathrm{SL}(2, \mathbb{C}))//\mathrm{SL}(2, \mathbb{C})$$

of this action provided by geometric invariant theory carries a natural structure as an affine algebraic subset of $\mathbb{C}^N$ defined over $\mathbb{Q}$ [MoS]; points can be interpreted as characters of representations $\Gamma \to \mathrm{SL}(2,\mathbb{C})$. The conjugacy class of a representation $\Gamma \to \mathrm{SU}(2)$ is determined by its character, and we have a natural inclusion

$$\mathcal{X}(\Gamma, \mathrm{SU}(2)) \subseteq \mathcal{X}(\Gamma, \mathrm{SL}(2,\mathbb{C})). \tag{14}$$

Trace functions. For $\gamma \in \Gamma$, the trace function

$$\begin{aligned} \mathrm{Hom}(\Gamma, \mathrm{SL}(2,\mathbb{C})) &\to \mathbb{C} \\ \rho \quad &\mapsto \mathrm{trace}(\rho(\gamma)) \end{aligned}$$

induces an algebraic function

$$I_\gamma : \mathcal{X}(\Gamma, \mathrm{SL}_2(\mathbb{C})) \to \mathbb{C}.$$

The ambient space $\mathbb{C}^N$. The embedding of $\mathcal{X}(\Gamma, \mathrm{SL}(2,\mathbb{C}))$ into $\mathbb{C}^N$ is realized by taking as coordinates the functions I_γ, where γ runs through the words of length at most three in the generators of Γ [GM].

Since traces of matrices in SU(2) are real, we have the embedding

$$\mathcal{X}(\Gamma, \mathrm{SU}(2)) \subset \mathcal{X}(\Gamma, \mathrm{SL}_2(\mathbb{C})) \cap \mathbb{R}^N$$

of $\mathcal{X}(\Gamma, \mathrm{SU}(2))$ into a real algebraic variety. One can show that $\mathcal{X}(\Gamma, \mathrm{SL}_2(\mathbb{C})) \cap \mathbb{R}^N = \mathcal{X}(\Gamma, \mathrm{SU}(2)) \cup \mathcal{X}(\Gamma, \mathrm{SL}_2(\mathbb{R}))$ [MoS]. The following lemma implies that $\mathcal{X}(\Gamma, \mathrm{SU}(2))$ is locally $\mathbb{R}$-algebraic as characters of non-abelian representations.

Lemma 9.1 ([Po1, lemme 5.25]). *Let* $[\rho] \in \mathcal{X}(\Gamma, \mathrm{SU}(2))$ *be the conjugacy class of a nonabelian representation. There exists an open neighborhood* $U \subset \mathbb{R}^N$ *of* $[\rho]$ *such that*:

$$\mathcal{X}(\Gamma, \mathrm{SU}(2)) \cap U = \mathcal{X}(\Gamma, \mathrm{SL}_2(\mathbb{C})) \cap \mathbb{R}^N \cap U.$$

9.2. *Lifts of holonomy representations into* $\mathrm{SU}(2) \times \mathrm{SU}(2)$ *and spin structures.* We recall that an element $(A, B) \in \mathrm{SU}(2) \times \mathrm{SU}(2) \cong \mathrm{Spin}(4)$ acts on $\mathrm{SU}(2) \cong S^3$ by

$$x \mapsto AxB^{-1}.$$

In particular, the kernel of the action is the order two subgroup $\{\pm(\mathrm{Id}, \mathrm{Id})\}$. The following lemma is classical:

Lemma 9.2. *The element* $(A, B) \in \mathrm{SU}(2) \times \mathrm{SU}(2)$ *acts on* $\mathrm{SU}(2) \cong S^3$ *as a rotation* (*i.e. has fixed points*) *if and only if* $\mathrm{trace}(A) = \mathrm{trace}(B)$. *Trace and rotation angle are related by*

$$\mathrm{trace}(A) = \pm 2 \cos\left(\frac{\alpha}{2}\right). \tag{15}$$

One can view SO(4) as the frame bundle on $S^3 = \mathrm{SO}(4)/\mathrm{SO}(3)$. The unique spin structure on S^3 is given by the canonical projection $\mathrm{Spin}(4) \to \mathrm{SO}(4)$. Given a spherical manifold N, not necessarily complete, with holonomy representation $\rho : \pi_1(N) \to \mathrm{SO}(4)$, the spin structures on N correspond to lifts of ρ to a representation into Spin(4). (The lift can be obtained from a spin structure by developing it, using the natural connection, onto Spin(4) and taking its holonomy.)

Let $p_i\colon \mathrm{SU}(2) \times \mathrm{SU}(2) \to \mathrm{SU}(2)$ denote the projection to the i-th factor, for $i = 1, 2$.

LEMMA 9.3. *Let $\phi : \pi_1(M) \to \mathrm{SU}(2) \times \mathrm{SU}(2)$ be the lift of the holonomy of a spherical cone manifold. Then both $p_1 \circ \phi$ and $p_2 \circ \phi$ are non-abelian unless Σ is a link and $M = X - \Sigma$ is Seifert fibred.*

Proof. Assume for instance that $p_1 \circ \phi$ is abelian. This means that the image of $p_1 \circ \phi$ is contained in a diagonalizable subgroup $\cong S^1$. Therefore ϕ preserves the corresponding Hopf fibration on S^3. It follows easily from this that Σ is a link and M is Seifert fibred (cf. [Po2, Lemma 9.1]). □

Let $\mu_1, \dots \mu_q$ be meridians for the singular edges and circles in Σ.

THEOREM 9.4 (Local parametrization). *Let $\mathcal{O}$ be a spherical orbifold such that $M = \mathcal{O} - \Sigma$ is not Seifert fibred. If $\phi : \pi_1(M) \to \mathrm{SU}(2) \times \mathrm{SU}(2)$ is a lift of the holonomy, then both $[p_1 \circ \phi]$ and $[p_2 \circ \phi]$ are smooth points of $\mathcal{X}(M, \mathrm{SU}(2))$. Moreover,*

$$(I_{\mu_1}, \dots, I_{\mu_q}) : \mathcal{X}(M, \mathrm{SU}(2)) \to \mathbb{R}^q$$

is a local diffeomorphism at both points $[p_1 \circ \phi]$ and $[p_2 \circ \phi]$.

This theorem is an infinitesimal rigidity result for spherical orbifolds, and its proof is postponed to the last subsection 9.5. It will be obtained from a cohomology computation, using the fact that spherical orbifolds are finitely covered by S^3.

9.3. *The deformation space of spherical structures.* The aim of this section is to prove Proposition 8.15. We adopt the notation of Section 8.4. Let μ_e denote the meridian of the edge e. We recall that $\mathcal{X}(M, \mathrm{SU}(2))$ is contained in the real algebraic set $V = \mathcal{X}(M, \mathrm{SL}_2(\mathbb{C})) \cap \mathbb{R}^N$; cf. (14). Let ϕ_0 be the lift to $\mathrm{SU}(2) \times \mathrm{SU}(2)$ of the holonomy representation of $\mathcal{O}'$ corresponding to the choice of spin structure on M. Its conjugacy class $[\phi_0]$ is contained in

$$\mathcal{X}(M, \mathrm{SU}(2) \times \mathrm{SU}(2)) = \mathcal{X}(M, \mathrm{SU}(2)) \times \mathcal{X}(M, \mathrm{SU}(2)) \subseteq V \times V.$$

According to Theorem 9.4, $\mathcal{X}(M, \mathrm{SU}(2) \times \mathrm{SU}(2))$ is locally bianalytic to $\mathbb{R}^{2q}$ at $[\phi_0]$ and, due to Lemma 9.1, is a neighborhood of $[\phi_0]$ in $V \times V$.

Let $p_1 : \mathrm{SU}(2) \times \mathrm{SU}(2) \to \mathrm{SU}(2)$ denote the projection to the first factor. Consider the algebraic subset

$$W = \left\{ ([\rho_1],[\rho_2]) \in V \times V \;\middle|\; \begin{array}{l} I_{\mu_i}([\rho_1]) = I_{\mu_i}([\rho_2]), \text{ for each meridian } \mu_i \\ I_{\mu_i}([\rho_1]) = I_{\mu_i}([p_1(\phi_0)]), \text{ for each meridian} \\ \hfill \mu_i \neq \mu_e \end{array} \right\}.$$

By Lemma 9.2, these equations are the algebraic conditions for a representation of $\pi_1(M)$ in $\mathrm{SU}(2) \times \mathrm{SU}(2)$ to be the lift of a representation in $\mathrm{SO}(4)$ with the properties:

– the images of the meridians are rotations;

– the rotation angles of all meridians are fixed except for the meridian μ_e of the edge e.

Let W_0 be the irreducible component of W containing $[\phi_0]$. By Theorem 9.4, W_0 is a real algebraic curve and $I_{\mu_e} \circ p_1$ restricted to W_0 is a smooth local parameter near $[\phi_0]$. In particular it is nonconstant on W_0.

A neighborhood of $[\phi_0]$ in W_0 can be lifted to a curve of representations $\pi_1(M) \to \mathrm{SU}(2) \times \mathrm{SU}(2)$ which are lifts of holonomies of spherical cone structures on $\mathcal{O}'$. The fact that the trace $I_{\mu_e} \circ p_1$ is a smooth local parameter on W_0 near $[\phi_0]$ implies that the cone angle at e is a (continuous) local parameter for the family of spherical cone structures near the orbifold structure. It takes values in a neighborhood of π.

We take $\mathcal{S}^{\pm}$ to be the connected component of the semi-algebraic set

$$\{([\rho_1],[\rho_2]) \in W_0 \mid 0 \leq \pm I_{\mu_e}([\rho_1]) \leq 2\cos(\pi/n_e)\}$$

that contains $[\phi_0]$. One of these two sets, say $\mathcal{S}^+$, contains representations arising from cone metrics with cone angle $< \pi$ at e.

Lemma 9.5. *All representations in $\mathcal{S}^+$ are lifts of holonomy representations for spherical cone structures on $\mathcal{O}'$ with cone angle at e in $[\frac{2\pi}{n_e}, \pi]$.*

Proof. Let A be the subset of representations in $\mathcal{S}^+$ that are such lifts. By our previous discussion, A contains a neighborhood of the endpoint $[\phi_0]$ of $\mathcal{S}^+$.

Openness of A. Lemmas 9.1 and 9.3 imply that perturbations of representations in A still take values in $\mathrm{SU}(2) \times \mathrm{SU}(2)$. Moreover, perturbations of holonomy representations are induced by perturbations of cone structures; cf. [CEG], [Po2].

Closedness of A. Consider a sequence of points $[\phi_n] \in A$ converging to $[\phi_\infty] \in \mathcal{S}^+$. We have to show that $[\phi_\infty] \in A$.

For $n \in \mathbb{N}$, let X_n be the spherical cone manifold with holonomy lift ϕ_n. The X_n are Alexandrov spaces with curvature ≥ 1 and therefore have diameter

$\leq \pi$. Hence if the sequence $(X_n)_{n\in\mathbb{N}}$ does not collapse then, up to a subsequence, it geometrically converges to a spherical cone manifold X_∞ with the same topological type and holonomy lift ϕ_∞. Thus $[\phi_\infty] \in A$ in this case.

Assume now that the sequence (X_n) collapses. In the spherical case, this is equivalent to $\mathrm{vol}(X_n) \to 0$, because

$$\mathrm{vol}(X_n) \leq 2\pi \ \mathrm{inj}(X_n).$$

This formula can be proved by using the Dirichlet domain centered at some point $x \in X_n$ of smallest injectivity radius, since this domain embeds isometrically in a lens in S^3 of width $2\,\mathrm{inj}(x)$.

Denote by α_n the cone angle of e for the cone structure X_n. We may assume that (α_n) converges and distinguish two cases.

Case 1. $\alpha_n \to \pi$. We will apply Schläfli's formula, relying on the algebraic structure of our deformation space. $\mathcal{S}^+$ is contained in the algebraic curve W_0. Hence, after passing to a subsequence, we may assume that the $[\phi_n]$ lie on an analytic path with endpoint $[\phi_\infty]$. Since the trace $I_{\mu_e} \circ p_1$ is a nonconstant analytic function on W_0, its critical points do not accumulate and we can assume that it is monotonic on this path. So the cone angle at e is monotonically increasing towards π. Schläfli's formula then implies that the volume of the cone structure *increases*, as we approach $[\phi_\infty]$. This contradicts collapse.

Case 2. $\alpha_n \to \alpha_\infty < \pi$. Since at least one of the singular vertices in $\mathcal{O}$ adjacent to e is platonic, each X_n has at least one vertex with two cone angles uniformly bounded away from π. Since (X_n) collapses by assumption, the thick vertex lemma (Lemma 5.10) implies that $\mathrm{diam}(X_n) \to 0$.

Now we rescale X_n by the (inverse of the) diameter. By application of the thick vertex lemma again, the rescaled sequence converges to a compact Euclidean cone manifold, which is a Euclidean structure on $\mathcal{O}$ with cone angles greater than or equal to the orbifold angles of $\mathcal{O}$.

SUBLEMMA 9.6. *This Euclidean cone structure on $\mathcal{O}$ has curvature ≤ 0 in the orbifold sense* (*i.e. it is locally the quotient of a* CAT(0) *space by a finite group of isometries*).

Proof. By Gromov's criterion [Bal], all we have to check is that the link of each point is CAT(1) in the orbifold sense (i.e. a quotient of a CAT(1) space by the isotropy group). The links are quotients of the unit sphere except at the points of e. The CAT(1) property for the links of interior points of e is clear because the cone angle is greater than or equal to the orbifold angle. It is clear for the same reason at a dihedral endpoint of e. (There may be at most one.) At platonic endpoints of e it follows from Proposition 9.10 if e has different endpoints, and from its addendum if e is a loop, to be proved in the next subsection. □

By the sublemma and from Haefliger's version of the Cartan-Hadamard theorem [Hae], $\pi_1(\mathcal{O})$ is infinite. We obtain a contradiction, because $\pi_1(\mathcal{O})$ is finite by Lemma 8.4. Thus (X_n) does not collapse. This finishes the proof of Lemma 9.5. □

LEMMA 9.7. *The map*

$$\alpha : \mathcal{S}^+ \longrightarrow \left[\frac{2\pi}{n_e}, \pi\right]$$

given by the cone angle at e is surjective.

Proof. Recall that $\mathcal{S}^+$ is a closed connected subset of the real algebraic curve W_0. Moreover, $\mathcal{S}^+$ is compact because it is contained in the subset of conjugacy classes of $\mathrm{SU}(2) \times \mathrm{SU}(2)$-valued representations. As an algebraic curve, W_0 is homeomorphic to a graph with finitely many vertices. It follows that $\mathcal{S}^+$ is a compact graph.

Consider the subset

$$\mathcal{S}_0^+ := \alpha^{-1}\left(\left(\frac{2\pi}{n_e}, \pi\right)\right) \subset \mathcal{S}^+.$$

The complement $\mathcal{S}^+ - \mathcal{S}_0^+$ is finite, because the nonconstant analytic trace function has discrete level sets.

Since $\mathcal{S}_0^+$ is locally algebraic, Sullivan's local Euler characteristic theorem [Sul] implies that vertices of $\mathcal{S}_0^+$ have even valency. It follows that $\mathcal{S}_0^+$ is a (noncompact) graph with an even number of ends. Recall that a neighborhood of $[\phi_0]$ in $\mathcal{S}^+$ is a curve starting in $[\phi_0]$, and therefore precisely one end of $\mathcal{S}_0^+$ converges to $[\phi_0]$. As a consequence, there exists another end of $\mathcal{S}_0^+$ converging to a point $[\phi_1] \in \mathcal{S}^+ - \mathcal{S}_0^+$ different from $[\phi_0]$. We have that

$$\alpha([\phi_1]) \in \left\{\frac{2\pi}{n_e}, \pi\right\}.$$

We are done if $\alpha([\phi_1]) = \frac{2\pi}{n_e}$. If $\alpha([\phi_1]) = \pi$, we obtain a contradiction from the following result.

THEOREM 9.8 (de Rham, cf. [Rot]). *A spherical structure on a closed orientable connected smooth* 3*-orbifold is unique up to isometry.*

That is, since a spin structure has been fixed on M, the nonconjugate representations ϕ_0 and ϕ_1 correspond to nonisometric spherical structures on the orbifold $\mathcal{O}'$. This concludes the proof of Lemma 9.7. □

Proposition 8.15 is a direct consequence of the results in this section.

9.4. *Certain spherical cone surfaces with the* CAT(1) *property.* The results of this subsection have been used in Sublemma 9.6 to show that a certain Euclidean cone structure on the orbifold $\mathcal{O}$ is a metric that satisfies locally the CAT(0) property, by showing that the links are CAT(1).

Let Λ be a spherical cone surface with cone angles $> 2\pi$. Such a surface has curvature ≤ 1 in the *local* sense. We discuss some examples when Λ has the CAT(1) property, i.e. satisfies *global* triangle comparison with *upper* curvature bound 1.

Due to a general criterion for piecewise spherical complexes, Λ is CAT(1) if and only if it contains no nonconstant closed geodesic with length $< 2\pi$; cf. [Gr2, §4.2 A and B]. An elegant way to check the CAT(1) property was discovered by Rivin in his thesis:

THEOREM 9.9 (Rivin, cf. [RH]). *The polar dual of a compact convex polyhedron in* $\mathbb{H}^3$ *is* CAT(1).

The result extends to ideal polyhedra; cf. [ChD, Thm. 4.1.1].

The *polar dual* or *Gauß image* $G(P)$ of a convex polyhedron P in $\mathbb{H}^3$ is constructed as follows, generalizing the Gauß map for convex polyhedra in Euclidean space. For every vertex v of P we take the set $G(v)$ of all outer unit normal vectors at v. Equipped with its natural metric as a subset of the unit tangent sphere, $G(v)$ is a spherical polygon. The sides of $G(v)$ correspond to the edges of P adjacent to v. If the vertices v_1 and v_2 are joined by an edge e, we glue the polygons $G(v_1)$ and $G(v_2)$ along their sides corresponding to e. The resulting spherical complex is $G(P)$, and it is easily seen to be a spherical cone surface with cone points of angles $> 2\pi$.

As an example, relevant in Sublemma 9.6, we determine the polar dual of the platonic solids. Let P be a regular polyhedron in $\mathbb{H}^3$. (P can be a tetrahedron with face angles $< \frac{\pi}{3}$, a cube with face angles $< \frac{\pi}{2}$, an octahedron with face angles $< \frac{\pi}{3}$, a dodecahedron with face angles $< \frac{3\pi}{5}$, or an icosahedron with face angles $< \frac{\pi}{3}$.) The isometry group $\mathrm{Isom}(P)$ acts on the dual $G(P)$ as a reflection group. Let $\Delta(P)$ be the quotient 2-orbifold. It is a triangle with reflector boundary equipped with a spherical metric. Each one of the three vertices of $\Delta(P)$ corresponds respectively to the vertices, edge midpoints and face centers of P. The angles of the triangle equal the orbifold angles at the first two vertices, but the third angle is bigger than the corresponding orbifold angle. In other words, in the orbifold sense, the metric is smooth everywhere except at the third vertex, and there it has concentrated negative curvature.

For instance, if P is a (possibly ideal) icosahedron with face angles $< \frac{\pi}{3}$, then $G(P)$ is a piecewise spherical dodecahedron composed of regular pentagons with angles $\in (\frac{2\pi}{3}, \pi]$. The orbifold $\Delta(P)$ has angles $\frac{\pi}{5}, \frac{\pi}{2}$ and third angle $\in (\frac{\pi}{3}, \frac{\pi}{2}]$. The various cases yield the following examples of piecewise spherical metrics on S^2 with the CAT(1) property.

PROPOSITION 9.10. *Let Δ be a spherical 2-orbifold which is topologically a triangle with reflector boundary. Suppose that Δ is equipped with a spherical metric so that the boundary is geodesic, all angles are $\leq \frac{\pi}{2}$, and the metric is smooth everywhere except at one vertex where it has concentrated negative curvature (in the orbifold sense).*

Then the pull-back of the metric to the universal covering orbifold $\widetilde{\Delta} \cong S^2$ satisfies the CAT(1) *property.*

Only the case when Δ is cyclic is not given by the platonic solids. But in this case, $\widetilde{\Delta}$ is the suspension of a circle with length $> 2\pi$, hence also CAT(1).

The next example is related.

ADDENDUM 9.11. *Let Δ be as in the proposition and with local isotropy groups D_2, D_3 and D_3 at the vertices. Suppose that the spherical metric has angle $\frac{\pi}{2}$ at the D_2-vertex and equal angles $\in (\frac{\pi}{3}, \frac{\pi}{2}]$ at the two vertices with D_3-isotropy. Then the same conclusion holds.*

Proof. By folding these orbifolds along their symmetry axis, one obtains index-two ramified coverings over the orbifolds with vertex isotropies (D_2, D_3, D_4). The angles of the quotient orbifolds satisfy the assumptions in Proposition 9.10. The quotients have the same universal cover, and the assertion follows from Proposition 9.10. $\square$

9.5. *Proof of the local parametrization theorem.*

Proof of Theorem 9.4. Let $\rho = p_i \circ \phi \colon \pi_1(M) \to \mathrm{SU}(2)$, for $i = 1$ or 2. The proof has three main steps. In Step 1 we show that the dimension of the Zariski tangent space of $\mathcal{X}(M, \mathrm{SU}(2))$ at $[\rho]$ equals the number of meridians. In Step 2 we prove that $[\rho]$ is a smooth point of $\mathcal{X}(M, \mathrm{SU}(2))$. In Step 3 we check that the differential forms $\{dI_{\mu_1}, \ldots, dI_{\mu_q}\}$ are a basis for the cotangent space.

Step 1. *Computation of the Zariski tangent space of $\mathcal{X}(M, \mathrm{SU}(2))$ at $[\rho]$.* The Zariski tangent space is given by the cohomology group $H^1(M; \mathrm{Ad} \circ \rho)$ where we work with coefficients in the Lie algebra su(2) twisted by the adjoint representation $\mathrm{Ad} \circ \rho$.

First compute the cohomology $H^*(\mathcal{O}; \mathrm{Ad} \circ \rho)$. Notice that $\mathrm{Ad} \circ \rho \colon \pi_1(M) \to \mathrm{Aut}(\mathrm{su}(2))$ factors through $\pi_1(M) \to \pi_1(\mathcal{O})$ because we may compose the holonomy representation of the spherical structure on $\mathcal{O}$ (which in general does not lift to Spin(4)) with the (first or second) projection $\mathrm{SO}(4) \to P\mathrm{SU}(2)$ and Ad:

$$\pi_1(\mathcal{O}) \longrightarrow \mathrm{SO}(4) \longrightarrow P\mathrm{SU}(2) \longrightarrow \mathrm{Aut}(\mathrm{su}(2)).$$

We use simplicial homology. Let K be a triangulation of the underlying space of $\mathcal{O}$ compatible with the branching locus and let $\tilde{K}$ be its lift to the universal cover $\tilde{\mathcal{O}} \cong S^3$. We consider the following chain and cochain complexes:

$$C_*(K;\mathrm{Ad}\circ\rho) = \mathrm{su}(2)\otimes_{\pi_1(\mathcal{O})} C_*(\tilde{K};\mathbb{Z}),$$
$$C^*(K;\mathrm{Ad}\circ\rho) = \mathrm{Hom}_{\pi_1(\mathcal{O})}(C_*(\tilde{K};\mathbb{Z}),\mathrm{su}(2)).$$

The homology of $C_*(K;\mathrm{Ad}\circ\rho)$ is denoted by $H_*(\mathcal{O};\mathrm{Ad}\circ\rho)$ and the cohomology of $C^*(K;\mathrm{Ad}\circ\rho)$ by $H^*(\mathcal{O};\mathrm{Ad}\circ\rho)$. From the differential point of view, $H^*(\mathcal{O};\mathrm{Ad}\circ\rho)$ is the cohomology of the su(2)-valued differential forms on S^3 which are $\pi_1(\mathcal{O})$-equivariant.

LEMMA 9.12. *$H_*(\mathcal{O};\mathrm{Ad}\circ\rho)\cong 0$.*

Proof. There is a canonical projection $C_*(S^3,\mathrm{su}(2))\to C_*(\mathcal{O},\mathrm{Ad}\circ\rho)$ and, since $\pi_1(\mathcal{O})$ is finite, an averaging map $C_*(\mathcal{O},\mathrm{Ad}\circ\rho)\to C_*(S^3,\mathrm{su}(2))$ which is a section for the projection. According to [Bro, Prop. 10.4 in Ch. 3], the induced map

$$H^*(\mathcal{O},\mathrm{Ad}\circ\rho)\longrightarrow \big(H^*(S^3,\mathrm{su}(2))\big)^{\pi_1(\mathcal{O})}\cong H^*(S^3,\mathbb{R})\otimes \mathrm{su}(2)^{\pi_1(\mathcal{O})}$$

is an isomorphism. The Lie algebra su(2) does not have nontrivial elements invariant by $\mathrm{Ad}\circ\rho$, because ρ is non-abelian (Lemma 9.3). $\square$

Let $\mathcal{N}(\Sigma)$ be a tubular neighborhood of Σ and let $N=\mathcal{O}-\mathcal{N}(\Sigma)$. We remark that N is compact and that the inclusion $N\subset M$ is a homotopy equivalence, so that $H_*(M;\mathrm{Ad}\circ\rho)\cong H_*(N;\mathrm{Ad}\circ\rho)$.

From the previous lemma and by application of Mayer-Vietoris to the pair $(N,\mathcal{N}(\Sigma))$, there is a natural isomorphism

$$H^*(N;\mathrm{Ad}\circ\rho)\oplus H^*(\mathcal{N}(\Sigma);\mathrm{Ad}\circ\rho)\cong H^*(\partial N;\mathrm{Ad}\circ\rho) \tag{16}$$

induced by inclusion.

Consider the following piece of the exact sequence of the pair $(N,\partial N)$:

$$0\longrightarrow H^1(N;\mathrm{Ad}\circ\rho)\xrightarrow{i^*} H^1(\partial N;\mathrm{Ad}\circ\rho)\xrightarrow{\delta} H^2(N,\partial N;\mathrm{Ad}\circ\rho). \tag{17}$$

The injectivity of i^* comes from (16). Poincaré duality implies that $H^1(N;\mathrm{Ad}\circ\rho)$ and $H^2(N,\partial N;\mathrm{Ad}\circ\rho)$ are dual, $H^1(\partial N;\mathrm{Ad}\circ\rho)$ is dual to itself, and moreover δ and i^* are dual maps. Hence δ is surjective and

$$\dim H_2(N,\partial N;\mathrm{Ad}\circ\rho)=\dim H_1(N;\mathrm{Ad}\circ\rho)=\frac{1}{2}\dim H_1(\partial N;\mathrm{Ad}\circ\rho). \tag{18}$$

PROPOSITION 9.13. *If Σ has q edges and circles, then*

$$\dim H^1(M;\mathrm{Ad}\circ\rho)=q.$$

Proof. We have to show that $\dim H^1(\partial N; \mathrm{Ad}\circ\rho) = 2q$. Since $\partial N = \partial\mathcal{N}(\Sigma)$ and the homology is the direct sum of the homology of the connected components, it suffices to compute $\dim H_1(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho)$ assuming that Σ is connected.

If Σ is a circle, then $\partial\mathcal{N}(\Sigma)$ is a torus and $\rho|_{\pi_1(\partial\mathcal{N}(\Sigma))}$ is abelian and nontrivial. Thus $H^0(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) \cong \mathrm{su}(2)^{\mathrm{Ad}\circ\rho} \cong \mathbb{R}$. By duality, $\dim H^2(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) = 1$. Since the torus has zero Euler characteristic, it follows that $\dim H^1(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) = 2$.

When Σ is a trivalent graph with v vertices, then it has $q = 3v/2$ edges. Note that the restriction $\rho|_{\pi_1(\partial\mathcal{N}(\Sigma))}$ is irreducible. (Namely, the holonomy lift of the neighborhood of a vertex takes values in the stabilizer of a point $\cong \mathrm{SU}(2)$ and has irreducible image. The stabilizer is a diagonal subgroup in $\mathrm{Spin}(4) \cong \mathrm{SU}(2)\times\mathrm{SU}(2)$ and projects isomorphically onto both factors.) Thus $H^0(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) \cong \mathrm{su}(2)^{\mathrm{Ad}\circ\rho} \cong 0$. By duality, $H^2(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) \cong 0$. Since $\chi(\partial\mathcal{N}(\Sigma)) = 2\chi(\Sigma) = 2(v-q) = -v = -\frac{2q}{3}$, we get $\dim H^1(\partial\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) = -\chi(\partial\mathcal{N}(\Sigma))\cdot\dim(\mathrm{su}(2)) = 2q$. $\square$

Step 2. *Smoothness of $\mathcal{X}(M, \mathrm{SU}(2))$ at $[\rho]$.* The cohomology group $H^1(M; \mathrm{Ad}\circ\rho)$ is naturally identified with the Zariski tangent space of $\mathcal{X}(M, \mathrm{SU}(2))$ at $[\rho]$. For an element in $H^1(M; \mathrm{Ad}\circ\rho)$, there is an infinite sequence of obstructions to be integrable (cf. [Dou], [Ab], [HPS], [Gol]). These obstructions are natural and live in the second cohomology group.

Note that $H^2(\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho) = 0$ because the orbifold $\mathcal{N}(\Sigma)$ is finitely covered by a manifold homotopy equivalent to a 1-dimensional complex. Hence (16) provides an isomorphism

$$H^2(M; \mathrm{Ad}\circ\rho) \cong H^2(\partial N; \mathrm{Ad}\circ\rho).$$

The obstructions to integrability vanish for ∂N because ∂N is a surface ([Gol] and [HPS]). Naturalness of the obstructions then implies that they also vanish for M. Thus every element in $H^1(M; \mathrm{Ad}\circ\rho)$ is formally integrable as a power series, and by a theorem of Artin [Ar] it is actually integrable. Since $H^1(M; \mathrm{Ad}\circ\rho)$ has dimension q, it follows that $[\rho]$ is a smooth point of $\mathcal{X}(M, \mathrm{SU}(2))$ with local dimension q.

Step 3. *A basis for $T^*_{[\rho]}\mathcal{X}(M, \mathrm{SU}(2))$.* By Step 2, there is a natural isomorphism

$$T^*_{[\rho]}\mathcal{X}(M, \mathrm{SU}(2)) \cong H_1(M; \mathrm{Ad}\circ\rho).$$

LEMMA 9.14. *The set $\{dI_{\mu_1}, \dots, dI_{\mu_q}\}$ is a basis for $H_1(M; \mathrm{Ad}\circ\rho)$.*

Proof. As before we identify $H_*(M; \mathrm{Ad}\circ\rho) \cong H_*(N; \mathrm{Ad}\circ\rho)$. Dual to (16), we have the isomorphism

$$H_1(\partial N; \mathrm{Ad}\circ\rho) \cong H_1(N; \mathrm{Ad}\circ\rho) \oplus H_1(\mathcal{N}(\Sigma); \mathrm{Ad}\circ\rho).$$

The map $\iota_* : H_1(\partial N; \mathrm{Ad} \circ \rho) \to H_1(\mathcal{N}(\Sigma); \mathrm{Ad} \circ \rho)$ of cotangent spaces is induced by the restriction map $\mathcal{X}(\mathcal{N}(\Sigma), \mathrm{SU}(2)) \to \mathcal{X}(\partial N, \mathrm{SU}(2))$. Since the meridians have finite order in $\pi_1(\mathcal{N}(\Sigma))$, the trace functions I_{μ_j} are constant on $\mathcal{X}(\mathcal{N}(\Sigma), \mathrm{SU}(2))$. Hence $\iota_*(dI_{\mu_j}) = 0$. Since $\dim H_1(N; \mathrm{Ad} \circ \rho) = q$ by Proposition 9.13, the proof reduces to the following lemma:

LEMMA 9.15. *The differential forms* $\{dI_{\mu_1}, \dots, dI_{\mu_q}\}$ *are linearly independent in* $H_1(\partial N; \mathrm{Ad} \circ \rho)$.

Proof. For each meridian μ_j, we construct a deformation of the restriction $\rho|_{\partial N}$ that is parametrized by the trace functions I_{μ_j} and leaves all other meridians constant. We proceed on each component, assuming that the deformation on the other components is trivial.

First we consider the case that μ is a meridian around a singular circle, so that the corresponding boundary component $T^2 \subset \partial N$ is a torus. We wish to deform $\rho|_{T^2}$. We choose λ so that λ and μ generate $\pi_1(T^2)$. The elements $\rho(\mu)$ and $\rho(\lambda)$ commute and can therefore be simultaneously diagonalized. Thus $\rho(\mu)$ and $\rho(\lambda)$ can be varied independently inside a circle subgroup of SU(2). The matrix $\rho(\mu)$ has eigenvalues $\pm e^{\pm i \frac{\alpha}{2}}$ and trace $\pm 2 \cos \frac{\alpha}{2}$ where α is the cone angle; cf. (15). Since $0 < \alpha \leq \pi$ there are variations of $\rho(\mu)$ with nonzero derivative of the trace function I_μ.

Now we deal with a surface of genus $g \geq 2$ in ∂N. We take a decomposition of this surface in pairs of pants, so that the curves of the decomposition are the meridians (each pair of pants P corresponds to a singular vertex and ∂P consists of meridians). The deformations we require are easily constructed if we prove that $\mathcal{X}(P, \mathrm{SU}(2))$ is locally parametrized by the trace functions of the components of ∂P. Choose generators a and b of $\pi_1(P)$, so that a, b and ab represent the three components of ∂P. Since a and b generate a free group, it is well known that

$$(I_a, I_b, I_{ab}) : \mathcal{X}(P, \mathrm{SL}_2(\mathbb{C})) \longrightarrow \mathbb{C}^3$$

is an isomorphism of algebraic varieties defined over $\mathbb{Q}$; see for instance [GM]. This isomorphism implies that, for each $\gamma \in \pi_1(P)$, I_γ is a polynomial on I_a, I_b, I_{ab} with coefficients in $\mathbb{Q}$. Thus

$$(I_a, I_b, I_{ab}) : \mathcal{X}(P, \mathrm{SL}_2(\mathbb{C})) \cap \mathbb{R}^N \longrightarrow \mathbb{R}^3$$

is also an isomorphism. The irreducibility of $\rho|_{\pi_1(P)}$ and Lemma 9.1 imply that $(I_a, I_b, I_{ab}) : \mathcal{X}(P, \mathrm{SU}(2)) \to \mathbb{R}^3$ is a local diffeomorphism at the conjugacy class of the restriction $\rho|_{\pi_1(P)}$. $\square$

This finishes the proof of Theorem 9.4. $\square$

10. The fibration theorem

Throughout this section, $\mathcal{O}$ denotes a closed orientable small 3-orbifold $\mathcal{O}$ and X a cone *metric* on $\mathcal{O}$ of constant curvature in $[-1, 0)$ and with cone angles less than or equal to the orbifold angles of $\mathcal{O}$. Since we assume that X is a cone metric on $\mathcal{O}$ and not only a cone structure, cf. Definition 6.3, the pairs $(|\mathcal{O}|, \Sigma)$ and (X, Σ_X) are homeomorphic, and in particular X is closed. The main result of this section is:

FIBRATION THEOREM (Theorem 6.14). *For $\omega, D_0 > 0$ there exists $\delta = \delta(\omega, D_0) > 0$ such that: If X has ω-thick links (cf. Definition 6.13), $\mathrm{diam}(X) \geq D_0$ and if X is δ-thin, then $\mathcal{O}$ is Seifert fibred.*

10.1. *Local Euclidean structures.* The local geometry of thin cone manifolds is modelled on noncompact Euclidean cone manifolds; cf. [ChG, part 2, Prop. 3.4] in the case of manifolds.

Recall that by Corollary 4.2 every noncompact Euclidean cone 3-manifold E^3 with cone angles $\leq \pi$ has a *soul* S.

LEMMA 10.1. *For every $\varepsilon > 0$ and $R > 1$, there exists $\delta_0 = \delta_0(\varepsilon, R, \omega) > 0$ such that: If $\delta < \delta_0$, X is a cone manifold of curvature in $[-1, 0)$ with ω-thick links, $\mathrm{diam}(X) \geq D_0$ and X is δ-thin, then each $x \in X$ has a neighborhood $U_x \subset X$, and a $(1+\varepsilon)$-bilipschitz homeomorphism*

$$f : U_x \to \mathcal{N}_{\nu_x}(S)$$

where $\mathcal{N}_{\nu_x}(S)$ is the normal cone fiber bundle, of radius $\nu_x \in (0, 1)$ depending on x, of the soul S of a noncompact Euclidean cone 3-manifold. In addition $\dim S = 1$ or 2, and

$$\max\big(d(f(x), S), \mathrm{diam}(S)\big) \leq \nu_x / R.$$

Proof. Assume that the assertion is false. Then there exist $\varepsilon > 0$, $R > 1$ and a sequence of cone manifolds X_n with diameter $\geq D_0$, curvature in $[-1, 0)$ and ω-thick links such that X_n is $\frac{1}{n}$-thin, and there exist points $x_n \in X_n$ for which the conclusion of the lemma does not hold.

The fact that X_n is $\frac{1}{n}$-thin and has ω-thick links implies that also the radii of embedded *singular* standard balls in X_n are $\leq r_n \to 0$. Let $\lambda_n > 0$ be the supremum of all radii r such that $B_r(x_n)$ is contained in a (smooth or singular) standard ball. We have that $\lambda_n \to 0$. The sequence of rescaled cone manifolds $(\frac{1}{\lambda_n} X_n, x_n)$ with base points subconverges to a limit space (E, x_∞). Observe that the balls $B_1(x_n) \subset \frac{1}{\lambda_n} X_n$ are uniformly thick, as shown by the following sublemma.

SUBLEMMA 10.2. *Assume that the cone 3-manifold X has curvature $k \in [-\kappa, \kappa]$, cone angles $\leq \pi$ and ω-thick links. Suppose that the distance ball*

$B_r(p) \subset X$ is contained in a standard ball. Then $B_r(p)$ is δ-thick with $\delta = \delta(\kappa, \omega, r) > 0$.

Proof. We may assume that X is a complete singular cone. The smooth points in X with injectivity radius $< \delta$ are contained in a tubular neighborhood of radius $\rho(\omega, \delta)$ around the singularity where $\lim_{\delta \to 0} \rho(\omega, \delta) = 0$, because X has ω-thick links. We choose $\delta < \frac{r}{2}$ sufficiently small such that $\rho(\omega, \delta) < \frac{r}{2}$. □

It follows with the compactness theorem (Corollary 3.22) that E is a 3-dimensional Euclidean cone manifold.

Since E is not compact, we can apply Theorem 4.1 which gives the classification of noncompact Euclidean cone 3-manifolds with cone angles $\leq \pi$. The space E is not a complete cone because balls around x_∞ with radii > 1 are not contained in a standard ball. Hence the soul S of E has dimension 1 or 2. Let $N \subset E$ be a tubular neighborhood around S with radius $\rho > R \cdot \mathrm{diam}(S \cup \{x_\infty\})$. We use now that the convergence is bilipschitz. For sufficiently large n, there exists a $(1 + \varepsilon)$-bilipschitz embedding $(N, x_\infty) \hookrightarrow (\frac{1}{\lambda_n} X_n, x_n)$, and hence $(\lambda_n N, x_\infty) \hookrightarrow (X_n, x_n)$. Hence x_n satisfies the conclusion of the lemma with $\nu_{x_n} = \lambda_n \rho \to 0$, a contradiction. □

We apply this lemma to each point of X with some constants $R > 1$, $\varepsilon > 0$ to be specified later. Consider the thickening

$$W_x := f^{-1}(\overline{\mathcal{N}_{\lambda \nu_x}(S)})$$

of the soul of U_x where $0 < \lambda < \frac{1}{R}$. We will also view W_x as a suborbifold of $\mathcal{O}$.

The local models E have 1- or 2-dimensional soul and therefore belong to the following list by Theorem 4.1:

- When S is 2-dimensional and orientable, then E is isometric to the product $S \times \mathbb{R}$. The possible Euclidean cone surfaces S are a torus T^2, a pillow $S^2(\pi, \pi, \pi, \pi)$, i.e. a Euclidean surface homeomorphic to S^2 with four cone points of angle π, and a Euclidean turnover $S^2(\alpha, \beta, \gamma)$ with cone angles α, β and γ satisfying $\alpha + \beta + \gamma = 2\pi$.

- When S is 2-dimensional but nonorientable (possibly with mirror boundary), then $E = \tilde{S} \times \mathbb{R}/\iota$, where $\tilde{S}$ is the orientable double covering of S and ι is an involution that preserves the product structure and reverses the orientation of each factor. Hence E is a twisted line bundle over S.

- When $\dim(S) = 1$, then either $S = S^1$ or S is an interval with mirror boundary (a quotient of S^1). In the former case, E is either a solid torus or a singular solid torus. In the latter, E is either a solid pillow or a singular solid pillow.

Not all of these possibilities can occur: W_x contains no turnover, because $\mathcal{O}$ is small and singular vertices of X have thick links (i.e. the addition of cone angles of edges adjacent to a singular vertex is $> 2\pi + \omega'$).

LEMMA 10.3. *If $\varepsilon = \varepsilon(\omega) > 0$ is small enough, then the soul S of the local model for W_x is neither a turnover nor the quotient of a turnover.*

Proof. Assume the contrary. By hypothesis, $\mathcal{O}$ is closed and small. Hence every turnover in $\mathcal{O}$ bounds a discal suborbifold. Using the bilipschitz homeomorphism from W_x to the local model, we see that the sum of cone angles of the turnover is close to 2π. This contradicts the ω-thickness of links. □

By looking at the remaining possibilities, we deduce:

COROLLARY 10.4. *Each W_x admits a Seifert fibration* (*in the orbifold sense*). *In particular, ∂W_x is a union of smooth tori and pillows.*

LEMMA 10.5. (i) *If W_x contains a singular vertex, then $\mathcal{O} - \operatorname{int}(W_x)$ is Haken.*

(ii) *Assume that $\mathcal{O}$ is cyclic. If $\varepsilon = \varepsilon(\omega) > 0$ is small enough, $R = R(\omega)$ is large enough and if $W_x \cap \Sigma \neq \emptyset$, then $\mathcal{O} - \operatorname{int}(W_x)$ is Haken.*

Proof. The main point is to prove that $\mathcal{O} - \operatorname{int}(W_x)$ is irreducible. For this we have to show that W_x is not contained in a discal suborbifold. Since the pairs $(\mathcal{O}, W_x)$ and $(\mathcal{O}, \bar{U}_x)$ are homeomorphic, this amounts to showing that U_x is not contained in a discal suborbifold, and we can use the metric properties of U_x.

If U_x contains a singular vertex, this vertex and at least one singular edge lie in the soul. Hence U_x contains an entire singular edge or loop and therefore cannot be included in a discal suborbifold. This proves irreducibility in case (i).

Now we proceed with case (ii). Suppose that $\mathcal{O}$ is of cyclic type, U_x meets the singular locus and is contained in the discal suborbifold Δ. Topologically, Δ is a singular ball with one axis a. Hence U_x cannot contain an entire singular edge. By looking at the possible local models we see that U_x contains at least two singular segments of length $> \nu_x$ whose midpoints m_1 and m_2 have distance between them $< \frac{\nu_x}{R}(1 + \varepsilon)$. By developing the smooth part of X into model space, and composing the developing map with the projection onto the axis fixed by the holonomy representation, we find a 1-Lipschitz function on X whose restriction to a is linear with slope 1. It follows that a is distance minimizing inside Δ and hence $d_\Delta(m_1, m_2) > \nu_x$, a contradiction. This finishes the proof of irreducibility in case (ii).

Since $\mathcal{O} - \operatorname{int}(W_x)$ has a boundary, all that remains to check for the Haken property is that there are no Euclidean or hyperbolic turnovers. This follows from the smallness of $\mathcal{O}$ in both cases (i) and (ii). □

10.2. *Covering by virtually abelian subsets.* In this section, we study general properties of coverings by virtually abelian subsets. The arguments in this section closely follow [Gr1] and [BoP].

We assign a special role to one of the subsets W_x along which we will cut $\mathcal{O}$ later on. Namely, we choose $x_0 \in X$ as follows: If $\mathcal{O}$ has singular vertices, we require that W_{x_0} contain a singular vertex and that its radius ν_{x_0} be almost maximal:

$$\nu_{x_0} \geq \frac{1}{1+\varepsilon} \sup\{\nu_x | W_x \cap \Sigma^{(0)} \neq \emptyset\}.$$

If $\mathcal{O}$ is cyclic, we make an analogous choice for W_{x_0} among all W_x that intersect the singular locus. We denote $W_0 = W_{x_0}$, $\mathcal{O}_0 = \mathcal{O} - \mathrm{int}(W_0)$, $\nu_0 = \nu_{x_0}$. In view of Lemma 10.5, $\mathcal{O}_0$ is Haken.

Definition 10.6. We say that a subset $S \subset \mathcal{O}$ is *virtually abelian in* $\mathcal{O}_0$ if, for each connected component Z of $S \cap \mathcal{O}_0$, the image of $\pi_1(Z) \to \pi_1(\mathcal{O}_0)$ in the fundamental group of the corresponding component of $\mathcal{O}_0$ is virtually abelian. Moreover, for $x \in X$ we define:

$$\mathrm{va}(x) = \sup\{r > 0 \mid B_r(x) \text{ is virtually abelian in } \mathcal{O}_0\}$$

and

$$r(x) = \inf\left(\frac{\mathrm{va}(x)}{8}, 1\right).$$

LEMMA 10.7. *Let $x, y \in X$. If $B_{r(x)}(x) \cap B_{r(y)}(y) \neq \emptyset$, then*

(a) $3/4 \leq r(x)/r(y) \leq 4/3$;

(b) $B_{r(x)}(x) \subset B_{4r(y)}(y)$.

Proof. We may assume that $r(x) \leq r(y)$, and moreover $r(x) = \frac{1}{8}\mathrm{va}(x) < 1$. By the triangle inequality,

$$\mathrm{va}(x) \geq \mathrm{va}(y) - r(x) - r(y),$$

so that

$$8r(x) = \mathrm{va}(x) \geq 8r(y) - r(x) - r(y) \geq 6r(y).$$

This shows part (a). Part (b) follows because $2r(x) + r(y) < 4r(y)$. □

LEMMA 10.8. *For R sufficiently large, $W_0 \subset B_{\frac{r(x_0)}{9}}(x_0)$.*

Proof. This follows because $\mathrm{va}(x_0) \geq \frac{1}{1+\varepsilon}\nu_{x_0}(1 - \frac{1}{R})$ and W_0 is contained in the ball of radius $3(1+\varepsilon)\frac{\nu_x}{R}$ around x_0. □

We proceed to construct coverings of X. We have already distinguished a point $x_0 \in W_0$. Consider sequences $\{x_0, x_1, \dots\}$ starting with x_0, such that:

(19) the balls $B_{\frac{1}{4}r(x_0)}(x_0), B_{\frac{1}{4}r(x_1)}(x_1), \dots$ are pairwise disjoint.

A sequence satisfying (19) is finite, by Lemma 10.7 and compactness of X.

Lemma 10.9. *If the sequence $\{x_0, x_1, \dots, x_p\}$ is maximal for property* (19), *then the balls $B_{\frac{2}{3}r(x_0)}(x_0), \dots, B_{\frac{2}{3}r(x_p)}(x_p)$ cover X.*

Proof. Let $x \in X$ be an arbitrary point. By maximality, there exists a point x_j such that $B_{\frac{1}{4}r(x)}(x) \cap B_{\frac{1}{4}r(x_j)}(x_j) \neq \emptyset$. By Lemma 10.7 we have $r(x) \leq \frac{4}{3}r(x_j)$ and $d(x, x_j) \leq \frac{1}{4}(r(x) + r(x_j)) \leq \frac{7}{12}r(x_j)$. Thus $x \in B_{\frac{2}{3}r(x_j)}(x_j)$. □

We fix now a sequence $x_0, x_1, \dots, x_p$ maximal for property (19) and consider the covering of X by the open sets

- $V_0 = B_{r(x_0)}(x_0)$ and
- $V_i = B_{r(x_i)}(x_i) - W_0$ for $i = 1, \dots, p$.

Lemmas 10.8 and 10.9 imply that the open sets $V_0, \dots, V_p$ cover X. We denote $r_i := r(x_i)$ and $B_i := B_{r(x_i)}(x_i)$.

Lemma 10.10. *There is a universal bound N on the number of balls B_i that can intersect a fixed ball B_k.*

Proof. For every ball B_i intersecting B_k, $B_i \subset B_{2r_i+r_k}(x_k) \subseteq B_{4r_k}(x_k)$. On the other hand, the points x_i are separated from each other, since $d(x_{i_1}, x_{i_2}) \geq \frac{1}{4}(r_{i_1} + r_{i_2}) \geq \frac{3}{8}r_k$. Thus the number of such x_i is bounded above by:

$$\frac{\operatorname{vol}(B_{4r_k}(x_k))}{\operatorname{vol}(B_{\frac{3}{16}r_k}(x_i))} \leq \frac{\operatorname{vol}(B_{8r_k}(x_i))}{\operatorname{vol}(B_{\frac{3}{16}r_k}(x_i))} \leq \frac{\mathrm{v}_\kappa(8r_k)}{\mathrm{v}_\kappa\left(\frac{3}{16}r_k\right)}.$$

Here $\mathrm{v}_\kappa(r)$ denotes the volume of the ball of radius r in the 3-space of constant curvature $\kappa \in [-1, 0)$, and the last inequality follows from Bishop-Gromov. Since $r_k \leq 1$, the ratio $\mathrm{v}_\kappa(8r_k)/\,\mathrm{v}_\kappa(\frac{3}{4}r_k)$ is bounded. □

In particular, the dimension of our covering is universally bounded by N. We want to decrease the dimension of the covering $\{V_0, \dots, V_p\}$ while keeping the properties that the covering sets are virtually abelian and only one of them meets W_0.

Using a partition of unity (ϕ_i) subordinate to (V_i) one can construct a map

$$f = \frac{1}{\sum_i \phi_i}(\phi_0, \dots, \phi_p) \colon X \to \Delta \subset \mathbb{R}^{p+1}, \tag{20}$$

where Δ is the unit simplex. The image of f is contained in the nerve of the covering $K \subset \Delta$, which has dimension $\leq N$. We start by controlling the Lipschitz constant of the map $f\colon X \to K$.

LEMMA 10.11. *There exists a constant $L_N > 0$ such that the partition of unity can be chosen so that the restriction $f|_{V_k}$ is $\frac{L_N}{r_k}$-Lipschitz.*

We first need the following geometric property of our covering:

SUBLEMMA 10.12. *Every $x \in X$ belongs to an open set V_k such that $d(x, \partial V_k) \geq r_k/3$.*

Proof. If $x \in B_{\frac{2}{3}r_0}(x_0)$ we choose $k = 0$. Suppose that $x \notin B_{\frac{2}{3}r_0}(x_0)$. Then there exists k with $x \in B_{\frac{2}{3}r_k}(x_k)$. The assertion is trivial if B_k and B_0 are disjoint. We assume therefore also that $B_k \cap B_0 \neq \emptyset$. Then by Lemma 10.8:

$$d(x, W_0) \geq d(x, x_0) - \frac{1}{9}r_0 \geq \frac{2}{3}r_0 - \frac{1}{9}r_0 \geq \frac{3}{4} \cdot \frac{5}{9}r_k > \frac{1}{3}r_k.$$

Hence $d(x, \partial V_k) \geq \frac{1}{3}r_k$. $\square$

Proof of Lemma 10.11. Let $\tau : [0,1] \to [0,1]$ be an auxiliary 4-Lipschitz function which vanishes in a neighborhood of 0 and satisfies $\tau|_{[\frac{1}{3},1]} \equiv 1$. We put $\phi_k := \tau(\frac{1}{r_k}d(\partial V_k, \cdot))$ on V_k and extend it trivially to X. Then ϕ_k is $\frac{4}{r_k}$-Lipschitz.

Let $x \in V_k$. Then at most $N+1$ functions ϕ_i are nonzero in x, and all of them have Lipschitz constant $\leq \frac{4}{3} \cdot \frac{4}{r_k}$. The claim follows since the functions

$$(x_0, \ldots, x_N) \mapsto \frac{x_k}{\sum_{i=0}^N x_i}$$

are Lipschitz on $\{x \in \mathbb{R}^{N+1} \mid x_i \geq 0 \; \forall i \wedge \sum_{i=0}^N x_i \geq 1\}$. $\square$

We now homotope f into the 3-skeleton $K^{(3)}$ by an inductive procedure while controlling the local Lipschitz constant.

LEMMA 10.13. *For $d \geq 4$ and $L > 0$, there exists $L' = L'(d, L) > 0$ such that the following is true*:

Suppose that $g : X \to K^{(d)}$ is a continuous map which is $\frac{L}{r_k}$-Lipschitz on V_k and has the property that the inverse image of the open star of the vertex $v_{V_k} \in K^{(0)}$ is contained in V_k. Then g can be homotoped to a map $\tilde{g} : X \to K^{(d-1)}$ with the same properties, L being replaced by L'.

Proof. It suffices to find a constant $\theta > 0$ such that every d-dimensional simplex $\sigma \subset K$ contains a point z at distance $\geq \theta$ from both $\partial\sigma$ and the image of g. To push g into the $(d-1)$-skeleton we compose it on σ with the central projection from z. This will increase the Lipschitz constant by a factor bounded in terms of d, and it reduces the inverse images of open stars of vertices.

If θ does not satisfy the desired property for some d-simplex σ, then $\mathrm{image}(g)\cap\mathrm{int}(\sigma)$ must contain a subset of at least $C(d)\cdot\frac{1}{\theta^d}$ points with pairwise distances $\geq \theta$. Let $A \subset X$ be a set of inverse images, one for each point. Let v_{V_k} be a vertex of σ. Then $A \subset V_k \subseteq B_k$. Since f is $\frac{L}{r_k}$-Lipschitz continuous on V_k, the points of A are separated by distance $\frac{1}{L}r_k\theta$. Since $r_k \leq 1$, volume comparison implies that A contains at most $C\cdot(\frac{L}{\theta})^3$ points. The inequality $C(d)\cdot\frac{1}{\theta^d} \leq C\cdot(\frac{L}{\theta})^3$ yields a positive lower bound $\theta_0(d, L)$ for θ. Hence any constant $\theta < \theta_0$ has the desired property. □

LEMMA 10.14. *For sufficiently small $\varepsilon > 0$ there exists a constant $C = C(\varepsilon) > 0$ such that*

$$\mathrm{vol}(V_i) \leq C\frac{1}{R}r_i^3$$

for all i.

Proof. We first show that W_0 does not enter too far into the other sets U_{x_i} given in Lemma 10.1.

SUBLEMMA 10.15. *There exists a constant $c = c(\varepsilon) > 0$ such that, if $R > 1$ is sufficiently large, then $d(x_i, W_0) \geq c\,\nu_{x_i}$ for all $i \neq 0$.*

Proof. By Lemma 10.8, $W_0 \subset B_{\frac{r_0}{9}}(x_0)$ and we obtain:

$$d(x_i, W_0) \geq d(x_i, x_0) - \frac{1}{9}r_0 \geq \frac{1}{4}r_0 - \frac{1}{9}r_0 > \frac{1}{8}r_0 \geq \frac{1}{64(1+\varepsilon)}\nu_{x_0}.$$

For the last estimate, we use the fact that $\mathrm{va}(x_0) \geq \frac{1}{1+\varepsilon}\nu_{x_0}$ and, since $\nu_{x_0} \leq 1$, $r_0 = \inf(\frac{\mathrm{va}(x_0)}{8}, 1) \geq \frac{1}{8(1+\varepsilon)}\nu_{x_0}$.

We now assume that W_0 intersects U_{x_i} because otherwise there is nothing to show. If $W_0 \subset U_{x_i}$ then, according to our choice of W_0, we can compare the radii ν_{x_0} and ν_{x_i} by $\nu_{x_0} \geq \frac{1}{1+\varepsilon}\nu_{x_i}$, and the assertion holds with $c < (8(1+\varepsilon))^{-2}$.

We are left with the case that $W_0 \not\subset U_{x_i}$ but intersects the ball of radius, say, $\frac{\nu_{x_i}}{4}$ around x_i. Then we can bound the ratio $\frac{\mathrm{diam}(W_0)}{\nu_{x_i}}$ from below by:

$$(1+\varepsilon)\frac{\nu_{x_i}}{R} + \frac{\nu_{x_i}}{4} + \mathrm{diam}(W_0) \geq \frac{\nu_{x_i}}{1+\varepsilon}.$$

By definition of W_0 we have $\mathrm{diam}(W_0) \leq (1+\varepsilon)\frac{2}{R}\nu_{x_0}$. Combining these estimates, we obtain a lower bound for $\frac{d(x_i,W_0)}{\nu_{x_i}}$, as claimed. □

The sublemma implies that $r_i \geq c\,\nu_{x_i}$. By the Bishop-Gromov inequality,

$$\mathrm{vol}(V_i) \leq \mathrm{vol}\left(B_{r_i}(x_i)\right) \leq \mathrm{vol}\left(B_{c\,\nu_{x_i}}(x_i)\right)\frac{\mathrm{v}_\kappa(r_i)}{\mathrm{v}_\kappa(c\,\nu_{x_i})} \leq \mathrm{vol}\left(B_{c\,\nu_{x_i}}(x_i)\right)c_1\,\frac{r_i^3}{\nu_{x_i}^3}$$

for some uniform $c_1 > 0$, where $\mathrm{v}_\kappa(r)$ denotes the volume of the ball of radius r in the space of constant curvature $\kappa \in [-1, 0)$. By the geometry of the local models, it follows that

$$\mathrm{vol}\big(B_{c\,\nu_{x_i}}(x_i)\big) \le \mathrm{vol}(U_{x_i}) \le c_2 \frac{\nu_{x_i}^3}{R}$$

for some $c_2 > 0$. Thus $\mathrm{vol}(V_i) \le c_1\, c_2\, \frac{r_i^3}{R}$. □

Now we can further homotope f into the 2-skeleton.

Proposition 10.16. *For suitable constants $\varepsilon > 0$ and $R > 1$, the map f in* (20) *is homotopic to a map*

$$\tilde{f} : X \to K^{(2)}$$

with the property that the inverse image of the open star of the vertex $v_{V_k} \in K^{(0)}$ is contained in V_k.

Proof. The inverse image under f of the open star of v_{V_k} is contained in V_k. Using Lemma 10.13 repeatedly, we can homotope f to a map $\hat{f} : X \to K^{(3)}$ which is locally Lipschitz and satisfies $\hat{f}^{-1}(\mathrm{star}(v_{V_k})) \subset V_k$. More precisely, there is a universal constant $\hat{L}$ such that $\hat{f}$ is $\frac{\hat{L}}{r_k}$-Lipschitz on V_k.

It suffices to show that no 3-simplex $\sigma \subset K$ is contained in the image of $\hat{f}$. The inverse image $\hat{f}^{-1}(\mathrm{int}(\sigma))$ lies in the intersection of sets V_j where v_{V_j} runs through the vertices of σ. Let V_k be one of these. With Lemma 10.14 it follows that

$$\mathrm{vol}(\mathrm{image}(\hat{f}) \cap \sigma) \le \mathrm{vol}(\hat{f}(V_k)) \le (\frac{\hat{L}}{r_k})^3\,\mathrm{vol}(V_k) \le C\hat{L}^3 \frac{1}{R}$$

with uniform constants C and $\hat{L}$. So, if R is large enough, $\mathrm{vol}(\mathrm{image}(\hat{f}) \cap \sigma) < \mathrm{vol}(\sigma)$. □

Note that f maps W_0 to the vertex v_{V_0} because W_0 intersects none of the sets V_j with $j \neq 0$. The proposition therefore implies the following properties which will be crucial below.

(i) $\tilde{f}(W_0) = \{v_{V_0}\}$.

(ii) For every vertex v of K, $\tilde{f}^{-1}(\mathrm{star}(v))$ is virtually abelian in $\mathcal{O}_0$.

10.3. *Vanishing of simplicial volume.* The orbifold $\mathcal{O}_0$ is Haken and therefore has a JSJ-splitting into Seifert and hyperbolic suborbifolds.

Proposition 10.17. *All components in the* JSJ *splitting of $\mathcal{O}_0$ are Seifert.*

Proof. Since the orbifold $\mathcal{O}_0$ is Haken, it is very good and there is a finite covering

$$p\colon N \to \mathcal{O}_0$$

by a manifold [McCMi]. The boundary ∂N is a union of tori. The JSJ splitting of $\mathcal{O}_0$ pulls back to the JSJ splitting of N. We have to show that no hyperbolic components occur in the JSJ splitting of N.

We may assume that the boundary of N is incompressible because otherwise N is a solid torus and the assertion holds. We construct a closed manifold $\bar{N}$ by Dehn filling on N as follows. Let $Y \subset N$ be a component of the JSJ splitting which meets the boundary, $Y \cap \partial N \neq \emptyset$. If Y is hyperbolic we choose, using the hyperbolic Dehn filling theorem, the Dehn fillings at the tori of $Y \cap \partial N$ such that the resulting manifold $\bar{Y}$ remains hyperbolic. If Y is Seifert, we fill in such a way that $\bar{Y}$ is Seifert and the components of $\partial Y - \partial N$ remain incompressible. This can be done because the base of the Seifert fibration of Y is neither an annulus nor a disc with zero or one cone point. The manifold $\bar{N}$ has a JSJ splitting along the same tori as N and with the same number of hyperbolic (and also Seifert) components.

It suffices to show that $\bar{N}$ has zero simplicial volume, because then [Gr1, §3.5] and [Kue] imply that $\bar{N}$ contains no hyperbolic component in its JSJ splitting. To this purpose we will apply Gromov's vanishing theorem; see [Gr1, §3.1], [Iva].

We compose $\tilde{f}$ with the projection p and extend the resulting map $N \to K^{(2)}$ to a map

$$h\colon \bar{N} \to K^{(2)}$$

by sending the filling solid tori to the vertex v_{V_0}. Note that h is continuous because $\tilde{f}(\partial\mathcal{O}_0) = \{v_{V_0}\}$. The inverse images under h of open stars of vertices are virtually abelian as subsets of $\bar{N}$, because they are already virtually abelian in N and the filling tori intersect only one of the subsets. These subsets yield an open covering of $\bar{N}$ with covering dimension ≤ 2. By Gromov's theorem, the simplicial volume of $\bar{N}$ vanishes. ◻

Conclusion of the proof of Theorem 6.14. Since $\mathcal{O}_0$ is graphed and $\mathcal{O}$ results from $\mathcal{O}_0$ by gluing in a Seifert orbifold it follows that $\mathcal{O}$ is graphed. Since $\mathcal{O}$ is moreover atoroidal, it must be Seifert. ◻

Laboratoire Émile Picard, CNRS UMR 5580, Université Paul Sabatier, Toulouse, France
E-mail address: boileau@picard.ups-tlse.fr

Mathematisches Institut, Universität München, München, Germany
E-mail address: leeb@mathematik.uni-muenchen.de

Departament de Matemàtiques, Universitat Autònoma de Barcelona, Spain
E-mail address: porti@mat.uab.es

References

[Ab] L. Ben Abdelghani, Espace des représentations du groupe d'un nœud dans un groupe de Lie, Thèse, Université de Bourgogne, 1998.

[Ar] M. Artin, On the solutions of analytic equations, *Invent. Math.* **5** (1968), 277–291.

[Bal] W. Ballmann, Singular spaces of nonpositive curvature, in *Sur les Groupes Hyperboliques d'après Mikhael Gromov* (E. Ghys, P. de la Harpe, ed.), Birkhäuser, Boston 1990, 189–201.

[BLP0] M. Boileau, B. Leeb, and J. Porti, Uniformization of small 3-orbifolds, *C.R. Acad. Sci. Paris Sér.* I *Math.* **332** (2001), 57–62.

[BMP] M. Boileau, S. Maillot, and J. Porti, Three-dimensional orbifolds and their geometric structures, *Panoramas et Synthèses* **15**, Soc. Math. de France, Paris, 2003.

[BoP] M. Boileau and J. Porti, *Geometrization of* 3*-Orbifolds of Cyclic Type*, *Astérisque* **272** (2001).

[BS1] F. Bonahon and L. C. Siebenmann, The characteristic tori splitting of irreducible compact three-orbifolds, *Math. Ann.* **278** (1987), 441–479.

[BS2] ———The classification of Seifert fibered three- manifolds, in *Low Dimensional Topology* (R. Fenn, ed.), *London Math. Soc. Lecture Notes Ser.* **95** (1985), Cambridge Univ. Press, Cambridge, 19–85.

[Bro] K. S. Brown, *Cohomology of Groups*, *GTM* **87**, Springer-Verlag, New York, 1982.

[BBI] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, *Grad Stud. in Math.* **33**, A.M.S., Providence, RI, 2001.

[BGP] Y. Burago, M. Gromov, and G. Perel'man, A. D. Alexandrov spaces with curvatures bounded below, *Russian Math. Surveys* **47** (1992), 1–58.

[CEG] R. D. Canary, D. B. A. Epstein, and P. Green, Notes on notes of Thurston, in *Analytical and Geometric Aspects of Hyperbolic Space* (D. B. A. Epstein, ed.), *London Math. Soc. Lecture Notes Ser*. **111** (1987), 3–92, Cambridge Univ. Press, Cambridge, 1987.

[ChD] R. Charney and M. Davis, The polar dual of a convex polyhedral set in hyperbolic space, *Michigan Math. J*. **42** (1995), 479–510.

[ChG] J. Cheeger and M. Gromov, Collapsing Riemannian manifolds while keeping their curvature bounded, part I: *J. Differential Geom*. **23** (1986), 309–346, part II: *J. Differential Geom*. **32** (1990), 269–298.

[CHK] D. Cooper, C. Hodgson, and S. Kerchkoff, *Three-dimensional Orbifolds and Cone-Manifolds*, *Math. Soc. of Japan Mem*. **5**, Math. Soc. of Japan, Tokyo, 2000.

[DaM] M. W. Davis and J. W. Morgan, Finite group actions on homotopy 3-spheres, in *The Smith Conjecture* (New York, 1979) (H. Bass, J. W. Morgan, eds.), *Pure Appl. Math*. **112**, Academic Press, Orlando, FL, 1984, 181–225.

[dRh] G. de Rham, Reidemeister's torsion invariant and rotations of S^n, in 1964 *Differential Analysis*, *Mumbay Colloq*., 27–36, Oxford Univ. Press, London, 1964.

[DD] W. Dicks and M. J. Dunwoody, *Groups Acting on Graphs*, Cambridge Univ. Press, Cambridge, 1989.

[Dou] A. Douady, Obstruction primaireè la déformation, in *Sém. Henri Cartan*, 1960/61, Exp. 4, Secrétariat mathématique, Paris, 1960/1961.

[Dun1] W. D. Dunbar, Hierarchies for 3-orbifolds, *Topology and its Appl*. **29** (1988), 267–283.

[Dun2] W. D. Dunbar, Classification of solvorbifolds in dimension three, in *Braids* (Santa Cruz, CA, 1986), *Contemp. Math.* **78**, 207–216, A.M.S., Providence, RI, 1988.

[DuM] W. D. Dunbar and R. G. Meyerhoff, Volumes of hyperbolic 3-orbifolds, *Indiana Univ. Math. J.* **43** (1994), 611–637.

[Fei] M. E. Feighn, Actions of finite groups on homotopy 3-spheres, *Trans. Amer. Math. Soc.* **284** (1984), 141–151.

[Gol] W. Goldman, The symplectic nature of fundamental groups of surfaces, *Adv. in Math.* **54** (1984), 200–225.

[GM] F. González-Acuña and J. M. Montesinos-Amibilia, On the character variety of group representations in $\mathrm{SL}(2, \mathbb{C})$ and $\mathrm{PSL}(2, \mathbb{C})$, *Math. Z.* **214** (1993), 627–652.

[Gr1] M. Gromov, Volume and bounded cohomology, *I.H.E.S. Publ. Math.* **56** (1983), 5–99.

[Gr2] ———, Hyperbolic groups, in *Essays in Group Theory* (S. Gersten, ed.), *MSRI Publ.* **8**, 75–263, Springer-Verlag, New York, 1987.

[GLP] M. Gromov, J. Lafontaine, and P. Pansu (eds.), *Structures métriques pour les variétés riemanniennes*, *Textes Math.* **1**, Cedic/Fernand Nathan, Paris 1981.

[Hae] A. Haefliger, Orbi-espaces, in *Sur les Groupes Hyperboliques d'après Mikhael Gromov* (E. Ghys, P. de la Harpe, eds.), 203–213, Birkhäuser, Boston, 1990.

[Ha1] R. S. Hamilton, Three-manifolds with positive Ricci curvature, *J. Differential Geom.* **17** (1982), 255–306.

[Ha2] ———, Four-manifolds with positive curvature operator, *J. Differential Geom.* **24** (1986), 153–179.

[HPS] M. Heusener, J. Porti, and E. Suárez, Deformations of reducible representations of 3-manifold groups into $\mathrm{SL}_2(\mathbb{C})$, *J. Reine Angew. Math.* **530** (2001), 191–227.

[Hod] C. Hodgson, Geometric structures on 3-dimensional orbifolds: Notes on Thurston's proof, preprint.

[Iva] N. V. Ivanov, Foundations of the theory of bounded cohomology, *J. Soviet Math.* **37** (1987), 1090–1115.

[JR] W. H. Jaco and J. H. Rubinstein, PL minimal surfaces in 3-manifolds, *J. Differential Geom.* **27** (1988), 493–524.

[Kap] M. Kapovich, *Hyperbolic Manifolds and Discrete Groups: Notes on Thurston's Hyperbolization*, *Progr. in Math.* **103**, Birkhäuser, Basel 2000.

[Koj] S. Kojima, Deformations of hyperbolic 3-cone-manifolds, *J. Differential Geom.* 49 (1998), 469–516.

[Kue] T. Kuessner, Relative simplicial volume, Ph.D. Thesis, Tübingen, 2001.

[KS] S. Kwasik and R. Schultz, Icosahedral group actions on $\mathbf{R}^3$, *Invent. Math.* **108** (1992), 385–402.

[LS] B. Leeb and P. Scott, A geometric characteristic splitting in all dimensions, *Comment. Math. Helv.* **75** (2000), 201–215.

[Mai] S. Maillot, Equivariant theorems for 3-orbifolds and geometrization of small 3-orbifolds with an infinite cyclic normal subgroup, *Trans. Amer. Math. Soc.* **355** (2003), 4595–4638.

[McCMi] D. McCullough and A. Miller, Manifold covers of 3-orbifolds with geometric pieces, *Topology and its Appl.* **31** (1989), 169–185.

[McM1] C. McMullen, Iteration on Teichmüller space, *Invent. Math.* **99** (1990), 425–454.

[McM2] C. McMullen, *Renormalization and 3-manifolds which Fiber over the Circle*, *Ann. of Math. Studies* **144**, Princeton Univ. Press, Princeton, NJ, 1996.

[MY1] W. H. Meeks and S.-T. Yau, Topology of three-dimensional manifolds and the embedding problems in minimal surface theory, *Ann. of Math.* **112** (1980), 441–484.

[MY2] ———, The equivariant Dehn's lemma and loop theorem, *Comment. Math. Helv.* **56** (1981), 225–239.

[MY3] ———, Group actions on $\mathbf{R}^3$, in *The Smith Conjecture* (H. Bass, J.W. Morgan, eds.), *Pure and Appl. Math.* **112**, 167–179, Academic Press, Orlando, FL, 1984.

[Mor] J. W. Morgan, *Actions de groupes finis sur* $\mathbf{S}^3$: *la conjecture de Smith*, *Sém. Bourbaki* **578**, *Lecture Notes in Math.* **901**, Springer-Verlag, New York, 1981.

[MB] J. W. Morgan and H. Bass, *The Smith Conjecture*, *Pure and Applied Math.* **112**, Academic Press, Orlando, FL, 1984.

[MoS] J. W. Morgan and P. B. Shalen, Valuations, trees, and degenerations of hyperbolic structures, I, *Ann. of Math.* **120** (1984), 401–476.

[Ot1] J.-P. Otal, Le théorème d'hyperbolisation pour les variétés fibrées de dimension trois, *Astérisque* **235** (1996).

[Ot2] ———, Thurston's hyperbolization of Haken manifolds, in *Surveys in Differential Geometry*, *Vol.* III (Cambridge, MA, 1996), 77–194, Internat. Press, Boston, MA, 1998.

[Po1] J. Porti, *Torsion de Reidemeister pour les Variétés Hyperboliques*, *Mem. Amer. Math. Soc.* **128**, Providence, RI, 1997.

[Po2] ———, Regenerating hyperbolic and spherical cone structures from Euclidean ones, *Topology* **37** (1998), 365–392.

[RH] I. Rivin and C. Hodgson, A characterization of compact convex polyhedra in hyperbolic 3-space, *Invent. Math.* **111** (1993), 77–111.

[Rot] M. Rothenberg, Torsion invariants and finite transformation groups, *Proc. Sympos. Pure Math.* **XXXII** (1978), 267–311.

[Sco] P. Scott, The geometries of 3-manifolds, *Bull. London Math. Soc.* **15** (1983), 401–487.

[SOK] T. Soma, K. Ohshika, and S. Kojima, Towards a proof of Thurston's geometrization theorem for orbifolds, *RIMS Kokyuroku* **568** (1985), 1–72.

[Sul] D. Sullivan, Combinatorial invariants of analytic spaces, in *Proc. of Liverpool Singularities—Symposium*, I (1969/70), 165–168, Springer-Verlag, New York, 1971.

[TY1] Y. Takeuchi and M. Yokoyama, The geometric realizations of the decompositions of 3-orbifold fundamental groups, *Topology Appl.* **95** (1999), 129–153.

[TY2] ———, PL-least area 2-orbifolds and its applications to 3-orbifolds, *Kyushu J. Math.* **55** (2001), 19–61.

[Thu1] W. P. Thurston, The geometry and topology of 3-manifolds, Princeton Math. Dept., 1979.

[Thu2] ———, Three dimensional manifolds, Kleinian groups and hyperbolic geometry, *Bull. Amer. Math. Soc.* **6** (1982), 357–381.

[Thu3] ———, Hyperbolic structures on 3-manifolds, I: Deformations of acylindrical manifolds, *Ann. of Math.* **124** (1986), 203–246.

[Thu4] ———, Hyperbolic structures on 3-manifolds, II: Surface groups and manifolds which fiber over S^1, preprint, 1986.

[Thu5] W. P. Thurston, Hyperbolic structures on 3-manifolds, III: Deformations of 3-manifolds with incompressible boundary, preprint, 1986.

[Thu6] ———, *Three-manifolds with symmetry*, preprint, 1982.

[Thu7] ———, *Three-dimensional Geometry and Topology* Vol. 1 (S. Levy, ed.), Princeton Univ. Press, Princeton, NJ, 1997.

[Zh1] Q. Zhou, 3-dimensional geometric cone structures, Ph.D. Thesis, University of California, L.A., 1989.

[Zh2] ———, The moduli space or hyperbolic cone structures, *J. Differential Geom.* **51** (1999), 517–550.

(Received March 4, 2002)