

ASTÉRIQUE ???

LOCAL COLLAPSING, ORBIFOLDS,
AND GEOMETRIZATION

Bruce Kleiner

John Lott

Société Mathématique de France 2014

Publié avec le concours du Centre National de la Recherche Scientifique

Bruce Kleiner

Courant Institute of Mathematical Sciences, 251 Mercer St., New York, NY 10012.

E-mail : `bkleiner@cims.nyu.edu`

John Lott

Department of Mathematics, University of California at Berkeley,
Berkeley, CA 94720.

E-mail : `lott@math.berkeley.edu`

2000 Mathematics Subject Classification. — ???

Key words and phrases. — Collapsing, Ricci flow, geometrization, orbifold.

LOCAL COLLAPSING, ORBIFOLDS, AND GEOMETRIZATION

Bruce Kleiner, John Lott

Abstract. — This volume has two papers, which can be read separately. The first paper concerns local collapsing in Riemannian geometry. We prove that a three-dimensional compact Riemannian manifold which is locally collapsed, with respect to a lower curvature bound, is a graph manifold. This theorem was stated by Perelman without proof and was used in his proof of the geometrization conjecture. The second paper is about the geometrization of orbifolds. A three-dimensional closed orientable orbifold, which has no bad suborbifolds, is known to have a geometric decomposition from work of Perelman in the manifold case, along with earlier work of Boileau-Leeb-Porti, Boileau-Maillot-Porti, Boileau-Porti, Cooper-Hodgson-Kerckhoff and Thurston. We give a new, logically independent, unified proof of the geometrization of orbifolds, using Ricci flow.

Résumé (???)— ???

CONTENTS

<i>Locally Collapsed 3-Manifolds</i>	7
1. Introduction	7
2. Notation and conventions	18
3. Preliminaries	21
4. Splittings, strainers, and adapted coordinates	27
5. Standing assumptions	39
6. The scale function τ	41
7. Stratification	45
8. The local geometry of the 2-stratum	46
9. Edge points and associated structure	48
10. The local geometry of the slim 1-stratum	57
11. The local geometry of the 0-stratum	59
12. Mapping into Euclidean space	62
13. Adjusting the map to Euclidean space	71
14. Extracting a good decomposition of M	80
15. Proof of Theorem ?? for closed manifolds	84
16. Manifolds with boundary	87
17. Application to the geometrization conjecture	90
18. Local collapsing without derivative bounds	91
19. Appendix A : Choosing ball covers	93
20. Appendix B : Cloudy submanifolds	94
21. Appendix C : An isotopy lemma	97
References	97
 <i>Geometrization of Three-Dimensional Orbifolds via Ricci Flow</i>	99
1. Introduction	99
2. Orbifold topology and geometry	102
3. Noncompact nonnegatively curved orbifolds	121
4. Riemannian compactness theorem for orbifolds	127

5. Ricci flow on orbifolds	128
6. κ -solutions	134
7. Ricci flow with surgery for orbifolds	140
8. Hyperbolic regions	147
9. Locally collapsed 3-orbifolds	150
10. Incompressibility of cuspidal cross-sections and proof of Theorem ??	157
11. Appendix A : Weak and strong graph orbifolds	163
References	172

LOCALLY COLLAPSED 3-MANIFOLDS

by

Bruce Kleiner & John Lott

Abstract. — We prove that a 3-dimensional compact Riemannian manifold which is locally collapsed, with respect to a lower curvature bound, is a graph manifold. This theorem was stated by Perelman and was used in his proof of the geometrization conjecture.

Résumé. —

1. Introduction

1.1. Overview. — In this paper we prove that a 3-dimensional Riemannian manifold which is locally collapsed, with respect to a lower curvature bound, is a graph manifold. This result was stated without proof by Perelman in [24, Theorem 7.4], where it was used to show that certain collapsed manifolds arising in his proof of the geometrization conjecture are graph manifolds. Our goal is to provide a proof of Perelman’s collapsing theorem which is streamlined, self-contained and accessible. Other proofs of Perelman’s theorem appear in [2, 5, 23, 30].

In the rest of this introduction we state the main result and describe some of the issues involved in proving it. We then give an outline of the proof. We finish by discussing the history of the problem.

2000 Mathematics Subject Classification. — ???

Key words and phrases. — ???

Research supported by NSF grants DMS-0903076 and DMS-1007508.

1.2. Statement of results. — We begin by defining an intrinsic local scale function for a Riemannian manifold.

Definition 1.1. — Let M be a complete Riemannian manifold. Given $p \in M$, the *curvature scale* R_p at p is defined as follows. If the connected component of M containing p has nonnegative sectional curvature then $R_p = \infty$. Otherwise, R_p is the (unique) number $r > 0$ such that the infimum of the sectional curvatures on $B(p, r)$ equals $-\frac{1}{r^2}$.

We need one more definition.

Definition 1.2. — Let M be a compact orientable 3-manifold (possibly with boundary). Give M an arbitrary Riemannian metric. We say that M is a *graph manifold* if there is a finite disjoint collection of embedded 2-tori $\{T_j\}$ in the interior of M such that each connected component of the metric closure of $M - \bigcup_j T_j$ is the total space of a circle bundle over a surface (generally with boundary).

For simplicity, in this introduction we state the main theorem in the case of closed manifolds. For the general case of manifolds with boundary, we refer the reader to Theorem 16.1.

Theorem 1.3 (cf. [24, Theorem 7.4]). — Let c_3 denote the volume of the unit ball in \mathbb{R}^3 and let $K \geq 10$ be a fixed integer. Fix a function $A : (0, \infty) \rightarrow (0, \infty)$. Then there is a $w_0 \in (0, c_3)$ such that the following holds.

Suppose that (M, g) is a closed orientable Riemannian 3-manifold. Assume in addition that for every $p \in M$,

- (1) $\text{vol}(B(p, R_p)) \leq w_0 R_p^3$ and
- (2) For every $w' \in [w_0, c_3]$, $k \in [0, K]$, and $r \leq R_p$ such that $\text{vol}(B(p, r)) \geq w' r^3$, the inequality

$$(1.4) \quad |\nabla^k \text{Rm}| \leq A(w') r^{-(k+2)}$$

holds in the ball $B(p, r)$.

Then M is a graph manifold.

1.3. Motivation. — Theorem 1.3, or more precisely the version for manifolds with boundary, is essentially the same as Perelman's [24, Theorem 7.4]. Either result can be used to complete the Ricci flow proof of Thurston's geometrization conjecture. We explain this in Section 17, following the presentation of Perelman's work in [21].

To give a brief explanation, let $(M, g(\cdot))$ be a Ricci flow with surgery whose initial manifold is compact, orientable and three-dimensional. Put $\widehat{g}(t) = \frac{g(t)}{t}$. Let M_t denote the time t manifold. (If t is a surgery time then we take M_t to be the postsurgery manifold.) For any $w > 0$, the Riemannian manifold $(M_t, \widehat{g}(t))$ has a decomposition into a w -thick part and a w -thin part. (Here the terms "thick" and

“thin” are suggested by the Margulis thick-thin decomposition but the definition is somewhat different. In the case of hyperbolic manifolds, the two notions are essentially equivalent.) As $t \rightarrow \infty$, the w -thick part of $(M_t, \widehat{g}(t))$ approaches the w -thick part of a complete finite-volume Riemannian manifold of constant curvature $-\frac{1}{4}$, whose cusps (if any) are incompressible in M_t . Theorem 1.3 implies that for large t , the w -thin part of M_t is a graph manifold. Since graph manifolds are known to have a geometric decomposition in the sense of Thurston, this proves the geometrization conjecture.

Independent of Ricci flow considerations, Theorem 1.3 fits into the program in Riemannian geometry of understanding which manifolds can collapse. The main geometric assumption in Theorem 1.3 is the first one, which is a local collapsing statement, as we discuss in the next subsection. The second assumption of Theorem 1.3 is more technical in nature. In the application to the geometrization conjecture, the validity of the second assumption essentially arises from the smoothing effect of the Ricci flow equation.

In fact, Theorem 1.3 holds without the second assumption. In order to prove this stronger result, one must use the highly nontrivial Stability Theorem of Perelman [19, 25]. As mentioned in [24], if one does make the second assumption then one can effectively replace the Stability Theorem by standard C^K -convergence of Riemannian manifolds. Our proof of Theorem 1.3 is set up so that it extends to a proof of the stronger theorem, without the second assumption, provided that one invokes the Stability Theorem in relevant places; see Sections 1.5.7 and 18.

1.4. Aspects of the proof. — The strategy in proving Theorem 1.3 is to first understand the local geometry and topology of the manifold M . One then glues these local descriptions together to give an explicit decomposition of M that shows it to be a graph manifold. This strategy is common to [5, 23, 30] and the present paper. In this subsection we describe the strategy in a bit more detail. Some of the new features of the present paper will be described more fully in Subsection 1.5.

1.4.1. An example. — The following simple example gives a useful illustration of the strategy of the proof.

Let $P \subset H^2$ be a compact convex polygonal domain in the two-dimensional hyperbolic space. Embedding H^2 in the four-dimensional hyperbolic space H^4 , let $N_s(P)$ be the metric s -neighborhood around P in H^4 . Take M to be the boundary $\partial N_s(P)$. If s is sufficiently small then one can check that the hypotheses of Theorem 1.3 are satisfied.

Consider the structure of M when s is small. There is a region $M^{\text{2-stratum}}$, lying at distance $\geq \text{const. } s$ from the boundary ∂P , which is the total space of a circle bundle. At scale comparable to s , a suitable neighborhood of a point in $M^{\text{2-stratum}}$ is nearly isometric to a product of a planar region with S^1 . There is also a region M^{edge} lying at distance $\leq \text{const. } s$ from an edge of P , but away from the vertices of P , which is

the total space of a 2-disk bundle. At scale comparable to s , a suitable neighborhood of a point in M^{edge} is nearly isometric to the product of an interval with a 2-disk. Finally, there is a region $M^{0\text{-stratum}}$ lying at distance $\leq \text{const.} \cdot s$ from the vertices of P . A connected component of $M^{0\text{-stratum}}$ is diffeomorphic to a 3-disk.

We can choose $M^{2\text{-stratum}}$, M^{edge} and $M^{0\text{-stratum}}$ so that there is a decomposition $M = M^{2\text{-stratum}} \cup M^{\text{edge}} \cup M^{0\text{-stratum}}$ with the property that on interfaces, fibration structures are compatible. Now $M^{\text{edge}} \cup M^{0\text{-stratum}}$ is a finite union of 3-disks and $D^2 \times I$'s, which is homeomorphic to a solid torus. Also, $M^{2\text{-stratum}}$ is a circle bundle over a 2-disk, *i.e.*, another solid torus, and $M^{2\text{-stratum}}$ intersects $M^{\text{edge}} \cup M^{0\text{-stratum}}$ in a 2-torus. So using this geometric decomposition, we recognize that M is a graph manifold. (In this case M is obviously diffeomorphic to S^3 , being the boundary of a convex set in H^4 , and so it is a graph manifold; the point is that one can recognize this using the geometric structure that comes from the local collapsing.)

1.4.2. Local collapsing. — The statement of Theorem 1.3 is in terms of a *local* lower curvature bound, as evidenced by the appearance of the curvature scale R_p . Assumption (1) of Theorem 1.3 can be considered to be a local collapsing statement. (This is in contrast to a global collapsing condition, where one assumes that the sectional curvatures are at least -1 and $\text{vol}(B(p, 1)) < \epsilon$ for every $p \in M$.) To clarify the local collapsing statement, we make one more definition.

Definition 1.5. — Let c_3 denote the volume of the Euclidean unit ball in \mathbb{R}^3 . Fix $\bar{w} \in (0, c_3)$. Given $p \in M$, the \bar{w} -volume scale at p is

$$(1.6) \quad r_p(\bar{w}) = \inf \{ r > 0 : \text{vol}(B(p, r)) = \bar{w} r^3 \}.$$

If there is no such r then we say that the \bar{w} -volume scale is infinite.

There are two ways to look at hypothesis (1) of Theorem 1.3, at the curvature scale or at the volume scale. Suppose first that we rescale the ball $B(p, R_p)$ to have radius one. Then the resulting ball will have sectional curvature bounded below by -1 and volume bounded above by w_0 . As w_0 will be small, we can say that on the curvature scale, the manifold is locally volume collapsed with respect to a lower curvature bound. On the other hand, suppose that we rescale $B(p, r_p(w_0))$ to have radius one. Let $B'(p, 1)$ denote the rescaled ball. Then $\text{vol}(B'(p, 1)) = w_0$. Hypothesis (1) of Theorem 1.3 implies that there is a big number \mathcal{R} so that the sectional curvature on the radius \mathcal{R} -ball $B'(p, \mathcal{R})$ (in the rescaled manifold) is bounded below by $-\frac{1}{\mathcal{R}^2}$. Using this, we deduce that on the volume scale, a large neighborhood of p is well approximated by a large region in a complete nonnegatively curved 3-manifold N_p . This gives a local model for the geometry of M . Furthermore, if w_0 is small then we can say that at the volume scale, the neighborhood of p is close in a coarse sense to a space of dimension less than three.

In order to prove Theorem 1.3, one must first choose on which scale to work. We could work on the curvature scale, or the volume scale, or some intermediate scale (as

is done in [5, 23, 30]). In this paper we will work consistently on the volume scale. This gives a uniform and simplifying approach.

1.4.3. Local structure. — At the volume scale, the local geometry of M is well approximated by that of a nonnegatively curved 3-manifold. (That we get a 3-manifold instead of a 3-dimensional Alexandrov space comes from the second assumption in Theorem 1.3.) The topology of nonnegatively curved 3-manifolds is known in the compact case by work of Hamilton [17, 18] and in the noncompact case by work of Cheeger-Gromoll [7]. In the latter case, the geometry is also well understood. Some relevant examples of such manifolds are:

- (1) $\mathbb{R}^2 \times S^1$,
- (2) $\mathbb{R} \times S^2$,
- (3) $\mathbb{R} \times \Sigma$, where Σ is a noncompact nonnegatively curved surface which is diffeomorphic to \mathbb{R}^2 and has a cylindrical end, and
- (4) $\mathbb{R} \times_{\mathbb{Z}_2} S^2$.

If a neighborhood of a point $p \in M$ is modeled by $\mathbb{R}^2 \times S^1$ at the volume scale then the length of the circle fiber is comparable to w_0 . Hence if w_0 is small then the neighborhood looks almost like a 2-plane. Similarly, if the neighborhood is modeled by $\mathbb{R} \times S^2$ then it looks almost like a line. On the other hand, if the neighborhood is modeled by $\mathbb{R} \times \Sigma$ then for small w_0 , the surface Σ looks almost like a half-line and the neighborhood looks almost like a half-plane. Finally, if the neighborhood is modeled by $\mathbb{R} \times_{\mathbb{Z}_2} S^2$ then it looks almost like a half-line.

1.4.4. Gluing. — The remaining issue is use the local geometry to deduce the global topology of M . This is a gluing issue, as the local models need to be glued together to obtain global information.

One must determine which local models should be glued together. We do this by means of a stratification of M . If p is a point in M then for $k \leq 2$, we say that p is a k -stratum point if on the volume scale, a large ball around p approximately splits off an \mathbb{R}^k -factor metrically, but not an \mathbb{R}^{k+1} -factor.

For $k \in \{1, 2\}$, neighborhoods of the k -stratum points will glue together in order to produce the total space of a fibration over a k -dimensional manifold. For example, neighborhoods of the 2-stratum points will glue together to form a circle bundle over a surface. Neighborhoods of the 0-stratum points play a somewhat different role. They will be inserted as “plugs”; for example, neighborhoods of the exceptional fibers in a Seifert fibration will arise in this way.

By considering how M is decomposed into these various subspaces that fiber, we will be able to show that M is a graph manifold.

1.5. Outline of the proof. — We now indicate the overall structure of the proof of Theorem 1.3. In this subsection we suppress parameters or denote them by const.

In the paper we will use some minimal facts about pointed Gromov-Hausdorff convergence and Alexandrov spaces, which are recalled in Section 3.

1.5.1. Modified volume scale. — The first step is to replace the volume scale by a slight modification of it. The motivation for this step is the fluctuation of the volume scale. Suppose that p and q are points in overlapping local models. As these local models are at the respective volume scales, there will be a problem in gluing the local models together if $r_q(\bar{w})$ differs wildly from $r_p(\bar{w})$. We need control on how the volume scale fluctuates on a ball of the form $B(p, \text{const} \cdot r_p(\bar{w}))$. We deal with this problem by replacing the volume scale $r_p(\bar{w})$ by a modified scale which has better properties. We assign a scale τ_p to each point $p \in M$ such that:

- (1) τ_p is much less than the curvature scale R_p .
- (2) The function $p \mapsto \tau_p$ is smooth and has Lipschitz constant $\Lambda \ll 1$.
- (3) The ball $B(p, \tau_p)$ has volume lying in the interval $[w'\tau_p^3, \bar{w}\tau_p^3]$, where $w' < \bar{w}$ are suitably chosen constants lying in the interval $[w_0, c_3]$.

The proof of the existence of the scale function $p \mapsto \tau_p$ follows readily from the local collapsing assumption, the Bishop-Gromov volume comparison theorem, and an argument similar to McShane's extension theorem for real-valued Lipschitz functions; see Section 6.

1.5.2. Implications of compactness. — Condition (1) above implies that the rescaled manifold $\frac{1}{\tau_p}M$, in the vicinity of p , has almost nonnegative curvature. Furthermore, condition (3) implies that it looks collapsed but not too collapsed, in the sense that the volume of the unit ball around p in the rescaled manifold $\frac{1}{\tau_p}M$ is small but not too small. Thus by working at the scale τ_p , we are able to retain the local collapsing assumption (in a somewhat weakened form) while gaining improved behavior of the scale function.

Next, the bounds (1.4) extend to give bounds on the derivatives of the curvature tensor of the form

$$(1.7) \quad |\nabla^k \text{Rm}| \leq A'(C, w')$$

for $0 \leq k \leq K$, when restricted to balls $B(p, C)$ in $\frac{1}{\tau_p}M$. Using (1.7) and standard compactness theorems for pointed Riemannian manifolds, we get:

- (4) For every $p \in M$, the rescaled pointed manifold $(\frac{1}{\tau(p)}M, p)$ is close in the pointed C^K -topology to a pointed nonnegatively curved C^K -smooth Riemannian 3-manifold (N_p, \star) .
- (5) For every $p \in M$, the pointed manifold $(\frac{1}{\tau(p)}M, p)$ is close in the pointed Gromov-Hausdorff topology to a pointed nonnegatively curved Alexandrov space (X_p, \star) of dimension at most 2.

1.5.3. *Stratification.* — The next step is to define a partition of M into k -stratum points, for $k \in \{0, 1, 2\}$. The partition is in terms of the number of \mathbb{R} -factors that approximately split off in $(\frac{1}{\tau(p)}M, p)$.

Let $0 < \beta_1 < \beta_2$ be new parameters. Working at scale τ_p , we classify points in M as follows (see Section 7):

- A point $p \in M$ lies in the **2-stratum** if $(\frac{1}{\tau(p)}M, p)$ is β_2 -close to $(\mathbb{R}^2, 0)$ in the pointed Gromov-Hausdorff topology.
- A point $p \in M$ lies in the **1-stratum**, if it does not lie in the 2-stratum, but $(\frac{1}{\tau(p)}M, p)$ is β_1 -close to $(\mathbb{R} \times Y_p, (0, \star_{Y_p}))$ in the pointed Gromov-Hausdorff topology, where Y_p is a point, a circle, an interval or a half-line, and \star_{Y_p} is a basepoint in Y_p .
- A point lies in the **0-stratum** if it does not lie in the k -stratum for $k \in \{1, 2\}$.

We now discuss the structure near points in the different strata in more detail, describing the model spaces X_p and N_p .

2-stratum points (Section 8). — If β_2 is small and $p \in M$ is a 2-stratum point then X_p is isometric to \mathbb{R}^2 , while N_p is isometric to a product $\mathbb{R}^2 \times S^1$ where the S^1 factor is small. Since the pointed rescaled manifold $(\frac{1}{\tau_p}M, p)$ is close to (N_p, \star) , we can transfer the projection map $N_p \simeq \mathbb{R}^2 \times S^1 \rightarrow \mathbb{R}^2$ to a map η_p defined on a large ball $B(p, C) \subset \frac{1}{\tau_p}M$, where it defines a circle fibration.

1-stratum points (Sections 9 and 10). — If β_1 is small and p is a 1-stratum point then $X_p = \mathbb{R} \times Y_p$, where Y_p is a point, a circle, an interval or a half-line. The C^K -smooth model space N_p will be an isometric product $\mathbb{R} \times \bar{N}_p$, where \bar{N}_p is a complete nonnegatively curved orientable surface. As in the 2-stratum case, we can transfer the projection map $N_p \simeq \mathbb{R} \times \bar{N}_p \rightarrow \mathbb{R}$ to a map η_p defined on a large ball $B(p, C) \subset \frac{1}{\tau_p}M$, where it defines a submersion.

We further classify the 1-stratum points according to the diameter of the cross-section Y_p . If the diameter of Y_p is not too large then we say that p lies in the *slim 1-stratum*. (The motivation for the terminology is that in this case $(\frac{1}{\tau_p}M, p)$ appears slim, being at moderate Gromov-Hausdorff distance from a line.) For slim 1-stratum points, the cross-section \bar{N}_p is diffeomorphic to S^2 or T^2 . Moreover, in this case the submersion η_p will be a fibration with fiber diffeomorphic to \bar{N}_p .

We also distinguish another type of 1-stratum point, the *edge points*. A 1-stratum point p is an edge point if $(X_p, \star) = (\mathbb{R} \times Y_p, \star)$ can be taken to be pointed isometric to a flat Euclidean half-plane whose basepoint lies on the edge. Roughly speaking, we show that near p , the set E' of edge points looks like a 1-dimensional manifold at scale τ_p . Furthermore, there is a smooth function $\eta_{E'}$ which behaves like the “distance to the edge” and which combines with η_p to yield “half-plane coordinates” for $\frac{1}{\tau_p}M$ near p . When restricted to an appropriate sublevel set of $\eta_{E'}$, the map η_p defines a fibration with fibers diffeomorphic to a compact surface with boundary F_p . Using the fact that for edge points $N_p = \mathbb{R} \times \bar{N}_p$ is Gromov-Hausdorff close to a

half-plane, one sees that the pointed surface (\bar{N}_p, \star) is Gromov-Hausdorff close to a pointed ray $([0, \infty), \star)$. This allows one to conclude that F_p is diffeomorphic to a closed 2-disk.

0-stratum points (Section 11). — We know by (4) that if p is a 0-stratum point, then $(\frac{1}{\mathfrak{r}_p}M, p)$ is C^K -close to a nonnegatively curved C^K -smooth 3-manifold N_p . The idea for analyzing the structure of M near a 0-stratum point p is to use the fact that nonnegatively curved manifolds look asymptotically like cones, and are diffeomorphic to any sufficiently large ball in them (centered at a fixed basepoint). More precisely, we find a scale r_p^0 with $\mathfrak{r}_p \leq r_p^0 \leq \text{const.} \cdot \mathfrak{r}_p$ so that:

- The pointed rescaled manifold $(\frac{1}{r_p^0}M, p)$ is close in the C^K -topology to a C^K -smooth nonnegatively curved 3-manifold (N'_p, \star) .
- The distance function d_p in $\frac{1}{r_p^0}M$ has no critical points in the metric annulus $A(p, \frac{1}{10}, 10) = \overline{B(p, 10)} - B(p, \frac{1}{10})$, and $B(p, 1) \subset \frac{1}{r_p^0}M$ is diffeomorphic to N'_p .
- The pointed space $(\frac{1}{r_p^0}M, p)$ is close in the pointed Gromov-Hausdorff topology to a Euclidean cone (in fact the Tits cone of N'_p).
- N'_p has at most one end.

The proof of the existence of the scale r_p^0 is based on the fact that nonnegatively curved manifolds are asymptotically conical, the critical point theory of Grove-Shiohama [16], and a compactness argument. Using the approximately conical structure, one obtains a smooth function η_p on $\frac{1}{r_p^0}M$ which, when restricted to the metric annulus $A(p, \frac{1}{10}, 10) \subset \frac{1}{r_p^0}M$, behaves like the radial function on a cone. In particular, for $t \in [\frac{1}{10}, 10]$, the sublevel sets $\eta_p^{-1}[0, t]$ are diffeomorphic to N'_p .

The soul theorem [7], together with Hamilton's classification of closed nonnegatively curved 3-manifolds [17, 18], implies that N'_p is diffeomorphic to one of the following: a manifold W/Γ where W is either S^3 , $S^2 \times S^1$ or T^3 equipped with a standard Riemannian metric and Γ is a finite group of isometries; $S^1 \times \mathbb{R}^2$; $S^2 \times \mathbb{R}$, $T^2 \times \mathbb{R}$; or a twisted line bundle over $\mathbb{R}P^2$ or the Klein bottle. Thus we know the possibilities for the topology of $B(p, 1) \subset \frac{1}{r_p^0}M$.

1.5.4. Compatibility of the local structures. — Having determined the local structure of M near each point, we examine how these local structures fit together on their overlap. For example, consider the slim 1-stratum points corresponding to an S^2 -fiber. A neighborhood of the set of such points looks like a union of cylindrical regions. If the axes of overlapping cylinders are very well-aligned then the process of gluing them together will be simplified. It turns out that such compatibility is automatic from our choice of stratification.

To see this, suppose that $p, q \in M$ are 2-stratum points with $B(p, \text{const.} \cdot \mathfrak{r}_p) \cap B(q, \text{const.} \cdot \mathfrak{r}_q) \neq \emptyset$. Then provided that Λ is small, we know that $\mathfrak{r}_p \approx \mathfrak{r}_q$. Suppose now that $z \in B(p, \text{const.} \cdot \mathfrak{r}_p) \cap B(q, \text{const.} \cdot \mathfrak{r}_q)$. We have two \mathbb{R}^2 -factors at z , coming from the approximate splittings at p and q . If the parameter β_2 is small then these

\mathbb{R}^2 -factors must align well at z . If not then we would get two misaligned \mathbb{R}^2 -factors at p , which would generate an approximate \mathbb{R}^3 -factor at p , contradicting the local collapsing assumption. Hence the maps η_p and η_q , which arose from approximate \mathbb{R}^2 -splittings, are nearly “aligned” with each other on their overlap, so that η_p and η_q are affine functions of each other, up to arbitrarily small C^1 -error.

Now fix β_2 . Let $p, q \in M$ be 1-stratum points. At any $z \in B(p, \text{const. } \tau_p) \cap B(q, \text{const. } \tau_q)$, there are two \mathbb{R} -factors, coming from the approximate \mathbb{R} -splittings at p and q . If β_1 is small then these two \mathbb{R} -factors must align well at z , or else we would get two misaligned \mathbb{R} -factors at p , contradicting the fact that p is not a 2-stratum point. Hence the functions η_p and η_q are also affine functions of each other, up to arbitrarily small C^1 -error.

One gets additional compatibility properties for pairs of points of different types. For example, if p lies in the 0-stratum and $q \in A(p, \frac{1}{10}, 10) \subset \frac{1}{\tau_p} M$ belongs to the 2-stratum then the radial function η_p , when appropriately rescaled, agrees with an affine function of η_q in $B(q, 10) \subset \frac{1}{\tau_q} M$ up to small C^1 -error.

1.5.5. Gluing the local pieces together (Sections 12–14). — To begin the gluing process, we select a separated collection of points of each type in M : $\{p_i\}_{i \in I_{2\text{-stratum}}}$, $\{p_i\}_{i \in I_{\text{slim}}}$, $\{p_i\}_{i \in I_{\text{edge}}}$, $\{p_i\}_{i \in I_{0\text{-stratum}}}$, so that

- $\bigcup_{i \in I_{2\text{-stratum}}} B(p_i, \text{const. } \tau_{p_i})$ covers the 2-stratum points,
- $\bigcup_{i \in I_{\text{slim}} \cup I_{\text{edge}}} B(p_i, \text{const. } \tau_{p_i})$ covers the 1-stratum points, and
- $\bigcup_{i \in I_{0\text{-stratum}}} B(p_i, \text{const. } r_{p_i}^0)$ covers the 0-stratum points.

Our next objective is to combine the η_{p_i} ’s so as to define global fibrations for each of the different types of points, and ensure that these fibrations are compatible on overlaps. To do this, we borrow an idea from the proof of the Whitney embedding theorem (as well as proofs of Gromov’s compactness theorem [14, Chapter 8.D], [20]): we define a smooth map $\mathcal{E}^0 : M \rightarrow H$ into a high-dimensional Euclidean space H . The components of \mathcal{E}^0 are functions of the η_{p_i} ’s, the edge function $\eta_{E'}$, and the scale function $p \mapsto \tau_p$, cutoff appropriately so that they define global smooth functions.

Due to the pairwise compatibility of the η_{p_i} ’s discussed above, it turns out that the image under \mathcal{E}^0 of $\bigcup_{i \in I_{2\text{-stratum}}} B(p_i, \text{const. } \tau_{p_i})$ is a subset $S \subset H$ which, when viewed at the right scale, is everywhere locally close (in the pointed Hausdorff sense) to a 2-dimensional affine subspace. We call such a set a *cloudy 2-manifold*. By an elementary argument, we show in Appendix B that a cloudy manifold of any dimension can be approximated by a core manifold W whose normal injectivity radius is controlled.

We adjust the map \mathcal{E}^0 by “pinching” it into the manifold core of S , thereby upgrading \mathcal{E}^0 to a new map \mathcal{E}^1 which is a circle fibration near the 2-stratum. The new map \mathcal{E}^1 is C^1 -close to \mathcal{E}^0 . We then perform similar adjustments near the edge points and slim 1-stratum points, to obtain a map $\mathcal{E} : M \rightarrow H$ which yields fibrations when

restricted to certain regions in M , see Section 13. For example, we obtain

- A D^2 -fibration of a region of M near the edge set E' ,
- S^2 or T^2 -fibrations of a region containing the slim stratum, and
- A surface fibration collaring (the boundary of) the region near 0-stratum points.

Furthermore, it is a feature built into the construction that where the fibered regions overlap, they do so in surfaces with boundary along which the two fibrations are compatible. For instance, the interface between the edge fibration and the 2-stratum fibration is a surface which inherits the same circle fibration from the edge fibration and the 2-stratum fibration. Similarly, the interface between the 2-stratum fibration and the slim 1-stratum fibration is a surface with boundary which inherits a circle fibration from the 2-stratum. See Proposition 14.1 for the properties of the fibrations.

1.5.6. Recognizing the graph manifold structure (Section 15). — At this stage of the argument, one has a decomposition of M into domains with disjoint interiors, where each domain is a compact 3-manifold with corners carrying a fibration of a specific kind, with compatibility of fibrations on overlaps. Using the topological classification of the fibers and the 0-stratum domains, one readily reads off the graph manifold structure. This completes the proof of Theorem 1.3.

1.5.7. Removing the bounds on derivatives of curvature (Section 18). — The proof of Theorem 1.3 uses the derivative bounds (1.4) only for C^K -precompactness results. In turn these are essentially used only to determine the topology of the 0-stratum balls and the fibers of the edge fibration. Without the derivative bounds (1.4), one can appeal to similar compactness arguments. However, one ends up with a sequence of pointed Riemannian manifolds $\{(M_k, \star_k)\}$ which converge in the pointed Gromov-Hausdorff topology to a pointed 3-dimensional nonnegatively curved Alexandrov space (M_∞, \star_∞) , rather than having C^K -convergence to a C^K -smooth limit. By invoking Perelman's Stability Theorem [19, 25], one can relate the topology of the limit space to those of the approximators. The only remaining step is to determine the topology of the nonnegatively curved Alexandrov spaces that arise as limits in this fashion. In the case of noncompact limits, this was done by Shioya-Yamaguchi [30]. In the compact case, it follows from Simon [31] or, alternatively, from the Ricci flow proof of the elliptization conjecture (using the finite time extinction results of Perelman and Colding-Minicozzi). For more details, we refer the reader to Section 18.

1.5.8. What's new in this paper. — The proofs of the collapsing theorems in [2, 5, 23, 29, 30], as well as the proof in this paper, all begin by comparing the local geometry at a certain scale with the geometry of a nonnegatively curved manifold, and then use this structure to deduce that one has a graph manifold. The paper [2] follows a rather different line from the other proofs, in that it uses the least amount of the information available from the nonnegatively curved models, and proceeds with a covering argument based on the theory of simplicial volume, as well as Thurston's

proof of the geometrization theorem for Haken manifolds. The papers [5, 23, 29, 30] and this paper have a common overall strategy, which is to use more of the theory of manifolds with nonnegative sectional curvature – Cheeger-Gromoll theory [7] and critical point theory [16] – to obtain a more refined version of the local models. Then the local models are spliced together to obtain a decomposition of the manifold into fibered regions from which one can recognize a graph manifold.

Overall, our proof uses a minimum of material beyond the theory of nonnegatively curved manifolds. It is essentially elementary in flavor. We now comment on some specific new points in our approach.

The scale function τ_p . — The existence of a scale function τ_p with the properties indicated in Section 1.5.2 makes it apparent that the theory of local collapsing is, at least philosophically, no different than the global version of collapsing.

We work consistently at the scale τ_p , which streamlines the argument. In particular, the structure theory of Alexandrov spaces, which enters if one works at the curvature scale, is largely eliminated. Also, in the selection argument, one considers ball covers where the radii are linked to the scale function τ , so one easily obtains bounds on the intersection multiplicity from the fact that the radii of intersecting balls are comparable (when the scale function $p \mapsto \tau_p$ has small Lipschitz constant). The technique of constructing a scale function with small Lipschitz constant could help in other geometric gluing problems.

The stratification. — Stratifications have a long history in geometric analysis, especially for singular spaces such as convex sets, minimal varieties, Alexandrov spaces, and Ricci limit spaces, where one typically looks at the number of \mathbb{R} -factors that split off in a tangent cone. The particular stratification that we use, based on the number of \mathbb{R} -factors that approximately split off in a manifold, was not used in collapsing theory before, to our knowledge. Its implications for achieving alignment may be useful in other settings.

The gluing procedure. — Passing from local models to global fibrations involves some kind of gluing process. Complications arise from the fact one has to construct a global base space for the fibration at the same time as one glues together the fibration maps; in addition, one has to make the fibrations from the different strata compatible. The most obvious approach is to add fibration patches inductively, by using small isotopies and the fact that on overlaps, the fibration maps are nearly affinely equivalent. Then one must perform further isotopies to make fibrations from the different strata compatible with one another. We find the gluing procedure used here to be more elegant; moreover, it produces fibrations which are automatically compatible.

Embeddings into a Euclidean space were used before to construct fibrations in a collapsing setting [12]. However, there is the important difference that in the earlier work the base of the fibration was already specified, and this base was embedded into

a Euclidean space. In contrast, in the present paper we must produce the base at the same time as the fibrations, so we produce it as a submanifold of the Euclidean space.

Cloudy manifolds. — The notion of cloudy manifolds, and the proof that they have a good manifold core, may be of independent interest. Cloudy manifolds are similar to objects that have been encountered before, in the work of Reifenberg [28] in geometric measure theory and also in [27]. However the clean elementary argument for the existence of a smooth core given in Appendix B, using the universal bundle and transversality, seems to be new.

1.5.9. A sketch of the history. — The theory of collapsing was first developed by Cheeger and Gromov [8, 9], assuming both upper and lower bounds on sectional curvature. Their work characterized the degeneration that can occur when one drops the injectivity radius bound in Gromov’s compactness theorem, generalizing Gromov’s theorem on almost flat manifolds [13]. The corresponding local collapsing structure was used by Anderson and Cheeger-Tian in work on Einstein manifolds [1, 10]. As far as we know, the first results on collapsing with a lower curvature bound were announced by Perelman in the early 90’s, as an application of the theory of Alexandrov spaces, in particular his Stability Theorem from [25] (see also [19]); however, these results were never published. Yamaguchi [33] established a fibration theorem for manifolds close to Riemannian manifolds, under a lower curvature bound. Shioya-Yamaguchi [29] studied collapsed 3-manifolds with a diameter bound and showed that they are graph manifolds, apart from an exceptional case. In [24], Perelman formulated without proof a theorem equivalent to our Theorem 1.3. A short time later, Shioya-Yamaguchi [30] proved that – apart from an exceptional case – sufficiently collapsed 3-manifolds are graph manifolds, this time without assuming a diameter bound. This result (or rather the localized version they discuss in their appendix) may be used in lieu of [24, Theorem 7.4] to complete the proof of the geometrization conjecture. Subsequently, Bessières-Besson-Boileau-Maillot-Porti [2] gave a different approach to the last part of the proof of the geometrization conjecture, which involves collapsing as well as refined results from 3-dimensional topology. Morgan-Tian [23] gave a proof of Perelman’s collapsing result along the lines of Shioya-Yamaguchi [30]. We also mention the paper [CG] by Cao-Ge which relies on more sophisticated Alexandrov space results.

Acknowledgements. — We thank Peter Scott for some references to the 3-manifold literature.

2. Notation and conventions

2.1. Parameters and constraints. — The rest of the paper develops a lengthy construction, many steps of which generate new constants; we will refer to these as *parameters*. Although the parameters remain fixed after being introduced, one should

view different sets of parameter values as defining different potential instances of the construction. This is necessary, because several arguments involve consideration of sequences of values for certain parameters, which one should associate with a sequence of distinct instances of the construction.

Many steps of the argument assert that certain statements hold provided that certain constraints on the parameters are satisfied. By convention, each time we refer to such a constraint, we will assume for the remainder of the paper that the inequalities in question are satisfied. Constraint functions will be denoted with a bar, e.g., $\beta_E < \bar{\beta}_E(\beta_1, \sigma)$ means that $\beta_E \in (0, \infty)$ satisfies an upper bound which is a function of β_1 and σ . By convention, all constraint functions take values in $(0, \infty)$.

At the end of the proof of Theorem 1.3, we will verify that the constraints on the various parameters can be imposed consistently. Fortunately, we do not have to carefully adjust each parameter in terms of the others; the constraints are rather of the form that one parameter is sufficiently small (or large) in terms of some others. Hence the only issue is the order in which the parameters are considered.

We follow Perelman's convention that a condition like $a > 0$ means that a should be considered to be a small parameter, while a condition like $A < \infty$ means that A should be considered to be a large parameter. This convention is only for expository purposes and may be ignored by a logically minded reader.

2.2. Notation. — We will use the following compact notation for cutoff functions with prescribed support. Let $\phi \in C^\infty(\mathbb{R})$ be a nonincreasing function so that $\phi|_{(-\infty, 0]} = 1$, $\phi|_{[1, \infty)} = 0$ and $\phi((0, 1)) \subset (0, 1)$. Given $a, b \in \mathbb{R}$ with $a < b$, we define $\Phi_{a,b} \in C^\infty(\mathbb{R})$ by

$$(2.1) \quad \Phi_{a,b}(x) = \phi(a + (b - a)x),$$

so that $\Phi_{a,b}|_{(-\infty, a]} = 1$ and $\Phi_{a,b}|_{[b, \infty)} = 0$. Given $a, b, c, d \in \mathbb{R}$ with $a < b < c < d$, we define $\Phi_{a,b,c,d} \in C^\infty(\mathbb{R})$ by

$$(2.2) \quad \Phi_{a,b,c,d}(x) = \phi_{-b,-a}(-x) \phi_{c,d}(x),$$

so that $\Phi_{a,b,c,d}|_{(-\infty, a]} = 0$, $\Phi_{a,b,c,d}|_{[b,c]} = 1$ and $\Phi_{a,b,c,d}|_{[d, \infty)} = 0$.

If X is a metric space and $0 < r \leq R$ then the annulus $A(x, r, R)$ is $\overline{B(x, R)} - B(x, r)$. The dimension of a metric space will always mean the Hausdorff dimension. For notation, if C is a metric cone with basepoint at the vertex \star then we will sometimes just write C for the pointed metric space (C, \star) . (Recall that a metric cone is a pointed metric space (Z, \star) , which is a union of rays leaving the basepoint \star , such that the union of any two such rays is isometric to the union of two rays leaving the origin in \mathbb{R}^2 .)

If Y is a subset of X and $t : Y \rightarrow (0, \infty)$ is a function then we write $N_t(Y)$ for the neighborhood of Y with variable thickness t : $N_t(Y) = \bigcup_{y \in Y} B(y, t(y))$.

If (X, d) is a metric space and $\lambda > 0$ then we write λX for the metric space $(X, \lambda d)$. For notational simplicity, we write $B(p, r) \subset \lambda X$ to denote the r -ball around p in the metric space λX .

Throughout the paper, a product of metric spaces $X_1 \times X_2$ will be endowed with the distance function given by the Pythagorean formula, *i.e.*, if $(x_1, x_2), (y_1, y_2) \in X_1 \times X_2$ then $d_{X_1 \times X_2}((x_1, x_2), (y_1, y_2)) = \sqrt{d_{X_1}^2(x_1, y_1) + d_{X_2}^2(x_2, y_2)}$.

2.3. Variables. — For the reader's convenience, we tabulate the variables in this paper, listed by the section in which they first appear.

- Section 1.2: R_p
- Section 1.4.2: $r_p(\cdot)$
- Section 6: $\Lambda, \sigma, \mathfrak{r}_p, \bar{w}, w'$
- Section 7.2: $\beta_1, \beta_2, \beta_3, \Delta$
- Section 8.1: ς_2 -stratum, η_p and ζ_p (for 2-stratum points)
- Section 8.2: \mathcal{M}
- Section 9.1: $\beta_E, \sigma_E, \beta_{E'}, \sigma_{E'}$
- Section 9.2: $d_{E'}, \rho_{E'}, \varsigma_{E'}$
- Section 9.3: $\varsigma_{\text{edge}}, \eta_p$ and ζ_p (for edge points)
- Section 9.6: $\zeta_{\text{edge}}, \zeta_{E'}$
- Section 10.1: $\varsigma_{\text{slim}}, \eta_p$ and ζ_p (for slim 1-stratum points)
- Section 11.1: $\Upsilon_0, \Upsilon'_0, \delta_0, r_p^0$
- Section 11.2: ς_0 -stratum, η_p and ζ_p (for 0-stratum points)
- Section 12.1: H, H_i, H'_i, H''_i, H_0 -stratum, $H_{\text{slim}}, H_{\text{edge}}, H_2$ -stratum, $Q_1, Q_2, Q_3, Q_4,$
 $\pi_i, \pi_i^\perp, \pi_{ij}, \pi_{H'_i}, \pi_{H''_i}, \mathcal{E}^0$
- Section 12.2: Ω_0
- Section 12.3: $A_1, \tilde{A}_1, S_1, \tilde{S}_1, \Omega_1, \Gamma_1, \Sigma_1, r_1, \Omega_1$
- Section 12.4: $A_2, \tilde{A}_2, S_2, \tilde{S}_2, \Omega_2, \Gamma_2, \Sigma_2, r_2, \Omega_2$
- Section 12.5: $A_3, \tilde{A}_3, S_3, \tilde{S}_3, \Omega_3, \Gamma_3, \Sigma_3, r_3, \Omega_3$
- Section 13: c_{adjust}
- Section 13.2: $W_1^0, \Xi_1, \psi_1, \Psi_1, \Omega'_1, \mathcal{E}^1, c_2$ -stratum
- Section 13.3: $W_2^0, \Xi_2, \psi_2, \Psi_2, \Omega'_2, \mathcal{E}^2, c_{\text{edge}}$
- Section 13.4: $W_3^0, \Xi_3, \psi_3, \Psi_3, \Omega'_3, \mathcal{E}^3, c_{\text{slim}}$
- Section 13.5: W_1, W_2, W_3
- Section 14.1: M^0 -stratum, M_1
- Section 14.2: $W'_3, U'_3, W''_3, M^{\text{slim}}, M_2$
- Section 14.3: $W'_2, U'_2, W''_2, M^{\text{edge}}, M_3$
- Section 14.4: W'_1, U'_1, W''_1, M^2 -stratum
- Section 15: $r_\partial, H_\partial, M_i^\partial$
- Appendix B: $S, \tilde{S}, r(\cdot), W$

3. Preliminaries

We refer to [3] for basics about length spaces and Alexandrov spaces.

3.1. Pointed Gromov-Hausdorff approximations

Definition 3.1. — Let (X, \star_X) be a pointed metric space. Given $\delta \in [0, \infty)$, two closed subspaces C_1 and C_2 are δ -close in the pointed Hausdorff sense if $C_1 \cap B(\star_X, \delta^{-1})$ and $C_2 \cap B(\star_X, \delta^{-1})$ have Hausdorff distance at most δ .

If X is complete and proper (*i.e.*, closed bounded sets are compact) then the corresponding pointed Hausdorff topology, on the set of closed subspaces of X , is compact and metrizable.

We now recall some definitions and basic results about the pointed Gromov-Hausdorff topology [3, Chapter 8.1].

Definition 3.2. — Let (X, \star_X) and (Y, \star_Y) be pointed metric spaces. Given $\delta \in [0, 1)$, a pointed map $f : (X, \star_X) \rightarrow (Y, \star_Y)$ is a δ -Gromov-Hausdorff approximation if for every $x_1, x_2 \in B(\star_X, \delta^{-1})$ and $y \in B(\star_Y, \delta^{-1} - \delta)$, we have

$$(3.3) \quad |d_Y(f(x_1), f(x_2)) - d_X(x_1, x_2)| \leq \delta \quad \text{and} \quad d_Y(y, f(B(\star_X, \delta^{-1}))) \leq \delta.$$

Two pointed metric spaces (X, \star_X) and (Y, \star_Y) are δ -close in the pointed Gromov-Hausdorff topology (or δ -close for short) if there is a δ -Gromov-Hausdorff approximation from (X, \star_X) to (Y, \star_Y) . We note that this does not define a metric space structure on the set of pointed metric spaces, but nevertheless defines a topology which happens to be metrizable. A sequence $\{(X_i, \star_i)\}_{i=1}^\infty$ of pointed metric spaces Gromov-Hausdorff converges to (Y, \star_Y) if there is a sequence $\{f_i : (X_i, \star_{X_i}) \rightarrow (Y, \star_Y)\}_{i=1}^\infty$ of δ_i -Gromov-Hausdorff approximations, where $\delta_i \rightarrow 0$. We will denote this by $(X_i, \star_{X_i}) \xrightarrow{\text{GH}} (Y, \star_Y)$.

Note that a δ -Gromov-Hausdorff approximation is a δ' -Gromov-Hausdorff approximation for every $\delta' \geq \delta$. A δ -Gromov-Hausdorff approximation f has a quasi-inverse $\hat{f} : (Y, \star_Y) \rightarrow (X, \star_X)$ constructed by saying that for $y \in B(\star_Y, \delta^{-1} - \delta)$, we choose some $x \in B(\star_X, \delta^{-1})$ with $d_Y(y, f(x)) \leq \delta$ and put $\hat{f}(y) = x$. There is a function $\delta' = \delta'(\delta) > 0$ with $\lim_{\delta \rightarrow 0} \delta' = 0$ so that if f is a δ -Gromov-Hausdorff approximation then \hat{f} is a δ' -Gromov-Hausdorff approximation and $\hat{f} \circ f$ (resp. $f \circ \hat{f}$) is δ' -close to the identity on $B(\star_X, (\delta')^{-1})$ (resp. $B(\star_Y, (\delta')^{-1})$). The condition $(X_i, \star_{X_i}) \xrightarrow{\text{GH}} (Y, \star_Y)$ is equivalent to the existence of a sequence $\{f_i : (Y, \star_Y) \rightarrow (X_i, \star_{X_i})\}_{i=1}^\infty$ (note the reversal of domain and target) of δ_i -Gromov-Hausdorff approximations, where $\delta_i \rightarrow 0$.

The relation of being δ -close is not symmetric. However, this does not create a problem because only the associated notion of convergence (*i.e.*, the topology) plays a role in our discussion.

The pointed Gromov-Hausdorff topology is a complete metrizable topology on the set of complete proper metric spaces (taken modulo pointed isometry). Hence we can talk about two such metric spaces being having distance at most δ from each other. There is a well-known criterion for a set of pointed metric spaces to be precompact in the pointed Gromov-Hausdorff topology [3, Theorem 8.1.10]. Complete proper *length spaces*, which are the main interest of this paper, form a closed subset of the set of complete proper metric spaces.

3.2. C^K -convergence

Definition 3.4. — Given $K \in \mathbb{Z}^+$, let (M_1, \star_{M_1}) and (M_2, \star_{M_2}) be complete pointed C^K -smooth Riemannian manifolds. (That is, the manifold transition maps are C^{K+1} and the metric in local coordinates is C^K). Given $\delta \in [0, \infty)$, a pointed C^{K+1} -smooth map $f : (M_1, \star_{M_1}) \rightarrow (M_2, \star_{M_2})$ is a δ - C^K approximation if it is a δ -Gromov-Hausdorff approximation and the C^K -norm of $f^*g_{M_2} - g_{M_1}$, computed on $B(\star_{M_1}, \delta^{-1})$, is bounded above by δ . Two C^K -smooth Riemannian manifolds (M_1, \star_{M_1}) and (M_2, \star_{M_2}) are δ - C^K close if there is a δ - C^K approximation from (M_1, \star_{M_1}) to (M_2, \star_{M_2}) .

In what follows, we will always take $K \geq 10$. We now state a C^K -precompactness result.

Lemma 3.5 (cf. [26, Chapter 10]). — Given $v, r > 0$, $n \in \mathbb{Z}^+$ and a function $A : (0, \infty) \rightarrow (0, \infty)$, the set of complete pointed C^{K+2} -smooth n -dimensional Riemannian manifolds (M, \star_M) such that

- (1) $\text{vol}(B(\star_M, r)) \geq v$ and
- (2) $|\nabla^k \text{Rm}| \leq A(R)$ on $B(\star_M, R)$, for all $0 \leq k \leq K$ and $R > 0$,

is precompact in the pointed C^K -topology.

The bounds on the derivatives of curvature in Lemma 3.5 give uniform C^{K+1} -bounds on the Riemannian metric in harmonic coordinates. One then obtains limit metrics which are C^K -smooth. One can get improved regularity but we will not need it.

3.3. Alexandrov spaces. — Recall that there is a notion of an Alexandrov space of curvature at least c , or equivalently a complete length space X having curvature bounded below by $c \in \mathbb{R}$ on an open set $U \subset X$ [3, Chapter 4]. In this paper we will only be concerned with Alexandrov spaces of finite Hausdorff dimension, so this will be assumed implicitly without further mention.

We will also have occasion to work with incomplete, but locally complete spaces. This situation typically arises when one has a metric space X where $X = B(p, r)$,

and every closed ball $\overline{B(p, r')}$ with $r' < r$ is complete. The version of Toponogov's theorem for Alexandrov spaces [3, Chapter 10.3], in which one deduces global triangle comparison inequality from local ones, also applies in the incomplete situation, provided that the geodesics arising in the proof lie in an *a priori* complete part of the space. In particular, if all sides of a geodesic triangle have length $< D$ then triangle comparison is valid provided that the closed balls of radius $2D$ centered at the vertices are complete.

We recall the notion of a strainer (cf. [3, Definition 10.8.9]).

Definition 3.6. — Given a point p in an Alexandrov space X of curvature at least c , an m -strainer at p of quality δ and scale r is a collection $\{(a_i, b_i)\}_{i=1}^m$ of pairs of points such that $d(p, a_i) = d(p, b_i) = r$ and in terms of comparison angles,

$$(3.7) \quad \begin{aligned} \tilde{Z}_p(a_i, b_i) &> \pi - \delta, \\ \tilde{Z}_p(a_i, a_j) &> \frac{\pi}{2} - \delta, \\ \tilde{Z}_p(a_i, b_j) &> \frac{\pi}{2} - \delta, \\ \tilde{Z}_p(b_i, b_j) &> \frac{\pi}{2} - \delta \end{aligned}$$

for all $i, j \in \{1, \dots, m\}$, $i \neq j$. Note that the comparison angles are defined using comparison triangles in the model space of constant curvature c .

For facts about strainers, we refer to [3, Chapter 10.8.2]. The Hausdorff dimension of X equals its strainer number, which is defined as follows.

Definition 3.8. — The *strainer number* of X is the supremum of numbers m such that there exists an m -strainer of quality $\frac{1}{100m}$ at some point and some scale.

By “dimension of X ” we will mean the Hausdorff dimension; this coincides with its topological dimension, although we will not need this fact. If (X, \star_X) is a pointed nonnegatively curved Alexandrov space then there is a pointed Gromov-Hausdorff limit $C_T X = \lim_{\lambda \rightarrow \infty} (\frac{1}{\lambda} X, \star_X)$ called the *Tits cone* of X . It is a nonnegatively curved Alexandrov space which is a metric cone, as defined in Subsection 2.2. We will consider Tits cones in the special case when X is a nonnegatively curved Riemannian manifold.

A *line* in a length space X is a curve $\gamma : \mathbb{R} \rightarrow X$ with the property that $d_X(\gamma(t), \gamma(t')) = |t - t'|$ for all $t, t' \in \mathbb{R}$. The splitting theorem (see [3, Chapter 10.5]) says that if a nonnegatively curved Alexandrov space X contains a line then it is isometric to $\mathbb{R} \times Y$ for some nonnegatively curved Alexandrov space Y .

If $\gamma : [0, T] \rightarrow X$ is a minimal geodesic in an Alexandrov space X , parametrized by arc-length, and $p \neq \gamma(0)$ is a point in X then the function $t \rightarrow d_p(\gamma(t))$ is right-differentiable and

$$(3.9) \quad \lim_{t \rightarrow 0^+} \frac{d_p(\gamma(t)) - d_p(\gamma(0))}{t} = -\cos \theta,$$

where θ is the minimal angle between γ and minimizing geodesics from $\gamma(0)$ to p [3, Corollary 4.5.7].

The proof of the next lemma is similar to that of [3, Theorem 10.7.2].

Lemma 3.10. — *Given $n \in \mathbb{Z}^+$, let $\{(X_i, \star_{X_i})\}_{i=1}^\infty$ be a sequence of complete pointed length spaces. Suppose that $c_i \rightarrow 0$ and $r_i \rightarrow \infty$ are positive sequences such that for each i , the ball $B(\star_{X_i}, r_i)$ has curvature bounded below by $-c_i$ and dimension bounded above by n . Then a subsequence of the (X_i, \star_{X_i}) 's converges in the pointed Gromov-Hausdorff topology to a pointed nonnegatively curved Alexandrov space of dimension at most n .*

3.4. Critical point theory. — We recall a few facts about critical point theory here, and refer the reader to [6, 16] for more information.

If M is a complete Riemannian manifold and $p \in M$, then a point $x \in M \setminus \{p\}$ is *noncritical* if there is a nonzero vector $v \in T_x M$ making an angle strictly larger than $\frac{\pi}{2}$ with the initial velocity of every minimizing segment from x to p . If there are no critical points in the set $d_p^{-1}(a, b)$ then the level sets $\{d_p^{-1}(t)\}_{t \in (a, b)}$, are pairwise isotopic Lipschitz hypersurfaces, and $d_p^{-1}(a, b)$ is diffeomorphic to a product, as in the usual Morse lemma for smooth functions. As with the traditional Morse lemma, the proof proceeds by constructing a smooth vector field ξ such that d_p has uniformly positive directional derivative in the direction ξ .

3.5. Topology of nonnegatively curved 3-manifolds. — In this subsection we describe the topology of certain nonnegatively curved manifolds. We start with 3-manifolds.

Lemma 3.11. — *Let M be a complete connected orientable 3-dimensional C^K -smooth Riemannian manifold with nonnegative sectional curvature. We have the following classification of the diffeomorphism type of M , based on the number of ends :*

- 0 ends: $S^1 \times S^2$, $S^1 \times_{\mathbb{Z}_2} S^2 = \mathbb{R}P^3 \# \mathbb{R}P^3$, T^3/Γ (where Γ is a finite subgroup of $\text{Isom}^+(T^3)$ which acts freely on T^3) or S^3/Γ (where Γ is a finite subgroup of $\text{SO}(4)$ that acts freely on S^3).
- 1 end: \mathbb{R}^3 , $S^1 \times \mathbb{R}^2$, $S^2 \times_{\mathbb{Z}_2} \mathbb{R} = \mathbb{R}P^3 - B^3$ or $T^2 \times_{\mathbb{Z}_2} \mathbb{R} =$ a flat \mathbb{R} -bundle over the Klein bottle.
- 2 ends: $S^2 \times \mathbb{R}$ or $T^2 \times \mathbb{R}$.

If M has two ends then it splits off an \mathbb{R} -factor isometrically.

Proof. — If M has no end then it is compact and the result follows for C^∞ -metrics from [18]. For C^K -smooth metrics, one could adapt the argument in [18] or alternatively use [31].

If M is noncompact then the Cheeger-Gromoll soul theorem says that M is diffeomorphic to the total space of a vector bundle over its soul, a closed lower-dimensional manifold with nonnegative sectional curvature [7]. (The proof in [7], which is for C^∞ -metrics, goes through without change for C^K -smooth metrics.) The possible dimensions of the soul are 0, 1 and 2. The possible topologies of M are listed in the lemma.

If M has two ends then it contains a line and the Toponogov splitting theorem implies that M splits off an \mathbb{R} -factor isometrically. \square

We now look at a pointed nonnegatively curved surface and describe the topology of a ball in it which is pointed Gromov-Hausdorff close to an interval.

Lemma 3.12. — *Suppose that (S, \star_S) is a pointed C^K -smooth nonnegatively curved complete orientable Riemannian 2-manifold. Let $\star_S \in S$ be a basepoint and suppose that the pointed ball $(B(\star_S, 10), \star_S)$ has pointed Gromov-Hausdorff distance at most δ from the pointed interval $([0, 10], 0)$.*

- (1) *Given $\theta > 0$ there is some $\bar{\delta}(\theta) > 0$ so that if $\delta < \bar{\delta}(\theta)$ then for every $x \in \overline{B(\star_S, 9)} - B(\star_S, 1)$ the set V_x of initial velocities of minimizing geodesic segments from x to \star_S has diameter bounded above by θ .*
- (2) *There is some $\bar{\delta} > 0$ so that if $\delta < \bar{\delta}$ then for every $r \in [1, 9]$ the ball $\overline{B(\star_S, r)}$ is homeomorphic to a closed 2-disk.*

Proof

(1). Choose a point x' with $d_S(\star_S, x') = 9.5$. Fix a minimizing geodesic γ' from x to x' and a minimizing geodesic γ'' from \star_S to x' . If γ is a minimizing geodesic from \star_S to x , consider the geodesic triangle with edges γ , γ' and γ'' . As

$$(3.13) \quad d(\star_S, x) + d(x, x') - d(\star_S, x'') \leq \text{const. } \delta,$$

triangle comparison implies that the angle at x between γ and γ' is bounded below by $\pi - a(\delta)$, where a is a positive monotonic function with $\lim_{\delta \rightarrow 0} a(\delta) = 0$. We take $\bar{\delta}$ so that $2a(\bar{\delta}) \leq \theta$.

(2). Suppose that $\delta < \bar{\delta}(\frac{\pi}{4})$. By critical point theory, the distance function $d_{\star_S} : A(\star_S, 1, 9) \rightarrow [1, 9]$ is a fibration with fibers diffeomorphic to a disjoint union of circles. In particular, the closed balls $\overline{B(\star_S, r)}$, for $r \in [1, 9]$, are pairwise homeomorphic. When $\delta \ll 1$, the fibers will be connected, since the diameter of $d_{\star_S}^{-1}(5)$ will be comparable to δ . Hence $\overline{B(\star_S, 1)}$ is homeomorphic to a surface with circle boundary.

Suppose that $\overline{B(\star_S, 1)}$ is not homeomorphic to a disk. A complete connected orientable nonnegatively curved surface is homeomorphic to S^2 , T^2 , \mathbb{R}^2 , or $S^1 \times \mathbb{R}$. By elementary topology, the only possibility is if S is homeomorphic to a 2-torus,

$\overline{B(\star_S, 1)}$ is homeomorphic to the complement of a 2-ball in S , and $S - B(\star_S, 2)$ is homeomorphic to a disk. In this case, S must be flat. However, the cylinder $A(\star_S, 1, 2)$ lifts to the universal cover of S , which is isometric to the flat \mathbb{R}^2 . If δ is sufficiently small then the flat \mathbb{R}^2 would contain a metric ball of radius $\frac{1}{10}$ which is Gromov-Hausdorff close to an interval, giving a contradiction. \square

3.6. Smoothing Lipschitz functions. — The technique of smoothing Lipschitz functions was introduced in Riemannian geometry by Grove and Shiohama [16].

If M is a Riemannian manifold and F is a Lipschitz function on M then the generalized gradient of F at $m \in M$ can be defined as follows. Given $\epsilon \in (0, \text{InjRad}_m)$, if $x \in B(m, \epsilon)$ is a point of differentiability of F then compute $\nabla_x F \in T_x M$ and parallel transport it along the minimizing geodesic to m . Take the closed convex hull of the vectors so obtained and then take the intersection as $\epsilon \rightarrow 0$. This gives a closed convex subset of $T_m M$, which is the generalized gradient of F at m [11]; we will denote this set by $\nabla_m^{\text{gen}} F$. The union $\bigcup_{m \in M} \nabla_m^{\text{gen}} F \subset TM$ will be denoted $\nabla^{\text{gen}} F$.

Lemma 3.14. — *Let M be a complete Riemannian manifold and let $\pi : TM \rightarrow M$ be the projection map. Suppose that $U \subset M$ is an open set, $C \subset U$ is a compact subset and S is an open fiberwise-convex subset of $\pi^{-1}(U)$.*

Then for every $\epsilon > 0$ and any Lipschitz function $F : M \rightarrow \mathbb{R}$ whose generalized gradient over U lies in S , there is a Lipschitz function $\widehat{F} : M \rightarrow \mathbb{R}$ such that:

- (1) \widehat{F} is C^∞ on an open set containing C .
- (2) The generalized gradient of \widehat{F} , over U , lies in S . (In particular, at every point in U where \widehat{F} is differentiable, the gradient lies in S .)
- (3) $|\widehat{F} - F|_\infty \leq \epsilon$.
- (4) $\widehat{F}|_{M-U} = F|_{M-U}$.

The proof of Lemma 3.14 proceeds by mollifying the Lipschitz function F , as in [16, Section 2]. We omit the details.

Corollary 3.15. — *Suppose that M is a compact Riemannian manifold. Given $K < \infty$ and $\epsilon > 0$, for any K -Lipschitz function F on M there is a $(K + \epsilon)$ -Lipschitz function $\widehat{F} \in C^\infty(M)$ with $|\widehat{F} - F|_\infty \leq \epsilon$.*

Proof. — Apply Lemma 3.14 with $C = U = M$, and $S = \{v \in TM : |v| < K + \epsilon\}$. \square

Corollary 3.16. — *For all $\epsilon > 0$ there is a $\theta > 0$ with the following property.*

Let M be a complete Riemannian manifold, let $Y \subset M$ be a closed subset and let $d_Y : M \rightarrow \mathbb{R}$ be the distance function from Y . Given $p \in M - Y$, let $V_p \subset T_p M$ be the set of initial velocities of minimizing geodesics from p to Y . Suppose that $U \subset M - Y$ is an open subset such that for all $p \in U$, one has $\text{diam}(V_p) < \theta$. Let C be a

compact subset of U . Then for every $\epsilon_1 > 0$ there is a Lipschitz function $\widehat{F} : M \rightarrow \mathbb{R}$ such that

- (1) \widehat{F} is smooth on a neighborhood of C .
- (2) $\|\widehat{F} - d_Y\|_\infty < \epsilon_1$.
- (3) $\widehat{F}|_{M-U} = d_Y|_{M-U}$.
- (4) For every $p \in C$, the angle between $-(\nabla\widehat{F})(p)$ and V_p is at most ϵ .
- (5) $\widehat{F} - d_Y$ is ϵ -Lipschitz.

Proof. — First, note that if $p \in M - Y$ is a point of differentiability of the distance function d_Y then $\nabla_p d_Y = -V_p$. Also, the assignment $x \mapsto V_x$ is semicontinuous in the sense that if $\{x_k\}_{k=1}^\infty$ is a sequence of points converging to x then by parallel transporting V_{x_k} radially to the fiber over x , we obtain a sequence $\{\bar{V}_{x_k}\}_{k=1}^\infty \subset T_x M$ which accumulates on a subset of V_x . It follows that the generalized derivative of d_Y at any point $p \in M - Y$ is precisely $-\text{Hull}(V_p)$, where $\text{Hull}(V_p)$ denotes the convex hull of V_p .

Put $S' = \bigcup_{p \in U} \text{Hull}(-V_p)$. Then S' is a relatively closed fiberwise-convex subset of $\pi^{-1}(U)$, with fibers of diameter less than θ . We can fatten S' slightly to form an open fiberwise-convex set $S \subset \pi^{-1}(U)$ which contains S' , with fibers of diameter less than 2θ .

Now take $\theta < \frac{\epsilon}{2}$ and apply Lemma 3.14 to $F = d_Y : M \rightarrow \mathbb{R}$, with S as in the preceding paragraph. The resulting function $\widehat{F} : M \rightarrow \mathbb{R}$ clearly satisfies (1)-(4). To see that (5) holds, note that if $p \in U$ is a point of differentiability of both \widehat{F} and F then $\nabla\widehat{F}(p)$ and $\nabla F(p)$ both lie in the fiber $S \cap T_p M$, which has diameter less than 2θ . Hence the gradient of the difference satisfies

$$(3.17) \quad \|\nabla(\widehat{F} - F)(p)\| = \|\nabla\widehat{F}(p) - \nabla F(p)\| < 2\theta < \epsilon.$$

Since \widehat{F} coincides with F outside U , this implies that $\widehat{F} - F$ is ϵ -Lipschitz. □

Remark 3.18. — When we apply Corollary 3.16, the hypotheses will be verified using triangle comparison.

4. Splittings, strainers, and adapted coordinates

This section is about the notion of a pointed metric space approximately splitting off an \mathbb{R}^k -factor. We first define an approximate \mathbb{R}^k -splitting, along with the notion of compatibility between an approximate \mathbb{R}^k -splitting and an approximate \mathbb{R}^j -splitting. We prove basic properties about approximate splittings. In the case of a pointed Alexandrov space, we show that having an approximate \mathbb{R}^k -splitting is equivalent to having a good k -strainer. We show that if there is not an approximate \mathbb{R}^{k+1} -splitting at a point p then any approximate \mathbb{R}^k -splitting at p is nearly-compatible with any approximate \mathbb{R}^j -splitting at p , for $j \leq k$.

We then introduce the notion of coordinates adapted to an approximate \mathbb{R}^k -splitting, in the setting of Riemannian manifolds with a lower curvature bound, proving existence and (approximate) uniqueness of such adapted coordinates.

4.1. Splittings. — We start with the notion of a splitting.

Definition 4.1. — A product structure on a metric space X is an isometry $\alpha : X \rightarrow X_1 \times X_2$. A k -splitting of X is a product structure $\alpha : X \rightarrow X_1 \times X_2$ where X_1 is isometric to \mathbb{R}^k . A *splitting* is a k -splitting for some k . Two k -splittings $\alpha : X \rightarrow X_1 \times X_2$ and $\beta : X \rightarrow Y_1 \times Y_2$ are *equivalent* if there are isometries $\phi_i : X_i \rightarrow Y_i$ such that $\beta = (\phi_1, \phi_2) \circ \alpha$.

In addition to equivalence of splittings, we can talk about compatibility of splittings.

Definition 4.2. — Suppose that $j \leq k$. A j -splitting $\alpha : X \rightarrow X_1 \times X_2$ is *compatible* with a k -splitting $\beta : X \rightarrow Y_1 \times Y_2$ if there is a j -splitting $\phi : Y_1 \rightarrow \mathbb{R}^j \times \mathbb{R}^{k-j}$ such that α is equivalent to the j -splitting given by the composition

$$(4.3) \quad X \xrightarrow{\beta} Y_1 \times Y_2 \xrightarrow{(\phi, \text{Id})} (\mathbb{R}^j \times \mathbb{R}^{k-j}) \times Y_2 \cong \mathbb{R}^j \times (\mathbb{R}^{k-j} \times Y_2).$$

Lemma 4.4

- (1) *Suppose $\alpha : X \rightarrow \mathbb{R}^k \times Y$ is a k -splitting of a metric space X , and $\beta : X \rightarrow \mathbb{R} \times Z$ is a 1-splitting. Then either β is compatible with α , or there is a 1-splitting $\gamma : Y \rightarrow \mathbb{R} \times W$ such that β is compatible with the induced splitting $X \rightarrow (\mathbb{R}^k \times \mathbb{R}) \times W$.*
- (2) *Any two splittings of a metric space are compatible with a third splitting.*

Before proving this, we need a sublemma.

Recall that a *line* in a metric space is a globally minimizing complete geodesic, *i.e.*, an isometrically embedded copy of \mathbb{R} . We will say that two lines are *parallel* if their union is isometric to the union of two parallel lines in \mathbb{R}^2 .

Sublemma 4.5

- (1) *A path $\gamma : \mathbb{R} \rightarrow X_1 \times X_2$ in a product is a constant speed geodesic if and only if the compositions $\pi_{X_i} \circ \gamma : \mathbb{R} \rightarrow X_i$ are constant speed geodesics.*
- (2) *Two lines in a metric space X are parallel if and only if they have constant speed parametrizations $\gamma_1 : \mathbb{R} \rightarrow X$ and $\gamma_2 : \mathbb{R} \rightarrow X$ such that $d^2(\gamma_1(s), \gamma_2(t))$ is a quadratic function of $(s - t)$.*
- (3) *If two lines γ_1, γ_2 in a product $\mathbb{R}^k \times X$ are parallel then either $\pi_X(\gamma_1), \pi_X(\gamma_2) \subset X$ are parallel lines, or they are both points.*
- (4) *Suppose \mathcal{L} is a collection of lines in a metric space X . If $\bigcup_{\gamma \in \mathcal{L}} \gamma = X$, and every pair $\gamma_1, \gamma_2 \in \mathcal{L}$ is parallel, then there is a 1-splitting $\alpha : X \rightarrow \mathbb{R} \times Y$ such that $\mathcal{L} = \{\alpha^{-1}(\mathbb{R} \times \{y\})\}_{y \in Y}$.*

Proof

(1). It follows from the Cauchy-Schwarz inequality that if $a, b, c \in X_1 \times X_2$ satisfy the triangle equation $d(a, c) = d(a, b) + d(b, c)$ then the same is true of their projections $a_1, b_1, c_1 \in X_1$ and $a_2, b_2, c_2 \in X_2$, and moreover $(d(a_1, b_1), d(a_2, b_2))$ and $(d(a_1, c_1), d(a_2, c_2))$ are linearly dependent in \mathbb{R}^2 . This implies (1).

(2). The parallel lines $y = 0$ and $y = a$ in \mathbb{R}^2 can be parametrized by $\gamma_1(s) = (s, 0)$ and $\gamma_2(t) = (t, a)$, with $d^2(\gamma_1(s), \gamma_2(t)) = (s - t)^2 + a^2$. Conversely, suppose that lines γ_1 and γ_2 in a metric space are such that $d^2(\gamma_1(s), \gamma_2(t))$ is quadratic in $(s - t)$. After affine changes of s and t , we can assume that $d^2(\gamma_1(s), \gamma_2(t)) = (s - t)^2 + a^2$ for some $a \in \mathbb{R}$. Then the union of γ_1 and γ_2 is isometric to the union of the lines $y = 0$ and $y = a$ in \mathbb{R}^2 .

(3). By (2), we may assume that for $i \in \{1, 2\}$ there are constant speed parametrizations $\gamma_i : \mathbb{R} \rightarrow \mathbb{R}^k \times X$, such that $d^2(\gamma_1(s), \gamma_2(t))$ is a quadratic function of $(s - t)$. The projections $\pi_{\mathbb{R}^k} \circ \gamma_i$ are constant speed geodesics in \mathbb{R}^k , and the quadratic function $d^2(\pi_{\mathbb{R}^k} \circ \gamma_1(s), \pi_{\mathbb{R}^k} \circ \gamma_2(t))$ is a function of $(s - t)$; otherwise $d^2(\pi_{\mathbb{R}^k} \circ \gamma_1(s), \pi_{\mathbb{R}^k} \circ \gamma_2(s))$ would be unbounded in s , which contradicts that $d^2(\gamma_1(s), \gamma_2(s))$ is constant in s . Therefore

$$(4.6) \quad d^2(\pi_X \circ \gamma_1(s), \pi_X \circ \gamma_2(t)) = d^2(\gamma_1(s), \gamma_2(t)) - d^2(\pi_{\mathbb{R}^k} \circ \gamma_1(s), \pi_{\mathbb{R}^k} \circ \gamma_2(t))$$

is a quadratic function of $(s - t)$. By (2) we conclude that $\pi_X \circ \gamma_1, \pi_X \circ \gamma_2$ are parallel.

(4). Let $\gamma : \mathbb{R} \rightarrow X$ be a unit speed parametrization of some line in \mathcal{L} , and let $b : X \rightarrow \mathbb{R}$ be the Busemann function $\lim_{t \rightarrow \infty} d(\gamma(t), \cdot) - t$. By assumption, the elements of \mathcal{L} partition X into the cosets of an equivalence relation; the quotient Y inherits a natural metric, namely the Hausdorff distance. The map $(b, \pi_Y) : X \rightarrow \mathbb{R} \times Y$ defines a 1-splitting – one verifies that it is an isometry using the fact that \mathcal{L} consists of parallel lines. \square

Proof of Lemma 4.4

(1). Consider the collection of lines $\mathcal{L}_\beta = \{\beta^{-1}(\mathbb{R} \times \{z\}) \mid z \in Z\}$. By Lemma 4.5 (3) it follows that $\pi_Y \circ \alpha(\mathcal{L}_\beta)$ consists of parallel lines, or consists entirely of points.

Case 1. $\pi_Y \circ \alpha(\mathcal{L}_\beta)$ consists of points. — In this case, Sublemma 4.5 implies that $\pi_{\mathbb{R}^k} \circ \alpha(\mathcal{L}_\beta)$ is a family of parallel lines in \mathbb{R}^k . Decomposing \mathbb{R}^k into a product $\mathbb{R}^k \simeq \mathbb{R}^1 \times \mathbb{R}^{k-1}$ in the direction defined by $\pi_{\mathbb{R}^k} \circ \alpha(\mathcal{L}_\beta)$, we obtain a 1-splitting of X which is easily seen to be equivalent to β .

Case 2. $\pi_Y \circ \alpha(\mathcal{L}_\beta)$ consists of parallel lines. — Since $\bigcup_{\gamma \in \pi_Y \circ \alpha(\mathcal{L}_\beta)} \gamma = Y$, by Lemma 4.5 (4), there is a 1-splitting $\gamma : Y \rightarrow \mathbb{R} \times W$ such that $\{\gamma^{-1}(\mathbb{R} \times \{w\}) \mid w \in W\} = \pi_Y \circ \alpha(\mathcal{L}_\beta)$. Letting $\alpha' : X \rightarrow (\mathbb{R}^k \times \mathbb{R}) \times W$ be the $(k + 1)$ -splitting given by $X \xrightarrow{\alpha} \mathbb{R}^k \times Y \rightarrow (\mathbb{R}^k \times \mathbb{R}) \times W$, we get that $\pi_W \circ \alpha'(\mathcal{L}_\beta)$ consists of points, so by Case 1 it follows that β is compatible with α' .

(2). Suppose that $\alpha : X \rightarrow \mathbb{R}^k \times Y$ and $\beta : X \rightarrow \mathbb{R}^l \times Z$ are splittings. We may apply part (1) to α and the 1-splitting obtained from the i^{th} coordinate direction of β , for successive values of $i \in \{1, \dots, l\}$. This will enlarge α to a splitting α' which is compatible with all of these 1-splittings, and clearly α' is then compatible with β . \square

4.2. Approximate splittings. — Next, we consider approximate splittings.

Definition 4.7. — Given $k \in \mathbb{Z}^{\geq 0}$ and $\delta \in [0, \infty)$, a (k, δ) -*splitting* of a pointed metric space (X, \star_X) is a δ -Gromov-Hausdorff approximation $(X, \star_X) \rightarrow (X_1, \star_{X_1}) \times (X_2, \star_{X_2})$, where (X_1, \star_{X_1}) is isometric to $(\mathbb{R}^k, \star_{\mathbb{R}^k})$. (We allow \mathbb{R}^k to have other basepoints than 0.)

There are “approximate” versions of equivalence and compatibility of splittings.

Definition 4.8. — Suppose that $\alpha : (X, \star_X) \rightarrow (X_1, \star_{X_1}) \times (X_2, \star_{X_2})$ is a (j, δ_1) -splitting and $\beta : (X, \star_X) \rightarrow (Y_1, \star_{Y_1}) \times (Y_2, \star_{Y_2})$ is an (k, δ_2) -splitting. Then

(1) α is ϵ -close to β if $j = k$ and there are ϵ -Gromov-Hausdorff approximations $\phi_i : (X_i, \star_{X_i}) \rightarrow (Y_i, \star_{Y_i})$ such that the composition $(\phi_1, \phi_2) \circ \alpha$ is ϵ -close to β , *i.e.*, agrees with β on $B(\star_X, \epsilon^{-1})$ up to error at most ϵ .

(2) α is ϵ -compatible with β if $j \leq k$ and there is a j -splitting $\gamma : (Y_1, \star_{Y_1}) \rightarrow (\mathbb{R}^j, \star_{\mathbb{R}^j}) \times (\mathbb{R}^{k-j}, \star_{\mathbb{R}^{k-j}})$ such that the (j, δ_2) -splitting defined by the composition

$$(4.9) \quad X \xrightarrow{\beta} Y_1 \times Y_2 \xrightarrow{(\gamma, \text{Id})} (\mathbb{R}^j \times \mathbb{R}^{k-j}) \times Y_2 \cong \mathbb{R}^j \times (\mathbb{R}^{k-j} \times Y_2)$$

is ϵ -close to α .

Lemma 4.10. — Given $\delta > 0$ and $C < \infty$, there is a $\delta' = \delta'(\delta, C) > 0$ with the following property. Suppose that (X, \star_X) is a complete pointed metric space with a (k, δ') -splitting α . Then for any $x \in B(\star_X, C)$, the pointed space (X, x) has a (k, δ) -splitting coming from a change of basepoint of α .

Proof. — In general, suppose that $f : (X, \star_X) \rightarrow (Y, \star_Y)$ is a δ' -Gromov-Hausdorff approximation. Given $x \in B(\star_X, C)$, consider x to be a new basepoint. Note that

$$(4.11) \quad d(\star_Y, f(x)) \leq d(\star_X, x) + \delta' \leq C + \delta'.$$

Suppose that δ satisfies

- (1) $\delta^{-1} \leq (\delta')^{-1} - C$,
- (2) $\delta^{-1} - \delta \leq (\delta')^{-1} - 2\delta' - C$ and
- (3) $\delta > 2\delta'$.

We claim that f is a δ -Gromov-Hausdorff between (X, x) and $(Y, f(x))$. To see this, first if $x_1, x_2 \in B(x, \delta^{-1})$ then $x_1, x_2 \in B(\star_X, (\delta')^{-1})$ and so

$$(4.12) \quad |d_Y(f(x_1), f(x_2)) - d_X(x_1, x_2)| \leq \delta' \leq \delta.$$

Next, if $y \in B(f(x), \delta^{-1} - \delta)$ then $y \in B(\star_Y, (\delta')^{-1} - \delta')$ and so there is some $\hat{x} \in B(\star_X, (\delta')^{-1})$ with $d(y, f(\hat{x})) \leq \delta'$. Now

$$(4.13) \quad d(x, \hat{x}) \leq d(f(x), f(\hat{x})) + \delta' \leq d(f(x), y) + d(y, f(\hat{x})) + \delta' \leq \delta^{-1} - \delta + 2\delta' < \delta^{-1},$$

which proves the claim.

The lemma now follows provided that we specialize to the case when Y splits off an \mathbb{R}^k -factor. \square

Lemma 4.14. — *Given $\delta > 0$ and $C < \infty$, there is a $\delta' = \delta'(\delta, C) > 0$ with the following property. Suppose that (X, \star_X) is a complete pointed metric space with a (k_1, δ') -splitting α_1 and a (k_2, δ') -splitting α_2 . Suppose that α_1 is δ' -compatible with α_2 . Given $x \in B(\star_X, C)$, let α'_1, α'_2 be the approximate splittings of (X, x) coming from a change of basepoint in α_1, α_2 . Then α'_1 and α'_2 are δ -compatible.*

Proof. — The proof is similar to that of Lemma 4.10. We omit the details. \square

4.3. Approximate splittings of Alexandrov spaces. — Recall the notion of a point in an Alexandrov space having a k -strainer of a certain size and quality; see Subsection 3.3. The next lemma shows that the notions of having a good strainer and having a good approximate \mathbb{R}^k -splitting are essentially equivalent for Alexandrov spaces.

Lemma 4.15

- (1) *Given $k \in \mathbb{Z}^+$ and $\delta > 0$, there is a $\delta' = \delta'(k, \delta) > 0$ with the following property. Suppose that (X, \star_X) is a complete pointed nonnegatively curved Alexandrov space with a (k, δ') -splitting. Then \star_X has a k -strainer of quality δ at a scale $\frac{1}{\delta}$.*
- (2) *Given $n \in \mathbb{Z}^+$ and $\delta > 0$, there is a $\delta' = \delta'(n, \delta) > 0$ with the following property. Suppose that (X, \star_X) is an complete pointed length space so that $B(\star_X, \frac{1}{\delta'})$ has curvature bounded below by $-\delta'$ and dimension bounded above by n . Suppose that for some $k \leq n$, \star_X has a k -strainer $\{p_i^\pm\}_{i=1}^k$ of quality δ' at a scale $\frac{1}{\delta'}$. Then (X, \star_X) has a (k, δ) -splitting $\phi : (X, \star_X) \rightarrow (\mathbb{R}^k \times X', (0, \star_{X'}))$ where the composition $\pi_{\mathbb{R}^k} \circ \phi$ has j^{th} component $d_X(p_j^+, \star_X) - d_X(p_j^+, \cdot)$.*

Proof. — The proof of (1) is immediate from the definitions.

Suppose that (2) were false. Then for each $i \in \mathbb{Z}^+$, there is an complete pointed length space (X_i, \star_{X_i}) so that

- (1) $B(\star_{X_i}, i)$ has dimension at most n ,
- (2) $B(\star_{X_i}, i)$ has curvature bounded below by $-\frac{1}{i}$ and
- (3) \star_{X_i} has a k -strainer of quality $\frac{1}{i}$ at a scale i but
- (4) If $\Phi_i : (X_i, \star_{X_i}) \rightarrow \mathbb{R}^k$ has j^{th} component defined as above, then Φ_i is not the \mathbb{R}^k part of a (k, δ) -splitting for any i .

After passing to a subsequence, we can assume that $\lim_{i \rightarrow \infty} (X_i, \star_{X_i}) = (X_\infty, \star_{X_\infty})$, for some pointed nonnegatively curved Alexandrov space $(X_\infty, \star_{X_\infty})$ of dimension at most n , the k -strainers yield k pairs $\{\gamma_j^\pm\}_{j=1}^k$ of opposite rays leaving \star_{X_∞} , the opposite rays γ_j^\pm fit together to form k orthogonal lines, and the j^{th} components of the Φ_i 's converge to the negative of the Busemann function of γ_j^+ . Using the Splitting Theorem [3, Chapter 10.5] it follows that X_∞ splits off an \mathbb{R}^k -factor. This gives a contradiction. \square

Lemma 4.16. — *Given $k \leq n \in \mathbb{Z}^+$, suppose that $\{(X_i, \star_{X_i})\}_{i=1}^\infty$ is a sequence of complete pointed length spaces and $\delta_i \rightarrow 0$ is a positive sequence such that*

- (1) *Each $B(\star_{X_i}, \frac{1}{\delta_i})$ has curvature bounded below by $-\delta_i$ and dimension bounded above by n .*
- (2) *Each (X_i, \star_{X_i}) has a (k, δ_i) -splitting.*
- (3) *$\lim_{i \rightarrow \infty} (X_i, \star_{X_i}) = (X_\infty, \star_{X_\infty})$ in the pointed Gromov-Hausdorff topology.*

Then $(X_\infty, \star_{X_\infty})$ is a nonnegatively curved Alexandrov space with a k -splitting.

Proof. — This follows from Lemma 4.15. \square

4.4. Compatibility of approximate splittings. — Next, we show that the nonexistence of an approximate $(k+1)$ -splitting implies that approximate j -splittings are approximately compatible with k -splittings for $j \leq k$.

Lemma 4.17. — *Given $j \leq k \leq n \in \mathbb{Z}^+$ and $\beta'_k, \beta_{k+1} > 0$, there are numbers $\delta = \delta(j, k, n, \beta'_k, \beta_{k+1}) > 0$. $\beta_j = \beta_j(j, k, n, \beta'_k, \beta_{k+1}) > 0$ and $\beta_k = \beta_k(j, k, n, \beta'_k, \beta_{k+1}) > 0$ with the following property. If (X, \star_X) is a complete pointed length space such that*

- (1) *The ball $B(\star_X, \delta^{-1})$ has curvature bounded below by $-\delta$ and dimension bounded above by n , and*
- (2) *(X, \star_X) does not admit a $(k+1, \beta_{k+1})$ -splitting*

then any (j, β_j) -splitting of (X, \star_X) is β'_k -compatible with any (k, β_k) -splitting.

Proof. — Suppose that the lemma is false. Then for some $j \leq k \leq n \in \mathbb{Z}^+$ and $\beta'_k, \beta_{k+1} > 0$, there are

- (1) A sequence $\{(X_i, \star_{X_i})\}_{i=1}^\infty$ of pointed complete length spaces,
- (2) A sequence $\{\alpha_i : (X_i, \star_{X_i}) \rightarrow (X_{1,i}, \star_{X_{1,i}}) \times (X_{2,i}, \star_{X_{2,i}})\}$ of (j, i^{-1}) -splittings and
- (3) A sequence $\{\bar{\alpha}_i : (X_i, \star_{X_i}) \rightarrow (Y_{1,i}, \star_{Y_{1,i}}) \times (Y_{2,i}, \star_{Y_{2,i}})\}$ of (k, i^{-1}) -splittings

such that

- (4) $B(\star_{X_i}, i^{-1})$ has curvature bounded below by $-i^{-1}$,
- (5) $B(\star_{X_i}, i^{-1})$ has dimension at most n ,
- (6) (X_i, \star_{X_i}) does not admit a $(k+1, \beta_{k+1})$ -splitting for any i and
- (7) α_i is not β'_k -compatible with $\bar{\alpha}_i$ for any i .

By (4), (5) and Lemma 3.10, after passing to a subsequence we can assume that there is a pointed nonnegatively curved Alexandrov space $(X_\infty, \star_{X_\infty})$, and a sequence $\{\Phi_i : (X_i, \star_{X_i}) \rightarrow (X_\infty, \star_{X_\infty})\}$ of i^{-1} -Gromov-Hausdorff approximations. In view of (2), (3) and Lemma 4.16, after passing to a further subsequence we can also assume that there is a pointed j -splitting $\alpha_\infty : (X_\infty, \star_{X_\infty}) \rightarrow (X_{\infty,1}, \star_{X_{\infty,1}}) \times (X_{\infty,2}, \star_{X_{\infty,2}})$ and a pointed k -splitting $\bar{\alpha}_\infty : (X_\infty, \star_{X_\infty}) \rightarrow (Y_{\infty,1}, \star_{Y_{\infty,1}}) \times (Y_{\infty,2}, \star_{Y_{\infty,2}})$ such that $\alpha_\infty \circ \Phi_i$ (respectively $\bar{\alpha}_\infty \circ \Phi_i$) is i^{-1} -close to α_i (respectively $\bar{\alpha}_i$). By Lemma 4.4 and the fact that $(X_\infty, \star_{X_\infty})$ does not admit a $(k+1)$ -splitting, we conclude that α_∞ is compatible with $\bar{\alpha}_\infty$. It follows that α_i is β'_k -compatible with $\bar{\alpha}_i$ for large i , which is a contradiction. \square

Remark 4.18. — Assumption (1) in Lemma 4.17 is probably not necessary but it allows us to give a simple proof, and it will be satisfied when we apply the lemma.

4.5. Overlapping cones. — In this subsection we prove a result about overlapping almost-conical regions that we will need later. We recall that a pointed metric space (X, \star) is a *metric cone* if it is a union of rays leaving the basepoint \star , and the union of any two rays γ_1, γ_2 leaving \star is isometric to the union of two rays $\bar{\gamma}_1, \bar{\gamma}_2 \subset \mathbb{R}^2$ leaving the origin $o \in \mathbb{R}^2$.

Lemma 4.19. — *If (X, \star_X) is a conical nonnegatively curved Alexandrov space and there is some $x \neq \star_X$ so that (X, x) is also a conical Alexandrov space then X has a 1-splitting such that the segment from \star_X to x is parallel to the \mathbb{R} -factor.*

Proof. — Let α be a segment joining \star_X to x . Since X is conical with respect to both \star_X and x , the segment α can be extended in both directions as a geodesic γ . The cone structure implies that γ is a line. The lemma now follows from the splitting theorem. \square

Lemma 4.20. — *Given $n \in \mathbb{Z}^+$ and $\delta > 0$, there is a $\delta' = \delta'(n, \delta) > 0$ with the following property. If*

- (X, \star_X) is a complete pointed length space,
- $x \in X$ has $d(\star_X, x) = 1$ and
- (X, \star_X) and (X, x) have pointed Gromov-Hausdorff distance less than δ' from conical nonnegatively curved Alexandrov spaces CY and CY' , respectively, of dimension at most n

then (X, x) has a $(1, \delta)$ -splitting.

Proof. — If the lemma were false, then there would be a $\delta > 0$, a positive sequence $\delta'_i \rightarrow 0$, a sequence of pointed complete length spaces $\{(X_i, \star_{X_i})\}_{i=1}^\infty$, and points $x_i \in X_i$ such that for every i :

- $d(\star_{X_i}, x_i) = 1$.

- (X_i, \star_{X_i}) and (X_i, x_i) have pointed Gromov-Hausdorff distance less than δ'_i from conical nonnegatively curved Alexandrov spaces CY_i and CY'_i , respectively, of dimension at most n .
- (X_i, x_i) does not have a $(1, \delta)$ -splitting.

After passing to a subsequence, we can assume that we have Gromov-Hausdorff limits $\lim_{i \rightarrow \infty} (X_i, \star_{X_i}) = (X_\infty, \star_{X_\infty})$, $\lim_{i \rightarrow \infty} x_i = x_\infty$ with $d(\star_{X_\infty}, x_\infty) = 1$, and both $(X_\infty, \star_{X_\infty})$ and (X_∞, x_∞) are conical nonnegatively curved Alexandrov spaces. By Lemma 4.19, (X_∞, x_∞) has a 1-splitting. This gives a contradiction. \square

4.6. Adapted coordinates. — In this subsection we discuss coordinate systems which arise from (k, δ) -splittings of Riemannian manifolds, in the presence of a lower curvature bound. The basic construction combines the standard construction of strainer coordinates [4] with the smoothing result of Corollary 3.16.

Definition 4.21. — Suppose $0 < \delta' \leq \delta$, and let α be a (k, δ') -splitting of a complete pointed Riemannian manifold (M, \star_M) . Let $\Phi : B(\star_M, \frac{1}{\delta}) \rightarrow \mathbb{R}^k$ be the composition $B(\star_M, \frac{1}{\delta}) \xrightarrow{\alpha} \mathbb{R}^k \times X_2 \rightarrow \mathbb{R}^k$. Then a map $\phi : (B(\star_M, 1), \star_M) \rightarrow (\mathbb{R}^k, \phi(\star_M))$ defines α -adapted coordinates of quality δ if

- (1) ϕ is smooth and $(1 + \delta)$ -Lipschitz.
- (2) The image of ϕ has Hausdorff distance at most δ from $B(\phi(\star_M), 1) \subset \mathbb{R}^k$.
- (3) For all $m \in B(\star_M, 1)$ and $m' \in B(\star_M, \frac{1}{\delta})$ with $d(m, m') > 1$, the (unit-length) initial velocity vector $v \in T_m M$ of any minimizing geodesic from m to m' satisfies

$$(4.22) \quad \left| D\phi(v) - \frac{\Phi(m') - \Phi(m)}{d(m, m')} \right| < \delta.$$

We will say that a map $\phi : (B(\star_M, 1), \star_M) \rightarrow (\mathbb{R}^k, 0)$ defines *adapted coordinates of quality δ* if there exists a (k, δ) -splitting α such that ϕ defines α -adapted coordinates of quality δ , as above. Likewise, (M, \star_M) admits *k -dimensional adapted coordinates of quality δ* if there is a map ϕ as above which defines adapted coordinates of quality δ .

We now give a sufficient condition for an approximate splitting to have good adapted coordinates.

Lemma 4.23 (Existence of adapted coordinates). — For all $n \in \mathbb{Z}^+$ and $\delta > 0$, there is a $\delta' = \delta'(n, \delta) > 0$ with the following property. Suppose that (M, \star_M) is an n -dimensional complete pointed Riemannian manifold with sectional curvature bounded below by $-\delta'^2$ on $B(\star_M, \frac{1}{\delta'})$, which has a (k, δ') -splitting α . Then there exist α -adapted coordinates of quality δ .

Proof. — The idea of the proof is to use the approximate splitting to construct a strainer and then use the strainer to construct Lipschitz-regular coordinates, which can be smoothed using Corollary 3.16.

Fix $n \in \mathbb{Z}^+$ and $\delta > 0$. Suppose that $\delta' < \delta$ and (M, \star_M) has a (k, δ') -splitting $\alpha : (M, \star_M) \rightarrow (\mathbb{R}^k, 0) \times (Y, \star_Y)$. Let $\{e_j\}_{j=1}^k$ be an orthonormal basis of \mathbb{R}^k . Given a parameter $s \in (\frac{1}{\delta}, \frac{1}{10\delta'})$, choose $p_{j\pm} \in M$ so that $\alpha(p_{j\pm})$ lies in the $10\delta'$ -neighborhood of $(\pm se_j, \star_Y)$.

Define $\phi_0 : B(\star_M, 1) \rightarrow \mathbb{R}^k$ by

$$(4.24) \quad \phi_0(m) = (d(\star_M, p_{1+}) - d(m, p_{1+}), \dots, d(\star_M, p_{k+}) - d(m, p_{k+})).$$

We will show that if s and δ' are chosen appropriately then we can smooth the component functions of ϕ_0 using Corollary 3.16, to obtain a map $\phi : B(\star_M, 1) \rightarrow (\mathbb{R}^k, 0)$ which defines α -adapted coordinates of quality δ . (Note that α is also a (k, δ) -splitting since $\delta' < \delta$.) We first estimate the left-hand side of (4.22). Recall that if m is a point of differentiability of $d_{p_{j+}}$ and $v \in T_m M$ is the initial vector of a unit-speed minimizing geodesic $\overline{mm'}$ then $Dd_{p_{j+}}(v) = -\cos(\tilde{Z}_m(m', p_{j+}))$.

Sublemma 4.25. — *There exists $\bar{s} = \bar{s}(n, \delta) > \frac{1}{\delta}$ so that for each $s \geq \bar{s}$, there is some $\bar{\delta}' = \bar{\delta}'(n, s, \delta) < \frac{1}{10s}$ such that if $\delta' < \bar{\delta}'$ then the following holds.*

Under the hypotheses of the lemma, suppose that $m \in B(\star_M, 1)$, $m' \in B(\star_M, \frac{1}{\delta})$ and $d(m, m') \geq 1$. Let $\overline{mp_{j+}}$ and $\overline{mm'}$ be minimizing geodesics. Then

$$(4.26) \quad \left| \cos(\tilde{Z}_m(m', p_{j+})) - \frac{d(m, p_{j+}) - d(m', p_{j+})}{d(m, m')} \right| < \delta.$$

Proof. — Suppose that the sublemma is not true. Then for each $\bar{s} > \frac{1}{\delta}$, one can find some $s \geq \bar{s}$ so that there are (see Figure 1):

- A positive sequence $\delta'_i \rightarrow 0$,
- A sequence $\{(M_i, \star_{M_i})\}_{i=1}^\infty$ of n -dimensional complete pointed Riemannian manifolds with sectional curvature bounded below by $-\delta_i'^2$ on $B(\star_{M_i}, \frac{1}{\delta'_i})$,
- A (k, δ'_i) -splitting α'_i of M_i and
- Points $m_i \in B(\star_{M_i}, 1)$ and $m'_i \in B(\star_{M_i}, \frac{1}{\delta'_i})$ with $d(m_i, m'_i) \geq 1$ so that
- For each i , the inequality (4.26) fails.




FIGURE 1

Using Lemma 4.16, after passing to a subsequence, we can assume that $\lim_{i \rightarrow \infty} (M_i, \star_{M_i}) = (X_\infty, \star_{X_\infty})$ in the pointed Gromov-Hausdorff topology for some pointed nonnegatively curved Alexandrov space $(X_\infty, \star_{X_\infty})$ which is an isometric product $\mathbb{R}^k \times Y_\infty$. After passing to a further subsequence, we can assume

- $\lim_{i \rightarrow \infty} m_i = x_\infty$ for some $x_\infty \in \overline{B(\star_{X_\infty}, 1)}$,
- $\lim_{i \rightarrow \infty} m'_i = x'_\infty$ for some $x'_\infty \in B(\star_{X_\infty}, \frac{1}{\delta})$ with $d(x_\infty, x'_\infty) \geq 1$,
- The segments $\overline{m_i m'_i}$ converge to a segment $\overline{x_\infty x'_\infty}$ and
- $\lim_{i \rightarrow \infty} p_{i, j+} = p_{\infty, j+}$ for some points $p_{\infty, j+} \in X_\infty$ in $\mathbb{R}^k \times \{\star_{Y_\infty}\}$ of distance s from \star_{X_∞} .

Then

- $\lim_{i \rightarrow \infty} \tilde{Z}_{m_i}(m'_i, p_{i, j+}) = \tilde{Z}_{x_\infty}(x'_\infty, p_{\infty, j+})$,
- $\lim_{i \rightarrow \infty} d(m_i, m'_i) = d(x_\infty, x'_\infty)$ and
- $\lim_{i \rightarrow \infty} d(m'_i, p_{i, j+}) = d(x'_\infty, p_{\infty, j+})$.

Now a straightforward verification shows that since we are in the case of an exact \mathbb{R}^k -splitting, there is a function $\eta = \eta(\delta, s)$ with $\lim_{s \rightarrow \infty} \eta(\delta, s) = 0$ so that

$$(4.27) \quad \left| \cos(\tilde{Z}_{x_\infty}(x'_\infty, p_{\infty, j+})) - \frac{d(x_\infty, p_{\infty, j+}) - d(x'_\infty, p_{\infty, j+})}{d(x_\infty, x'_\infty)} \right| < \eta.$$

Taking s large enough gives a contradiction, thereby proving the sublemma. \square

Returning to the proof of the lemma, with $s \geq \bar{s}$, define ϕ_0 as in (4.24). We have shown that if δ' is sufficiently small then ϕ_0 satisfies (4.22) with δ replaced by $\frac{1}{2}\delta$. By a similar contradiction argument, one can show that if δ' is sufficiently small, as a function of n , s and δ , then ϕ_0 is a $(1 + \frac{1}{2}\delta)$ -Lipschitz map whose image is a δ -Hausdorff approximation to $B(0, 1) \subset \mathbb{R}^k$.

If it were not for problems with cutpoints which could cause ϕ_0 to be nonsmooth, then we could take $\phi = \phi_0$. In general, we claim that if s is large enough then we can apply Lemma 3.16 in order to smooth $d_{p_{j+}}$ on $B(\star_M, 1)$, and thereby construct ϕ from ϕ_0 . To see this, note that for any $\epsilon > 0$, by making s sufficiently large, we can arrange that for any $m \in B(\star_M, \frac{3}{\delta})$, the comparison angle $\tilde{Z}_m(p_{j+}, p_{j-})$ is as close to π as we wish, and hence by triangle comparison the hypotheses of Corollary 3.16 will hold with $Y = p_{j+}$, $C = \overline{B(\star_M, \frac{2}{\delta})} \subset U = B(\star_M, \frac{3}{\delta})$, and $\theta = \theta(\epsilon)$ as in the statement of Corollary 3.16. \square

We now show that under certain conditions, the adapted coordinates associated to an approximate splitting are essentially unique.

Lemma 4.28 (Uniqueness of adapted coordinates). — *Given $1 \leq k \leq n \in \mathbb{Z}^+$ and $\epsilon > 0$, there is an $\epsilon' = \epsilon'(n, \epsilon) > 0$ with the following property. Suppose that*

- (1) (M, \star_M) is an n -dimensional complete pointed Riemannian manifold with sectional curvature bounded below by $-(\epsilon')^2$ on $B(\star_M, \frac{1}{\epsilon'})$.
- (2) $\alpha : (M, \star_M) \rightarrow (\mathbb{R}^k \times Z, (0, \star_Z))$ is a (k, ϵ') -splitting of (M, \star_M) .

(3) $\phi_1 : (B(\star_M, 1), \star_M) \rightarrow (\mathbb{R}^k, 0)$ is an α -adapted coordinate on $B(\star_M, 1)$ of quality ϵ' .

(4) Either

(a) $\phi_2 : (B(\star_M, 1), \star_M) \rightarrow (\mathbb{R}^k, 0)$ is an α -adapted coordinate on $B(\star_M, 1)$ of quality ϵ' , or

(b) ϕ_2 has $(1 + \epsilon')$ -Lipschitz components and the following holds :

For every $m \in B(\star_M, 1)$ and every $j \in \{1, \dots, k\}$, there is an $m'_j \in B(\star_M, (\epsilon')^{-1})$ with $d(m'_j, m) > 1$ satisfying (4.22) (with $\phi \rightsquigarrow \phi_2$), such that $(\pi_{\mathbb{R}^k} \circ \alpha)(m'_j)$ lies in the ϵ' -neighborhood of the line $(\pi_{\mathbb{R}^k} \circ \alpha)(m) + \mathbb{R}e_j$, and $(\pi_Z \circ \alpha)(m'_j)$ lies in the ϵ' -ball centered at $(\pi_Z \circ \alpha)(m)$.

Then $\|\phi_1 - \phi_2\|_{C^1} \leq \epsilon$ on $B(\star_M, 1)$.

Proof. — We first give the proof when ϕ_2 is also an α -adapted coordinate on $B(\star_M, 1)$ of quality ϵ' .

Let $\Phi : (M, \star_M) \rightarrow \mathbb{R}^k$ be the composition $(M, \star_M) \xrightarrow{\alpha} \mathbb{R}^k \times Z \rightarrow \mathbb{R}^k$.

Given $\epsilon_1 > 0$, if ϵ' is sufficiently small, then by choosing points $\{p_{i,\pm}\}_{i=1}^k$ in M with $d(\alpha(p_{i,\pm}), (\pm\epsilon_1^{-1}e_i, \star_Z)) \leq \epsilon_1$, we obtain a strainer of quality comparable to ϵ_1 at scale ϵ_1^{-1} . Given $m \in B(\star_M, 1)$, let γ_i be a unit speed minimizing geodesic from m to $p_{i,+}$, let $v_i \in T_m M$ be the initial velocity of γ_i , and let m_i be the point on γ_i with $d(m_i, m) = 2$. Given $\epsilon_2 > 0$, if ϵ_1 is sufficiently small then using triangle comparison, we get

$$(4.29) \quad |\langle v_i, v_j \rangle - \delta_{ij}| < \epsilon_2, \quad |(\Phi(m_i) - \Phi(m)) - 2e_i| < \epsilon_2.$$

for all $1 \leq i, j \leq k$. Applying (4.22) with $m' = m_i$ gives

$$(4.30) \quad \max(|D\phi_1(v_i) - e_i|, |D\phi_2(v_i) - e_i|) < \epsilon_2.$$

Finally, given $\epsilon_3 > 0$, since ϕ_1, ϕ_2 are $(1 + \epsilon')$ -Lipschitz, if ϵ' and ϵ_2 are small enough then we can assume that if v is a unit vector and $v \perp \text{span}(v_1, \dots, v_k)$ then $\max(|D\phi_1(v)|, |D\phi_2(v)|) < \epsilon_3$. So for any $\epsilon_4 > 0$, if ϵ_3 is sufficiently small then the operator norm of $D\phi_1 - D\phi_2$ is bounded above by ϵ_4 on $B(\star_M, 1)$.

Since $\phi_1(\star_M) - \phi_2(\star_M) = 0$, we can integrate the inequality $\|D\phi_1 - D\phi_2\| \leq \epsilon_4$ along minimizing curves in $B(\star_M, 1)$ to conclude that $\|\phi_1 - \phi_2\|_{C^0} \leq 2\epsilon_4$ on $B(\star_M, 1)$. Since ϵ_4 is arbitrary, the lemma follows in this case.

Now suppose ϕ_2 satisfies instead the second condition in (4). If $m \in B(\star_M, 1)$, $j \in \{1, \dots, k\}$, m'_j is as in (4), and v_j is the initial velocity of a unit speed geodesic from m to m'_j , then (4.22) implies that $(D(\phi_2)_j)(v_j)$ is close to 1 when ϵ' is small. Since the j^{th} component $(\phi_2)_j$ is $(1 + \epsilon')$ -Lipschitz, this implies $\nabla(\phi_2)_j$ is close to v_j when ϵ' is small. Applying the reasoning of the above paragraphs to ϕ_1 , we get that $\nabla(\phi_1)_j$ is also close to v_j when ϵ' is small. Hence $|D\phi_1 - D\phi_2|$ is small when ϵ' is small, and integrating within $B(\star_M, 1)$ as before, the lemma follows. \square

Finally, we show that approximate compatibility of two approximate splittings leads to an approximate compatibility of their associated adapted coordinates.

Lemma 4.31. — *Given $1 \leq j \leq k \leq n \in \mathbb{Z}^+$ and $\epsilon > 0$, there is an $\epsilon' = \epsilon'(n, \epsilon) > 0$ with the following property. Suppose that*

- (1) (M, \star_M) is an n -dimensional complete pointed Riemannian manifold with sectional curvature bounded below by $-(\epsilon')^2$ on $B(\star_M, \frac{1}{\epsilon'})$.
- (2) α_1 is a (j, ϵ') -splitting of (M, \star_M) and α_2 is a (k, ϵ') -splitting of (M, \star_M) .
- (3) α_1 is ϵ' -compatible with α_2 .
- (4) $\phi_1 : (M, \star_M) \rightarrow (\mathbb{R}^j, 0)$ and $\phi_2 : (M, \star_M) \rightarrow (\mathbb{R}^k, 0)$ are adapted coordinates on $B(\star_M, 1)$ of quality ϵ' , associated to α_1 and α_2 , respectively.

Then there exists a map $T : \mathbb{R}^k \rightarrow \mathbb{R}^j$, which is a composition of an isometry with an orthogonal projection, such that $\|\phi_1 - T \circ \phi_2\|_{C^1} \leq \epsilon$ on $B(\star_M, 1)$.

Proof. — Let $\alpha_1 : M \rightarrow \mathbb{R}^j \times Z_1$ and $\alpha_2 : M \rightarrow \mathbb{R}^k \times Z_2$ be the approximate splittings. Let Φ_1 be the composition $B(\star_M, (\epsilon')^{-1}) \xrightarrow{\alpha_1} \mathbb{R}^j \times Z_1 \rightarrow \mathbb{R}^j$ and Φ_2 be the composition $B(\star_M, (\epsilon')^{-1}) \xrightarrow{\alpha_2} \mathbb{R}^k \times Z_2 \rightarrow \mathbb{R}^k$.

By (3), there is a j -splitting $\gamma : \mathbb{R}^k \rightarrow \mathbb{R}^j \times \mathbb{R}^{k-j}$ and a pair of ϵ' -Gromov-Hausdorff approximations $\xi_1 : \mathbb{R}^j \rightarrow \mathbb{R}^j$, $\xi_2 : \mathbb{R}^{n-j} \times Z_2 \rightarrow Z_1$ such that the map $\widehat{\alpha}_2$ given by the composition

$$(4.32) \quad M \xrightarrow{\alpha_2} \mathbb{R}^k \times Z_2 \xrightarrow{(\gamma, \text{Id}_{Z_2})} \mathbb{R}^j \times \mathbb{R}^{k-j} \times Z_2 \xrightarrow{(\xi_1, \xi_2)} \mathbb{R}^j \times Z_1$$

agrees with the map α_1 on the ball $B(\star_M, (\epsilon')^{-1})$ up to error at most ϵ' . Since Gromov-Hausdorff approximations $(\mathbb{R}^j, 0) \rightarrow (\mathbb{R}^j, 0)$ are close to isometries, for all $\epsilon_1 > 0$, there will be a map $T : \mathbb{R}^k \rightarrow \mathbb{R}^j$ (a composition of an isometry and a projection) which agrees with the composition

$$(4.33) \quad \mathbb{R}^k \xrightarrow{\gamma} \mathbb{R}^j \times \mathbb{R}^{k-j} \xrightarrow{\xi_1} \mathbb{R}^j$$

up to error at most ϵ_1 on the ball $B(\star_M, \epsilon_1^{-1})$, provided ϵ' is sufficiently small. Thus for all $\epsilon_2 > 0$, the composition

$$(4.34) \quad M \xrightarrow{\alpha_2} \mathbb{R}^k \times Z_2 \longrightarrow \mathbb{R}^k \xrightarrow{T} \mathbb{R}^j$$

agrees with Φ_1 up to error at most ϵ_2 on $B(\star_M, \epsilon_2^{-1})$, provided ϵ_1 and ϵ' are sufficiently small. Using Definition 4.21 for the approximate splitting α_2 and applying T , it follows that for all $\epsilon_3 > 0$, if ϵ_2 is sufficiently small then we are ensured that $T \circ \phi_2$ defines α_1 -adapted coordinates on $B(\star_M, 1)$ of quality ϵ_3 . By Lemma 4.28 (using the first criterion in part (4) of Lemma 4.28), it follows that if ϵ_3 is sufficiently small then $\|\phi_1 - T \circ \phi_2\|_{C^1} < \epsilon$. \square

Remark 4.35. — In Definition 4.21 we defined adapted coordinates on the unit ball $B(\star_M, 1)$. By a rescaling, we can define adapted coordinates on a ball of any specified size, and the results of this section will remain valid.

5. Standing assumptions

We now start on the proof of Theorem 1.3 in the case of closed manifolds. The proof is by contradiction.

The next lemma states that if we can get a contradiction from a certain “Standing Assumption” then we have proven Theorem 1.3. We recall from the definition of the volume scale in Definition 1.5 that if $w \leq w'$ then $r_p(w) \geq r_p(w')$.

Lemma 5.1. — *If Theorem 1.3 is false then we can satisfy the following Standing Assumption, for an appropriate choice of A' .*

Standing Assumption 5.2. — *Let $K \geq 10$ be a fixed integer, and let $A' : (0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ be a function.*

We assume that $\{(M^\alpha, g^\alpha)\}_{\alpha=1}^\infty$ is a sequence of connected closed Riemannian 3-manifolds such that

- (1) *For all $p \in M^\alpha$, the ratio $\frac{R_p}{r_p(1/\alpha)}$ of the curvature scale at p to the $\frac{1}{\alpha}$ -volume scale at p is bounded below by α .*
- (2) *For all $p \in M^\alpha$ and $w' \in [\frac{1}{\alpha}, c_3)$, let $r_p(w')$ denote the w' -volume scale at p . Then for each integer $k \in [0, K]$ and each $C \in (0, \alpha)$, we have $|\nabla^k \text{Rm}| \leq A'(C, w') r_p(w')^{-(k+2)}$ on $B(p, Cr_p(w'))$.*
- (3) *Each M^α fails to be a graph manifold.*

Proof. — If Theorem 1.3 is false then for every $\alpha \in \mathbb{Z}^+$, there is a manifold (M^α, g^α) which satisfies the hypotheses of Theorem 1.3 with the parameter w_0 of the theorem set to $w_0^\alpha = \frac{1}{8\alpha^4}$, but M^α is not a graph manifold.

We claim first that for every $p^\alpha \in M^\alpha$, we have $r_{p^\alpha}(1/\alpha) < R_{p^\alpha}$. If not then for some $p^\alpha \in M^\alpha$, we have $r_{p^\alpha}(1/\alpha) \geq R_{p^\alpha}$. From the definition of $r_{p^\alpha}(1/\alpha)$, it follows that

$$(5.3) \quad \text{vol}(B(p^\alpha, R_{p^\alpha})) \geq \frac{1}{\alpha} R_{p^\alpha}^3 > \frac{1}{8\alpha^4} R_{p^\alpha}^3,$$

which contradicts our choice of w_0^α .

Thus $r_{p^\alpha}(1/\alpha) < R_{p^\alpha}$. Then

$$(5.4) \quad \frac{1}{\alpha} (r_{p^\alpha}(1/\alpha))^3 = \text{vol}(B(p^\alpha, r_{p^\alpha}(1/\alpha))) \leq \text{vol}(B(p^\alpha, R_{p^\alpha})) \leq \frac{1}{8\alpha^4} R_{p^\alpha}^3,$$

so $\frac{R_{p^\alpha}}{r_{p^\alpha}(1/\alpha)} \geq 2\alpha$. This shows that $\{(M^\alpha, g^\alpha)\}_{\alpha=1}^\infty$ satisfies condition (1) of Standing Assumption 5.2.

To see that condition (2) of Standing Assumption 5.2 holds, for an appropriate choice of A' , we first note that it suffices to just consider $C \in [1, \alpha)$, since a derivative

bound on a bigger ball implies a derivative bound on a smaller ball. For $\tilde{w}' \in [\frac{1}{\alpha}, c_3)$, we have

$$(5.5) \quad Cr_{p^\alpha}(\tilde{w}') \leq \alpha r_{p^\alpha}(1/\alpha) \leq R_p.$$

Now

$$(5.6) \quad \begin{aligned} \text{vol}(B(p^\alpha, Cr_{p^\alpha}(\tilde{w}'))) &\geq \text{vol}(B(p^\alpha, r_{p^\alpha}(\tilde{w}'))) = \tilde{w}'(r_{p^\alpha}(\tilde{w}'))^3 \\ &= C^{-3}\tilde{w}'(Cr_{p^\alpha}(\tilde{w}'))^3. \end{aligned}$$

Put $w' = C^{-3}\tilde{w}'$. Then

$$(5.7) \quad w_0^\alpha = \frac{1}{8\alpha^4} \leq w' < c_3.$$

Hypothesis (2) of Theorem 1.3 implies that

$$(5.8) \quad |\nabla^k \text{Rm}| \leq A(w') (Cr_{p^\alpha}(\tilde{w}'))^{-(k+2)}$$

on $B(p^\alpha, Cr_{p^\alpha}(\tilde{w}'))$. Hence condition (2) of Standing Assumption 5.2 will be satisfied, for $C \in [1, \alpha)$, if we take

$$(5.9) \quad A'(C, \tilde{w}') = \max_{0 \leq k \leq K} A(C^{-3}\tilde{w}') C^{-(k+2)}. \quad \square$$

Standing Assumption 5.2 will remain in force until Section 16, where we consider manifolds with boundary. We will eventually get a contradiction to Standing Assumption 5.2.

For the sake of notational brevity, we will usually suppress the superscript α ; thus M will refer to M^α . By convention, each of the statements made in the proof is to be interpreted as being valid provided α is sufficiently large, whether or not this qualification appears explicitly.

Remark 5.10. — The condition $K \geq 10$ in Standing Assumption 5.2 is clearly not optimal but it is good enough for the application to the geometrization conjecture.

Remark 5.11. — We note that for fixed $\hat{w} \in (0, c_3)$, conditions (1) and (2) of Standing Assumption 5.2 imply that for large α , the following holds for all $p \in M^\alpha$:

- (1) $\frac{R_p}{r_p(\hat{w})} \geq \alpha$ and
- (2) For each integer $k \in [0, K]$ and each $C \in (0, \alpha)$, we have $|\nabla^k \text{Rm}| \leq A'(C, \hat{w}) r_p(\hat{w})^{-(k+2)}$ on $B(p, Cr_p(\hat{w}))$.

Since in addition $\text{vol}(B(p, r_p(\hat{w}))) = \hat{w}(r_p(\hat{w}))^3$, we have all of the ingredients to extract convergent subsequences, at the \hat{w} -volume scale, with smooth pointed limits that are nonnegatively curved. This is how the hypotheses of Standing Assumption 5.2 will enter into finding a contradiction. In effect, \hat{w} will eventually become a judiciously chosen constant.

6. The scale function τ

In this section we introduce a smooth scale function $\tau : M \rightarrow (0, \infty)$ which will be used throughout the rest of the paper. This function is like a volume scale in the sense that one has lower bounds on volume at scale τ , which enables one to appeal to C^K -precompactness arguments. The advantage of τ over a volume scale is that τ can be arranged to have small Lipschitz constant, which is technically useful.

We will use the following lemma to construct slowly varying functions subject to *a priori* upper and lower bounds.

Lemma 6.1. — *Suppose X is a metric space, $C \in (0, \infty)$, and $l, u : X \rightarrow (0, \infty)$ are functions. Then there is a C -Lipschitz function $r : X \rightarrow (0, \infty)$ satisfying $l \leq r \leq u$ if and only if*

$$(6.2) \quad l(p) - Cd(p, q) \leq u(q)$$

for all $p, q \in X$.

Proof. — Clearly if such an r exists then (6.2) must hold.

Conversely, suppose that (6.2) holds and define $r : X \rightarrow (0, \infty)$ by

$$(6.3) \quad r(q) = \sup\{l(p) - Cd(p, q) \mid p \in X\}.$$

Then $l \leq r \leq u$. For $q, q' \in X$, since $l(p) - Cd(p, q) \geq l(p) - Cd(p, q') - Cd(q, q')$, we obtain $r(q) \geq r(q') - Cd(q, q')$, from which it follows that r is C -Lipschitz. \square

Recall that c_3 is the volume of the unit ball in \mathbb{R}^3 . Let $\Lambda > 0$ and $\bar{w} \in (0, c_3)$ be new parameters, and put

$$(6.4) \quad w' = \frac{\bar{w}}{2(1 + 2\Lambda^{-1})^3}.$$

Recall the notion of the volume scale $r_p(\bar{w})$ from Definition 1.5.

Corollary 6.5. — *There is a smooth Λ -Lipschitz function $\tau : M \rightarrow (0, \infty)$ such that for every $p \in M$, we have*

$$(6.6) \quad \frac{1}{2} r_p(\bar{w}) \leq \tau(p) \leq 2r_p(w').$$

Proof. — We let $l : M \rightarrow (0, \infty)$ be the \bar{w} -volume scale, and $u : M \rightarrow (0, \infty)$ be the w' -volume scale. We first verify (6.2) with parameter $C = \frac{\Lambda}{2}$. To argue by contradiction, suppose that for some $p, q \in M$ we have $l(p) - \frac{1}{2}\Lambda d(p, q) > u(q)$. In particular, $d(p, q) < \frac{2l(p)}{\Lambda}$ and $u(q) < l(p)$. There are inclusions $B(p, l(p)) \subset B(q, (1 + 2\Lambda^{-1})l(p)) \subset B(p, (1 + 4\Lambda^{-1})l(p))$. Then

$$(6.7) \quad \text{vol}(B(q, (1 + 2\Lambda^{-1})l(p))) \geq \text{vol}(B(p, l(p))) = \bar{w}l^3(p) = 2w'((1 + 2\Lambda^{-1})l(p))^3.$$

For any $c > 0$, if α is sufficiently large then the sectional curvature on $B(p, (1 + 4\Lambda^{-1})l(p))$, and hence on $B(q, (1 + 2\Lambda^{-1})l(p))$, is bounded below by $-c^2 l(p)^{-2}$. As $u(q) < l(p) < (1 + 2\Lambda^{-1})l(p)$, the Bishop-Gromov inequality implies that

$$(6.8) \quad \frac{w' u(q)^3}{\int_0^{u(q)/l(p)} \sinh^2(cr) dr} = \frac{\text{vol}(B(q, u(q)))}{\int_0^{u(q)/l(p)} \sinh^2(cr) dr} \geq \frac{\text{vol}(B(q, (1 + 2\Lambda^{-1})l(p)))}{\int_0^{1+2\Lambda^{-1}} \sinh^2(cr) dr} \\ \geq \frac{2w' ((1 + 2\Lambda^{-1})l(p))^3}{\int_0^{1+2\Lambda^{-1}} \sinh^2(cr) dr}$$

or

$$(6.9) \quad \frac{c^2 \left(\frac{u(q)}{l(p)}\right)^3}{\int_0^{u(q)/l(p)} \sinh^2(cr) dr} \geq \frac{2c^2(1 + 2\Lambda^{-1})^3}{\int_0^{1+2\Lambda^{-1}} \sinh^2(cr) dr}.$$

As the function $x \mapsto \frac{x^2}{3} \frac{x^3}{\int_0^x \sinh^2(cr) dr}$ tends uniformly to 1 as $c \rightarrow 0$, for $x \in (0, 3]$, taking c small gives a contradiction. (Note the factor of 2 on the right-hand side of (6.9).)

By Lemma 6.1, there is a $\frac{\Lambda}{2}$ -Lipschitz function r on M satisfying $l \leq r \leq u$. The corollary now follows from Corollary 3.15. \square

We will write \mathfrak{r}_p for $\mathfrak{r}(p)$. Recall our convention that the index α in the sequence $\{M^\alpha\}_{\alpha=1}^\infty$ has been suppressed, and that all statements are to be interpreted as being valid provided α is sufficiently large. The next lemma shows C^K -precompactness at scale \mathfrak{r} .

Lemma 6.10

- (1) *There is a constant $\widehat{w} = \widehat{w}(w') > 0$ such that $\text{vol}(B(p, \mathfrak{r}_p)) \geq \widehat{w}(\mathfrak{r}_p)^3$ for every $p \in M$.*
- (2) *For every $p \in M$, $C < \infty$ and $k \in [0, K]$, we have*

$$(6.11) \quad |\nabla^k \text{Rm}| \leq 2^{k+2} A'(C, w') \mathfrak{r}_p^{-(k+2)} \quad \text{on the ball } B\left(p, \frac{1}{2} C \mathfrak{r}_p\right).$$

- (3) *Given $\epsilon > 0$, for sufficiently large α and for every $p \in M^\alpha$, the rescaled pointed manifold $(\frac{1}{\mathfrak{r}_p} M^\alpha, p)$ is ϵ -close in the pointed C^K -topology to a complete non-negatively curved C^K -smooth Riemannian 3-manifold. Moreover, this manifold belongs to a family which is compact in the pointed C^K -topology.*

Proof

(1). As $\frac{1}{2} \mathfrak{r}_{p^\alpha} \leq r_{p^\alpha}(w')$, the Bishop-Gromov inequality gives

$$(6.12) \quad \frac{\text{vol}(B(p^\alpha, \frac{1}{2} \mathfrak{r}_{p^\alpha}))}{\int_0^{\frac{\mathfrak{r}_{p^\alpha}}{2r_{p^\alpha}(w')}} \sinh^2(r) dr} \geq \frac{\text{vol}(B(p^\alpha, r_{p^\alpha}(w')))}{\int_0^1 \sinh^2(r) dr} = \frac{w'(r_{p^\alpha}(w'))^3}{\int_0^1 \sinh^2(r) dr},$$

or

$$(6.13) \quad \frac{\text{vol}(B(p^\alpha, \frac{1}{2} \mathfrak{r}_{p^\alpha}))}{(\frac{1}{2} \mathfrak{r}_{p^\alpha})^3} \geq \frac{w'}{\int_0^1 \sinh^2(r) dr} \frac{\int_0^{\frac{\mathfrak{r}_{p^\alpha}}{2r_{p^\alpha}(w')}} \sinh^2(r) dr}{\left(\frac{\mathfrak{r}_{p^\alpha}}{2r_{p^\alpha}(w')}\right)^3} \geq \frac{w'}{3 \int_0^1 \sinh^2(r) dr}.$$

Thus

$$(6.14) \quad \text{vol}(B(p^\alpha, \mathfrak{r}_{p^\alpha})) \geq \text{vol}(B(p^\alpha, \mathfrak{r}_{p^\alpha}/2)) \geq \frac{w'}{24 \int_0^1 \sinh^2(r) dr} (\mathfrak{r}_{p^\alpha})^3,$$

which gives (1).

(2). From hypothesis (2) of Standing Assumption 5.2, for each $C < \alpha$ and $k \in [0, K]$ we have

$$(6.15) \quad |\nabla^k \text{Rm}| \leq A'(C, w') r_{p^\alpha}(w')^{-(k+2)} \leq 2^{k+2} A'(C, w') \mathfrak{r}_{p^\alpha}^{-(k+2)}$$

on $B(p^\alpha, Cr_{p^\alpha}(w')) \supset B(p^\alpha, \frac{1}{2}C\mathfrak{r}_{p^\alpha})$.

(3). If not, then for some $\epsilon > 0$, after passing to a subsequence, for every α we could find $p^\alpha \in M^\alpha$ such that $(\frac{1}{\mathfrak{r}_{p^\alpha}}M^\alpha, p^\alpha)$ has distance at least ϵ in the C^K -topology from a complete nonnegatively curved 3-manifold.

(1) and (2) imply that after passing to a subsequence, the sequence $\{(\frac{1}{\mathfrak{r}_{p^\alpha}}M^\alpha, p^\alpha)\}$ converges in the pointed C^K -topology to a complete Riemannian 3-manifold. But since the ratio $\frac{R_p}{\mathfrak{r}_p}$ tends uniformly to infinity as $\alpha \rightarrow \infty$, the limit manifold has nonnegative curvature, which is a contradiction. \square

We now extend Lemma 6.10 to provide C^K -splittings.

Lemma 6.16. — *Given $\epsilon > 0$ and $0 \leq j \leq 3$, provided $\delta < \bar{\delta}(\epsilon, w')$ the following holds. If $p \in M$, and $\phi : (\frac{1}{\mathfrak{r}_p}M, p) \rightarrow (\mathbb{R}^j \times X, (0, \star_X))$ is a (j, δ) -splitting, then ϕ is ϵ -close to a (j, ϵ) -splitting $\hat{\phi} : (\frac{1}{\mathfrak{r}_p}M, p) \rightarrow (\mathbb{R}^j \times \hat{X}, (0, \star_{\hat{X}}))$, where \hat{X} is a complete nonnegatively curved C^K -smooth Riemannian $(3 - j)$ -manifold, and $\hat{\phi}$ is ϵ -close to an isometry on the ball $B(p, \epsilon^{-1}) \subset \frac{1}{\mathfrak{r}_p}M$, in the C^{K+1} -topology.*

Proof. — Suppose not. Then for some $\epsilon > 0$, after passing to a subsequence we can assume that there are points $p^\alpha \in M^\alpha$ so that $(\frac{1}{\tau_{p^\alpha}}M, p^\alpha)$ admits a (j, α^{-1}) -splitting ϕ_j , but the conclusion of the lemma fails.

By Lemma 6.10, a subsequence of $\{(\frac{1}{\tau_{p^\alpha}}M^\alpha, p^\alpha)\}_{\alpha=1}^\infty$ converges in the pointed C^K -topology to a complete pointed nonnegatively curved C^K -smooth 3-dimensional Riemannian manifold (M^∞, p^∞) , such that ϕ_j converges to a $(j, 0)$ -splitting ϕ_∞ of (M^∞, p^∞) . This is a contradiction. \square

Remark 6.17. — If we only assume condition (1) of Assumption 5.2 then the proof of Lemma 6.16 yields the following weaker conclusion: \widehat{X} is a (nonnegatively curved) $(3 - j)$ -dimensional Alexandrov space, and $\widehat{\phi}$ is a homeomorphism on $B(p, \epsilon^{-1})$. See Section 18 for more discussion.

Let $\sigma > 0$ be a new parameter. In the next lemma, we show that if the parameter \bar{w} is small then the pointed 3-manifold $(\frac{1}{\tau_p}M, p)$ is Gromov-Hausdorff close to something of lower dimension.

Lemma 6.18. — *Under the constraint $\bar{w} < \bar{w}(\sigma, \Lambda)$, the following holds. For every $p \in M$, the pointed space $(\frac{1}{\tau_p}M, p)$ is σ -close in the pointed Gromov-Hausdorff metric to a nonnegatively curved Alexandrov space of dimension at most 2.*

Proof. — Suppose that the lemma is not true. Then for some $\sigma, \Lambda > 0$, there is a sequence $\bar{w}_i \rightarrow 0$ and for each i , a sequence $\{(M^{\alpha(i,j)}, p^{\alpha(i,j)})\}_{j=1}^\infty$ so that for each j , $(\frac{1}{\tau_{p^{\alpha(i,j)}}}M^{\alpha(i,j)}, p^{\alpha(i,j)})$ has pointed Gromov-Hausdorff distance at least σ from any nonnegatively curved Alexandrov space of dimension at most 2.

Given i , as $j \rightarrow \infty$ the curvature scale at $p^{\alpha(i,j)}$ divided by $r_{p^{\alpha(i,j)}}(w')$ goes to infinity. Hence the curvature scale at $p^{\alpha(i,j)}$ divided by $\tau_{p^{\alpha(i,j)}}$ also goes to infinity. Thus we can find some $j = j(i)$ so that the curvature scale at $p^{\alpha(i,j(i))}$ is at least $i \tau_{p^{\alpha(i,j(i))}}$. We relabel $M^{\alpha(i,j(i))}$ as M^i and $p^{\alpha(i,j(i))}$ as p^i . Thus we have a sequence $\{(M^i, p^i)\}_{i=1}^\infty$ so that for each i , $(\frac{1}{\tau_{p^i}}M^i, p^i)$ has pointed Gromov-Hausdorff distance at least σ from any nonnegatively curved Alexandrov space of dimension at most 2, and the curvature scale at p^i is at least $i \tau_{p^i}$. In particular, a subsequence of the $(\frac{1}{\tau_{p^i}}M^i, p^i)$'s converges in the pointed Gromov-Hausdorff topology to a nonnegatively curved Alexandrov space (X, x) , necessarily of dimension 3. Hence there is a uniform positive lower bound on $\frac{\text{vol}(B(p^i, 2\tau_{p^i}))}{(2\tau_{p^i})^3}$.

As $r_{p^i}(\bar{w}_i) \leq 2\tau_{p^i}$, the Bishop-Gromov inequality implies that

$$(6.19) \quad \frac{\bar{w}_i (r_{p^i}(\bar{w}_i))^3}{\int_0^{\frac{r_{p^i}(\bar{w}_i)}{2\tau_{p^i}}} \sinh^2(r) dr} = \frac{\text{vol}(B(p^i, r_{p^i}(\bar{w}_i)))}{\int_0^{\frac{r_{p^i}(\bar{w}_i)}{2\tau_{p^i}}} \sinh^2(r) dr} \geq \frac{\text{vol}(B(p^i, 2\tau_{p^i}))}{\int_0^1 \sinh^2(r) dr}.$$

That is,

$$(6.20) \quad \frac{\text{vol}(B(p^i, 2\tau_{p^i}))}{(2\tau_{p^i})^3} \leq \bar{w}_i \left(\int_0^1 \sinh^2(r) dr \right) \frac{\left(\frac{r_{p^i}(\bar{w}_i)}{2\tau_{p^i}} \right)^3}{\int_0^{\frac{r_{p^i}(\bar{w}_i)}{2\tau_{p^i}}} \sinh^2(r) dr} \\ \leq 3 \bar{w}_i \int_0^1 \sinh^2(r) dr.$$

Since $\bar{w}_i \rightarrow 0$, we obtain a contradiction. □

As explained in Section 2, we will assume henceforth that the constraint

$$(6.21) \quad \bar{w} < \bar{w}(\sigma, \Lambda)$$

is satisfied.

7. Stratification

In this section we define a rough stratification of a Riemannian 3-manifold, based on the maximal dimension of a Euclidean factor of an approximate splitting at a given point.

7.1. Motivation. — Recall that in a metric polyhedron P , there is a natural metrically defined filtration $P_0 \subset P_1 \subset \dots \subset P$, where $P_k \subset P$ is the set of points $p \in P$ that do not have a neighborhood that isometrically splits off a factor of \mathbb{R}^{k+1} . The associated strata $\{P_k - P_{k-1}\}$ are manifolds of dimension k . An approximate version of this kind of filtration/stratification will be used in the proof of Theorem 1.3.

For the proof of Theorem 1.3, we use a stratification of M so that if $p \in M$ lies in the k -stratum then there is a metrically defined fibration of an approximate ball centered at p , over an open subset of \mathbb{R}^k . We now give a rough description of the strata; a precise definition will be given shortly.

2-stratum. — Here $(\frac{1}{\tau_p}M, p)$ is close to splitting off an \mathbb{R}^2 -factor and, due to the collapsing assumption, it is Gromov-Hausdorff close to \mathbb{R}^2 . One gets a circle fibration over an open subset of \mathbb{R}^2 .

1-stratum. — Here $(\frac{1}{\tau_p}M, p)$ is not close to splitting off an \mathbb{R}^2 -factor, but is close to splitting off an \mathbb{R} -factor. These points fall into two types: those where $(\frac{1}{\tau_p}M, p)$ looks like a half-plane, and those where it look like the product of \mathbb{R} with a space with controlled diameter. One gets a fibration over an open subset of \mathbb{R} , whose fiber is D^2 , S^2 , or T^2 .

0-stratum. — Here $(\frac{1}{\tau_p}M, p)$ is not close to splitting off an \mathbb{R} -factor. We will show that for some radius r comparable to τ_p , $(\frac{1}{r}M, p)$ is Gromov-Hausdorff close to the Tits cone $C_T N$ of some complete nonnegatively curved 3-manifold N with at most one end, and the ball $B(p, r) \subset M$ is diffeomorphic to N . The possibilities for the

topology of N are: $S^1 \times B^2$, B^3 , $\mathbb{R}P^3 - D^3$, a twisted interval bundle over the Klein bottle, $S^1 \times S^2$, $\mathbb{R}P^3 \# \mathbb{R}P^3$, a spherical space form S^3/Γ and a isometric quotient T^3/Γ of the 3-torus.

7.2. The k -stratum points. — To define the stratification precisely, we introduce the additional parameters $0 < \beta_1 < \beta_2 < \beta_3$, and put $\beta_0 = 0$. Recall that the parameter σ has already been introduced in Section 6.

Definition 7.1. — A point $p \in M$ belongs to the k -stratum, $k \in \{0, 1, 2, 3\}$, if $(\frac{1}{\tau_p}M, p)$ admits a (k, β_k) -splitting, but does not admit a (j, β_j) -splitting for any $j > k$.

Note that every pointed space has a $(0, 0)$ -splitting, so every $p \in M$ belongs to the k -stratum for some $k \in \{0, 1, 2, 3\}$.

Lemma 7.2. — Under the constraints $\beta_3 < \bar{\beta}_3$ and $\sigma < \bar{\sigma}$ there are no 3-stratum points.

Proof. — Let $c > 0$ be the minimal distance, in the pointed Gromov-Hausdorff metric, between $(\mathbb{R}^3, 0)$ and a nonnegatively curved Alexandrov space of dimension at most 2. Taking $\bar{\beta}_3 = \bar{\sigma} = \frac{c}{4}$, the lemma follows from Lemma 6.18. \square

Let $\Delta \in (\beta_2^{-1}, \infty)$ be a new parameter.

Lemma 7.3. — Under the constraint $\Delta > \bar{\Delta}(\beta_2)$, if $p \in M$ has a 2-strainer of size $\frac{\Delta}{100}\tau_p$ and quality $\frac{1}{\Delta}$ at p , then $(\frac{1}{\tau_p}M, p)$ has a $(2, \frac{1}{2}\beta_2)$ -splitting $\frac{1}{\tau_p}M \rightarrow \mathbb{R}^2$. In particular p is in the 2-stratum.

Proof. — This follows from Lemma 4.15. \square

Definition 7.4. — A 1-stratum point $p \in M$ is in the *slim 1-stratum* if there is a $(1, \beta_1)$ -splitting $(\frac{1}{\tau_p}M, p) \rightarrow (\mathbb{R} \times X, (0, \star_X))$ where $\text{diam}(X) \leq 10^3\Delta$.

8. The local geometry of the 2-stratum

In the next few sections, we examine the local geometry near points of different type, introducing adapted coordinates and certain associated cutoff functions.

In this section we consider the 2-stratum points. Along with the adapted coordinates and cutoff functions, we discuss the local topology and a selection process to get a ball covering of the 2-stratum points.

8.1. Adapted coordinates, cutoff functions and local topology near 2-stratum points. — Let p denote a point in the 2-stratum and let $\phi_p : (\frac{1}{r_p}M, p) \rightarrow \mathbb{R}^2 \times (X, \star_X)$ be a $(2, \beta_2)$ -splitting.

Lemma 8.1. — *Under the constraints $\beta_2 < \bar{\beta}_2$ and $\sigma < \bar{\sigma}$, the factor (X, \star_X) has diameter < 1 .*

Proof. — If not then we could find a subsequence $\{M^{\alpha_j}\}$ of the M^{α} 's, and $p_j \in M^{\alpha_j}$, such that with $\beta_2 = \sigma = \frac{1}{j}$, the map $\phi_{p_j} : (\frac{1}{r_{p_j}}M^{\alpha_j}, p_j) \rightarrow (\mathbb{R}^2 \times X_j, (0, \star_{X_j}))$ violates the conclusion of the lemma. We then pass to a pointed Gromov-Hausdorff sublimit (M_∞, p_∞) , which will be a nonnegatively curved Alexandrov space of dimension at most 2, and extract a limiting 2-splitting $\phi_\infty : (M_\infty, p_\infty) \rightarrow \mathbb{R}^2 \times X_\infty$. The only possibility is that $\dim(M_\infty) = 2$, ϕ_∞ is an isometry and X_∞ is a point. This contradicts the diameter assumption. \square

Let $\varsigma_{2\text{-stratum}} > 0$ be a new parameter.

Lemma 8.2. — *Under the constraint $\beta_2 < \bar{\beta}_2(\varsigma_{2\text{-stratum}})$, there is a ϕ_p -adapted coordinate η_p of quality $\varsigma_{2\text{-stratum}}$ on $B(p, 200) \subset (\frac{1}{r_p}M, p)$*

Proof. — This follows from Lemma 4.23 (see also Remark 4.35). \square

Definition 8.3. — Let ζ_p be the smooth function on M which is the extension by zero of $\Phi_{8,9} \circ |\eta_p|$. (See Section 2 for the definition of $\Phi_{a,b}$.)

Lemma 8.4. — *Under the constraints $\beta_2 < \bar{\beta}_2$, $\varsigma_{2\text{-stratum}} < \bar{\varsigma}_{2\text{-stratum}}$ and $\sigma < \bar{\sigma}$, the restriction of η_p to $\eta_p^{-1}(B(0, 100))$ is a fibration with fiber S^1 . In particular, for all $R \in (0, 100)$, $|\eta_p|^{-1}[0, R]$ is diffeomorphic to $S^1 \times \overline{B(0, R)}$.*

Proof. — For small β_2 and $\varsigma_{2\text{-stratum}}$, the map $\eta_p : \frac{1}{r_p}M \supset B(p, 200) \rightarrow \mathbb{R}^2$ is a submersion; this follows by applying (4.22) to an appropriate 2-strainer at $x \in B(p, 200)$ furnished by the $(2, \beta_2)$ -splitting ϕ_p .

By Lemma 8.1, if $\varsigma_{2\text{-stratum}} < \bar{\varsigma}_{2\text{-stratum}}$ then $\eta_p^{-1}(B(0, 100)) \subset B(p, 102) \subset \frac{1}{r_p}M$. Therefore if $K \subset B(0, 100)$ is compact then $\eta_p^{-1}(K)$ is a closed subset of $\overline{B(p, 102)} \subset \frac{1}{r_p}M$, and hence compact. It follows that $\eta_p|_{\eta_p^{-1}(B(0, 100))} : \eta_p^{-1}(B(0, 100)) \rightarrow B(0, 100)$ is a proper map. Thus $\eta_p|_{\eta_p^{-1}(B(0, 100))}$ is a trivial fiber bundle with compact 1-dimensional fibers.

Since $\eta_p^{-1}(0)$ has diameter at most 2 by Lemma 8.1 (assuming $\varsigma_{2\text{-stratum}} < \bar{\varsigma}_{2\text{-stratum}}$), it follows that any two points in the fiber $\eta_p^{-1}(0)$ can be joined by a path in $\eta_p^{-1}(B(0, 100))$. Now the triviality of the bundle implies that the fibers are connected, *i.e.*, diffeomorphic to S^1 . \square

8.2. Selection of 2-stratum balls. — Let \mathcal{M} be a new parameter, which will become a bound on intersection multiplicity of balls. The corresponding bound $\overline{\mathcal{M}}$ will describe how big \mathcal{M} has to be taken in order for various assertions to be valid.

Let $\{p_i\}_{i \in I_{2\text{-stratum}}}$ be a maximal set of 2-stratum points with the property that the collection $\{B(p_i, \frac{1}{3}\mathfrak{r}_{p_i})\}_{i \in I_{2\text{-stratum}}}$ is disjoint. We write ζ_i for ζ_{p_i} .

Lemma 8.5. — *Under the constraints $\mathcal{M} > \overline{\mathcal{M}}$ and $\Lambda < \overline{\Lambda}$,*

- (1) $\bigcup_{i \in I_{2\text{-stratum}}} B(p_i, \mathfrak{r}_{p_i})$ contains all 2-stratum points.
- (2) The intersection multiplicity of the collection $\{\text{supp}(\zeta_i)\}_{i \in I_{2\text{-stratum}}}$ is bounded by \mathcal{M} .

Proof

(1). Assume $1 + \frac{2}{3}\Lambda < 1.01$. If p is a 2-stratum point, there is an $i \in I_{2\text{-stratum}}$ such that $B(p, \frac{1}{3}\mathfrak{r}_p) \cap B(p_i, \frac{1}{3}\mathfrak{r}_{p_i}) \neq \emptyset$ for some $i \in I_{2\text{-stratum}}$. Then $\frac{\mathfrak{r}_p}{\mathfrak{r}_{p_i}} \in [.9, 1.1]$, and $p \in B(p_i, \mathfrak{r}_{p_i})$.

(2). From the definition of ζ_i , if $\varsigma_{2\text{-stratum}}$ is sufficiently small then we are ensured that $\text{supp}(\zeta_i) \subset B(p_i, 10\mathfrak{r}_{p_i})$.

Suppose that for some $p \in M$, we have $p \in \bigcap_{j=1}^N B(p_{i_j}, 10\mathfrak{r}_{p_{i_j}})$ for distinct i_j 's. We relabel so that $B(p_{i_1}, \mathfrak{r}_{p_{i_1}})$ has the smallest volume among the $B(p_{i_j}, \mathfrak{r}_{p_{i_j}})$'s.

If 10Λ is sufficiently small then we can assume that for all j , $\frac{1}{2} \leq \frac{\mathfrak{r}_{p_{i_j}}}{\mathfrak{r}_{p_{i_1}}} \leq 2$. Hence the N disjoint balls $\{B(p_{i_j}, \frac{1}{3}\mathfrak{r}_{p_{i_j}})\}_{j=1}^N$ lie in $B(p_{i_1}, 100\mathfrak{r}_{p_{i_1}})$ and by Bishop-Gromov volume comparison

$$(8.6) \quad N \leq \frac{\text{vol}(B(p_{i_1}, 100\mathfrak{r}_{p_{i_1}}))}{\text{vol}(B(p_{i_1}, \frac{1}{3}\mathfrak{r}_{p_{i_1}}))} \leq \frac{\int_0^{100} \sinh^2(r) dr}{\int_0^{\frac{1}{3}} \sinh^2(r) dr}.$$

This proves the lemma. □

9. Edge points and associated structure

In this section we study points $p \in M$ where the pair (M, p) looks like a half-plane with a basepoint lying on the edge. Such points define an edge set E . We show that any 1-stratum point, which is not a slim 1-stratum point, is not far from E .

As a technical tool, we also introduce an approximate edge set E' consisting of points where the edge structure is of slightly lower quality than that of E . The set E' will fill in the boundary edges of the approximate half-plane regions around points in E . We construct a smoothed distance function from E' , along with an associated cutoff function.

We describe the local topology around points in E and choose a ball covering of E .

9.1. Edge points. — We begin with a general lemma about 1-stratum points.

Lemma 9.1. — *Given $\epsilon > 0$, if $\beta_1 < \bar{\beta}_1(\epsilon)$ and $\sigma < \bar{\sigma}(\epsilon)$ then the following holds. If $(\frac{1}{\tau_p}M, p)$ has a $(1, \beta_1)$ -splitting then there is a $(1, \epsilon)$ -splitting $(\frac{1}{\tau_p}M, p) \rightarrow (\mathbb{R} \times Y, (0, \star_Y))$, where Y is an Alexandrov space with $\dim(Y) \leq 1$.*

Proof. — This is similar to the proof of Lemma 8.1. If the lemma were false then we could find a sequence $\alpha_j \rightarrow \infty$ so that taking $\beta_1 = j^{-1}$ and $\sigma = j^{-1}$, for every j there would be $p_j \in M^{\alpha_j}$ and a $(1, j^{-1})$ -splitting of $(\frac{1}{\tau_{p_j}}M^{\alpha_j}, p_j)$, but no $(1, \epsilon)$ -splitting as asserted. Passing to a subsequence, we obtain a pointed Gromov-Hausdorff limit (M_∞, p_∞) , and the $(1, j^{-1})$ -splittings converge to a 1-splitting $\phi_\infty : M_\infty \rightarrow \mathbb{R} \times Y$. It follows from Lemma 6.18 that $\dim M_\infty \leq 2$, and hence $\dim Y \leq 1$. This implies that for large j , we can find arbitrarily good splittings $(\frac{1}{\tau_{p_j}}M^{\alpha_j}, p_j) \rightarrow (\mathbb{R} \times Y_j, (0, \star_{Y_j}))$ where Y_j is an Alexandrov space with $\dim(Y_j) \leq 1$. This is a contradiction. \square

Let $0 < \beta_E < \beta_{E'}$ and $0 < \sigma_E < \sigma_{E'}$ be new parameters.

Definition 9.2. — A point $p \in M$ is an (s, t) -edge point if there is a $(1, s)$ -splitting

$$(9.3) \quad F_p : \left(\frac{1}{\tau_p}M, p \right) \longrightarrow (\mathbb{R} \times Y, (0, \star_Y))$$

and a t -pointed-Gromov-Hausdorff approximation

$$(9.4) \quad G_p : (Y, \star_Y) \longrightarrow ([0, C], 0),$$

with $C \geq 200\Delta$. Given F_p and G_p , we put

$$(9.5) \quad Q_p = (\text{Id} \times G_p) \circ F_p : \left(\frac{1}{\tau_p}M, p \right) \longrightarrow (\mathbb{R} \times [0, C], (0, 0)).$$

We let E denote the set of (β_E, σ_E) -edge points, and E' denote the set of $(\beta_{E'}, \sigma_{E'})$ -edge points. Note that $E \subset E'$. We will often refer to elements of E as *edge points*.

We emphasize that in the definition above, Q_p maps the basepoint $p \in M$ to $(0, 0) \in \mathbb{R} \times [0, C]$.

Lemma 9.6. — *Under the constraints $\beta_{E'} < \bar{\beta}_{E'}$, $\sigma_{E'} < \bar{\sigma}_{E'}$ and $\beta_2 < \bar{\beta}_2$, no element $p \in E'$ can be a 2-stratum point.*

Proof. — By Lemma 8.1, if p is a 2-stratum point and $\phi_p : (\frac{1}{\tau_p}M, p) \rightarrow (\mathbb{R}^2 \times X, (0, \star_X))$ is a $(2, \beta_2)$ -splitting then $\text{diam } X < 1$. Thus if $\beta_{E'}$, $\sigma_{E'}$ and β_2 are all small then a large pointed ball in $(\mathbb{R}^2, 0)$ has pointed Gromov-Hausdorff distance less than two from a large pointed ball in $(\mathbb{R} \times [0, C], (0, 0))$, which is a contradiction. \square

We now show that in a neighborhood of $p \in E$, the set E' looks like the border of a half-plane.

Lemma 9.7. — Given $\epsilon > 0$, if $\beta_{E'} < \bar{\beta}_{E'}(\epsilon, \Delta)$, $\sigma_{E'} < \bar{\sigma}_{E'}(\epsilon, \Delta)$, $\beta_E < \bar{\beta}_E(\beta_{E'}, \sigma_{E'})$, $\sigma_E < \bar{\sigma}_E(\beta_{E'}, \sigma_{E'})$ and $\Lambda < \bar{\Lambda}(\epsilon, \Delta)$ then the following holds.

For $p \in E$, if Q_p is as in Definition 9.2 and $\widehat{Q}_p : (\mathbb{R} \times [0, C], (0, 0)) \rightarrow (\frac{1}{\mathfrak{r}_p}M, p)$ is a quasi-inverse for Q_p (see Subsection 3.1), then $\widehat{Q}_p([-100\Delta, 100\Delta] \times \{0\})$ is $\frac{\epsilon}{2}$ -Hausdorff close to $E' \cap Q_p^{-1}([-100\Delta, 100\Delta] \times [0, 100\Delta])$.

Proof. — Suppose the lemma were false. Then for some $\epsilon > 0$, there would be sequences $\alpha_i \rightarrow \infty$, $s_i \rightarrow 0$ and $\Lambda_i \rightarrow 0$ so that for each i ,

- (1) The scale function \mathfrak{r} of M^{α_i} has Lipschitz constant bounded above by Λ_i , and
- (2) There is an (s_i^2, s_i^2) -edge point $p_i \in M^{\alpha_i}$ such that $\widehat{Q}_{p_i}([-100\Delta, 100\Delta] \times \{0\})$ is not $\frac{\epsilon}{2}$ -Hausdorff close to $E'_i \cap Q_{p_i}^{-1}([-100\Delta, 100\Delta] \times [0, 100\Delta])$, where E'_i denotes the set of (s_i, s_i) -edge points in M^{α_i} .

After passing to a subsequence, we can assume that $\lim_{i \rightarrow \infty} (\frac{1}{\mathfrak{r}_{p_i}}M^{\alpha_i}, p_i) = (X^\infty, p_\infty)$ for some pointed nonnegatively curved Alexandrov space (X^∞, p_∞) . We can also assume that $\lim_{i \rightarrow \infty} \widehat{Q}_{p_i}|_{[-200\Delta, 200\Delta] \times [0, 200\Delta]}$ exists and is an isometric embedding $\widehat{Q}_\infty : [-200\Delta, 200\Delta] \times [0, 200\Delta] \rightarrow X^\infty$, with $\widehat{Q}_\infty(0, 0) = p_\infty$. Then

$$(9.8) \quad \lim_{i \rightarrow \infty} \widehat{Q}_{p_i}([-100\Delta, 100\Delta] \times \{0\}) = \widehat{Q}_\infty([-100\Delta, 100\Delta] \times \{0\}).$$

However, since $s_i \rightarrow 0$ and $\Lambda_i \rightarrow 0$, it follows that

$$(9.9) \quad \lim_{i \rightarrow \infty} E'_i \cap Q_{p_i}^{-1}([-100\Delta, 100\Delta] \times [0, 100\Delta]) = \widehat{Q}_\infty([-100\Delta, 100\Delta] \times \{0\}).$$

Hence for large i , $\widehat{Q}_{p_i}([-100\Delta, 100\Delta] \times \{0\})$ is $\frac{\epsilon}{2}$ -Hausdorff close to $E'_i \cap Q_{p_i}^{-1}([-100\Delta, 100\Delta] \times [0, 100\Delta])$. This is a contradiction. \square

The first part of the next lemma says that 1-stratum points are either slim 1-stratum points or lie not too far from an edge point. The second part says that E is coarsely dense in E' .

Lemma 9.10. — Under the constraints $\beta_{E'} < \bar{\beta}_{E'}(\Delta)$, $\sigma_{E'} < \bar{\sigma}_{E'}(\Delta)$, $\beta_E < \bar{\beta}_E(\beta_{E'}, \sigma_{E'})$, $\sigma_E < \bar{\sigma}_E(\beta_{E'}, \sigma_{E'})$, $\beta_1 < \bar{\beta}_1(\Delta, \beta_E)$, $\sigma < \bar{\sigma}(\Delta, \sigma_E)$ and $\Lambda < \bar{\Lambda}(\Delta)$, the following holds.

- (1) For every 1-stratum point p which is not in the slim 1-stratum, there is some $q \in E$ with $p \in B(q, \Delta \mathfrak{r}_q)$.
- (2) For every 1-stratum point p which is not in the slim 1-stratum and for every $p' \in E' \cap B(p, 10\Delta \mathfrak{r}_p)$, there is some $q \in E$ with $p' \in B(q, \mathfrak{r}_q)$. See Figure 2 below.

Proof. — Let $\epsilon > 0$ be a constant that will be adjusted during the proof. Let p be a 1-stratum point which is not in the slim 1-stratum.

By Lemma 9.1, if $\beta_1 < \bar{\beta}_1(\Delta, \epsilon)$ and $\sigma < \bar{\sigma}(\Delta, \epsilon)$ then there is a $(1, \epsilon)$ -splitting $F : (\frac{1}{\mathfrak{r}_p}M, p) \rightarrow (\mathbb{R} \times Y, (0, \star_Y))$, where Y is a nonnegatively curved Alexandrov space of dimension at most one.

• p




FIGURE 2

Sublemma 9.11. — $\text{diam}(Y) \geq 500\Delta$.

Proof. — Suppose that $\text{diam}(Y) < 500\Delta$. Let $\phi : (\frac{1}{r_p}M, p) \rightarrow (\mathbb{R} \times X, (0, \star_X))$ be a $(1, \beta_1)$ -splitting.

Since p belongs to the 1-stratum, $(\frac{1}{r_p}M, p)$ does not admit a $(2, \beta_2)$ -splitting. By Lemma 4.17, it follows that if $\epsilon < \bar{\epsilon}(\Delta)$ and $\beta_1 < \bar{\beta}_1(\Delta)$ then there is a $\frac{1}{10^3\Delta}$ -Gromov-Hausdorff approximation $(X, \star_X) \rightarrow (Y, \star_Y)$. Since $Y \subset B(\star_Y, 500\Delta)$, we conclude that the metric annulus $A(\star_X, 600\Delta, 900\Delta) \subset X$ is empty. But then the image of the ball $B(p, \beta_1^{-1}) \subset \frac{1}{r_p}M$ under the composition $B(p, \beta_1^{-1}) \xrightarrow{\phi} \mathbb{R} \times X \xrightarrow{\pi_X} X$ will be contained in $B(\star_X, 600\Delta) \subset X$ (because the inverse image of $B(\star_X, 600\Delta)$ under $\pi_X \circ \phi$ is open and closed in the connected set $B(p, \beta_1^{-1})$). Thus $\phi : (\frac{1}{r_p}M, p) \rightarrow (\mathbb{R} \times B(\star_X, 600\Delta), (0, \star_X))$ is a $(1, \beta_1)$ -splitting, and p is in the slim 1-stratum. This is a contradiction. \square

Proof of Lemma 9.10 continued. — Suppose first that Y is a circle. If ϵ is sufficiently small then there is a 2-strainer of size $\frac{\Delta}{100}r_p$ and quality $\frac{1}{\Delta}$ at p . By the choice of Δ (see Section 7), p is a 2-stratum point, which is a contradiction.

Hence up to isometry, Y is an interval $[0, C]$ with $C > 500\Delta$. If $\star_Y \in (\frac{\Delta}{10}, C - \frac{\Delta}{10})$ then the same argument as in the preceding paragraph shows that \star_Y is a 2-stratum point provided ϵ is sufficiently small. Hence $\star_Y \in [0, \frac{\Delta}{10}]$ or $\star_Y \in [C - \frac{\Delta}{10}, C]$. In the second case, we can flip $[0, C]$ around its midpoint to reduce to the first case. So

we can assume that $\star_Y \in [0, \frac{\Delta}{10}]$. Let \widehat{F} be a quasi-inverse of F and put $q = \widehat{F}(0, 0)$. If $\Lambda\Delta$ is sufficiently small then we can ensure that $.9 \leq \frac{\tau_p}{\tau_q} \leq 1.1$. From Lemma 4.10, if β_1 is sufficiently small, relative to β_E , then q has a $(1, \beta_E)$ -splitting. If in addition ϵ is sufficiently small, relative to σ_E , then q is guaranteed to be in E . Then $d(p, q) \leq \frac{1}{2}\Delta\tau_p < \Delta\tau_q$.

To prove the second part of the lemma, Lemma 9.7 implies that if $p' \in E' \cap B(p, 10\Delta\tau_p)$ then we can assume that p' lies within distance $\frac{1}{2}\tau_p$ from $\widehat{F}([-100\Delta, 100\Delta] \times \{0\})$. (This is not a constraint on the present parameter ϵ .) Choose $q = \widehat{F}(x, 0)$ for some $x \in [-100\Delta, 100\Delta]$ so that $d(p', q) \leq \frac{1}{2}\tau_p$. From Lemma 4.10, if β_1 is sufficiently small, relative to β_E , then q has a $(1, \beta_E)$ -splitting. If in addition ϵ is sufficiently small, relative to σ_E , then q is guaranteed to be in E . If $\Lambda\Delta$ is sufficiently small then $d(p', q) \leq \tau_q$. This proves the lemma. \square

9.2. Regularization of the distance function $d_{E'}$. — Let $d_{E'}$ be the distance function from E' . We will apply the smoothing results from Section 3.6 to $d_{E'}$. We will see that the resulting smoothing of the distance function from E' defines part of a good coordinate in a collar region near E .

Let $\varsigma_{E'} > 0$ be a new parameter.

Lemma 9.12. — *Under the constraints $\beta_{E'} < \overline{\beta}_{E'}(\Delta, \varsigma_{E'})$ and $\sigma_{E'} < \overline{\sigma}_{E'}(\Delta, \varsigma_{E'})$ there is a function $\rho_{E'} : M \rightarrow [0, \infty)$ such that if $\eta_{E'} = \frac{\rho_{E'}}{\tau}$ then:*

(1) *We have*

$$(9.13) \quad \left| \frac{\rho_{E'}}{\tau} - \frac{d_{E'}}{\tau} \right| \leq \varsigma_{E'}.$$

(2) *In the set $\eta_{E'}^{-1}[\frac{\Delta}{10}, 10\Delta] \cap (\frac{d_{E'}}{\tau})^{-1}[0, 50\Delta]$, the function $\rho_{E'}$ is smooth and its gradient lies in the $\varsigma_{E'}$ -neighborhood of the generalized gradient of $d_{E'}$.*

(3) *$\rho_{E'} - d_{E'}$ is $\varsigma_{E'}$ -Lipschitz.*

Proof. — Let $\epsilon_1 \in (0, \infty)$ and $\theta \in (0, \pi)$ be constants, to be determined during the proof.

Put

$$(9.14) \quad C = \left(\frac{d_{E'}}{\tau} \right)^{-1} \left[\frac{\Delta}{20}, 20\Delta \right] \cap \left(\frac{d_E}{\tau} \right)^{-1} [0, 50\Delta].$$

If $x \in C$ and $\Lambda < \overline{\Lambda}(\Delta)$ then there exists a $p \in E$ such that $x \in B(p, 60\Delta) \subset \frac{1}{\tau_p}M$. By Lemma 9.7, provided that $\beta_{E'} < \overline{\beta}_{E'}(\epsilon_1, \Delta)$, $\sigma_{E'} < \overline{\sigma}_{E'}(\epsilon_1, \Delta)$, $\beta_E < \overline{\beta}_E(\beta_{E'}, \sigma_{E'})$, $\sigma_E < \overline{\sigma}_E(\beta_{E'}, \sigma_{E'})$, and $\Lambda < \overline{\Lambda}(\epsilon_1, \Delta)$, there is a $y \in M$ such that in $\frac{1}{\tau_p}M$,

$$(9.15) \quad |d(y, x) - d_{E'}(x)| < \epsilon_1, \quad |d(y, E') - 2d_{E'}(x)| < \epsilon_1.$$




FIGURE 3

By triangle comparison, involving triangles whose vertices are at x , y and points in E' whose distance to x is almost infimal, it follows that if $\epsilon_1 < \bar{\tau}_1(\theta, \Delta)$ then $\text{diam}(V_x) < \theta$, where V_x is the set of initial velocities of minimizing geodesic segments from x to E' . See Figure 3.

Applying Corollary 3.16, if $\theta < \bar{\theta}(\varsigma_{E'})$ then we obtain a function $\rho_{E'} : M \rightarrow [0, \infty)$ such that

- (1) $\rho_{E'}$ is smooth in a neighborhood of C .
- (2) $\left\| \frac{\rho_{E'}}{\tau} - \frac{d_{E'}}{\tau} \right\| < \varsigma_{E'}$.
- (3) For every $x \in C$, the gradient of $\rho_{E'}$ lies in the $\varsigma_{E'}$ -neighborhood of the generalized gradient of $d_{E'}$.
- (4) $\rho_{E'} - d_{E'}$ is $\varsigma_{E'}$ -Lipschitz.

If $\varsigma_{E'} < \frac{\Delta}{20}$ then

$$(9.16) \quad \eta_{E'}^{-1} \left[\frac{\Delta}{10}, 10\Delta \right] \cap \left(\frac{d_E}{\tau} \right)^{-1} [0, 50\Delta] \subset C,$$

so the lemma follows. □

9.3. Adapted coordinates tangent to the edge. — In this subsection, $p \in E$ will denote an edge point and Q_p will denote a map as in (9.5).

Let $\varsigma_{\text{edge}} > 0$ be a new parameter. Applying Lemma 4.23, we get:

Lemma 9.17. — *Under the constraint $\beta_E < \bar{\beta}_E(\Delta, \varsigma_{\text{edge}})$, there is a Q_p -adapted coordinate*

$$(9.18) \quad \eta_p : \left(\frac{1}{\mathfrak{r}_p} M, p \right) \supset B(p, 100\Delta) \longrightarrow \mathbb{R}$$

of quality ς_{edge} .

We define a global function $\zeta_p : M \rightarrow [0, 1]$ by extending

$$(9.19) \quad (\Phi_{-9\Delta, -8\Delta, 8\Delta, 9\Delta} \circ \eta_p) \cdot (\Phi_{8\Delta, 9\Delta} \circ \eta_{E'}) : B(p, 100\Delta) \longrightarrow [0, 1]$$

by zero.

Lemma 9.20. — *The following holds:*

- (1) ζ_p is smooth.
- (2) *Under the constraints $\beta_2 < \bar{\beta}_2(\varsigma_2\text{-stratum})$, $\Lambda < \bar{\Lambda}(\varsigma_2\text{-stratum}, \Delta)$, $\beta_{E'} < \bar{\beta}_{E'}(\varsigma_2\text{-stratum}, \Delta)$, $\sigma_{E'} < \bar{\sigma}_{E'}(\varsigma_2\text{-stratum}, \Delta)$, $\beta_E < \bar{\beta}_E(\beta_2, \beta_{E'}, \sigma_{E'}, \varsigma_2\text{-stratum})$, $\sigma_E < \bar{\sigma}_E(\beta_2, \beta_{E'}, \sigma_{E'}, \varsigma_2\text{-stratum})$, $\varsigma_{E'} < \bar{\varsigma}_{E'}(\varsigma_2\text{-stratum})$ and $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\varsigma_2\text{-stratum})$, if $x \in (\eta_p, \eta_{E'})^{-1}([-10\Delta, 10\Delta] \times [\frac{1}{10}\Delta, 10\Delta])$ then x is a 2-stratum point, and there is a $(2, \beta_2)$ -splitting $\phi : (\frac{1}{\mathfrak{r}_x} M, x) \rightarrow \mathbb{R}^2$ such that $(\eta_i, \eta_{E'}) : (\frac{1}{\mathfrak{r}_x} M, x) \rightarrow (\mathbb{R}^2, \phi(x))$ defines ϕ -adapted coordinates of quality $\varsigma_2\text{-stratum}$ on the ball $B(x, 100) \subset \frac{1}{\mathfrak{r}_x} M$.*

Proof

- (1). This follows from Lemma 9.12.

(2). Let $\epsilon_1, \dots, \epsilon_5 > 0$ be constants, to be chosen at the end of the proof. For $i \in \{1, 2\}$ choose points $x_i^\pm \in M$ such that $Q_p(x_i^\pm) \in B(Q_p(x) \pm \frac{\Delta}{20} e_i, \sigma_E) \subset \mathbb{R} \times [0, C]$. Provided that $\beta_E < \bar{\beta}_E(\beta_2, \Delta, \epsilon_1)$, $\sigma_E < \bar{\sigma}_E(\beta_2, \Delta, \epsilon_1)$ and $\Lambda < \bar{\Lambda}(\Delta)$, the tuple $\{x_i^\pm\}_{i=1}^2$ will be a 2-strainer at x of quality ϵ_1 , and scale at least $\frac{\Delta}{30}$ in $\frac{1}{\mathfrak{r}_x} M$. Therefore, if $\epsilon_1 < \bar{\epsilon}_1(\beta_2)$ then x will be a 2-stratum point, with a $(2, \beta_2)$ -splitting $\phi : (\frac{1}{\mathfrak{r}_x} M, x) \rightarrow (\mathbb{R}^2, 0)$ given by strainer coordinates as in Lemma 4.15.

Suppose that y is a point in $\frac{1}{\mathfrak{r}_x} M$ with $d(y, x) < \epsilon_2 \cdot \frac{\Delta}{20}$, and $z \in E'$ is a point with $d(y, z) \leq d(y, E') + \epsilon_2 \Delta$; see Figure 4. Then by Lemma 9.7, if $\beta_{E'} < \bar{\beta}_{E'}(\epsilon_3, \Delta)$, $\sigma_{E'} < \bar{\sigma}_{E'}(\epsilon_3, \Delta)$, $\beta_E < \bar{\beta}_E(\epsilon_3, \beta_{E'}, \sigma_{E'})$, $\sigma_E < \bar{\sigma}_E(\epsilon_3, \beta_{E'}, \sigma_{E'})$, $\Lambda < \bar{\Lambda}(\epsilon_3, \Delta)$ and $\epsilon_2 < \bar{\epsilon}_2(\epsilon_3)$ then the comparison angles $\tilde{\angle}_y(x_1^\pm, z)$, $\tilde{\angle}_y(x_2^\pm, z)$ will satisfy $|\tilde{\angle}_y(x_1^\pm, z) - \frac{\pi}{2}| < \epsilon_3$ and $|\tilde{\angle}_y(x_2^\pm, z) - \pi| < \epsilon_3$. If γ_i^\pm is a minimizing segment from y to x_i^\pm , and γ_z is a minimizing segment from y to z , it follows that $|\angle_y(\gamma_1^\pm, \gamma_z) - \frac{\pi}{2}| < \epsilon_4$ and $|\angle_y(\gamma_2^\pm, \gamma_z) - \pi| < \epsilon_4$, provided that $\epsilon_i < \bar{\epsilon}_i(\epsilon_4)$ for $i \leq 3$. Therefore $|D\eta_{E'}((\gamma_1^\pm)'(0))| < \epsilon_5$ and $|D\eta_{E'}((\gamma_2^\pm)'(0)) - 1| < \epsilon_5$, provided that $\epsilon_4 < \bar{\epsilon}_4(\epsilon_5)$ and $\varsigma_{E'} < \bar{\varsigma}_{E'}(\epsilon_5)$. Likewise, $|D\eta_p((\gamma_1^\pm)'(0)) - 1| < \epsilon_5$ and




FIGURE 4

$|D\eta_p((\gamma_2^\pm)'(0))| < \epsilon_5$, provided that $\beta_E < \bar{\beta}_E(\epsilon_5)$ and $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\epsilon_5)$. It follows from Lemma 4.28 that $(\eta_p, \eta_{E'})$ defines ϕ -adapted coordinates of quality $\varsigma_{2\text{-stratum}}$ on $B(x, 100) \subset \frac{1}{\epsilon_x}M$, provided that $\epsilon_5 < \bar{\epsilon}_5(\varsigma_{2\text{-stratum}})$ and $\beta_2 < \bar{\beta}_2(\varsigma_{2\text{-stratum}})$.

We may fix the constants in the order $\epsilon_5, \dots, \epsilon_1$. The lemma follows. \square

9.4. The topology of the edge region. — In this subsection we determine the topology of a suitable neighborhood of an edge point $p \in E$.

Lemma 9.21. — *Under the constraints $\beta_{E'} < \bar{\beta}_{E'}(\Delta)$, $\sigma_{E'} < \bar{\sigma}_{E'}(\Delta)$, $\beta_E < \bar{\beta}_E(\beta_{E'}, \sigma_{E'}, w')$, $\sigma_E < \bar{\sigma}_E(\beta_{E'}, \sigma_{E'})$, $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\Delta)$, $\varsigma_{E'} < \bar{\varsigma}_{E'}(\Delta)$, $\Lambda < \bar{\Lambda}(\Delta)$ and $\sigma < \bar{\sigma}(\Delta)$, the map η_p restricted to $(\eta_p, \eta_{E'})^{-1}((-4\Delta, 4\Delta) \times (-\infty, 4\Delta])$ is a fibration with fiber diffeomorphic to the closed 2-disk D^2 .*

Proof. — Let $\epsilon > 0$ be a constant which will be internal to this proof.

By Lemma 6.16, if $\beta_E < \bar{\beta}_E(\epsilon, \Delta, w')$ then the map F_p of Definition 9.2 is ϵ -close to a $(1, \epsilon)$ -splitting $\phi : (\frac{1}{\epsilon_p}M, p) \rightarrow (\mathbb{R} \times Z, (0, \star_Z))$, where Z is a complete pointed nonnegatively curved C^K -smooth surface, and ϕ is ϵ -close in the C^{K+1} -topology to an isometry between the ball $B(p, 1000\Delta) \subset (\frac{1}{\epsilon_p}M, p)$ and its image in $(\mathbb{R} \times Z, (0, \star_Z))$. If in addition $\sigma_E < \bar{\sigma}_E(\Delta)$ then the pointed ball $(B(\star_Z, 10\Delta), \star_Z) \subset (Z, \star_Z)$ will have pointed-Gromov-Hausdorff distance at most δ from the pointed interval $([0, 10], 0)$, where δ is the parameter of Lemma 3.12. By Lemma 3.12, we conclude that $B(\star_Z, \Delta)$ is homeomorphic to the closed 2-disk.

Put $Y = \mathbb{R} \times \{\star_Z\} \subset \mathbb{R} \times Z$ and let $d_Y : \mathbb{R} \times Z \rightarrow \mathbb{R}$ be the distance to Y . By Lemma 3.12, if ϵ is sufficiently small then for every $x \in \mathbb{R} \times Z$ with $d_Y(x) \in [\Delta, 9\Delta]$, the set V_x of initial velocities of minimizing segments from x to Y has small diameter; moreover V_x is orthogonal to the \mathbb{R} -factor of $\mathbb{R} \times Z$. Thus we may apply Lemma 3.16 to find a smoothing ρ_Y of d_Y , where $\|\rho_Y - d_Y\|_\infty$ is small, and in $d_Y^{-1}(\Delta, 9\Delta)$ the gradient of ρ_Y is close to the generalized gradient of d_Y .

Note that by Lemma 9.7, we may assume that $\phi^{-1}(\mathbb{R} \times \{\star_Z\}) \cap B(p, 50\Delta)$ is Hausdorff close to $E' \cap B(p, 50\Delta)$. Since ϕ is C^{K+1} -close to an isometry, the generalized gradient of $d_Y \circ \phi$ will be close to the generalized gradient of $d_{E'}$ in the region $(\eta_p, \eta_{E'})^{-1}(-5\Delta, 5\Delta) \times (2\Delta, 5\Delta)$, where the gradients are taken with respect to the metric on $\frac{1}{v_p}M$. (One may see this by applying a compactness argument to conclude that minimizing geodesics to E' in this region are mapped by ϕ to be C^1 -close to minimizing geodesics to Y .) Hence if $\zeta_{E'}$ is small then $\eta_{E'}$ and $\rho_Y \circ \phi$ will be C^1 -close in the region $(\eta_p, \eta_{E'})^{-1}(-5\Delta, 5\Delta) \times (2\Delta, 5\Delta)$. Similarly, if β_E and ζ_{edge} are small then η_p will be C^1 -close to $\pi_Z \circ \phi$ in the region $(\eta_p, \eta_{E'})^{-1}(-5\Delta, 5\Delta) \times (-\infty, 5\Delta)$.

For $t \in [0, 1]$, define a map $f^t : (\eta_p, \eta_{E'})^{-1}((-5\Delta, 5\Delta) \times (-\infty, 5\Delta)) \rightarrow \mathbb{R}^2$ by

$$(9.22) \quad f^t = (t\eta_p + (1-t)\pi_Z \circ \phi, t\eta_{E'} + (1-t)\rho_Y \circ \phi).$$

Let $F : (\eta_p, \eta_{E'})^{-1}((-5\Delta, 5\Delta) \times (-\infty, 5\Delta)) \times [0, 1] \rightarrow \mathbb{R}^2$ be the map with slices $\{f^t\}$. In view of the C^1 -closeness discussed above, we may now apply Lemma 21.1 to conclude that $(\eta_p, \eta_{E'})^{-1}(\{0\} \times (-\infty, 4\Delta])$ is diffeomorphic to $(\pi_Z, \rho_Y)^{-1}(\{0\} \times (-\infty, 4\Delta])$, which is a closed 2-disk.

Finally, we claim that the restriction of η_p to $(\eta_p, \eta_{E'})^{-1}(-4\Delta, 4\Delta) \times (-\infty, 4\Delta]$ yields a proper submersion to $(-4\Delta, 4\Delta)$, and is therefore a fibration. The properness follows from the fact that $(\eta_p, \eta_{E'})^{-1}((-4\Delta, 4\Delta) \times (-\infty, 4\Delta])$ is contained in a compact subset of the domain of $(\eta_p, \eta_{E'})$. The fact that it is a submersion follows from the nonvanishing of $D\eta_p$, and the linear independence of $\{D\eta_p, D\eta_{E'}\}$ at points with $(\eta_p, \eta_{E'}) \in (-4\Delta, 4\Delta) \times \{4\Delta\}$. \square

Remark 9.23. — Given w' , we take β_E very small in the proof of Lemma 9.21 in order to get a very good 1-splitting. On the other hand, we just have to take σ_E , and hence σ , small enough to apply Lemma 3.12; the parameter δ of Lemma 3.12 is independent of w' . This will be important for the order in which we choose the parameters.

9.5. Selection of edge balls. — Let $\{p_i\}_{i \in I_{\text{edge}}}$ be a maximal set of edge points with the property that the collection $\{B(p_i, \frac{1}{3}\Delta r_{p_i})\}_{i \in I_{\text{edge}}}$ is disjoint. We write ζ_i for ζ_{p_i} .

Lemma 9.24. — Under the constraints $\mathcal{M} > \overline{\mathcal{M}}$ and $\Lambda < \overline{\Lambda}(\Delta)$,

- $\bigcup_{i \in I_{\text{edge}}} B(p_i, \Delta r_{p_i})$ contains E .
- The intersection multiplicity of the collection $\{\text{supp}(\zeta_i)\}_{i \in I_{\text{edge}}}$ is bounded by \mathcal{M} .

Proof. — We omit the proof, as it is similar to the proof of Lemma 8.5. \square

We now give a useful covering of the 1-stratum points.

Lemma 9.25. — *Under the constraint $\Lambda < \overline{\Lambda}(\Delta)$, any 1-stratum point lies in the slim 1-stratum or lies in $\bigcup_{i \in I_{\text{edge}}} B(p_i, 3\Delta\tau_{p_i})$.*

Proof. — From Lemma 9.10, if a 1-stratum point does not lie in the slim 1-stratum then it lies in $B(p, \Delta\tau_p)$ for some $p \in E$. By Lemma 9.24 we have $p \in B(p_i, \Delta\tau_{p_i})$ for some $i \in I_{\text{edge}}$. If $\Lambda\Delta$ is sufficiently small then we can assume that $.9 \leq \frac{\tau_p}{\tau_{p_i}} \leq 1.1$. The lemma follows. \square

The next lemma will be used later for the interface between the slim stratum and the edge stratum.

Lemma 9.26. — *Under the constraints $\beta_E < \overline{\beta}_E(\Delta, \beta_2)$, $\varsigma_{\text{edge}} < \overline{\varsigma}_{\text{edge}}(\Delta, \beta_2)$ and $\Lambda < \overline{\Lambda}(\Delta)$, the following holds. Suppose for some $i \in I_{\text{edge}}$ and $q \in B(p_i, 10\Delta\tau_{p_i})$ we have*

$$(9.27) \quad \eta_{E'}(q) < 5\Delta, \quad |\eta_{p_i}(q)| < 5\Delta.$$

Then either p_i belongs to the slim 1-stratum, or there is a $j \in I_{\text{edge}}$ such that $q \in B(p_j, 10\Delta\tau_{p_j})$ and $|\eta_{p_j}(q)| < 2\Delta$.

Proof. — We may assume that p_i does not belong to the slim 1-stratum.

From the definition of η_i and Lemma 9.7, provided that $\beta_{E'} < \overline{\beta}_{E'}(\Delta)$, $\beta_E < \overline{\beta}_E(\beta_{E'})$, $\Lambda < \overline{\Lambda}(\Delta)$ and $\varsigma_{\text{edge}} < \overline{\varsigma}_{\text{edge}}(\Delta)$, there will be a $q' \in E' \cap B(p_i, 10\Delta\tau_{p_i})$ such that $|\eta_{p_i}(q') - \eta_{p_i}(q)| < \frac{1}{10}\Delta$. Since p_i is not in the slim 1-stratum, by Lemma 9.10 (2) there is a $p \in E$ such that $q' \in B(p, \tau_p)$, and by Lemma 9.24 we have $p \in B(p_j, \Delta\tau_{p_j})$ for some $j \in I_{\text{edge}}$. If $\Lambda\Delta$ is small then we will have $|\eta_{p_j}(q')| < 1.5\Delta$. Lemma 4.31 now implies that if $\beta_{\text{edge}} < \overline{\beta}_{\text{edge}}(\Delta, \beta_2)$ and $\varsigma_{\text{edge}} < \overline{\varsigma}_{\text{edge}}(\Delta, \beta_2)$ then $|\eta_{p_j}(q') - \eta_{p_j}(q)| < \frac{1}{2}\Delta$. Thus $|\eta_{p_j}(q)| < 2\Delta$. \square

9.6. Additional cutoff functions. — We define two additional cutoff functions for later use:

$$(9.28) \quad \zeta_{\text{edge}} = 1 - \Phi_{\frac{1}{2}, 1} \circ \left(\sum_{i \in I_{\text{edge}}} \zeta_i \right)$$

and

$$(9.29) \quad \zeta_{E'} = \left(\Phi_{\frac{2}{10}\Delta, \frac{3}{10}\Delta, 8\Delta, 9\Delta} \circ \eta_{E'} \right) \cdot \zeta_{\text{edge}}.$$

10. The local geometry of the slim 1-stratum

In this section we consider the slim 1-stratum points. Along with the adapted coordinates and cutoff functions, we discuss the local topology and a selection process to get a ball covering of the slim 1-stratum points.

10.1. Adapted coordinates, cutoff functions and local topology near slim 1-stratum points. — In this subsection, we let p denote a point in the slim 1-stratum, and $\phi_p : (\frac{1}{\tau_p}M, p) \rightarrow (\mathbb{R} \times X, (0, \star_X))$ be a $(1, \beta_1)$ -splitting, with $\text{diam}(X) \leq 10^3\Delta$. Let $\varsigma_{\text{slim}} > 0$ be a new parameter.

Lemma 10.1. — *Under the constraint $\beta_1 < \overline{\beta}_1(\Delta, \varsigma_{\text{slim}})$:*

- *There is a ϕ_p -adapted coordinate η_p of quality ς_{slim} on $B(p, 10^6\Delta) \subset (\frac{1}{\tau_p}M, p)$.*
- *The cutoff function*

$$(10.2) \quad (\Phi_{-9 \cdot 10^5\Delta, -8 \cdot 10^5\Delta, 8 \cdot 10^5\Delta, 9 \cdot 10^5\Delta}) \circ \eta_p$$

extends by zero to a smooth function ζ_p on M .

Proof. — This follows from Lemma 4.23 (see also Remark 4.35). □

Let η_p and ζ_p be as in Lemma 10.1.

Lemma 10.3. — *Under the constraints $\beta_1 < \overline{\beta}_1(\varsigma_{\text{slim}}, \Delta, w')$, $\varsigma_{\text{slim}} < \overline{\varsigma}_{\text{slim}}(\Delta)$, then $\eta_p^{-1}\{0\}$ is diffeomorphic to S^2 or T^2 .*

Proof. — From Lemma 6.16, if $\beta_1 < \overline{\beta}_1(\Delta, w')$ then close to ϕ_p , there is a $(1, \beta_1)$ -splitting $\phi : (\frac{1}{\tau_p}M, p) \rightarrow (\mathbb{R} \times Z, (0, \star_Z))$ for some complete pointed non-negatively curved C^K -smooth surface (Z, \star_Z) , with the map being C^{K+1} -close to an isometry on $B(p, 10^6\Delta)$. From Definition 7.4, the diameters of the Z -fibers are at most $10^4\Delta$. In particular, since Z is compact and M is orientable, Z must be diffeomorphic to S^2 or T^2 . Furthermore, we may assume that for any pair of points $m, m' \in M$ with $m \in B(p, 10^6\Delta) \subset \frac{1}{\tau_p}M$, $d(m, m') \in [2, 10]$, and $\pi_Z(\phi(m)) = \pi_Z(\phi(m'))$, the initial velocity v of a minimizing segment γ from m to m' maps under ϕ_* to a vector almost tangent to the \mathbb{R} -factor of $\mathbb{R} \times Z$.

As ϕ is close to ϕ_p , we may assume that η_p is a ϕ -adapted coordinate of quality $2\varsigma_{\text{slim}}$. If $\varsigma_{\text{slim}} < \overline{\varsigma}_{\text{slim}}(\Delta)$, then we may apply the estimate from the preceding paragraph, and the definition of adapted coordinates (specifically (4.22)), to conclude that η_p is C^1 -close to the composition $\frac{1}{\tau_p}M \xrightarrow{\phi} \mathbb{R} \times Z \rightarrow \mathbb{R}$ on the ball $B(p, 10^6\Delta)$. The lemma now follows from Lemma 21.3. □

10.2. Selection of slim 1-stratum balls. — Let $\{p_i\}_{i \in I_{\text{slim}}}$ be a maximal set of slim stratum points with the property that the collection $\{B(p_i, \frac{1}{3}\Delta\tau_{p_i})\}_{i \in I_{\text{slim}}}$ is disjoint. We write ζ_i for ζ_{p_i} .

Lemma 10.4. — *Under the constraints $\mathcal{M} > \overline{\mathcal{M}}$ and $\Lambda < \overline{\Lambda}(\Delta)$,*

- $\bigcup_{i \in I_{\text{slim}}} B(p_i, \Delta\tau_{p_i})$ *contains all slim stratum points.*
- *The intersection multiplicity of the collection $\{\text{supp}(\zeta_i)\}_{i \in I_{\text{slim}}}$ is bounded by \mathcal{M} .*

We omit the proof, as it is similar to the proof of Lemma 8.5.

11. The local geometry of the 0-stratum

Thus far, points in the 0-stratum have been defined by a process of elimination (they are points that are neither 2-stratum points nor 1-stratum points) rather than by the presence of some particular geometric structure. We now discuss their geometry. We show in Lemma 11.1 that M has conical structure near every point – not just the 0-stratum points – provided one looks at an appropriate scale larger than \mathfrak{r} . We then use this to define radial and cutoff functions near 0-stratum points.

Let $\delta_0 > 0$ and $\Upsilon_0, \Upsilon'_0, \tau_0 > 1$ be new parameters.

11.1. The Good Annulus Lemma. — We now show that for every point p in M , there is a scale at which a neighborhood of p is well approximated by a model geometry in two different ways: by a nonnegatively curved 3-manifold in the pointed C^K -topology, and by the Tits cone of this manifold in the pointed Gromov-Hausdorff topology.

Lemma 11.1. — *Under the constraint $\Upsilon'_0 > \overline{\Upsilon}'_0(\delta_0, \Upsilon_0, w')$, if $p \in M$ then there exists $r_p^0 \in [\Upsilon_0 \mathfrak{r}_p, \Upsilon'_0 \mathfrak{r}_p]$ and a complete 3-dimensional nonnegatively curved C^K -smooth Riemannian manifold N_p such that:*

- (1) $(\frac{1}{r_p^0} M, p)$ is δ_0 -close in the pointed Gromov-Hausdorff topology to the Tits cone $C_T N_p$ of N_p .
- (2) The ball $B(p, r_p^0) \subset M$ is diffeomorphic to N_p .
- (3) The distance function from p has no critical points in the annulus $A(p, \frac{r_p^0}{100}, r_p^0)$.

Proof. — Suppose that conclusion (1) does not hold. Then for each j , if we take $\Upsilon'_0 = j\Upsilon_0$, it is not true that conclusion (1) holds for sufficiently large α . Hence we can find a sequence $\alpha_j \rightarrow \infty$ so that for each j , $(M^{\alpha_j}, p_{\alpha_j})$ provides a counterexample with $\Upsilon'_0 = j\Upsilon_0$.

For convenience of notation, we relabel $(M^{\alpha_j}, p_{\alpha_j})$ as (M_j, p_j) and write \mathfrak{r}_j for $\mathfrak{r}_{p_{\alpha_j}}$. Then by assumption, for each $r_j^0 \in [\Upsilon_0 \mathfrak{r}_j, j\Upsilon_0 \mathfrak{r}_j]$ there is no 3-dimensional nonnegatively curved Riemannian manifold N_j such that conclusion (1) holds.

Assumption 5.2 implies that a subsequence of $\{(\frac{1}{\mathfrak{r}_j} M_j, p_j)\}_{j=1}^\infty$, which we relabel as $\{(\frac{1}{\mathfrak{r}_j} M_j, p_j)\}_{j=1}^\infty$, converges in the pointed C^K -topology to a pointed 3-dimensional nonnegatively curved C^K -smooth Riemannian manifold (N, p_∞) . Now N is asymptotically conical. That is, there is some $R > 0$ so that if $R' > R$ then $(\frac{1}{R'} N, p_\infty)$ is $\frac{\delta_0}{2}$ -close in the pointed Gromov-Hausdorff topology to the Tits cone $C_T N$.

By critical point theory, large open balls in N are diffeomorphic to N itself. Hence we can find $R' > 10^3 \max(\Upsilon_0, R)$ so that for any $R'' \in (\frac{1}{2}R', 2R')$, there are no critical points of the distance function from p_∞ in $A(p_\infty, \frac{R''}{10^3}, 10R'') \subset N$, and the ball $B(p_\infty, R'')$ is diffeomorphic to N . In view of the convergence $(\frac{1}{\mathfrak{r}_j} M_j, p_j) \rightarrow (N, p_\infty)$ in the pointed C^K -topology, it follows that for large j there are no critical points of the

distance function in $A(p_j, \frac{R'\tau_j}{100}, R''\tau_j) \subset N_j$, and $B(p_j, R''\tau_j) \subset M_j$ is diffeomorphic to $B(p_\infty, R'') \subset N$. Taking $r_j^0 = R'\tau_j$ gives a contradiction. \square

Remark 11.2. — If we take the parameter σ of Lemma 6.18 to be small then we can additionally conclude that $C_T N_p$ is pointed Gromov-Hausdorff close to a conical nonnegatively curved Alexandrov space of dimension at most two.

11.2. The radial function near a 0-stratum point. — For every $p \in M$, we apply Lemma 11.1 to get a scale $r_p^0 \in [\Upsilon_0\tau_p, \Upsilon'_0\tau_p]$ for which the conclusion of Lemma 11.1 holds. In particular, $(\frac{1}{r_p^0}M, p)$ is δ_0 -close in the pointed Gromov-Hausdorff topology to the Tits cone $C_T N_p$ of a nonnegatively curved 3-manifold N_p .

Let d_p be the distance function from p in $(\frac{1}{r_p^0}M, p)$. Let $\varsigma_{0\text{-stratum}} > 0$ be a new parameter.

Lemma 11.3. — *Under the constraint $\delta_0 < \bar{\delta}_0(\varsigma_{0\text{-stratum}})$, there is a function $\eta_p : \frac{1}{r_p^0}M \rightarrow [0, \infty)$ such that:*

- (1) η_p is smooth on $A(p, \frac{1}{10}, 10) \subset \frac{1}{r_p^0}M$.
- (2) $\|\eta_p - d_p\|_\infty < \varsigma_{0\text{-stratum}}$.
- (3) $\eta_p - d_p : \frac{1}{r_p^0}M \rightarrow [0, \infty)$ is $\varsigma_{0\text{-stratum}}$ -Lipschitz.
- (4) η_p is smooth and has no critical points in $\eta_p^{-1}([\frac{2}{10}, 2])$, and for every $\rho \in [\frac{2}{10}, 2]$, the sublevel set $\eta_p^{-1}([0, \rho])$ is diffeomorphic to either the closed disk bundle in the normal bundle νS of the soul $S \subset N_p$, if N_p is noncompact, or to N_p itself when N_p is compact.
- (5) The composition $\Phi_{\frac{2}{10}, \frac{3}{10}, \frac{8}{10}, \frac{9}{10}} \circ \eta_p$ extends by zero to a smooth cutoff function $\zeta_p : M \rightarrow [0, 1]$.

Proof. — We apply Lemma 3.16 with $Y = \{p\}$, $U = A(p, \frac{1}{20}, 20)$ and $C = \overline{A(p, \frac{1}{10}, 10)}$. To verify the hypotheses of Lemma 3.16, suppose that $q \in U$. From Lemma 11.1, for any $\mu > 0$, there is an $\bar{\delta}_0 = \bar{\delta}_0(\mu)$ so that if $\delta_0 < \bar{\delta}_0$ then we can find some $q' \in M$ with $d(p, q') = 2d(p, q)$ and $d(q, q') \geq (1 - \mu)d(p, q)$. Fix a minimizing geodesic γ_1 from q to q' . By triangle comparison, for any $\theta > 0$, if μ is sufficiently small then we can ensure that for any minimizing geodesic γ from q to p , the angle between $\gamma'(0)$ and $\gamma_1'(0)$ is at least $\pi - \frac{\theta}{2}$. Parts (1), (2) and (3) of the lemma now follow from Lemma 3.16.

(4). Using the same proof as the “Morse lemma” for distance functions, one gets a smooth vector field ξ in $A(p, \frac{1}{10}, 10)$, such that ξd_p and $\xi \eta_p$ are both close to 1. Using the flow of ξ , if $\varsigma_{0\text{-stratum}}$ is sufficiently small then for every $\rho \in [\frac{2}{10}, 2]$, the sublevel sets $d_p^{-1}([0, \rho])$ and $\eta_p^{-1}([0, \rho])$ are homeomorphic. Let \bar{N} be the closed disk bundle in νS . Then $\text{int}(\bar{N}) \stackrel{\text{homeo}}{\simeq} N_p \stackrel{\text{homeo}}{\simeq} \text{int}(d_p^{-1}([0, \rho])) \stackrel{\text{homeo}}{\simeq} \text{int}(\eta_p^{-1}([0, \rho]))$. Since two compact orientable 3-manifolds with boundary are homeomorphic provided that their interiors are homeomorphic, we have $\bar{N} \stackrel{\text{homeo}}{\simeq} \eta_p^{-1}([0, \rho])$. (This may be readily

deduced from the fact that if S is a closed orientable surface then any smooth embedding $S \rightarrow S \times \mathbb{R}$, which is also a homotopy equivalence, is isotopic to the fiber $S \times \{0\}$, as follows from the Schoenflies theorem when $S = S^2$ and from [32] when $\text{genus}(S) > 0$.)

(5) follows from the fact that the composition $\Phi_{\frac{2}{10}, \frac{3}{10}, \frac{8}{10}, \frac{9}{10}} \circ \eta_p$ is compactly supported in the annulus $A(p, \frac{1}{10}, 10)$. \square

Remark 11.4. — One may avoid the Schoenflies and Stallings theorems in the proof of Lemma 11.3 (4). If M is a complete noncompact nonnegatively curved manifold, and $p \in M$, then the distance function d_p has no critical points outside $B(p, r_0)$ for some $r_0 \in (0, \infty)$. In fact, for every $r > r_0$, the closed ball $\overline{B(p, r)}$ is isotopic, by an isotopy with arbitrarily small tracks, to a compact domain with smooth boundary D ; moreover, the smooth isotopy class $[D]$ is canonical and independent of $r \in (r_0, \infty)$. (These assertions are true in general for noncritical sublevel sets of proper distance functions. They are proved by showing that one may smooth d_p near $S(p, r)$ without introducing critical points.) The proof of the soul theorem actually shows that the isotopy class $[D]$ is the same as that of a closed smooth tubular neighborhood of the soul, which is diffeomorphic to the unit normal bundle of the soul.

11.3. Selecting the 0-stratum balls. — The next lemma has a statement about an adapted coordinate for the radial splitting in an annular region of a 0-stratum ball. We use the parameter ς_{slim} for the quality of this splitting, even though there is no *a priori* relationship to slim 1-stratum points. Our use of this parameter will simplify the later parameter ordering.

Lemma 11.5. — *Under the constraints $\delta_0 < \bar{\delta}_0(\beta_1, \varsigma_{\text{slim}})$, $\Upsilon_0 > \bar{\Upsilon}_0(\beta_1)$, $\beta_1 < \bar{\beta}_1(\varsigma_{\text{slim}})$ and $\varsigma_{0\text{-stratum}} < \bar{\varsigma}_{0\text{-stratum}}(\varsigma_{\text{slim}})$, there is a finite collection $\{p_i\}_{i \in I_{0\text{-stratum}}}$ of points in M so that*

- (1) *The balls $\{B(p_i, r_{p_i}^0)\}_{i \in I_{0\text{-stratum}}}$ are disjoint.*
- (2) *If $q \in B(p_i, 10r_{p_i}^0)$, for some $i \in I_{0\text{-stratum}}$, then $r_q^0 \leq 20r_{p_i}^0$ and $\frac{r_{p_i}^0}{r_q} \geq \frac{1}{20}\Upsilon_0$.*
- (3) *For each i , every $q \in A(p_i, \frac{1}{10}r_{p_i}^0, 10r_{p_i}^0)$ belongs to the 1-stratum or 2-stratum, and there is a $(1, \beta_1)$ -splitting of $(\frac{1}{r_q}M, q)$, for which $\frac{r_{p_i}^0}{r_q} \eta_{p_i}$ is an adapted coordinate of quality ς_{slim} .*
- (4) $\bigcup_{i \in I_{0\text{-stratum}}} B(p_i, \frac{1}{10}r_{p_i}^0)$ *contains all the 0-stratum points.*
- (5) *For each $i \in I_{0\text{-stratum}}$, the manifold N_{p_i} has at most one end.*

Proof

(1). Let $V_0 \subset M$ be the set of points $p \in M$ such that the ball $B(p, r_p^0)$ contains a 0-stratum point. We partially order V_0 by declaring that $p_1 \prec p_2$ if and only if $(2r_{p_1}^0 < r_{p_2}^0$ and $B(p_1, r_{p_1}^0) \subset B(p_2, r_{p_2}^0)$). Note that every chain in the poset (V_0, \prec) has an upper bound, since $r_p^0 < \Upsilon'_0 r_p$ is bounded above. Let $V \subset V_0$ be the subset of

elements which are maximal with respect to \prec , and apply Lemma 19.1 with $\mathcal{R}_p = r_p^0$ to get the finite disjoint collection of balls $\{B(p_i, r_{p_i}^0)\}_{i \in I_{0\text{-stratum}}}$. Thus (1) holds.

(2). If $q \in B(p_i, 10r_{p_i}^0)$ then $r_q^0 \leq 20r_{p_i}^0$, for otherwise we would have $q \in V_0$ and $p_i \prec q$, contradicting the maximality of p_i . Thus $\frac{r_{p_i}^0}{r_q^0} \geq \frac{r_q^0}{20r_{p_i}^0} \geq \frac{1}{20}\Upsilon_0$.

(3). Suppose that $i \in I_{0\text{-stratum}}$ and $q \in A(p_i, \frac{1}{10}r_{p_i}^0, 10r_{p_i}^0)$. Recall that $(\frac{1}{r_{p_i}^0}M, p_i)$ is δ_0 -close in the pointed Gromov-Hausdorff topology to the Tits cone $C_T N_{p_i}$. If the Tits cone $C_T N_{p_i}$ were a single point then $\text{diam}(M)$ would be bounded above by $\delta_0 r_{p_i}^0$; taking $\delta_0 < \frac{1}{10}$ we get $q \in B(p_i, \frac{1}{10}r_{p_i}^0)$, which is a contradiction. Therefore $C_T N_{p_i}$ is not a point. It follows that there is a 1-strainer at q of scale comparable to $r_{p_i}^0$ and quality comparable to δ_0 , where one of the strainer points is p_i . By (2), if $\Upsilon_0 > \overline{\Upsilon}_0(\beta_1)$ and $\delta_0 < \overline{\delta}(\beta_1)$ then Lemma 4.15 implies there is a $(1, \beta_1)$ -splitting $\alpha : (\frac{1}{r_q^0}M, q) \rightarrow (\mathbb{R} \times X, (0, \star_X))$, where the first component is given by $d_{p_i} - d_{p_i}(q)$. In particular q is a 1-stratum point or a 2-stratum point. By Lemma 11.3, the smooth radial function η_{p_i} is $\varsigma_{0\text{-stratum}}$ -Lipschitz close to d_{p_i} . Lemma 4.28 implies that if $\varsigma_{0\text{-stratum}} < \overline{\varsigma}_{0\text{-stratum}}(\varsigma_{\text{slim}})$ then we are ensured that η_{p_i} is an α -adapted coordinate of quality ς_{slim} . Hence (3) holds.

(4). If q is in the 0-stratum then $q \in V_0$, so $q \prec \bar{q}$ for some $\bar{q} \in V$. By Lemma 19.1, for some $i \in I_{0\text{-stratum}}$ we have $B(\bar{q}, r_{\bar{q}}^0) \cap B(p_i, r_{p_i}^0) \neq \emptyset$ and $r_{\bar{q}}^0 \leq 2r_{p_i}^0$. Therefore $q \in B(p_i, 5r_{p_i}^0)$ and by (3), we have $q \in B(p_i, \frac{1}{10}r_{p_i}^0)$.

(5). Let $\epsilon > 0$ be a new constant. Suppose that N_{p_i} has more than one end. Then $C_T N_{p_i} \simeq \mathbb{R}$. If $\delta_0 < \overline{\delta}_0(\epsilon)$ then every point $q \in B(p_i, 1) \subset \frac{1}{r_{p_i}^0}M$ will have a strainer of quality ϵ and scale ϵ^{-1} . By (2) and Lemma 4.15, if $\epsilon < \overline{\epsilon}(\beta_1)$ and $\Upsilon_0 > \overline{\Upsilon}_0(\beta_1)$ then there is a $(1, \beta_1)$ splitting of $(\frac{1}{r_q^0}M, q)$. Thus every point in $B(p_i, r_{p_i}^0)$ is in the 1-stratum or 2-stratum. This contradicts the definition of V , and hence N_{p_i} has at most one end. \square

12. Mapping into Euclidean space

12.1. The definition of the map $\mathcal{E}^0 : M \rightarrow H$. — We will now use the ball collections defined in Sections 8-11, and the geometrically defined functions discussed in earlier sections, to construct a smooth map $\mathcal{E}^0 : M \rightarrow H = \bigoplus_{i \in I} H_i$, where

- $I = I_{\mathfrak{r}} \cup I_{E'} \cup I_{0\text{-stratum}} \cup I_{\text{slim}} \cup I_{\text{edge}} \cup I_{2\text{-stratum}}$, where the two index sets $I_{\mathfrak{r}}$ and $I_{E'}$ are singletons $I_{\mathfrak{r}} = \{\mathfrak{r}\}$ and $I_{E'} = \{E'\}$ respectively,
- H_i is a copy of \mathbb{R} when $i = \mathfrak{r}$,
- H_i is a copy of $\mathbb{R} \oplus \mathbb{R}$ when $i \in I_{E'} \cup I_{0\text{-stratum}} \cup I_{\text{slim}} \cup I_{\text{edge}}$, and
- H_i is a copy of $\mathbb{R}^2 \oplus \mathbb{R}$ when $i \in I_{2\text{-stratum}}$.

We also put

- $H_{0\text{-stratum}} = \bigoplus_{i \in I_{0\text{-stratum}}} H_i$,
- $H_{\text{slim}} = \bigoplus_{i \in I_{\text{slim}}} H_i$,
- $H_{\text{edge}} = \bigoplus_{i \in I_{\text{edge}}} H_i$,

- $H_{2\text{-stratum}} = \bigoplus_{i \in I_{2\text{-stratum}}} H_i,$
- $Q_1 = H,$
- $Q_2 = H_{0\text{-stratum}} \oplus H_{\text{slim}} \oplus H_{\text{edge}},$
- $Q_3 = H_{0\text{-stratum}} \oplus H_{\text{slim}},$
- $Q_4 = H_{0\text{-stratum}},$ and
- $\pi_{i,j} : Q_i \rightarrow Q_j, \pi_i = \pi_{1,i} : H \rightarrow Q_i, \pi_i^\perp : H \rightarrow Q_i^\perp$ are the orthogonal projections, for $1 \leq i \leq j \leq 4.$

If $x \in Q_j,$ we denote the projection to a summand H_i by $\pi_{H_i}(x) = x_i.$ When $i \neq \tau,$ we write $H_i = H'_i \oplus H''_i \cong \mathbb{R}^{k_i} \oplus \mathbb{R},$ where $k_i \in \{1, 2\},$ and we denote the decomposition of $x_i \in H_i$ into its components by $x_i = (x'_i, x''_i) \in H'_i \oplus H''_i.$ We denote orthogonal projection onto H'_i and H''_i by $\pi_{H'_i}$ and $\pi_{H''_i},$ respectively.

In Sections 8-11, we defined adapted coordinates $\eta_p,$ and cutoff functions ζ_p corresponding to points $p \in M$ of different types. If $\{p_i\}$ is a collection of points used to define a ball cover, as in Sections 8-11, then we write η_i for η_{p_i} and ζ_i for $\zeta_{p_i}.$ Recall that we also defined $\eta_{E'}$ and $\zeta_{E'}$ in Sections 9.2 and 9.6, respectively. For $i \in I \setminus \{\tau\},$ we will also define a new scale parameter $R_i,$ as follows:

- If $i \in I_{0\text{-stratum}}$ we put $R_i = r_{p_i}^0,$ where $r_{p_i}^0$ is as in Lemma 11.5;
- If $i \in I_{\text{slim}} \cup I_{\text{edge}} \cup I_{2\text{-stratum}},$ then $R_i = r_{p_i};$
- If $i = E',$ then $R_i = r;$ note that unlike in the other cases, R_i is not a constant.

The component $\mathcal{E}_i^0 : M \rightarrow H_i$ of the map $\mathcal{E}^0 : M \rightarrow H$ is defined to be τ when $i = \tau,$ and

$$(12.1) \quad (R_i \eta_i \zeta_i, R_i \zeta_i)$$

otherwise.

In the remainder of this section we prepare for the adjustment procedure in Section 13 by examining the behavior of \mathcal{E}^0 near the different strata.

12.2. The image of $\mathcal{E}^0.$ — Before proceeding, we make some observations about the image of $\mathcal{E}^0,$ to facilitate the choice of cutoff functions. Let $x = \mathcal{E}^0(p) \in H.$ Then the components of x satisfy the following inequalities:

$$(12.2) \quad x_\tau > 0$$

and for every $i \in I_{0\text{-stratum}} \cup I_{\text{slim}} \cup I_{\text{edge}} \cup I_{2\text{-stratum}},$

$$(12.3) \quad x''_i \in [0, R_i] \quad \text{and} \quad |x'_i| \leq c_i x''_i,$$

where

$$(12.4) \quad c_i = \begin{cases} 9\Delta & \text{when } i \in I_{E'}, \\ \frac{9}{10} & \text{when } i \in I_{0\text{-stratum}}, \\ 10^5 \Delta & \text{when } i \in I_{\text{slim}}, \\ 9\Delta & \text{when } i \in I_{\text{edge}}, \\ 9 & \text{when } i \in I_{2\text{-stratum}}. \end{cases}$$

Lemma 12.5. — Under the constraint $\Lambda \leq \bar{\Lambda}(\mathcal{M})$, there is a number $\Omega_0 = \Omega_0(\mathcal{M})$ so that for all $p \in M$, $|D\mathcal{E}_p^0| \leq \Omega_0$.

Proof. — This follows from the definition of \mathcal{E}^0 . \square

12.3. Structure of \mathcal{E}^0 near the 2-stratum. — Put

$$(12.6) \quad \tilde{A}_1 = \bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| \leq 8\}, \quad A_1 = \bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| \leq 7\}.$$

We refer to Definition 20.1 for the definition of a cloudy manifold. We will see that on a scale which is sufficiently small compared with \mathfrak{r} , the pair $(\tilde{S}_1, S_1) = (\mathcal{E}^0(\tilde{A}_1), \mathcal{E}^0(A_1)) \subset H$ is a cloudy 2-manifold. In brief, this is because, on a scale small compared with \mathfrak{r} , near any point in A_1 the map \mathcal{E}^0 is well approximated in the C^1 topology by an affine function of η_i , for some $i \in I_{2\text{-stratum}}$.

Let $\Sigma_1, \Gamma_1 > 0$ be new parameters. Define $r_1 : \tilde{S}_1 \rightarrow (0, \infty)$ by putting $r_1(x) = \Sigma_1 \mathfrak{r}_p$ for some $p \in (\mathcal{E}^0)^{-1}(x) \cap \tilde{A}_1$.

Lemma 12.7. — There is a constant $\Omega_1 = \Omega_1(\mathcal{M})$ so that under the constraints $\Sigma_1 < \bar{\Sigma}_1(\Gamma_1, \mathcal{M})$, $\beta_2 < \bar{\beta}_2(\Gamma_1, \Sigma_1, \mathcal{M})$, $\varsigma_{2\text{-stratum}} < \bar{\varsigma}_{2\text{-stratum}}(\Gamma_1, \Sigma_1, \mathcal{M})$, $\beta_E < \bar{\beta}_E(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\sigma_E < \bar{\sigma}_E(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\varsigma_{E'} < \bar{\varsigma}_{E'}(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\beta_1 < \bar{\beta}_1(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\varsigma_{\text{slim}} < \bar{\varsigma}_{\text{slim}}(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\varsigma_{0\text{-stratum}} < \bar{\varsigma}_{0\text{-stratum}}(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, $\Upsilon_0 \geq \bar{\Upsilon}_0(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$ and $\Lambda < \bar{\Lambda}(\Gamma_1, \Sigma_1, \Delta, \mathcal{M})$, the following holds.

- (1) The triple (\tilde{S}_1, S_1, r_1) is a $(2, \Gamma_1)$ cloudy 2-manifold.
- (2) The affine subspaces $\{A_x\}_{x \in S_1}$ inherent in the definition of the cloudy 2-manifold can be chosen to have the following property. Pick $p \in A_1$ and put $x = \mathcal{E}^0(p) \in S_1$. Let $A_x^0 \subset H$ be the linear subspace parallel to A_x (i.e., $A_x = A_x^0 + x$) and let $\pi_{A_x^0} : H \rightarrow A_x^0$ denote orthogonal projection onto A_x^0 . Then

$$(12.8) \quad \|D\mathcal{E}_p^0 - \pi_{A_x^0} \circ D\mathcal{E}_p^0\| < \Gamma_1,$$

and

$$(12.9) \quad \Omega_1^{-1} \|v\| \leq \|\pi_{A_x^0} \circ D\mathcal{E}^0(v)\| \leq \Omega_1 \|v\|$$

for every $v \in T_p M$ which is orthogonal to $\ker(\pi_{A_x^0} \circ D\mathcal{E}_p^0)$.

- (3) Given $i \in I_{2\text{-stratum}}$, there is a smooth map $\widehat{\mathcal{E}}_i^0 : (\overline{B(0, 8)} \subset \mathbb{R}^2) \rightarrow (H'_i)^\perp$ such that

$$(12.10) \quad \|D\widehat{\mathcal{E}}_i^0\| \leq \Omega_1 R_i$$

and on the subset $\{|\eta_i| \leq 8\} \subset \frac{1}{R_i} M$, we have

$$(12.11) \quad \left\| \frac{1}{R_i} \mathcal{E}^0 - \left(\eta_i, \frac{1}{R_i} \widehat{\mathcal{E}}_i^0 \circ \eta_i \right) \right\|_{C^1} < \Gamma_1.$$

Furthermore, if $x \in S_1$ then there are some $i \in I_{2\text{-stratum}}$ and $p \in \{|\eta_i| \leq 7\}$ such that $x = \mathcal{E}^0(p)$ and $A_x^0 = \text{Im} \left(I, \frac{1}{R_i} (D\widehat{\mathcal{E}}_i^0)_{\eta_i(p)} \right)$.

The parameters $\epsilon_1, \epsilon_2 > 0$ will be internal to this subsection, which is devoted to the proof of Lemma 12.7. Until further notice, the index i will denote a fixed element of $I_{2\text{-stratum}}$.

Put $J = \{j \in I_{E'} \cup I_{0\text{-stratum}} \cup I_{\text{slim}} \cup I_{\text{edge}} \cup I_{2\text{-stratum}} \mid \text{supp } \zeta_j \cap B(p_i, 10R_i) \neq \emptyset\}$.

Sublemma 12.12. — *Under the constraints $\beta_2 < \overline{\beta}_2(\epsilon_1)$, $\varsigma_{2\text{-stratum}} < \overline{\varsigma}_{2\text{-stratum}}(\epsilon_1)$, $\beta_E < \overline{\beta}_E(\epsilon_1, \Delta)$, $\varsigma_{\text{edge}} < \overline{\varsigma}_{\text{edge}}(\epsilon_1, \Delta)$, $\varsigma_{E'} < \overline{\varsigma}_E(\epsilon_1, \Delta)$, $\beta_1 < \overline{\beta}_1(\epsilon_1, \Delta)$, $\varsigma_{\text{slim}} < \overline{\varsigma}_{\text{slim}}(\epsilon_1, \Delta)$, $\varsigma_{0\text{-stratum}} < \overline{\varsigma}_{0\text{-stratum}}(\epsilon_1, \Delta)$ and $\Lambda < \overline{\Lambda}(\epsilon_1, \Delta)$, the following holds.*

For each $j \in J$, there is a map $T_{ij} : \mathbb{R}^2 \rightarrow \mathbb{R}^{k_j}$ which is a composition of an isometry and an orthogonal projection, such that on the ball $B(p_i, 10) \subset \frac{1}{R_i}M$, the map η_j is defined and satisfies

$$(12.13) \quad \left\| \frac{R_j}{R_i} \eta_j - (T_{ij} \circ \eta_i) \right\|_{C^1} < \epsilon_1.$$

Proof. — As we are assuming the hypotheses of Lemma 7.2, there are no 3-stratum points.

Suppose first that $j \in I_{2\text{-stratum}}$. Then $d(p_j, p_i) \leq 10(R_i + R_j)$. If Λ is sufficiently small then we can assume that $\frac{R_j}{R_i}$ is arbitrarily close to 1, so in particular $d(p_j, p_i) \leq 40R_j$. By Lemma 4.10, if β_2 is sufficiently small then the $(2, \beta_2)$ -splitting of $(\frac{1}{R_j}M, p_j)$ gives an arbitrarily good 2-splitting of $(\frac{1}{R_j}M, p_i)$. By Lemma 4.17, if β_2 is sufficiently small then this splitting of $(\frac{1}{R_j}M, p_i)$ is compatible, to an arbitrarily degree of closeness, with the $(2, \beta_2)$ -splitting of $(\frac{1}{R_i}M, p_i)$ coming from the fact that p_i is a 2-stratum point. Hence in this case, if β_2 and $\varsigma_{2\text{-stratum}}$ are sufficiently small (as functions of ϵ_1) then the sublemma follows from Lemma 4.31, along with Remark 4.35.

If $j \in I_{\text{edge}} \cup I_{\text{slim}}$ then $d(p_i, p_j) \leq 10R_i + 10^5 \Delta R_j$. We now have an approximate 1-splitting at p_j , which gives an approximate 1-splitting at p_i . As before, if β_2 , $\varsigma_{2\text{-stratum}}$, β_1 , ς_{edge} , ς_{slim} , and Λ are sufficiently small (as functions of ϵ_1 and Δ) then we can apply Lemmas 4.17 and 4.31 to deduce the conclusion of the sublemma. Note that in this case, we have to allow Λ to depend on Δ .

If $j \in I_{E'}$ then since $\text{supp } \zeta_{E'} \cap B(p_i, 10R_i) \neq \emptyset$, we know that $\eta_{E'}(q) \in [\frac{2}{10}\Delta, 9\Delta]$ for some $q \in B(p_i, 10R_i)$. As $\Delta \gg 10$, it follows from Lemma 9.12 that if Λ is sufficiently small then $B(p_i, 10R_i) \subset \eta_{E'}^{-1}([\frac{2}{10}\Delta, 10\Delta])$. From the definition of E' , it follows that if $\beta_{E'}$ and Λ are sufficiently small then there is a 1-splitting at p_i of arbitrarily good quality, coming from the $[0, C]$ -factor in Definition 9.2. As before, if $\beta_{E'}$, Λ , β_2 , $\varsigma_{2\text{-stratum}}$, $\beta_{E'}$ and ς_E are sufficiently small (as functions of ϵ_1 and Δ) then we can apply Lemmas 4.17 and 4.31 to deduce the conclusion of the sublemma.

If $j \in I_{0\text{-stratum}}$ then since $\text{supp } \zeta_j \cap B(p_i, 10R_i) \neq \emptyset$, we know that $\eta_j(q) \in [\frac{2}{10}, \frac{9}{10}]$ for some $q \in B(p_i, 10R_i)$. From Lemma 11.5, $\frac{r_{2i}^0}{R_i} \geq \frac{1}{20} \Upsilon_0$. Hence we may

assume that $B(p_i, 10R_i) \subset A(p_j, \frac{1}{10}r_{p_j}^0, r_{p_j}^0)$. Lemma 11.5 also gives a $(1, \beta_1)$ -splitting of $(\frac{R_i}{R_j}M, p_i)$. If β_1 and β_2 are sufficiently small then by Lemma 4.17, this 1-splitting is compatible with the $(2, \beta_2)$ -splitting of $(\frac{1}{R_i}M, p_i)$ to an arbitrary degree of closeness. As before, if $\beta_1, \beta_2, \varsigma_{2\text{-stratum}}$ and $\bar{\varsigma}_{0\text{-stratum}}$ are sufficiently small (as functions of ϵ_1 and Δ) then the sublemma follows from Lemma 4.31. \square

We retain the hypotheses of Sublemma 12.12.

For $j \in J$, the cutoff function ζ_j is a function of the $\eta_{j'}$'s for $j' \in J$, *i.e.*, there is a smooth function $\Phi_j \in C_c^\infty(\mathbb{R}^J)$ such that $\zeta_j(\cdot) = \Phi_j(\{\eta_{j'}(\cdot)\}_{j' \in J})$. (Note from (9.29) that $\zeta_{E'}$ depends on $\eta_{E'}$ and $\{\zeta_k\}_{k \in I_{\text{edge}}}$.) The H_j -component of \mathcal{E}^0 , after dividing by R_i , can be written as

$$(12.14) \quad \frac{1}{R_i} \mathcal{E}_j^0 = \left(\frac{R_j}{R_i} \eta_j \zeta_j, \frac{R_j}{R_i} \zeta_j \right) = \left(\frac{R_j}{R_i} \eta_j \cdot (\Phi_j \circ \{\eta_{j'}\}_{j' \in J}), \frac{R_j}{R_i} \Phi_j \circ \{\eta_{j'}\}_{j' \in J} \right).$$

Let $\mathcal{F}^0 : \mathbb{R}^2 \rightarrow H$ be the map so that the H_j -component of $\mathcal{F}^0 \circ \eta_i$, for $j \in J$, is obtained from the preceding formula by replacing each occurrence of η_j with the approximation $\frac{R_i}{R_j}(T_{ij} \circ \eta_i)$, *i.e.*,

$$(12.15) \quad \frac{1}{R_i} \mathcal{F}_j^0(u) = \left(T_{ij}(u) \cdot \left(\Phi_j \left(\left\{ \frac{R_i}{R_{j'}} T_{ij'}(u) \right\}_{j' \in J} \right) \right), \frac{R_j}{R_i} \Phi_j \left(\left\{ \frac{R_i}{R_{j'}} T_{ij'}(u) \right\}_{j' \in J} \right) \right),$$

whose H_τ -component is the constant function R_i , and whose other components vanish. That is,

$$(12.16) \quad \frac{1}{R_i} \mathcal{F}_j^0 \circ \eta_i = \left((T_{ij} \circ \eta_i) \cdot \left(\Phi_j \left(\left\{ \frac{R_i}{R_{j'}} T_{ij'} \circ \eta_i \right\}_{j' \in J} \right) \right), \frac{R_j}{R_i} \Phi_j \left(\left\{ \frac{R_i}{R_{j'}} T_{ij'} \circ \eta_i \right\}_{j' \in J} \right) \right).$$

Sublemma 12.17. — *Under the constraints $\epsilon_1 \leq \bar{\epsilon}_1(\epsilon_2, \mathcal{M})$, $\Upsilon_0 \geq \bar{\Upsilon}_0(\epsilon_2, \mathcal{M})$ and $\Lambda \leq \bar{\Lambda}(\epsilon_2, \mathcal{M})$,*

$$(12.18) \quad \left\| \frac{1}{R_i} \mathcal{E}^0 - \frac{1}{R_i} \mathcal{F}^0 \circ \eta_i \right\|_{C^1} < \epsilon_2$$

on $B(p_i, 10) \subset \frac{1}{R_i}M$.

Proof. — First note that $\mathcal{E}_\tau(p_i) = \mathcal{F}_\tau^0(p_i) = R_i$ and the \mathcal{E}_τ -component of \mathcal{E} has Lipschitz constant Λ , so it suffices to control the remaining components. For $j \in J$ and $j \in I_{E'} \cup I_{\text{slim}} \cup I_{\text{edge}} \cup I_{2\text{-stratum}}$, if Λ is sufficiently small then we can assume that $\frac{R_i}{R_j}$ is arbitrarily close to one. Then the H_j -component of $\frac{1}{R_i} \mathcal{E}^0 - \frac{1}{R_i} \mathcal{F}^0 \circ \eta_i$ can be estimated in C^1 -norm by using (12.13) to estimate $\eta_{j'}$, plugging this into (12.14) and applying the chain rule. In applying the chain rule, we use the fact that the functions Φ_j have explicit bounds on their derivatives of order up to 2.

If $j \in J \cap I_{0\text{-stratum}}$ then the only relevant argument of Φ_j is when $j' = j$. Hence in this case we can write

$$(12.19) \quad \frac{1}{R_i} \mathcal{E}_j^0 = \left(\frac{R_j}{R_i} \eta_j \cdot \Phi_j(\eta_j), \frac{R_j}{R_i} \Phi_j(\eta_j) \right)$$

and

$$(12.20) \quad \frac{1}{R_i} \mathcal{F}_j^0 \circ \eta_i = \left((T_{ij} \circ \eta_i) \cdot \left(\Phi_j \left(\frac{R_i}{R_j} T_{ij} \circ \eta_i \right) \right), \frac{R_j}{R_i} \Phi_j \left(\frac{R_i}{R_j} T_{ij} \circ \eta_i \right) \right).$$

From part (2) of Lemma 11.5, $\frac{R_i}{R_j} \leq \frac{20}{\Upsilon_0}$. Then the H_j -component of $\frac{1}{R_i} \mathcal{E}^0 - \frac{1}{R_i} \mathcal{F}^0 \circ \eta_i$ can be estimated in C^1 -norm by using (12.13). Note when we use the chain rule to estimate the second component of $\frac{1}{R_i} \mathcal{E}^0 - \frac{1}{R_i} \mathcal{F}^0 \circ \eta_i$, namely $\frac{R_j}{R_i} (\Phi_j(\eta_j) - \Phi_j(\frac{R_i}{R_j} T_{ij} \circ \eta_i))$, we differentiate Φ_j and this brings down a factor of $\frac{R_i}{R_j}$ when estimating norms. \square

Sublemma 12.21. — Given $\Sigma \in (0, \frac{1}{10})$, suppose that $|\eta_i(p)| < 8$ for some $p \in M$. Put $x = \mathcal{E}^0(p)$. For any $q \in M$, if $\mathcal{E}^0(q) \in B(x, \Sigma R_i)$ then $|\eta_i(p) - \eta_i(q)| < 20\Sigma$.

Proof. — We know that $\zeta_i(p) = 1$. By hypothesis, $|\mathcal{E}^0(p) - \mathcal{E}^0(q)| < \Sigma R_i$. In particular, $|\zeta_i(p) - \zeta_i(q)| < \Sigma$ and $|\zeta_i(p)\eta_i(p) - \zeta_i(q)\eta_i(q)| < \Sigma$. Then

$$(12.22) \quad \begin{aligned} |\eta_i(p) - \eta_i(q)| &= \frac{1}{\zeta_i(q)} |\zeta_i(q)\eta_i(p) - \zeta_i(q)\eta_i(q)| \\ &\leq \frac{1}{\zeta_i(q)} [|\zeta_i(p)\eta_i(p) - \zeta_i(q)\eta_i(q)| + |\zeta_i(p) - \zeta_i(q)| |\eta_i(p)|] \\ &\leq \frac{10\Sigma}{1 - \Sigma} \leq 20\Sigma. \end{aligned}$$

This proves the sublemma. \square

We now prove Lemma 12.7. We no longer fix $i \in I_{2\text{-stratum}}$. Given $x \in S_1$, choose $p \in A_1$ and $i \in I_{2\text{-stratum}}$ so that $\mathcal{E}^0(p) = x$ and $|\eta_i(p)| \leq 8$. Put $A_x^0 = \text{Im}(d\mathcal{F}_{\eta_i(p)}^0)$, a 2-plane in H , and let $A_x = x + A_x^0$ be the corresponding affine subspace through x . We first show that under the constraints $\Sigma_1 \leq \bar{\Sigma}_1(\Gamma_1, \mathcal{M})$, $\epsilon_2 \leq \bar{\epsilon}_2(\Gamma_1, \mathcal{M})$ and $\Lambda \leq \bar{\Lambda}(\Gamma_1, \mathcal{M})$, the triple (\tilde{S}_1, S_1, r_1) is a $(2, \Gamma_1)$ cloudy 2-manifold.

We verify condition (1) of Definition 20.2. Pick $x, y \in \tilde{S}_1$, and choose $p \in (\mathcal{E}^0)^{-1}(x) \cap \bigcup_{i \in I_{2\text{-stratum}}} |\eta_i|^{-1}[0, 8)$ (respectively $q \in (\mathcal{E}^0)^{-1}(y) \cap \bigcup_{i \in I_{2\text{-stratum}}} |\eta_i|^{-1}[0, 8)$) satisfying $r_1(x) = \Sigma_1 \mathbf{r}_p$ (respectively $r_1(y) = \Sigma_1 \mathbf{r}_q$).

We can assume that $\Lambda < \frac{1}{100}$. Suppose first that $d(p, q) \leq \frac{\mathbf{r}_p}{\Lambda}$. Then since \mathbf{r} is Λ -Lipschitz, we get $|\mathbf{r}_p - \mathbf{r}_q| \leq \mathbf{r}_p$, so in this case

$$(12.23) \quad |r_1(x) - r_1(y)| = \Sigma_1 |\mathbf{r}_p - \mathbf{r}_q| \leq \Sigma_1 \mathbf{r}_p = r_1(x).$$

Now suppose that $d(p, q) \geq 20\mathfrak{r}_p$. We claim that if Λ is sufficiently small then this implies that $d(p, q) \geq 19\mathfrak{r}_q$ as well. Suppose not. Then $20\mathfrak{r}_p \leq d(p, q) \leq 19\mathfrak{r}_q$, so $\frac{\mathfrak{r}_p}{\mathfrak{r}_q} \leq \frac{19}{20}$. On the other hand, since $|\mathfrak{r}_q - \mathfrak{r}_p| \leq \Lambda d(p, q)$, we also know that $\mathfrak{r}_q - \mathfrak{r}_p \leq \Lambda d(p, q) \leq 19\Lambda\mathfrak{r}_q$, so $\frac{\mathfrak{r}_p}{\mathfrak{r}_q} \geq 1 - 19\Lambda$. If Λ is sufficiently small then this is a contradiction.

Thus there are $i, j \in I_{2\text{-stratum}}$ such that $p \in |\eta_i|^{-1}[0, 8)$, $q \in |\eta_j|^{-1}[0, 8)$, $\zeta_i(p) = 1 = \zeta_j(q)$ and $\zeta_i(q) = 0 = \zeta_j(p)$. Then

$$(12.24) \quad \begin{aligned} |x - y| &= |\mathcal{E}^0(p) - \mathcal{E}^0(q)| \geq \max(\mathfrak{r}_{p_i}|\zeta_i(p) - \zeta_i(q)|, \mathfrak{r}_{p_j}|\zeta_j(p) - \zeta_j(q)|) \\ &= \max(\mathfrak{r}_{p_i}, \mathfrak{r}_{p_j}) \geq \frac{1}{2} \max(\mathfrak{r}_p, \mathfrak{r}_q) = \frac{\max(r_1(x), r_1(y))}{2\Sigma_1}. \end{aligned}$$

So $|r_1(x) - r_1(y)| \leq |x - y|$ provided $\Sigma_1 \leq \frac{1}{4}$. Thus condition (1) of Definition 20.2 will be satisfied.

We now verify condition (2) of Definition 20.2. Given $x \in S_1$, let $i \in I_{2\text{-stratum}}$ and $p \in M$ be such that $\mathcal{E}^0(p) = x$ and $|\eta_i(p)| \leq 7$. Taking $\Sigma = \frac{1}{100}$ in Sublemma 12.21, we have $\text{Im}(\mathcal{E}^0) \cap B(x, \frac{R_i}{100}) \subset \text{Im}(\mathcal{E}^0|_{|\eta_i|^{-1}[0, 7.2)})$. Thus we can restrict attention to the action of \mathcal{E}^0 on $|\eta_i|^{-1}[0, 7.2)$. Now $\text{Im}(\mathcal{F}^0|_{B(0, 7.2)})$ is the restriction to $B(0, 7.2)$ of the graph of a function $G_i^0 : H_i' \rightarrow (H_i')^\perp$, since $T_{ii} = \text{Id}$ and $\zeta_i|_{B(0, 7.2)} = 1$. Furthermore, in view of the universality of the functions $\{\Phi_j\}_{j \in J}$ and the bound on the cardinality of J , there are uniform C^1 -estimates on G_i^0 . Hence we can find Σ_1 (as a function of Γ_1 and \mathcal{M}) to ensure that $(\frac{1}{r_1(x)} \text{Im}(\mathcal{F}^0|_{B(0, 7.2)}), x)$ is $\frac{\Gamma_1}{2}$ -close in the pointed Hausdorff topology to $x + \text{Im}(d\mathcal{F}_p^0)$. Finally, if the parameter ϵ_2 of Sublemma 12.17 is sufficiently small then we can ensure that $(\frac{1}{r_1(x)} \text{Im}(\mathcal{E}^0), x)$ is Γ_1 -close in the pointed Hausdorff topology to $x + \text{Im}(D\mathcal{F}_p^0)$. Thus condition (2) of Definition 20.2 will be satisfied.

To finish the proof of Lemma 12.7, equation (12.8) is clearly satisfied if the parameter ϵ_2 of Sublemma 12.17 is sufficiently small. Equation (12.9) is equivalent to upper and lower bounds on the eigenvalues of the matrix $(\pi_{A_x^0} \circ D\mathcal{E}_p^0)(\pi_{A_x^0} \circ D\mathcal{E}_p^0)^*$, which acts on the two-dimensional space A_x^0 . In view of Sublemma 12.17 and the definition of A_x , it is sufficient to show upper and lower bounds on the eigenvalues of $D\mathcal{F}_{\eta_i(p)}^0(D\mathcal{F}_{\eta_i(p)}^0)^*$ acting on A_x^0 . In terms of the function G_i^0 , these are the same as the eigenvalues of $I_2 + ((DG_i^0)_{\eta_i(p)})^*(DG_i^0)_{\eta_i(p)}$, acting on \mathbb{R}^2 . The eigenvalues are clearly bounded below by one. In view of the C^1 -bounds on G_i^0 , there is an upper bound on the eigenvalues in terms of $\dim(H)$, which in turn is bounded above in terms of \mathcal{M} . This shows equation (12.9).

Finally, given $i \in I_{2\text{-stratum}}$, we can write $\frac{1}{R_i}\mathcal{F}^0$ on $\overline{B(0, 8)} \subset \mathbb{R}^2$ in the form $\frac{1}{R_i}\mathcal{F}^0 = (I, \frac{1}{R_i}\widehat{\mathcal{E}}_i^0)$ with respect to the orthogonal decomposition $H = H_i' \oplus (H_i')^\perp$. (Recall that \mathcal{F}^0 is defined in reference to the given value of i .) We use this to define $\widehat{\mathcal{E}}_i^0$. Equation (12.11) is a consequence of Sublemma 12.17. The last statement of Lemma 12.7 follows from the definition of A_x^0 .

12.4. Structure of \mathcal{E}^0 near the edge stratum. — Recall that $Q_2 = H_{0\text{-stratum}} \oplus H_{\text{slim}} \oplus H_{\text{edge}}$, and $\pi_2 : H \rightarrow Q_2$ is the orthogonal projection.

Put

$$(12.25) \quad \tilde{A}_2 = \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| \leq 8\Delta, \eta_{E'} \leq 8\Delta\}, \quad A_2 = \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| \leq 7\Delta, \eta_{E'} \leq 7\Delta\}$$

and

$$(12.26) \quad \tilde{S}_2 = (\pi_2 \circ \mathcal{E}^0)(\tilde{A}_2), \quad S_2 = (\pi_2 \circ \mathcal{E}^0)(A_2).$$

Let $\Sigma_2, \Gamma_2 > 0$ be new parameters. Define $r_2 : \tilde{S}_2 \rightarrow (0, \infty)$ by putting $r_2(x) = \Sigma_2 \tau_p$ for some $p \in (\pi_2 \circ \mathcal{E}^0)^{-1}(x) \cap \tilde{A}_2$.

The analog of Lemma 12.7 for the region near edge points is:

Lemma 12.27. — *There is a constant $\Omega_2 = \Omega_2(\mathcal{M})$ so that under the constraints $\Sigma_2 < \bar{\Sigma}_2(\Gamma_2, \mathcal{M})$, $\beta_E < \bar{\beta}_E(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$, $\sigma_E < \bar{\sigma}_E(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$, $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$, $\beta_1 < \bar{\beta}_1(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$, $\varsigma_{\text{slim}} < \bar{\varsigma}_{\text{slim}}(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$, $\varsigma_{0\text{-stratum}} < \bar{\varsigma}_{0\text{-stratum}}(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$ and $\Lambda < \bar{\Lambda}(\Gamma_2, \Sigma_2, \beta_2, \Delta, \mathcal{M})$, the following holds.*

- (1) *The triple (\tilde{S}_2, S_2, r_2) is a $(2, \Gamma_2)$ cloudy 1-manifold.*
- (2) *The affine subspaces $\{A_x\}_{x \in S_2}$ inherent in the definition of the cloudy 1-manifold can be chosen to have the following property. Pick $p \in A_2$ and put $x = (\pi_2 \circ \mathcal{E}^0)(p) \in S_2$. Let $A_x^0 \subset Q_2$ be the linear subspace parallel to A_x (i.e., $A_x = A_x^0 + x$) and let $\pi_{A_x^0} : H \rightarrow A_x^0$ denote orthogonal projection onto A_x^0 . Then*

$$(12.28) \quad \|D(\pi_2 \circ \mathcal{E}^0)_p - \pi_{A_x^0} \circ D(\pi_2 \circ \mathcal{E}^0)_p\| < \Gamma_2$$

and

$$(12.29) \quad \Omega_2^{-1} \|v\| \leq \|(\pi_{A_x^0} \circ D(\pi_2 \circ \mathcal{E}^0))(v)\| \leq \Omega_2 \|v\|$$

for every $v \in T_p M$ which is orthogonal to $\ker(\pi_{A_x^0} \circ D(\pi_2 \circ \mathcal{E}^0)_p)$.

- (3) *Given $i \in I_{\text{edge}}$, there is a smooth map $\widehat{\mathcal{E}}_i^0 : (\overline{B(0, 8\Delta)} \subset \mathbb{R}) \rightarrow (H_i^0)^\perp \cap Q_2$ such that*

$$(12.30) \quad \|D\widehat{\mathcal{E}}_i^0\| \leq \Omega_2 R_i$$

and on the subset $\{|\eta_i| \leq 8\Delta, \eta_{E'} \leq 8\Delta\}$, we have

$$(12.31) \quad \left\| \frac{1}{R_i} \pi_2 \circ \mathcal{E}^0 - \left(\eta_i, \frac{1}{R_i} \widehat{\mathcal{E}}_i^0 \circ \eta_i \right) \right\|_{C^1} < \Gamma_2.$$

Furthermore, if $x \in S_2$ then there are some $i \in I_{\text{edge}}$ and $p \in \{|\eta_i| \leq 7\Delta, \eta_{E'} \leq 7\Delta\}$ such that $x = (\pi_2 \circ \mathcal{E}^0)(p)$ and $A_x^0 = \text{Im} \left(I, \frac{1}{R_i} (D\widehat{\mathcal{E}}_i^0)_{\eta_i(p)} \right)$.

We omit the proof as it is similar to the proof of Lemma 12.7.

12.5. Structure of \mathcal{E}^0 near the slim 1-stratum. — Recall that $Q_3 = H_{0\text{-stratum}} \oplus H_{\text{slim}}$, and $\pi_3 : H \rightarrow Q_3$ is the orthogonal projection.

Put

$$(12.32) \quad \tilde{A}_3 = \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| \leq 8 \cdot 10^5 \Delta\}, \quad A_3 = \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| \leq 7 \cdot 10^5 \Delta\}$$

and

$$(12.33) \quad \tilde{S}_3 = (\pi_3 \circ \mathcal{E}^0)(\tilde{A}_3), \quad S_3 = (\pi_3 \circ \mathcal{E}^0)(A_3).$$

Let $\Sigma_3, \Gamma_3 > 0$ be new parameters. Define $r_3 : \tilde{S}_3 \rightarrow (0, \infty)$ by putting $r_3(x) = \Sigma_3 \mathfrak{r}_p$ for some $p \in (\pi_3 \circ \mathcal{E}^0)^{-1}(x) \cap \tilde{A}_3$.

The analog of Lemma 12.7 for the slim 1-stratum points is:

Lemma 12.34. — *There is a constant $\Omega_3 = \Omega_3(\mathcal{M})$ so that under the constraints $\Sigma_3 < \bar{\Sigma}_3(\Gamma_3, \mathcal{M})$, $\beta_E < \bar{\beta}_E(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$, $\sigma_E < \bar{\sigma}_E(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$, $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$, $\beta_1 < \bar{\beta}_1(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$, $\varsigma_{\text{slim}} < \bar{\varsigma}_{\text{slim}}(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$, $\varsigma_{0\text{-stratum}} < \bar{\varsigma}_{0\text{-stratum}}(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$ and $\Lambda < \bar{\Lambda}(\Gamma_3, \Sigma_3, \beta_2, \Delta, \mathcal{M})$, the following holds.*

- (1) *The triple (\tilde{S}_3, S_3, r_3) is a $(2, \Gamma_3)$ cloudy 1-manifold.*
- (2) *The affine subspaces $\{A_x\}_{x \in S_3}$ inherent in the definition of the cloudy 1-manifold can be chosen to have the following property. Pick $p \in A_3$ and put $x = (\pi_3 \circ \mathcal{E}^0)(p) \in S_3$. Let $A_x^0 \subset Q_3$ be the linear subspace parallel to A_x (i.e., $A_x = A_x^0 + x$) and let $\pi_{A_x^0} : H \rightarrow A_x^0$ denote orthogonal projection onto A_x^0 . Then*

$$(12.35) \quad \|D(\pi_3 \circ \mathcal{E}^0)_p - \pi_{A_x^0} \circ D(\pi_3 \circ \mathcal{E}^0)_p\| < \Gamma_3$$

and

$$(12.36) \quad \Omega_3^{-1} \|v\| \leq \|\pi_{A_x^0} \circ D(\pi_3 \circ \mathcal{E}^0)(v)\| \leq \Omega_3 \|v\|$$

for every $v \in T_p M$ which is orthogonal to $\ker(\pi_{A_x^0} \circ D(\pi_3 \circ \mathcal{E}^0)_p)$.

- (3) *Given $i \in I_{\text{slim}}$, there is a smooth map $\widehat{\mathcal{E}}_i^0 : (\overline{B(0, 8 \cdot 10^5 \Delta)} \subset \mathbb{R}) \rightarrow (H'_i)^\perp \cap Q_3$ such that*

$$(12.37) \quad \|D\widehat{\mathcal{E}}_i^0\| \leq \Omega_3 R_i$$

and on the subset $\{|\eta_i| \leq 8 \cdot 10^5 \Delta\}$, we have

$$(12.38) \quad \left\| \frac{1}{R_i} \pi_3 \circ \mathcal{E}^0 - \left(\eta_i, \frac{1}{R_i} \widehat{\mathcal{E}}_i^0 \circ \eta_i \right) \right\|_{C^1} < \Gamma_3.$$

Furthermore, if $x \in S_3$ then there are some $i \in I_{\text{slim}}$ and $p \in \{|\eta_i| \leq 7 \cdot 10^5 \Delta\}$ such that $x = (\pi_3 \circ \mathcal{E}^0)(p)$ and $A_x^0 = \text{Im} \left(I, \frac{1}{R_i} (D\widehat{\mathcal{E}}_i^0)_{\eta_i(p)} \right)$.

We omit the proof as it is similar to the proof of Lemma 12.7.

12.6. Structure of \mathcal{E}^0 near the 0-stratum. — The only information we will need near the 0-stratum is:

Lemma 12.39. — *For $i \in I_{0\text{-stratum}}$, the only nonzero component of the map $\pi_4 \circ \mathcal{E}^0 : M \rightarrow Q_4 = H_{0\text{-stratum}}$ in the region $\{\eta_i \in [\frac{3}{10}, \frac{8}{10}]\}$ is \mathcal{E}_i^0 , where it coincides with $(R_i \eta_i, R_i)$.*

13. Adjusting the map to Euclidean space

The main result of this section is the following proposition, which asserts that it is possible to adjust \mathcal{E}^0 slightly, to get a new map \mathcal{E} which is a submersion in different parts of M . In Section 14 this structure will yield compatible fibrations of different parts of M .

Let $c_{\text{adjust}} > 0$ be a parameter.

Proposition 13.1. — *Under the constraints imposed in this and prior sections, there is a smooth map $\mathcal{E} : M \rightarrow H$ with the following properties:*

(1) *For every $p \in M$,*

$$(13.2) \quad \|\mathcal{E}(p) - \mathcal{E}^0(p)\| < c_{\text{adjust}} \mathfrak{t}(p) \quad \text{and} \quad \|D\mathcal{E}_p - D\mathcal{E}_p^0\| < c_{\text{adjust}}.$$

(2) *For $j \in \{1, 2, 3\}$ the restriction of $\pi_j \circ \mathcal{E} : M \rightarrow Q_j$ to the region $U_j \subset M$ is a submersion to a k_j -manifold $W_j \subset Q_j$, where*

$$(13.3) \quad \begin{aligned} U_1 &= \bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| < 5\}, \\ U_2 &= \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 5\Delta, \eta_{E'} < 5\Delta\}, \\ U_3 &= \bigcup_{i \in I_{\text{slim}}} \{|\eta_i| < 5 \cdot 10^5 \Delta\} \end{aligned}$$

and $k_1 = 2, k_2 = k_3 = 1$.

We will use the following additional parameters in this section: $c_{2\text{-stratum}}, c_{\text{edge}}, c_{\text{slim}} > 0$ and $\Xi_i > 0$ for $i \in \{1, 2, 3\}$.

13.1. Overview of the proof of Proposition 13.1. — In certain regions of M , the map \mathcal{E}^0 defined in the previous section, as well as its composition with projection onto certain summands of H , behaves like a “rough fibration”. As indicated in the overview in Section 1.5, the next step is to modify the map \mathcal{E}^0 so as to promote these rough fibrations to honest fibrations, in such a way that they are compatible on their overlap. We will do this by producing a sequence of maps $\mathcal{E}^j : M \rightarrow H$, for $j \in \{1, 2, 3\}$, which are successive adjustments of the map \mathcal{E}^0 .

To construct the map \mathcal{E}^j from \mathcal{E}^{j-1} , $j \in \{1, 2, 3\}$, we will use the following procedure. We consider the orthogonal splitting $H = Q_j \oplus Q_j^\perp$ of H , and let

$\pi_j = \pi_{1,j} : H \rightarrow Q_j$, $\pi_j^\perp : H \rightarrow Q_j^\perp$ be the orthogonal projections. In Section 12 we considered a pair of subsets (\widetilde{A}_j, A_j) in M whose image (\widetilde{S}_j, S_j) under the composition $\pi_j \circ \mathcal{E}^{j-1}$ is a cloudy k_j -manifold in Q_j , in the sense of Definition 20.1 of Appendix B. We think of the restriction of \mathcal{E}^{j-1} to A_j as defining a “rough submersion” over the cloudy k_j -manifold (\widetilde{S}_j, S_j) . By Lemma 20.2, there is a k_j -dimensional manifold $W_j \subset Q_j$ near (\widetilde{S}_j, S_j) and a projection map P_j onto W_j , defined in a neighborhood \widehat{W}_j of W_j . Hence we have a well-defined map

$$(13.4) \quad H \supset \widehat{W}_j \times Q_j^\perp \xrightarrow{(P_j \circ \pi_j, \pi_j^\perp)} Q_j \oplus Q_j^\perp = H.$$

Then using a partition of unity, we blend the composition $(P_j \circ \pi_j, \pi_j^\perp) \circ \mathcal{E}^{j-1}$ with $\mathcal{E}^{j-1} : M \rightarrow H$ to obtain $\mathcal{E}^j : M \rightarrow H$. In fact, \mathcal{E}^j will be the postcomposition of \mathcal{E}^{j-1} with a map from H to itself.

We draw attention to two key features of the construction. First, in passing from \mathcal{E}^{j-1} to \mathcal{E}^j , we do not change it much. More precisely, at a point $p \in M$, we have $|\mathcal{E}^{j-1}(p) - \mathcal{E}^j(p)| < \text{const. } \mathbf{r}_p$ and $|D\mathcal{E}_p^{j-1} - D\mathcal{E}_p^j| < \text{const.}$ for some small constants. Second, the passage from \mathcal{E}^j to \mathcal{E}^{j-1} respects the submersions defined by \mathcal{E}^{j-1} .

13.2. Adjusting the map near the 2-stratum. — Our first adjustment step involves the 2-stratum.

We take $Q_1 = H$, $Q_1^\perp = \{0\}$, and we let $\widetilde{A}_1, A_1, \widetilde{S}_1, S_1$ and $r_1 : \widetilde{S}_1 \rightarrow (0, \infty)$ be as in Section 12.3.

Thus $(\widetilde{S}_1, S_1, r_1)$ is a $(2, \Gamma_1)$ cloudy 2-manifold by Lemma 12.7. By Lemma 20.2, there is a 2-manifold $W_1^0 \subset H$ so that the conclusion of Lemma 20.2 holds, where the parameter ϵ in the lemma is given by $\Xi_1 = \Xi_1(\Gamma_1)$. (We remark that W_1^0 will not be the same as the W_1 of Proposition 13.1, due to subsequent adjustments.) In particular, there is a well-defined nearest point projection

$$(13.5) \quad P_1 : N_{r_1}(S_1) = \widehat{W}_1 \rightarrow W_1^0,$$

where we are using the notation for variable thickness neighborhoods from Section 3.

We now define a certain cutoff function.

Lemma 13.6. — *There is a smooth function $\psi_1 : H \rightarrow [0, 1]$ with the following properties:*

$$(1) \quad \begin{aligned} (13.7) \quad & \psi_1 \circ \mathcal{E}^0 \equiv 1 \text{ in } \bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| < 6\} \text{ and} \\ & \psi_1 \circ \mathcal{E}^0 \equiv 0 \text{ outside } \bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| < 7\}. \end{aligned}$$

$$(2) \quad \text{supp}(\psi_1) \cap \text{im}(\mathcal{E}^0) \subset \widehat{W}_1.$$

(3) *There is a constant $\Omega'_1 = \Omega'_1(\mathcal{M})$ such that*

$$(13.8) \quad |(d\psi_1)_x| < \Omega'_1 x_\tau^{-1}$$

for all $x \in \text{im}(\mathcal{E}^0)$.

Proof. — Let $\psi_1 : H \rightarrow [0, 1]$ be given by

$$(13.9) \quad \psi_1(x) = 1 - \Phi_{\frac{1}{2},1} \left(\sum_{\{i \in I_{2\text{-stratum}} \mid x''_i > 0\}} \Phi_{6,6.5} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) \right) \right).$$

For each $i \in I_{2\text{-stratum}}$, the function $x \mapsto \Phi_{6,6.5} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) \right)$ is well-defined and smooth in the set $\{x''_i > 0\}$, with support contained in the set $\{x''_i \geq \frac{1}{2}R_i\}$; so extending it by zero defines a smooth function on H . Hence ψ_1 is smooth.

To prove part (1), suppose that $i \in I_{2\text{-stratum}}$ and $|\eta_i(p)| < 6$. Then $\zeta_i(p) = 1$ and $|\eta_i| < 6$. Putting $x = \mathcal{E}^0(p)$, we have

$$(13.10) \quad x_i = (x'_i, x''_i) = \mathcal{E}^0_i(p) = (R_i \zeta_i(p) \eta_i(p), R_i \zeta_i(p)),$$

so $x''_i = R_i$ and $\frac{|x'_i|}{x''_i} \in [0, 6)$. Hence

$$(13.11) \quad \Phi_{6,6.5} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) \right) = 1,$$

so $\psi_1(x) = 1$.

Suppose now that $|\eta_i(p)| \geq 7$ for every $i \in I_{2\text{-stratum}}$. Putting $x = \mathcal{E}^0(p)$, for each $i \in I_{2\text{-stratum}}$ we claim that

$$(13.12) \quad \Phi_{6,6.5} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) \right) = 0;$$

otherwise we would have $|x'_i| < 6.5 x''_i$ and $x''_i \geq R_i/2$, which contradicts our assumption on p . It follows that $\psi_1(x) = 0$. This proves part (1).

To prove part (2), suppose $x = \mathcal{E}^0(p)$ and $\psi_1(x) > 0$. Then from part (1), $|\eta_i(p)| < 7$ for some $i \in I_{2\text{-stratum}}$. Therefore, $p \in A_1$ and $x \in \mathcal{E}^0(A_1) = S_1 \subset \widehat{W}_1$, so part (2) follows.

To prove part (3), suppose that $x = \mathcal{E}^0(p)$. If $x''_i > 0$ then $\zeta_i(p) > 0$, so the number of such indices $i \in I_{2\text{-stratum}}$ is bounded by the multiplicity of the 2-stratum cover; for the remaining indices $j \in I_{2\text{-stratum}}$, the quantity $1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_j}{R_j} \right)$ vanishes near x . Thus by the chain rule, it suffices to bound the differential of

$$(13.13) \quad \Phi_{6,6.5} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) \right)$$

for each $i \in I_{2\text{-stratum}}$ for which $x''_i > 0$. But the differential is nonzero only when $\frac{|x'_i|}{x''_i} \leq 6.5$ and $\frac{x''_i}{R_i} \geq \frac{1}{2}$. In this case, R_i will be comparable to x_τ and the estimate (13.8) follows easily. \square

Define $\Psi_1 : H \rightarrow H$ by $\Psi_1(x) = x$ if $x \notin \widehat{W}_1$ and

$$(13.14) \quad \Psi_1(x) = \psi_1(x)P_1(x) + (1 - \psi_1(x))x$$

otherwise. Put $\mathcal{E}^1 = \Psi_1 \circ \mathcal{E}^0$.

Lemma 13.15. — *Under the constraints $\Sigma_1 < \overline{\Sigma}_1(\Omega_1, c_{2\text{-stratum}})$, $\Gamma_1 < \overline{\Gamma}_1(\Omega_1, c_{2\text{-stratum}})$ and $\Xi_1 < \overline{\Xi}_1(c_{2\text{-stratum}})$, we have:*

(1) \mathcal{E}^1 is smooth.

(2) For all $p \in M$,

$$(13.16) \quad \|\mathcal{E}^1(p) - \mathcal{E}^0(p)\| < c_{2\text{-stratum}} \mathfrak{r}(p) \quad \text{and} \quad \|D\mathcal{E}_p^1 - D\mathcal{E}_p^0\| < c_{2\text{-stratum}}.$$

(3) The restriction of \mathcal{E}^1 to $\bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| < 6\}$ is a submersion to W_1^0 .

Proof. — That \mathcal{E}^1 is smooth follows from part (2) of Lemma 13.6.

Given $p \in M$, put $x = \mathcal{E}^0(p)$. We have

$$(13.17) \quad \mathcal{E}^1(p) - \mathcal{E}^0(p) = \psi_1(x) (P_1(x) - x).$$

Now $|\psi_1(x)| \leq 1$. From Lemma 20.2 (1), $|P_1(x) - x| \leq \Xi_1 r_1(x)$. From Sublemma 12.21, we can assume that $r_1(x) \leq 10\mathfrak{r}_p$. This gives the first equation in (13.16).

Next,

$$(13.18) \quad \begin{aligned} D\mathcal{E}_p^1 - D\mathcal{E}_p^0 &= (D\psi_1)_x (P_1(x) - x) + \psi_1(x) ((DP_1)_x \circ D\mathcal{E}_p^0 - D\mathcal{E}_p^0) \\ &= (D\psi_1)_x (P_1(x) - x) + \psi_1(x) ((DP_1)_x - \pi_{A_x^0}) \circ D\mathcal{E}_p^0 + \\ &\quad \psi_1(x) (\pi_{A_x^0} \circ D\mathcal{E}_p^0 - D\mathcal{E}_p^0). \end{aligned}$$

Equation (13.8) gives a bound on $|(D\psi_1)_x|$. Lemma 20.2 (1) gives a bound on $|P_1(x) - x|$. Lemma 20.2 (7) gives a bound on $|(DP_1)_x - \pi_{A_x^0}|$. Lemma 12.5 gives a bound on $|D\mathcal{E}_p^0|$. Equation (12.8) gives a bound on $|\pi_{A_x^0} \circ D\mathcal{E}_p^0 - D\mathcal{E}_p^0|$. The second equation in (13.16) follows from these estimates.

Finally, the restriction of \mathcal{E}^1 to $\bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| < 6\}$ equals $P_1 \circ \mathcal{E}^0$. For $p \in \bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| < 6\}$, put $x = \mathcal{E}^0(p)$. Then

$$(13.19) \quad D(P_1 \circ \mathcal{E}^0)_p = \pi_{A_x^0} \circ d\mathcal{E}_p^0 + ((DP_1)_x - \pi_{A_x^0}) \circ D\mathcal{E}_p^0.$$

Using (12.9) and Lemma 20.2 (7), if Ξ_1 is sufficiently small then $D(P_1 \circ \mathcal{E}^0)_p$ maps onto $(TW_1^0)_{P_1(x)}$. This proves the lemma. \square

13.3. Adjusting the map near the edge points. — Our second adjustment step involves the region near the edge points.

Recall that $Q_2 = H_{0\text{-stratum}} \oplus H_{\text{slim}} \oplus H_{\text{edge}}$ and $\pi_2 : H \rightarrow Q_2$ is orthogonal projection. We let $\tilde{A}_2, A_2, \tilde{S}_2, S_2$ and $r_2 : \tilde{S}_2 \rightarrow (0, \infty)$ be as in Section 12.4.

Thus (\tilde{S}_2, S_2, r_2) is a $(2, \Gamma_2)$ cloudy 1-manifold by Lemma 12.27. By Lemma 20.2, there is a 1-manifold $W_2^0 \subset Q_2$ so that the conclusion of Lemma 20.2 holds, where the parameter ϵ in the lemma is given by $\Xi_2 = \Xi_2(\Gamma_2)$. (We remark that W_2^0 will

not be the same as the W_2 of Proposition 13.1, due to subsequent adjustments.) In particular, there is a well-defined nearest point projection

$$(13.20) \quad P_2 : N_{r_2}(S_2) = \widehat{W}_2 \rightarrow W_2^0,$$

where we are using the notation for variable thickness neighborhoods from Section 3.

Lemma 13.21. — *Under the constraint $c_{2\text{-stratum}} < \bar{c}_{2\text{-stratum}}$, there is a smooth function $\psi_2 : \{x_\tau > 0\} \rightarrow [0, 1]$ with the following properties:*

$$(1) \quad (13.22) \quad \psi_2 \circ \mathcal{E}^1 \equiv 1 \text{ in } \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6\Delta, \eta_{E'} < 6\Delta\} \text{ and}$$

$$\psi_2 \circ \mathcal{E}^1 \equiv 0 \text{ outside } \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 7\Delta, \eta_{E'} < 7\Delta\}.$$

$$(2) \quad \text{supp}(\psi_2) \cap \text{im}(\mathcal{E}^1) \subset \widehat{W}_2 \times Q_2^\perp.$$

(3) *There is a constant $\Omega'_2 = \Omega'_2(\mathcal{M})$ such that*

$$(13.23) \quad |(D\psi_2)_x| < \Omega'_2 x_\tau^{-1}$$

for all $x \in \text{im}(\mathcal{E}^1)$.

Proof. — If the parameter $c_{2\text{-stratum}}$ is sufficiently small then $\mathcal{E}^1(p) \in \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6\Delta, \eta_{E'} < 6\Delta\}$ implies that $\mathcal{E}^0(p) \in \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6.1\Delta, \eta_{E'} < 6.1\Delta\}$, and $\mathcal{E}^1(p) \notin \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 7\Delta, \eta_{E'} < 7\Delta\}$ implies that $\mathcal{E}^0(p) \notin \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6.9\Delta, \eta_{E'} < 6.9\Delta\}$.

In analogy to (9.28), put

$$(13.24) \quad z_{\text{edge}} = 1 - \Phi_{\frac{1}{2}, 1} \left(\sum_{i \in I_{\text{edge}}} \frac{x''_i}{R_i} \right).$$

Define $\psi_2 : \{x_\tau > 0\} \rightarrow [0, 1]$ by

$$(13.25) \quad \psi_2(x) = 1 - \Phi_{\frac{1}{2}, 1} \left(\sum_{\{i \in I_{\text{edge}} \mid x''_i > 0\}} \Phi_{6.1\Delta, 6.5\Delta} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2}, 1} \left(\frac{x''_i}{R_i} \right) \right) \right. \\ \left. \left[\left(1 - \Phi_{\frac{1}{4}, \frac{1}{2}} \left(\frac{x''_{E'}}{x_\tau} \right) \right) \Phi_{6.1\Delta, 6.5\Delta} \left(\frac{|x'_{E'}|}{x''_{E'}} \right) + 10 \left(\frac{x''_i}{R_i} z_{\text{edge}} - \frac{x''_{E'}}{x_\tau} \right) \right] \right).$$

It is easy to see that ψ_2 is smooth.

To prove part (1), it is enough to show that

$$(13.26) \quad \psi_2 \circ \mathcal{E}^0 \equiv 1 \text{ in } \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6.1\Delta, \eta_{E'} < 6.1\Delta\} \text{ and}$$

$$\psi_2 \circ \mathcal{E}^0 \equiv 0 \text{ outside } \bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6.9\Delta, \eta_{E'} < 6.9\Delta\}.$$

Suppose that $i \in I_{\text{edge}}$, $|\eta_i(p)| < 6.1\Delta$ and $\eta_{E'}(p) < 6.1\Delta$. Put $x = \mathcal{E}^0(p)$. Recall that $x''_i = R_i \zeta_i(p)$, where ζ_i is given in (9.19) with $p \rightsquigarrow p_i$, and $x''_{E'} = \tau_p \zeta_{E'}(p)$,

where $\zeta_{E'}$ is the expression in (9.29). Hence

$$(13.27) \quad \begin{aligned} \frac{x''_i}{R_i} &= \zeta_i(p) = 1, \\ 1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) &= 1, \\ \Phi_{6.1\Delta,6.5\Delta} \left(\frac{|x''_i|}{x''_i} \right) &= \Phi_{6.1\Delta,6.5\Delta} (|\eta_i(p)|) = 1. \end{aligned}$$

If $\frac{x''_{E'}}{x_\tau} = \zeta_{E'}(p) \geq \frac{1}{2}$ then

$$(13.28) \quad \begin{aligned} 1 - \Phi_{\frac{1}{4},\frac{1}{2}} \left(\frac{x''_{E'}}{x_\tau} \right) &= 1, \\ \Phi_{6.1\Delta,6.5\Delta} \left(\frac{|x'_{E'}|}{x''_{E'}} \right) &= \Phi_{6.1\Delta,6.5\Delta} (|\eta_{E'}(p)|) = 1, \\ \frac{x''_i}{R_i} z_{\text{edge}} - \frac{x''_{E'}}{x_\tau} &= \zeta_i(p) \zeta_{\text{edge}}(p) - \zeta_{E'}(p) = \zeta_{\text{edge}}(p) - \zeta_{E'}(p) \geq 0. \end{aligned}$$

If $\frac{x''_{E'}}{x_\tau} = \zeta_{E'}(p) < \frac{1}{2}$ then $\left(1 - \Phi_{\frac{1}{4},\frac{1}{2}} \left(\frac{x''_{E'}}{x_\tau} \right)\right) \Phi_{6.1\Delta,6.5\Delta} \left(\frac{|x'_{E'}|}{x''_{E'}} \right) \geq 0$ and

$$(13.29) \quad \frac{x''_i}{R_i} z_{\text{edge}} - \frac{x''_{E'}}{x_\tau} = \zeta_i(p) \zeta_{\text{edge}}(p) - \zeta_{E'}(p) = 1 - \zeta_{E'}(p) \geq \frac{1}{2}.$$

In either case, the argument of $\Phi_{\frac{1}{2},1}$ in (13.25) is bounded below by one and so $\psi_2(x) = 1$.

Now suppose that for all $i \in I_{\text{edge}}$, either $\zeta_i(p) = 0$, or $\zeta_i(p) > 0$ and $|\eta_i(p)| \geq 6.9\Delta$, or $\zeta_i(p) > 0$ and $|\eta_i(p)| < 6.9\Delta$ and $\eta_{E'}(p) \geq 6.9\Delta$. If $\zeta_i(p) = 0$, or $\zeta_i(p) > 0$ and $|\eta_i(p)| \geq 6.9\Delta$, then

$$(13.30) \quad \Phi_{6.1\Delta,6.5\Delta} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2},1} \left(\frac{x''_i}{R_i} \right) \right) = \Phi_{6.1\Delta,6.5\Delta} (|\eta_i(p)|) \cdot \left(1 - \Phi_{\frac{1}{2},1} (\zeta_i(p)) \right) = 0.$$

If $|\eta_i(p)| < 6.9\Delta$ and $\eta_{E'}(p) \geq 6.9\Delta$ then

$$(13.31) \quad \left(1 - \Phi_{\frac{1}{4},\frac{1}{2}} \left(\frac{x''_{E'}}{x_\tau} \right) \right) \Phi_{6.1\Delta,6.5\Delta} \left(\frac{|x'_{E'}|}{x''_{E'}} \right) = \left(1 - \Phi_{\frac{1}{4},\frac{1}{2}} (\zeta_{E'}(p)) \right) \cdot \Phi_{6.1\Delta,6.5\Delta} (|\eta_{E'}(p)|) = 0.$$

and

$$(13.32) \quad \begin{aligned} \frac{x''_i}{R_i} z_{\text{edge}} - \frac{x''_{E'}}{x_\tau} &= \zeta_i(p) \zeta_{\text{edge}}(p) - \zeta_{E'}(p) \\ &= \Phi_{8\Delta,9\Delta}(\eta_{E'}(p)) \cdot \zeta_{\text{edge}}(p) - \Phi_{\frac{2}{10}\Delta, \frac{3}{10}\Delta, 8\Delta, 9\Delta}(\eta_{E'}(p)) \cdot \zeta_{\text{edge}}(p) = 0. \end{aligned}$$

Hence $\psi_2(x) = 0$.

This proves part (1) of the lemma.

The proof of the rest of the lemma is similar to that of Lemma 13.6. \square

We can assume that $\widehat{W}_2 \subset \{x_\tau > 0\}$. Define $\Psi_2 : \{x_\tau > 0\} \rightarrow \{x_\tau > 0\}$ by $\Psi_2(x) = x$ if $\pi_2(x) \notin \widehat{W}_2$ and

$$(13.33) \quad \Psi_2(x) = (\psi_2(x)P_2(\pi_2(x)) + (1 - \psi_2(x))\pi_2(x), \pi_2^\perp(x))$$

otherwise. Put $\mathcal{E}^2 = \Psi_2 \circ \mathcal{E}^1$.

Lemma 13.34. — *Under the constraints $\Sigma_2 < \overline{\Sigma}_2(\Omega_2, c_{\text{edge}})$, $\Gamma_2 < \overline{\Gamma}_2(\Omega_2, c_{\text{edge}})$, $\Xi_2 < \overline{\Xi}_2(c_{\text{edge}})$ and $c_{2\text{-stratum}} < \overline{c}_{2\text{-stratum}}(c_{\text{edge}})$, we have:*

(1) \mathcal{E}^2 is smooth.

(2) For all $p \in M$,

$$(13.35) \quad \|\mathcal{E}^2(p) - \mathcal{E}^0(p)\| < c_{\text{edge}} \mathfrak{r}(p) \quad \text{and} \quad \|D\mathcal{E}_p^2 - D\mathcal{E}_p^0\| < c_{\text{edge}}.$$

(3) The restriction of $\pi_2 \circ \mathcal{E}^2$ to $\bigcup_{i \in I_{\text{edge}}} \{|\eta_i| < 6\Delta, \eta_{E'} < 6\Delta\}$ is a submersion to W_2^0 .

Proof. — As in the proof of Lemma 13.15, \mathcal{E}^2 is smooth and we can ensure that

$$(13.36) \quad \|\mathcal{E}^2(p) - \mathcal{E}^1(p)\| < \frac{1}{2}c_{\text{edge}} \mathfrak{r}(p) \quad \text{and} \quad \|D\mathcal{E}_p^2 - D\mathcal{E}_p^1\| < \frac{1}{2}c_{\text{edge}}.$$

Along with (13.16), part (2) of the lemma follows.

The proof of part (3) is similar to that of Lemma 13.15 (3). We omit the details. \square

13.4. Adjusting the map near the slim 1-stratum. — Our third adjustment step involves the slim stratum.

Recall that $Q_3 = H_{0\text{-stratum}} \oplus H_{\text{slim}}$ and $\pi_3 : H \rightarrow Q_3$ is orthogonal projection. We let $\tilde{A}_3, A_3, \tilde{S}_3, S_3$ and $r_3 : \tilde{S}_3 \rightarrow (0, \infty)$ be as in Section 12.5.

Thus (\tilde{S}_3, S_3, r_3) is a $(2, \Gamma_3)$ cloudy 1-manifold by Lemma 12.34. By Lemma 20.2, there is a 1-manifold $W_3^0 \subset Q_3$ so that the conclusion of Lemma 20.2 holds, where the parameter ϵ in the lemma is given by $\Xi_3 = \Xi_3(\Gamma_3)$. In particular, there is a well-defined nearest point projection

$$(13.37) \quad P_3 : N_{r_3}(S_3) = \widehat{W}_3 \rightarrow W_3^0,$$

where we are using the notation for variable thickness neighborhoods from Section 3.

Lemma 13.38. — *Under the constraint $c_{\text{edge}} < \overline{c}_{\text{edge}}$, there is a smooth function $\psi_3 : H \rightarrow [0, 1]$ with the following properties:*

$$(1) \quad (13.39) \quad \begin{aligned} \psi_3 \circ \mathcal{E}^2 &\equiv 1 \text{ in } \bigcup_{i \in I_{\text{slim}}} \{|\eta_i| < 6 \cdot 10^5 \Delta\} \text{ and} \\ \psi_3 \circ \mathcal{E}^2 &\equiv 0 \text{ outside } \bigcup_{i \in I_{\text{slim}}} \{|\eta_i| < 7 \cdot 10^5 \Delta\}. \end{aligned}$$

$$(2) \quad \text{supp}(\psi_3) \cap \text{im}(\mathcal{E}^2) \subset \widehat{W}_3 \times Q_3^\perp.$$

(3) There is a constant $\Omega'_3 = \Omega'_3(\mathcal{M})$ such that

$$(13.40) \quad |(D\psi_3)_x| < \Omega'_3 x_\tau^{-1}$$

for all $x \in \text{im}(\mathcal{E}^2)$.

Proof. — Let $\psi_3 : H \rightarrow [0, 1]$ be given by

$$(13.41) \quad \psi_3(x) = 1 - \Phi_{\frac{1}{2}, 1} \left(\sum_{\{i \in I_{\text{slim}} \mid x''_i > 0\}} \Phi_{6.1 \cdot 10^5 \Delta, 6.5 \cdot 10^5 \Delta} \left(\frac{|x'_i|}{x''_i} \right) \cdot \left(1 - \Phi_{\frac{1}{2}, 1} \left(\frac{x''_i}{R_i} \right) \right) \right).$$

The rest of the proof is similar to that of Lemma 13.21. We omit the details. \square

Define $\Psi_3 : H \rightarrow H$ by $\Psi_3(x) = x$ if $\pi_3(x) \notin \widehat{W}_3$ and

$$(13.42) \quad \Psi_3(x) = (\psi_3(x)P_3(\pi_3(x)) + (1 - \psi_3(x))\pi_3(x), \pi_3^\perp(x))$$

otherwise. Put $\mathcal{E}^3 = \Psi_3 \circ \mathcal{E}^2$.

Lemma 13.43. — Under the constraints $\Sigma_3 < \overline{\Sigma}_3(\Omega_3, c_{\text{slim}})$, $\Gamma_3 < \overline{\Gamma}_3(\Omega_3, c_{\text{slim}})$, $\Xi_3 < \overline{\Xi}_3(c_{\text{slim}})$ and $c_{\text{edge}} < \overline{c}_{\text{edge}}(c_{\text{slim}})$, we have:

(1) \mathcal{E}^3 is smooth.

(2) For all $p \in M$,

$$(13.44) \quad \|\mathcal{E}^3(p) - \mathcal{E}^0(p)\| < c_{\text{slim}} \mathfrak{r}(p) \quad \text{and} \quad \|D\mathcal{E}_p^3 - D\mathcal{E}_p^0\| < c_{\text{slim}}.$$

(3) The restriction of $\pi_3 \circ \mathcal{E}^3$ to $\bigcup_{i \in I_{\text{slim}}} \{|\eta_i| < 6 \cdot 10^5 \Delta\}$ is a submersion to W_3^0 .

Proof. — The proof is similar to that of Lemma 13.34. We omit the details. \square

13.5. Proof of Proposition 13.1. — Note from (13.42) that Ψ_3 can be factored as $\Psi_3^{Q_2} \times I_{Q_2^\perp}$ for some $\Psi_3^{Q_2} : Q_2 \rightarrow Q_2$. In particular, $\pi_2 \circ \Psi_3 = \Psi_3^{Q_2} \circ \pi_2$.

Put $\mathcal{E} = \mathcal{E}^3$, $c_{\text{adjust}} = c_{\text{slim}}$ and

$$(13.45) \quad \begin{aligned} W_1 &= (\Psi_3 \circ \Psi_2)(W_1^0) \cap \bigcup_{i \in I_{2\text{-stratum}}} \{y \in H : y''_i > .9R_i, |y'_i| < 5.5R_i\}, \\ W_2 &= \Psi_3^{Q_2}(W_2^0) \cap \bigcup_{i \in I_{\text{edge}}} \{y \in Q_2 : y''_i > .9R_i, |y'_i| < 5.5\Delta R_i, y_\tau > 0, y_{E'} < 5.5\Delta y_\tau\}, \\ W_3 &= W_3^0 \cap \bigcup_{i \in I_{\text{slim}}} \{y \in Q_3 : y''_i > .9R_i, |y'_i| < 5.5 \cdot 10^5 \Delta R_i\}. \end{aligned}$$

The smoothness of \mathcal{E} follows from part (1) of Lemma 13.43. Part (1) of Proposition 13.1 follows from part (2) of Lemma 13.43.

Lemma 13.46. — W_i is a k_i -manifold.

Proof. — We will show that W_1 is a 2-manifold; the proofs for W_2 and W_3 are similar.

Choose $x \in W_1$. For some $i \in I_{2\text{-stratum}}$, we have $x''_i > .9R_i$ and $|x'_i| < 5.5R_i$. Putting

$$(13.47) \quad V_i = W_1 \cap \{y \in H : y''_i > .9R_i, |y'_i| < 5.5R_i\}$$

gives a neighborhood of x in W_1 . As $(\pi_{H'_i}, \pi_{H''_i}) \circ (\Psi_3 \circ \Psi_2) = (\pi_{H'_i}, \pi_{H''_i})$, it follows that V_i is the image, under $\Psi_3 \circ \Psi_2$, of the 2-manifold

$$(13.48) \quad V_i^0 = W_1^0 \cap \{y \in H : y''_i > .9R_i, |y'_i| < 5.5R_i\}.$$

If we can show that $\pi_{H'_i}$ maps V_i^0 diffeomorphically to its image in H'_i then V_i^0 will be a graph over a domain in H'_i , and the same will be true for V_i .

In view of (13.16) and the definition of \mathcal{E}^0 , if $c_{2\text{-stratum}}$ is sufficiently small then we are ensured that $V_i^0 = \mathcal{E}^1(\{|\eta_i| < 7\}) \cap \{y \in H : |y'_i| < 5.5R_i\}$. From Lemma 12.7 (3), Lemma 20.2 (3) and Lemma 20.2 (5), if Ξ_1 is sufficiently small then we are ensured that $\pi_{H'_i}$ restricts to a proper surjective local diffeomorphism from V_i^0 to $B(0, 5.5R_i) \subset H'_i$. Hence V_i^0 is a proper covering space of $B(0, 5.5R_i) \subset H'_i$ and so consists of a finite number of connected components, each mapping diffeomorphically under π'_i to $B(0, 5.5R_i) \subset H'_i$. It remains to show that there is only one connected component.

If V_i^0 has more than one connected component then there are $y_1, y_2 \in V_i^0 \cap \pi_{H'_i}^{-1}(0)$ with $y_1 \neq y_2$. We can write $y_1 = \mathcal{E}^1(p_1)$ and $y_2 = \mathcal{E}^1(p_2)$ for some $p_1, p_2 \in \{|\eta_i| < 7\}$. We claim that there is a smooth path γ in M from p_1 to p_2 so that $\mathcal{E}^1 \circ \gamma$ lies within $B(y_1, \frac{1}{10}R_i)$. To see this, we first note that if Γ_1 and $c_{2\text{-stratum}}$ are sufficiently small then Lemma 12.7 (3) and (13.16) ensure that $|\eta_i(p_1)| \ll 1$ and $|\eta_i(p_2)| \ll 1$, as otherwise we would contradict the assumption that $(y_1)'_i = (y_2)'_i = 0$. Let $\hat{\gamma}$ be a straight line from $\eta_i(p_1)$ to $\eta_i(p_2)$. Relative to the fiber bundle structure defined by η_i (see Lemma 8.4), let γ_1 be a lift of $\hat{\gamma}$, with initial point p_1 . Let γ_2 be a curve in the S^1 -fiber containing p_2 , going from the endpoint of γ_1 to p_2 . Let γ be a smooth concatenation of γ_1 and γ_2 . Then $\eta_i \circ \gamma$ lies in a ball whose diameter is much smaller than one. If Γ_1 and $c_{2\text{-stratum}}$ are sufficiently small then Lemma 12.7 (3) and (13.16) ensure that $\mathcal{E}^1 \circ \gamma$ lies in a ball whose diameter is much smaller than R_i .

On the other hand, since p_1 and p_2 lie in different connected components of V_i^0 , any curve in W_1^0 from p_1 to p_2 must go from p_1 to $\{y \in H : |y'_i| = R_i\}$. This is a contradiction.

Thus V_i^0 is connected and W_1 is a manifold. □

Recall the definition of U_1 from Proposition 13.1. By Lemma 13.15 (3), the restriction of \mathcal{E}^1 to U_1 is a submersion from U_1 to W_1^0 . From Lemma 12.7 (3) and (13.44), if Γ_1 and c_{slim} are sufficiently small then $\mathcal{E} = \Psi_3 \circ \Psi_2 \circ \mathcal{E}^1$ maps U_1 to $W_1 \subset (\Psi_3 \circ \Psi_2)(W_1^0)$. To see that it is a submersion, suppose that $|\eta_i(p)| < 5$ for some $i \in I_{2\text{-stratum}}$. Put $x^0 = \mathcal{E}^0(p)$ and $x = \mathcal{E}(p)$. Note that $x'_i = (x_0)'_i$. From Lemma 12.7 (3) and Lemma 20.2 (3), if Ξ_1 is sufficiently small then we are ensured that $(D\pi_{H'_i})_{x^0} \circ D\mathcal{E}^0_p$

maps onto $T_{(x^0)_i} H'_i \cong \mathbb{R}^2$. Then $(D\pi_{H'_i})_x \circ D\mathcal{E}_p = (D\pi_{H'_i})_x \circ D(\Psi_3 \circ \Psi_2)_{x^0} \circ D\mathcal{E}_p^0 = (D\pi_{H'_i})_{x^0} \circ D\mathcal{E}_p^0$ maps onto $T_{x'_i} H'_i \cong \mathbb{R}^2$. Thus $D\mathcal{E}_p$ must map $T_p M$ onto $T_x W_1$, showing that \mathcal{E} is a submersion near p .

Next, by Lemma 13.34 (3), the restriction of $\pi_2 \circ \mathcal{E}^2$ to U_2 is a submersion from U_2 to W_2^0 . Lemma 12.27 (3) and (13.44) imply that if Γ_2 and c_{slim} are sufficiently small then $\pi_2 \circ \mathcal{E} = \pi_2 \circ \Psi_3 \circ \mathcal{E}^2 = \Psi_3^{Q_3} \circ \pi_2 \circ \mathcal{E}^2$ maps U_2 to $W_2 \subset \Psi_3^{Q_3}(W_2^0)$. By a similar argument to the preceding paragraph, the restriction of $\pi_2 \circ \mathcal{E}$ to U_2 is a submersion to W_2 .

Finally, by Lemma 13.43 (3), the restriction of $\pi_3 \circ \mathcal{E} = \pi_3 \circ \mathcal{E}^3$ to U_3 is a submersion to $W_3 = W_3^0$. This proves Proposition 13.1.

14. Extracting a good decomposition of M

In this section we will use the map \mathcal{E} to find a decomposition of M into fibered pieces which are compatible along the intersections:

Proposition 14.1. — *There is a decomposition*

$$(14.2) \quad M = M^{0\text{-stratum}} \cup M^{\text{slim}} \cup M^{\text{edge}} \cup M^{2\text{-stratum}}$$

into compact domains with disjoint interiors, where each connected component of M^{slim} , M^{edge} , or $M^{2\text{-stratum}}$ may be endowed with a fibration structure, such that:

- (1) $M^{0\text{-stratum}}$ and M^{slim} are domains with smooth boundary, while M^{edge} and $M^{2\text{-stratum}}$ are smooth manifolds with corners, each point of which has a neighborhood diffeomorphic to $\mathbb{R}^{3-k} \times [0, \infty)^k$ for some $k \leq 2$.
- (2) Connected components of $M^{0\text{-stratum}}$ are diffeomorphic to one of the following: $S^1 \times S^2$, $S^1 \times_{\mathbb{Z}_2} S^2 = \mathbb{R}P^3 \# \mathbb{R}P^3$, T^3/Γ (where Γ is a finite subgroup of $\text{Isom}^+(T^3)$ which acts freely on T^3), S^3/Γ (where Γ is a finite subgroup of $\text{Isom}^+(S^3)$ which acts freely on S^3), a solid torus $S^1 \times D^2$, a twisted line bundle $S^2 \times_{\mathbb{Z}_2} I$ over $\mathbb{R}P^2$, or a twisted line bundle $T^2 \times_{\mathbb{Z}_2} I$ over a Klein bottle.
- (3) The components of M^{slim} have a fibration with S^2 -fibers or T^2 -fibers.
- (4) Components of M^{edge} are diffeomorphic (as manifolds with corners) to a solid torus $S^1 \times D^2$ or $I \times D^2$, and have a fibration with D^2 fibers.
- (5) $M^{2\text{-stratum}}$ is a smooth domain with corners with a smooth S^1 -fibration; in particular the S^1 -fibration is compatible with any corners.
- (6) Each fiber of the fibration $M^{\text{edge}} \rightarrow B^{\text{edge}}$, lying over a boundary point of the base B^{edge} , is contained in the boundary of $M^{0\text{-stratum}}$ or the boundary of M^{slim} .
- (7) The part of ∂M^{edge} which carries an induced S^1 -fibration is contained in $M^{2\text{-stratum}}$, and the S^1 -fibration induced from M^{edge} agrees with the one inherited from $M^{2\text{-stratum}}$.

To prove the proposition, we show that the submersions identified in Proposition 13.1 become fibrations, when restricted to appropriate subsets. Using this, we remove fibered regions around successive strata in the following order: 0-stratum, slim stratum, the edge region and the 2-stratum. The compatibility of the fibrations is automatic from the compatibility of the various projection maps π_j , for $j \in \{1, 2, 3, 4\}$.

14.1. The definition of $M^{0\text{-stratum}}$. — For each $i \in I_{0\text{-stratum}}$, put

$$(14.3) \quad M_i^{0\text{-stratum}} = B(p_i, .35R_i) \cup \mathcal{E}^{-1} \left\{ x \in H : x''_i \geq .9R_i, \frac{x'_i}{x''_i} \leq \frac{4}{10} \right\}.$$

Lemma 14.4. — *Under the constraints $\varsigma_{0\text{-stratum}} < \bar{\varsigma}_{0\text{-stratum}}$ and $c_{\text{adjust}} < \bar{c}_{\text{adjust}}$, $\{M_i^{0\text{-stratum}}\}_{i \in I_{0\text{-stratum}}}$ is a disjoint collection and each $M_i^{0\text{-stratum}}$ is a compact manifold with boundary, which is diffeomorphic to one of the possibilities in Proposition 14.1 (2).*

Proof. — Note that

$$(14.5) \quad (\mathcal{E}^0)^{-1} \left\{ x \in H : x''_i \geq .9R_i, \frac{x'_i}{x''_i} \leq \frac{4}{10} \right\} = \{p \in M : \zeta_i(p) \geq .9, \eta_i(p) \leq .4\}.$$

In particular, if $\varsigma_{0\text{-stratum}}$ is sufficiently small then this set contains $A(p_i, .31R_i, .39R_i)$ and is contained in $A(p_i, .29R_i, .41R_i)$. Then if c_{adjust} is sufficiently small, $\mathcal{E}^{-1}\{x \in H \mid x''_i \geq .9R_i, \frac{x'_i}{x''_i} \leq \frac{4}{10}\}$ contains $A(p_i, .32R_i, .38R_i)$ and is contained in $A(p_i, .28R_i, .42R_i)$.

In particular, $B(p_i, .38R_i) \subset M_i^{0\text{-stratum}} \subset B(p_i, .42R_i)$. It now follows from Lemma 11.5 that $\{M_i^{0\text{-stratum}}\}_{i \in I_{0\text{-stratum}}}$ are disjoint.

To characterize the topology of $M_i^{0\text{-stratum}}$, if c_{adjust} is sufficiently small then we can find a smooth function $f^0 : M \rightarrow \mathbb{R}$ such that

1. If $p \in A(p_i, .3R_i, .5R_i)$ and $x = \mathcal{E}(p)$ then $f^0(p) = \frac{x'_i}{x''_i}$.
2. If $p \in B(p_i, .35R_i)$ then $f^0(p) \leq .39$.
3. If $p \notin B(p_i, .5R_i)$ then $f^0(p) \geq .41$.

Put $f^1 = \eta_i$ and define $F : M \times [0, 1] \rightarrow \mathbb{R}$ by $F(p, t) = (1 - t)f^0(p) + tf^1(p)$. Put $f^t(p) = F(p, t)$ and $X = (-\infty, .4]$. If c_{adjust} and $\varsigma_{0\text{-stratum}}$ are sufficiently small then Lemma 11.1 implies that for each $t \in [0, 1]$, f^t is transverse to $\partial X = \{.4\}$. By Lemma 21.1, $M_i^{0\text{-stratum}} = (f^0)^{-1}(X)$ is diffeomorphic to $(f^1)^{-1}(X)$. By Lemma 11.3, the latter is diffeomorphic to one of the possibilities in Proposition 14.1 (2). This proves the lemma. \square

We let $M^{0\text{-stratum}} = \bigcup_{i \in I_{0\text{-stratum}}} M_i^{0\text{-stratum}}$, and put $M_1 = M \setminus \text{int}(M^{0\text{-stratum}})$. Thus $M^{0\text{-stratum}}$ and M_1 are smooth compact manifolds with boundary.

14.2. The definition of M^{slim} . — We first truncate W_3 . Put

$$(14.6) \quad W'_3 = W_3 \cap \bigcup_{i \in I_{\text{slim}}} \left\{ x \in Q_3 \mid x''_i > .9R_i, \left| \frac{x'_i}{x''_i} \right| < 4 \cdot 10^5 \Delta \right\}$$

and define $U'_3 = (\pi_3 \circ \mathcal{E})^{-1}(W'_3)$.

Lemma 14.7. — *Under the constraints $c_{\text{slim}} < \bar{c}_{\text{slim}}(\Delta)$ and $c_{\text{adjust}} < \bar{c}_{\text{adjust}}$, we have*

- (1) $\bigcup_{i \in I_{\text{slim}}} \{|\eta_i| \leq 3.5 \cdot 10^5 \Delta\} \subset U'_3 \subset U_3$, where U_3 is as in Proposition 13.1.
- (2) *The restriction of $\pi_3 \circ \mathcal{E}$ to U'_3 gives a proper submersion to W'_3 . In particular, it is a fibration.*
- (3) *The fibers of $\pi_3 \circ \mathcal{E} : U'_3 \rightarrow W'_3$ are diffeomorphic to S^2 or T^2 .*
- (4) M_1 intersects U'_3 in a submanifold with boundary which is a union of fibers of $\pi_3 \circ \mathcal{E} : U'_3 \rightarrow W'_3$.

Proof. — For a given $i \in I_{\text{slim}}$, suppose that $p \in M$ satisfies $|\eta_i(p)| \leq 3.5 \cdot 10^5 \Delta$. Putting $y = (\pi_3 \circ \mathcal{E}^0)(p) \in Q_3$, we have $y''_i = R_i$ and $|\frac{y''_i}{y'_i}| \leq 3.5 \cdot 10^5 \Delta$. Hence if c_{adjust} is small enough then since $\Delta \gg 1$, we are ensured that, putting $x = (\pi_3 \circ \mathcal{E})(p) \in Q_3$, we have $x''_i > .9R_i$ and $|\frac{x''_i}{x'_i}| < 4 \cdot 10^5 \Delta$. As $p \in U_3$, Proposition 13.1 implies that $x \in W_3$. Hence $\bigcup_{i \in I_{\text{slim}}} \{|\eta_i| \leq 3.5 \cdot 10^5 \Delta\} \subset U'_3$.

Now suppose that $p \in U'_3$. Putting $x = (\pi_3 \circ \mathcal{E})(p)$, for some $i \in I_{\text{slim}}$ we have $x''_i > .9R_i$ and $|\frac{x''_i}{x'_i}| < 4 \cdot 10^5 \Delta$. If c_{adjust} is small enough then we are ensured that, putting $y = (\pi_3 \circ \mathcal{E}^0)(p)$, we have $y''_i \geq .8R_i$ and $|\frac{y''_i}{y'_i}| \leq 4.5 \cdot 10^5 \Delta$. Hence $|\eta_i(p)| \leq 4.5 \cdot 10^5 \Delta$. This shows that $U'_3 \subset U_3$, proving part (1) of the lemma.

By Proposition 13.1, $\pi_3 \circ \mathcal{E}$ is a submersion from U_3 to W_3 . Hence it restricts to a surjective submersion on U'_3 .

Suppose that K is a compact subset of W'_3 . Then $(\pi_3 \circ \mathcal{E})^{-1}(K)$ is a closed subset of M which is contained in $\bar{U}_3 = \bigcup_{i \in I_{\text{slim}}} \{|\eta_i| \leq 5 \cdot 10^5 \Delta\}$. As $\{p_i\}_{i \in I_{\text{slim}}}$ are in the slim 1-stratum, it follows from the definition of adapted coordinates that $\{|\eta_i| \leq 5 \cdot 10^5 \Delta\}$ is a compact subset of M ; cf. the proof of Lemma 10.3. Thus the restriction of $\pi_3 \circ \mathcal{E}$ to U'_3 is a proper submersion. This proves part (2) of the lemma.

To prove part (3) of the lemma, given $x \in W'_3$, suppose that $p \in U'_3$ satisfies $(\pi_3 \circ \mathcal{E})(p) = x$. Choose $i \in I_{\text{slim}}$ so that $|\eta_i(p)| \leq 4.5 \cdot 10^5 \Delta$. If c_{adjust} is sufficiently small then by looking at the components in H_i , one sees that for any $p' \in U'_3$ satisfying $(\pi_3 \circ \mathcal{E})(p') = x$, we have $p' \in \{|\eta_i| < 5 \cdot 10^5 \Delta\}$. Thus to determine the topology of the fiber, we can just consider the restriction of $\pi_3 \circ \mathcal{E}$ to $\{|\eta_i| < 5 \cdot 10^5 \Delta\}$.

Let $\pi_{H'_i} : Q_3 \rightarrow H'_i$ be orthogonal projection and put $X = \pi_{H'_i}(x) \in H'_i$. As the restriction of $\pi_{H'_i} \circ \pi_3 \circ \mathcal{E}^0$ to $\{|\eta_i| < 5 \cdot 10^5 \Delta\}$ equals η_i , it follows that $\pi_{H'_i} \circ \pi_3 \circ \mathcal{E}^0$ is transverse there to X . By Lemma 10.3, $\{|\eta_i| < 5 \cdot 10^5 \Delta\} \cap (\pi_{H'_i} \circ \pi_3 \circ \mathcal{E}^0)^{-1}(X)$ is diffeomorphic to S^2 or T^2 .

Consider the restriction of $(\pi_{H'_i} \circ \pi_3 \circ \mathcal{E})$ to $\{|\eta_i| < 5 \cdot 10^5 \Delta\}$. Proposition 13.1 and Lemma 21.3 imply that if c_{adjust} is sufficiently small then the fiber $\{|\eta_i| < 5 \cdot 10^5 \Delta\} \cap (\pi_{H'_i} \circ \pi_3 \circ \mathcal{E})^{-1}(X)$ is diffeomorphic to S^2 or T^2 . In particular, it is connected. Now $(\pi_{H'_i} \circ \pi_3 \circ \mathcal{E})^{-1}(X)$ is the preimage, under $\pi_3 \circ \mathcal{E} : U'_3 \rightarrow W'_3$, of the preimage of X under $\pi_{H'_i} : W'_3 \rightarrow H'_i$. From connectedness of the fiber, the preimage of X under $\pi_{H'_i} : W'_3 \rightarrow H'_i$ must just be x . Hence $(\pi_3 \circ \mathcal{E})^{-1}(x)$ is diffeomorphic to S^2 or T^2 . This proves part (3) of the lemma.

To prove part (4) of the lemma, given $j \in I_{0\text{-stratum}}$, suppose that $p \in \partial M_j^{0\text{-stratum}}$. If $x = \mathcal{E}(p)$ then $x'_j \geq .9R_j$ and $x'_j = .4x''_j$. Suppose that $p \in U'_3$. If $q \in U'_3$ is a point in the same fiber of $\pi_3 \circ \mathcal{E} : U'_3 \rightarrow W'_3$ as p , put $y = \mathcal{E}(q) \in H$. As $\pi_3(x) = \pi_3(y)$, we have $y''_j \geq .9R_j$ and $y'_j = .4y''_j$. Thus $q \in \partial M_j^{0\text{-stratum}}$. Hence $\partial M_j^{0\text{-stratum}}$ is a union of fibers of $\pi_3 \circ \mathcal{E} : U'_3 \rightarrow W'_3$. In fact, since $\partial M_j^{0\text{-stratum}}$ is a connected 2-manifold, it is a single fiber of $\pi_3 \circ \mathcal{E}$. This proves part (4) of the lemma. \square

Let $W''_3 \subset W'_3$ be a compact 1-dimensional manifold with boundary such that $(\pi_3 \circ \mathcal{E})^{-1}(W''_3)$ contains $\bigcup_{i \in I_{\text{slim}}} \{|\eta_i| \leq 3.5 \cdot 10^5 \Delta\}$, and put $M^{\text{slim}} = M_1 \cap (\pi_3 \circ \mathcal{E})^{-1}(W''_3)$. We endow M^{slim} with the fibration induced by $\pi_3 \circ \mathcal{E}$.

Put $M_2 = M_1 \setminus \text{int}(M^{\text{slim}})$.

14.3. The definition of M^{edge} . — We first truncate W_2 . Put

$$(14.8) \quad W'_2 = W_2 \cap \bigcup_{i \in I_{\text{edge}}} \left\{ x \in Q_2 \mid x'_i \geq .9R_i, \left| \frac{x'_i}{x''_i} \right| < 4\Delta \right\}$$

and

$$(14.9) \quad U'_2 = (\pi_2 \circ \mathcal{E})^{-1}(W'_2) \cap \left(\{\eta_{E'} \leq .35\Delta\} \cup \mathcal{E}^{-1}\{x \in H \mid x_\tau > 0, \frac{x_{E'}}{x_\tau} \leq 4\Delta\} \right).$$

Lemma 14.10. — *Under the constraints $\Lambda < \bar{\Lambda}(\Delta)$, $\varsigma_{\text{edge}} < \bar{\varsigma}_{\text{edge}}(\Delta)$ and $c_{\text{adjust}} < \bar{c}_{\text{adjust}}$, we have*

- (1) $\bigcup_{i \in I_{\text{edge}}} \{|\eta_i| \leq 3.5\Delta, |\eta_{E'}| \leq 3.5\Delta\} \subset U'_2 \subset U_2$, where U_2 is as in Proposition 13.1.
- (2) The restriction of $\pi_2 \circ \mathcal{E}$ to U'_2 gives a proper submersion to W'_2 . In particular, it is a fibration.
- (3) The fibers of $\pi_2 \circ \mathcal{E} : U'_2 \rightarrow W'_2$ are diffeomorphic to D^2 .
- (4) M_2 intersects U'_2 in a submanifold with corners which is a union of fibers of $\pi_2 \circ \mathcal{E} : U'_2 \rightarrow W'_2$.

Proof. — The proof is similar to that of Lemmas 14.4 and 14.7. We omit the details. \square

Lemma 14.11. — *Under the constraint $c_{\text{adjust}} < \bar{c}_{\text{adjust}}$, $M_2 \cap U'_2$ is compact.*

Proof. — Suppose that $M_2 \cap U'_2$ is not compact. As M is compact, there is a sequence $\{q^k\}_{k=1}^\infty \subset M_2 \cap U'_2$ with a limit $q \in M$, for which $q \notin M_2 \cap U'_2$. Put $y = \mathcal{E}(q)$.

Since M_2 is closed we have $q \in M_2$ and so $q \notin U'_2$. Since $y_\tau > 0$ (assuming c_{adjust} is sufficiently small) we also have $q \in \{\eta_{E'} \leq .35\Delta\} \cup \mathcal{E}^{-1}\{x \in H \mid x_\tau > 0, \frac{x_{E'}}{x_\tau} \leq 4\Delta\}$.

We know that $\pi_2(y) \in \overline{W'_2}$. As $q \notin U'_2$, it must be that $\pi_2(y) \notin W'_2$. Then for some $i \in I_{\text{edge}}$, we have $y''_i \geq .9R_i$ and $\left| \frac{y'_i}{y''_i} \right| = 4\Delta$. Now p_i cannot be a slim 1-stratum point, as otherwise the preceding truncation step would force $B(p_i, 1000\Delta R_i) \cap M_2 = \emptyset$, which contradicts the facts that $q \in M_2$ and $d(p_i, q) < 10\Delta R_i$.

Lemma 9.26 now implies that there is a $j \in I_{\text{edge}}$ such that $|\eta_j(q)| < 2\Delta$. If c_{adjust} is sufficiently small then we are ensured that $y_j'' \geq .9R_j$ and $|\frac{y_j}{y_j''}| < 3\Delta$. Thus $\pi_2(y) \in W_2'$ and so $q \in U_2'$, which is a contradiction. \square

We put $M^{\text{edge}} = U_2' \cap M_2$ and $W_2'' = (\pi_2 \circ \mathcal{E})(M^{\text{edge}})$. We endow M^{edge} with the fibration induced by $\pi_2 \circ \mathcal{E}$.

Put $M_3 = M_2 \setminus \text{int}(M^{\text{edge}})$.

14.4. The definition of $M^{2\text{-stratum}}$. — We first truncate W_1 . Put

$$(14.12) \quad W_1' = W_1 \cap \bigcup_{i \in I_{2\text{-stratum}}} \left\{ x \in H \mid x_i'' > .9, \left| \frac{x_i'}{x_i''} \right| < 4 \right\}$$

and define $U_1' = \mathcal{E}^{-1}(W_1')$.

Lemma 14.13. — *Under the constraints $\varsigma_{2\text{-stratum}} < \bar{\varsigma}_{2\text{-stratum}}$ and $c_{\text{adjust}} < \bar{c}_{\text{adjust}}$, we have*

- (1) $\bigcup_{i \in I_{2\text{-stratum}}} \{|\eta_i| \leq 3.5\} \subset U_1' \subset U_1$, where U_1 is as in Proposition 13.1.
- (2) The restriction of \mathcal{E} to U_1' gives a proper submersion to W_1' . In particular, it is a fibration.
- (3) The fibers of $M^{2\text{-stratum}}$ are circles.
- (4) M_3 is contained in U_1' , and is a submanifold with corners which is a union of fibers of $\mathcal{E}|_{U_1'} : U_1' \rightarrow W_1'$.

Proof. — The proof is similar to that of Lemma 14.7. We omit the details. \square

We put $M^{2\text{-stratum}} = M_3$, and endow it with the fibration $\mathcal{E}|_{M^{2\text{-stratum}}} : M^{2\text{-stratum}} \rightarrow \mathcal{E}(M^{2\text{-stratum}})$.

14.5. The proof of Proposition 14.1. — Proposition 14.1 now follows from combining the results in this section.

15. Proof of Theorem 16.1 for closed manifolds

Recall that we are trying to get a contradiction to Standing Assumption 5.2. As before, we let M denote M^α for large α . Then M satisfies the conclusion of Proposition 14.1. To get a contradiction, we will show that M is a graph manifold.

We recall the definition of a graph manifold from Definition 1.2. It is obvious that boundary components of graph manifolds are tori. It is also obvious that if we glue two graph manifolds along boundary components then the result is a graph manifold, provided that it is orientable. In addition, the connected sum of two graph manifolds is a graph manifold. For more information about graph manifolds, we refer to [22, Chapter 2.4]

15.1. M is a graph manifold. — Each connected component of $M^{0\text{-stratum}}$ has boundary either \emptyset , S^2 or T^2 . If there is a connected component of $M^{0\text{-stratum}}$ with empty boundary then M is diffeomorphic to $S^1 \times S^2$, $S^1 \times_{\mathbb{Z}_2} S^2 = \mathbb{R}P^3 \# \mathbb{R}P^3$, T^3/Γ (where Γ is a finite subgroup of $\text{Isom}^+(T^3)$ which acts freely on T^3) or S^3/Γ (where Γ is a finite subgroup of $\text{Isom}^+(S^3)$ which acts freely on S^3). In any case M is a graph manifold. So we can assume that each connected component of $M^{0\text{-stratum}}$ has nonempty boundary.

Each connected component of M^{slim} fibers over S^1 or I . If it fibers over S^1 then M is diffeomorphic to $S^1 \times S^2$ or the total space of a T^2 -bundle over S^1 . In either case, M is a graph manifold. Hence we can assume that each connected component of M^{slim} is diffeomorphic to $I \times S^2$ or $I \times T^2$.

Lemma 15.1. — *Let $M_i^{0\text{-stratum}}$ be a connected component of $M^{0\text{-stratum}}$. If $M_i^{0\text{-stratum}} \cap M^{\text{slim}} \neq \emptyset$ then $\partial M_i^{0\text{-stratum}}$ is a boundary component of a connected component of M^{slim} . If $M_i^{0\text{-stratum}} \cap M^{\text{slim}} = \emptyset$ then we can write $\partial M_i^{0\text{-stratum}} = A_i \cup B_i$ where*

- (1) $A_i = M_i^{0\text{-stratum}} \cap M^{\text{edge}}$ is a disjoint union of 2-disks,
- (2) $B_i = M_i^{0\text{-stratum}} \cap M^{2\text{-stratum}}$ is the total space of a circle bundle and
- (3) $A_i \cap B_i = \partial A_i \cap \partial B_i$ is a union of circle fibers.

Furthermore, if $\partial M_i^{0\text{-stratum}}$ is a 2-torus then $A_i = \emptyset$, while if $\partial M_i^{0\text{-stratum}}$ is a 2-sphere then A_i consists of exactly two 2-disks.

Proof. — Proposition 14.1 implies all but the last sentence of the lemma. The statement about A_i follows from an Euler characteristic argument. □

Lemma 15.2. — *Let M_i^{slim} be a connected component of M^{slim} . Let Y_i be one of the connected components of $\partial M_i^{\text{slim}}$. If $Y_i \cap M^{0\text{-stratum}} \neq \emptyset$ then $Y_i = \partial M_i^{0\text{-stratum}}$ for some connected component $M_i^{0\text{-stratum}}$ of $M^{0\text{-stratum}}$.*

If $Y_i \cap M^{0\text{-stratum}} = \emptyset$ then we can write $Y_i = A_i \cup B_i$ where

- (1) $A_i = Y_i \cap M^{\text{edge}}$ is a disjoint union of 2-disks,
- (2) $B_i = Y_i \cap M^{2\text{-stratum}}$ is the total space of a circle bundle and
- (3) $A_i \cap B_i = \partial A_i \cap \partial B_i$ is a union of circle fibers.

Furthermore, if Y_i is a 2-torus then $A_i = \emptyset$, while if Y_i is a 2-sphere then A_i consists of exactly two 2-disks.

Proof. — The proof is similar to that of Lemma 15.1. We omit the details. □

Hereafter we can assume that there is a disjoint union $M^{0\text{-stratum}} = M_{S^2}^{0\text{-stratum}} \cup M_{T^2}^{0\text{-stratum}}$, based on what the boundaries of the connected components are. Similarly, each fiber of M^{slim} is S^2 or T^2 , so there is a disjoint union $M^{\text{slim}} = M_{S^2}^{\text{slim}} \cup M_{T^2}^{\text{slim}}$.

It follows from Lemmas 15.1 and 15.2 that each connected component of $M_{T^2}^{0\text{-stratum}} \cup M_{T^2}^{\text{slim}}$ is diffeomorphic to

1. A connected component of $M_{T^2}^{0\text{-stratum}}$,
2. The gluing of two connected components of $M_{T^2}^{0\text{-stratum}}$ along a 2-torus, or
3. $I \times T^2$.

In case 1, the connected component is diffeomorphic to $S^1 \times D^2$ or the total space of a twisted interval bundle over a Klein bottle. In any case, we can say that $M_{T^2}^{0\text{-stratum}} \cup M_{T^2}^{\text{slim}}$ is a graph manifold. Put $X_1 = M - \text{int}(M_{T^2}^{0\text{-stratum}} \cup M_{T^2}^{\text{slim}})$. To show that M is a graph manifold, it suffices to show that X_1 is a graph manifold. Note that $X_1 = M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}} \cup M^{\text{edge}} \cup M^{2\text{-stratum}}$.

Suppose that $M_i^{0\text{-stratum}}$ is a connected component of $M_{S^2}^{0\text{-stratum}}$. From Proposition 14.1, $M_i^{0\text{-stratum}}$ is diffeomorphic to D^3 or $\mathbb{R}P^3 \# D^3$. If $M_i^{0\text{-stratum}}$ is diffeomorphic to $\mathbb{R}P^3 \# D^3$, let Z_i be the result of replacing $M_i^{0\text{-stratum}}$ in X_1 by D^3 . Then X_1 is diffeomorphic to $\mathbb{R}P^3 \# Z_i$. As $\mathbb{R}P^3$ is a graph manifold, if Z_i is a graph manifold then X_1 is a graph manifold. Hence without loss of generality, we can assume that each connected component of $M_{S^2}^{0\text{-stratum}}$ is diffeomorphic to a 3-disk.

From Lemmas 15.1 and 15.2, each connected component of $M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}}$ is diffeomorphic to

1. D^3 ,
2. $I \times S^2$ or
3. S^3 , the result of attaching two connected components of $M_{S^2}^{0\text{-stratum}}$ by a connected component $I \times S^2$ of $M_{S^2}^{\text{slim}}$.

In case 3, X_1 is diffeomorphic to a graph manifold. In case 2, if Z is a connected component of $M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}}$ which is diffeomorphic to $I \times S^2$ then we can do surgery along $\{\frac{1}{2}\} \times S^2 \subset X_1$ to replace $I \times S^2 \subset X_1$ by a union of two 3-disks. Let X_2 be the result of performing the surgery. Then X_1 is recovered from X_2 by either taking a connected sum of two connected components of X_2 or by taking a connected sum of X_2 with $S^1 \times S^2$. In either case, if X_2 is a graph manifold then X_1 is a graph manifold. Hence without loss of generality, we can assume that each connected component of $(M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}}) \subset X_1$ is diffeomorphic to D^3 .

Some connected components of M^{edge} may fiber over S^1 . If Z is such a connected component then it is diffeomorphic to $S^1 \times D^2$. If $X_1 - \text{int}(Z)$ is a graph manifold then X_1 is a graph manifold. Hence without loss of generality, we can assume that each connected component of M^{edge} is diffeomorphic to $I \times D^2$.

Let G be a graph (*i.e.*, 1-dimensional CW-complex) whose vertices correspond to connected components of $M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}}$, and whose edges correspond to connected components of M^{edge} joining such ‘‘vertex’’ components. From Lemmas 15.1 and 15.2, each vertex of G has degree two. Again from Lemmas 15.1 and 15.2, we can label the connected components of $M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}} \cup M^{\text{edge}}$ by connected components of G . It follows that each connected component of $M_{S^2}^{0\text{-stratum}} \cup M_{S^2}^{\text{slim}} \cup M^{\text{edge}}$ is diffeomorphic to $S^1 \times D^2$.

We have now shown that X_1 is the result of gluing a disjoint collection of $S^1 \times D^2$'s to $M^{2\text{-stratum}}$, with each gluing being performed between the boundary of a $S^1 \times D^2$ factor and a toral boundary component of $M^{2\text{-stratum}}$. As $M^{2\text{-stratum}}$ is the total space of a circle bundle, it is a graph manifold. Thus X_1 is a graph manifold. Hence we have shown:

Proposition 15.3. — *Under the constraints imposed in the earlier sections, M is a graph manifold.*

15.2. Satisfying the constraints. — We now verify that it is possible to simultaneously satisfy all the constraints that appeared in the construction.

We indicate a partial ordering of the parameters which is respected by all the constraints appearing in the paper. This means that every constraint on a given parameter is an upper (or lower) bound given as a function of other parameters which are strictly smaller in the partial order. Consequently, all constraints can be satisfied simultaneously, since we may choose values for parameters starting with those parameters which are minimal with respect to the partial order, and proceeding upward.

$$(15.4) \quad \{\mathcal{M}, \beta_3\} \prec \{c_{\text{slim}}, \Omega_i, \Omega'_i\} \prec \Gamma_3 \prec \{\Sigma_3, \Xi_3\} \prec c_{\text{edge}} \prec \Gamma_2 \prec \{\Sigma_2, \Xi_2\} \prec \\ c_{2\text{-stratum}} \prec \Gamma_1 \prec \{\Sigma_1, \Xi_1\} \prec c_{2\text{-stratum}} \prec \beta_2 \prec \Delta \prec \{c_{\text{edge}}, c_{E'}, c_{\text{slim}}\} \prec \\ c_{0\text{-stratum}} \prec \{\beta_{E'}, \sigma_{E'}\} \prec \sigma_E \prec \{\sigma, \Lambda\} \prec \bar{w} \prec w' \prec \beta_E \prec \beta_1 \prec \{\Upsilon_0, \delta_0\} \prec \Upsilon'_0.$$

This proves Theorem 1.3.

16. Manifolds with boundary

In this section we consider manifolds with boundary. Since our principal application is to the geometrization conjecture, we will only deal with manifolds whose boundary components have a nearly cuspidal collar. We recall that a *hyperbolic cusp* is a complete manifold with boundary diffeomorphic to $T^2 \times [0, \infty)$, which is isometric to the quotient of a horoball by an isometric \mathbb{Z}^2 -action. More explicitly, a cusp is isometric to a quotient of the upper half space $\mathbb{R}^2 \times [0, \infty) \subset \mathbb{R}^3$, with the metric $dz^2 + e^{-z}(dx^2 + dy^2)$, by a rank-2 group of horizontal translations. (For application to the geometrization conjecture, we take the cusp to have constant sectional curvature $-\frac{1}{4}$.)

Theorem 16.1. — *Let $K \geq 10$ be a fixed integer. Fix a function $A : (0, \infty) \rightarrow (0, \infty)$. Then there is some $w_0 \in (0, c_3)$ such that the following holds.*

Suppose that (M, g) is a compact connected orientable Riemannian 3-manifold with boundary. Assume in addition that

- (1) *The diameters of the connected components of ∂M are bounded above by w_0 .*

- (2) For each component X of ∂M , there is a hyperbolic cusp \mathcal{H}_X with boundary $\partial\mathcal{H}_X$, along with a C^{K+1} -embedding of pairs $e : (N_{100}(\partial\mathcal{H}_X), \partial\mathcal{H}_X) \rightarrow (M, X)$ which is w_0 -close to an isometry.
- (3) For every $p \in M$ with $d(p, \partial M) \geq 10$, we have $\text{vol}(B(p, R_p)) \leq w_0 R_p^3$.
- (4) For every $p \in M$, $w' \in [w_0, c_3]$, $k \in [0, K]$, and $r \leq R_p$ such that $\text{vol}(B(p, r)) \geq w' r^3$, the inequality

$$(16.2) \quad |\nabla^k \text{Rm}| \leq A(w') r^{-(k+2)}$$

holds in the ball $B(p, r)$.

Then M is a graph manifold.

In order to prove Theorem 16.1, we make the following assumption.

Standing Assumption 16.3. — Let $K \geq 10$ be a fixed integer and let $A' : (0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ be a function.

We assume that $\{(M^\alpha, g^\alpha)\}_{\alpha=1}^\infty$ is a sequence of connected closed Riemannian 3-manifolds such that

- (1) The diameters of the connected components of ∂M^α are bounded above by $\frac{1}{\alpha}$.
- (2) For each component X^α of ∂M^α , there is a hyperbolic cusp \mathcal{H}_{X^α} along with a C^{K+1} -embedding of pairs $e : (N_{100}(\partial\mathcal{H}_{X^\alpha}), \partial\mathcal{H}_{X^\alpha}) \rightarrow (M^\alpha, X^\alpha)$ which is $\frac{1}{\alpha}$ -close to an isometry.
- (3) For all $p \in M^\alpha$ with $d(p, \partial M^\alpha) \geq 10$, the ratio $\frac{R_p}{r_p(1/\alpha)}$ of the curvature scale at p to the $\frac{1}{\alpha}$ -volume scale at p is bounded below by α .
- (4) For all $p \in M^\alpha$ and $w' \in [\frac{1}{\alpha}, c_3]$, let $r_p(w')$ denote the w' -volume scale at p . Then for each integer $k \in [0, K]$ and each $C \in (0, \alpha)$, we have $|\nabla^k \text{Rm}| \leq A'(C, w') r_p(w')^{-(k+2)}$ on $B(p, Cr_p(w'))$.
- (5) Each M^α fails to be a graph manifold.

As in Lemma 5.1, to prove Theorem 16.1 it suffices to get a contradiction from Standing Assumption 16.3. As before, we let M denote the manifold M^α for large α . The argument to get a contradiction from Standing Assumption 16.3 is a slight modification of the argument in the closed case, the main difference being the appearance of a new family of points – those lying in a collared region near the boundary.

We will use the same set of the parameters as in the case of closed manifolds, with an additional parameter r_∂ . It will be placed at the end of the partial ordering in (15.4), after Υ'_0 .

Let $\{\partial_i M\}_{i \in I_\partial}$ be the collection of boundary components of M , and let $e_i : (N_{100}(\partial\mathcal{H}_i), \partial\mathcal{H}_i) \rightarrow (M, \partial_i M)$ be the embedding from Standing Assumption 16.3. Note that the restriction of e_i to $\partial\mathcal{H}_i$ is a diffeomorphism. Put $b_i = d_{\partial\mathcal{H}_i} \in C^\infty(\mathcal{H}_i)$. Let η_i be a slight smoothing of $b_i \circ e_i^{-1}$ on $(b_i \circ e_i^{-1})^{-1}(1, 99)$, as in Lemma 3.14.

Lemma 16.4. — We may assume that for all $p \in \eta_i^{-1}(5, 95)$,

- (1) The curvature scale satisfies $R_p \in (1, 3)$.
- (2) $\mathfrak{r}_p < r_\partial$.
- (3) There is a $(1, \beta_1)$ -splitting of $(\frac{1}{\mathfrak{r}_p}M, p)$ for which $\frac{1}{\mathfrak{r}_p}\eta_i$ is an adapted coordinate of quality ζ_{slim} .

Proof. — In view of the quality of the embedding e_i , it suffices to check the claim on the constant-curvature space $b_i^{-1}(4, 96)$. The diameter of $\partial\mathcal{H}_i$ can be assumed to be arbitrarily small by taking α to be large enough. The Riemannian metric on \mathcal{H}_i has the form $dz^2 + e^{-z}g_{T^2}$ for a flat metric g_{T^2} on T^2 , with $z \in [0, 100)$. The lemma follows from elementary estimates. \square

We now select 2-stratum balls, edge balls, slim 1-stratum balls and 0-balls as in the closed case, except with the restriction that the center points p_i all satisfy $d(p_i, \partial M) \geq 10$.

Given $i \in I_\partial$, let B_i be the connected component of $M - \eta_i^{-1}(90)$ containing $\partial_i M$.

Lemma 16.5. — If $r_\partial < \bar{r}_\partial(\Upsilon'_0)$ then M is diffeomorphic to $I \times T^2$ or $\{B_i\}_{i \in I_\partial} \cup \{B(p_i, r_{p_i}^0)\}_{i \in I_{0\text{-stratum}}}$ is a disjoint collection of open sets.

Proof. — We can assume that M is not diffeomorphic to $I \times T^2$.

Suppose first that $B_i \cap B_j \neq \emptyset$ for some $i, j \in I_\partial$ with $i \neq j$. Then $\eta_i^{-1}(5, 90)$ must intersect $\eta_j^{-1}(5, 90)$. It follows easily that $M = N_{10}(B_i) \cup N_{10}(B_j)$ is diffeomorphic to $I \times T^2$, which is a contradiction. Thus $B_i \cap B_j = \emptyset$.

Next, suppose that $B_i \cap B(p_j, r_{p_j}^0) \neq \emptyset$ for some $i \in I_\partial$ and $j \in I_{0\text{-stratum}}$. If $r_\partial < \bar{r}_\partial(\Upsilon'_0)$ then by Lemma 16.4 we will have $\Upsilon'_0 \mathfrak{r}_{p_j} < \frac{1}{100}$. Hence $r_{p_j}^0 < \frac{1}{100}$ and the triangle inequality implies that $p_j \in \eta_i^{-1}(5, 95)$. However, from Lemma 16.4 (3), this contradicts the fact that p_j is a 0-stratum point.

Finally, if $i, j \in I_{0\text{-stratum}}$ and $i \neq j$ then $B(p_i, r_{p_i}^0) \cap B(p_j, r_{p_j}^0) = \emptyset$ from Lemma 11.5 (1). \square

Hereafter we assume that M is not diffeomorphic to $I \times T^2$, which is already a graph manifold.

For each $i \in I_\partial$, let H_i be a copy of \mathbb{R}^2 . Put $H_\partial = \bigoplus_{i \in I_\partial} H_i$. We also put

- $Q_1 = H \oplus H_\partial$,
- $Q_2 = H_{0\text{-stratum}} \oplus H_{\text{slim}} \oplus H_{\text{edge}} \oplus H_\partial$,
- $Q_3 = H_{0\text{-stratum}} \oplus H_{\text{slim}} \oplus H_\partial$,
- $Q_4 = H_{0\text{-stratum}} \oplus H_\partial$.

For $i \in I_\partial$, let $\zeta_i \in C^\infty(M)$ be the extension by zero of $\Phi_{20,30,80,90} \circ \eta_i$ to M . Define $\mathcal{E}_i^0 : M \rightarrow H_i$ by $\mathcal{E}_i^0(p) = (\eta_i(p)\zeta_i(p), \zeta_i(p))$. We now go through Sections 12 and 13,

treating H_∂ in parallel to $H_{0\text{-stratum}}$. Next, in analogy to (14.3), for each $i \in I_\partial$ we put

$$(16.6) \quad M_i^\partial = N_{35}(\partial_i M) \cup \mathcal{E}^{-1} \left\{ x \in H : x''_i \geq .9, \frac{x'_i}{x''_i} \leq 40 \right\}.$$

Then M_i^∂ is diffeomorphic to $I \times T^2$. We now go through the argument of Section 15, treating each M_i^∂ as if it were an element of $M_{T^2}^{0\text{-stratum}}$ without a core. As in Section 15, we conclude that M is a graph manifold. This proves Theorem 16.1.

17. Application to the geometrization conjecture

We now use the terminology of [21] and [24]. Let $(M, g(\cdot))$ be a Ricci flow with surgery whose initial 3-manifold is compact. We normalize the metric by putting $\widehat{g}(t) = \frac{g(t)}{t}$. Let $(M_t, \widehat{g}(t))$ be the time- t manifold. (If t is a surgery time then we take M_t to be the post-surgery manifold.) We recall that the w -thin part $M^-(w, t)$ of M_t is defined to be the set of points $p \in M_t$ so that either $R_p = \infty$ or $\text{vol}(B(p, R_p)) < wR_p^3$. The w -thick part $M^+(w, t)$ of M_t is $M_t - M^-(w, t)$.

The following theorem is proved in [24, Section 7.3]; see also [21, Proposition 90.1].

Theorem 17.1. — [24] *There is a finite collection $\{(H_i, x_i)\}_{i=1}^k$ of pointed complete finite-volume Riemannian 3-manifolds with constant sectional curvature $-\frac{1}{4}$ and, for large t , a decreasing function $\beta(t)$ tending to zero and a family of maps*

$$(17.2) \quad f_t : \bigsqcup_{i=1}^k H_i \supset \bigsqcup_{i=1}^k B\left(x_i, \frac{1}{\beta(t)}\right) \longrightarrow M_t$$

such that

- (1) f_t is $\beta(t)$ -close to being an isometry.
- (2) The image of f_t contains $M^+(\beta(t), t)$.
- (3) The image under f_t of a cuspidal torus of $\{H_i\}_{i=1}^k$ is incompressible in M_t .

Given a sequence $t^\alpha \rightarrow \infty$, let Y^α be the truncation of $\bigsqcup_{i=1}^k H_i$ obtained by removing horoballs at distance approximately $\frac{1}{2\beta(t^\alpha)}$ from the basepoints x_i . Put $M^\alpha = M_{t^\alpha} - f_{t^\alpha}(Y_{t^\alpha})$.

Theorem 17.3. — [24] *For large α , M^α is a graph manifold.*

Proof. — We check that the hypotheses of Theorem 16.1 are satisfied for large α . Conditions (1) and (2) of Theorem 16.1 follow from the almost-isometric embedding of $\bigsqcup_{i=1}^k (B(x_i, \frac{1}{\beta(t^\alpha)}) - B(x_i, \frac{1}{2\beta(t^\alpha)})) \subset \bigsqcup_{i=1}^k H_i$ in M^α .

Next, Theorem 17.1 says that for any $\bar{w} > 0$, for large α the \bar{w} -thick part of M_{t^α} has already been removed in forming M^α . Thus Condition (3) of Theorem 16.1 holds.

From Ricci flow arguments, for each $w' \in (0, c_3)$ there are $\bar{r}(w') > 0$ and $K_k(w') < \infty$ so that for large α the following holds: for every $p \in M^\alpha$, $w' \in (0, c_3)$, $k \in [0, K]$ and

$r \leq \min(R_p, \bar{r}(w'))$, the inequality $|\nabla^k \text{Rm}| \leq K_k(w')r^{-(k+2)}$ holds in the ball $B(p, r)$ [21, Lemma 92.13]. Hence to verify Condition (4) of Theorem 16.1, at least for large α , we must show that if $p \in M^\alpha$ then the conditions $r \leq R_p$ and $\text{vol}(B(p, r)) \geq w'r^3$ imply that $r \leq \bar{r}(w')$.

Suppose not, *i.e.*, we have $\bar{r}(w') < r \leq R_p$. Then $\text{Rm}|_{B(p,r)} \geq -\frac{1}{r^2}$. Using the fact that $\text{vol}(B(p, r)) \geq w'r^3$, the Bishop-Gromov inequality gives an inequality of the form $\text{vol}(B(p, \bar{r}(w'))) \geq w''\bar{r}^3(w')$ for some $w'' = w''(w') > 0$.

We also have $\text{Rm}|_{B(p,\bar{r}(w'))} \geq -\frac{1}{\bar{r}^2(w')}$. Then from [25, Lemma 7.2] or [21, Lemma 88.1], for large α we can assume that the sectional curvatures on $B(p, \bar{r}(w'))$ are arbitrarily close to $-\frac{1}{4}$. In particular, $R_p \leq 5$. Then

$$(17.4) \quad \text{vol}(B(p, R_p)) \geq \text{vol}(B(p, r)) \geq w'r^3 = w' \left(\frac{r}{R_p}\right)^3 R_p^3 \geq w' \left(\frac{\bar{r}(w')}{5}\right)^3 R_p^3.$$

If α is sufficiently large then we conclude that $p \in f_{t^\alpha}(Y_{t^\alpha})$, which is a contradiction.

We now take $A(w')$ to be a number so that Condition (4) of Theorem 16.1 holds for all M^α . From the preceding discussion, there is a finite such number. Then for large α , all of the hypotheses of Theorem 16.1 hold. The theorem follows. \square

Theorems 17.1 and 17.3, along with the description of how M_t changes under surgery [24, Section 3], [21, Lemma 73.4], imply Thurston's geometrization conjecture.

18. Local collapsing without derivative bounds

In this section, we explain how one can remove the bounds on derivatives of curvature from the hypotheses of Theorem 1.3, to obtain:

Theorem 18.1. — *There exists a $w_0 \in (0, c_3)$ such that if M is a closed, orientable, Riemannian 3-manifold satisfying*

$$(18.2) \quad \text{vol}(B(p, R_p)) < w_0 R_p^3$$

for every $p \in M$, then M is a graph manifold.

The bounds on the derivatives of curvature are only used to obtain pointed C^K -limits of sequences at the (modified) volume scale. This occurs in Lemmas 9.21 and 11.1. We explain how to adapt the statements and proofs.

Modifications in Lemma 9.21. — The statement of the Lemma does not require modification. In the proof, the map ϕ will be a Gromov-Hausdorff approximation rather than a C^{K+1} -map close to an isometry, and Z will be a complete 2-dimensional nonnegatively curved Alexandrov space. As critical point theory for functions works the same way for Alexandrov spaces as for Riemannian manifolds, and 2-dimensional Alexandrov spaces are topological manifolds, the statement and proof of Lemma 3.12

remain valid for 2-dimensional Alexandrov spaces. The main difference in the proof of Lemma 9.21 is the method for verifying the fiber topology. For this, we use:

Theorem 18.3 (Linear local contractibility [15]). — *For every $w \in (0, \infty)$ and every positive integer n , there exist $r_0 \in (0, \infty)$ and $C \in (1, \infty)$ with the following property. If $B(p, 1)$ is a unit ball with compact closure in a Riemannian n -manifold, $\text{Rm}|_{B(p,1)} \geq -1$ and $\text{vol}(B(p, 1)) \geq w$ then the inclusion $B(p, r) \rightarrow B(p, Cr)$ is null-homotopic for every $r \in (0, r_0)$.*

This uniform contractibility may be used to promote a Gromov-Hausdorff approximation f_0 to a nearby continuous map f : one first restricts f_0 to the 0-skeleton of a fine triangulation, and then extends it inductively to higher skeleta simplex by simplex, using the controlled contractibility radius.

Lemma 18.4. — *With notation from the proof of Lemma 9.21, the fiber $F = \eta_p^{-1}(\{0\})$ is homotopy equivalent to $B(\star_Z, 4\Delta) \subset Z$.*

Proof. — Let $\widehat{\phi}$ be a quasi-inverse to the Gromov-Hausdorff approximation ϕ .

To produce a map $F \rightarrow B(\star_Z, 4\Delta)$ we take $\pi_Z \circ \phi|_F$, promote it to a continuous map as above, and then use the absence of critical points of d_Y near $S(\star_Z, 4\Delta)$ to homotope this to a map taking values in $B(\star_Z, 4\Delta)$.

To get the map $B(\star_Z, 4\Delta) \rightarrow F$, we apply the above procedure to promote $\widehat{\phi}|_F$ to a nearby continuous map $B(\star_Z, 4\Delta) \rightarrow \frac{1}{\tau_p}M$. Then using the fibration structures defined by η_p and $(\eta_p, \eta_{E'})$, we may perturb this to a map taking values in F .

The compositions of these maps are close to the identity maps; using a relative version of the approximation procedure one shows that these are homotopic to identity maps. \square

Thus we conclude that the fiber is a contractible compact 2-manifold with boundary, so it is a 2-disk.

Modifications in Lemma 11.1. — In the statement of the lemma, N_p is a 3-dimensional nonnegatively curved Alexandrov space instead of a nonnegatively curved Riemannian manifold, and “diffeomorphism” is replaced by “homeomorphism”.

In the proof, the pointed C^K -convergence is replaced by pointed Gromov-Hausdorff convergence to a 3-dimensional nonnegatively curved Alexandrov space N ; otherwise, we retain the notation from the proof. We need:

Theorem 18.5 (The Stability Theorem [19, 25]). — *Suppose $\{(M_k, \star_k)\}$ is a sequence of Riemannian n -manifolds, such that the sectional curvature is bounded below by a (k -independent) function of the distance to the basepoint \star_k . Let X be an n -dimensional Alexandrov space with curvature bounded below, and assume that $\phi_k : (X, \star_\infty) \rightarrow (M_k, \star_k)$ is a δ_k -pointed Gromov Hausdorff approximation, where $\delta_k \rightarrow 0$. Then for every $R \in (0, \infty)$, $\epsilon \in (0, \infty)$, and every sufficiently large k , there is a pointed map*

$\psi_k : (B(\star_\infty, R + \epsilon), \star_\infty) \rightarrow (M_k, \star_k)$ which is a homeomorphism onto an open subset containing $B(\star_k, R)$, where $d_{C^0}(\psi_k, \phi_k|_{B(\star_\infty, R+\epsilon)}) < \epsilon$.

Using critical point theory as before, we get that the limiting Alexandrov space N is homeomorphic to the balls $B(p_\infty, R'')$ for $R'' \in (\frac{1}{2}R', 2R')$, and there are no critical points for d_{p_∞} or d_{p_j} in the respective annuli $A(p_\infty, \frac{R''}{10^3}, 10R'') \subset N$ and $A(p_j, \frac{R''}{10^3}, 10R'') \subset \frac{1}{\tau_{p_j}}M_j$. The Stability Theorem produces a homeomorphism ψ from the closed ball $\overline{B(p_\infty, R'')} \subset M$ to a subset close to the ball $B(p_j, R'') \subset \frac{1}{\tau_{p_j}}M_j$; in particular, restricting ψ to the sphere $S(p_\infty, R'')$ we obtain a Gromov-Hausdorff approximation from the surface $S(p_\infty, R'')$ to the surface $S(p_j, R'')$. Appealing to uniform contractibility (Theorem 18.3), and using homotopies guaranteed by the absence of critical points we get that $\psi|_{S(p_\infty, R'')}$ is close to a homotopy equivalence. As in the proof of Theorem 11.3, we conclude that $\psi(\overline{B(p_\infty, R'')})$ is isotopic to $\overline{B(p_j, R'')}$.

Finally, we appeal to the classification of complete, noncompact, orientable, nonnegatively curved Alexandrov spaces N , when N is a noncompact topological 3-manifold, from [29] to conclude that the list of possible topological types is the same as in the smooth case.

Theorem 18.6 (Shioya-Yamaguchi [29]). — *If X is a noncompact, orientable, 3-dimensional nonnegatively curved Alexandrov space which is a topological manifold, then X is homeomorphic to one of the following: \mathbb{R}^3 , $S^1 \times \mathbb{R}^2$, $S^2 \times \mathbb{R}$, $T^2 \times \mathbb{R}$, or a twisted line bundle over $\mathbb{R}P^2$ or the Klein bottle.*

When N is compact, we may apply the main theorem of [31] to see that the topological classification is the same as in the smooth case. Alternatively, using the splitting theorem, one may reduce to the case when N has finite fundamental group and use the elliptization conjecture (now a theorem via Ricci flow due to finite extinction time results).

Remark 18.7. — Theorem 18.1 implies the collapsing result stated in the appendix of [30]. Note that Theorem 18.1 is strictly stronger, since the curvature scale need not be small compared to the diameter. However, we remark that the argument of [30] also gives the stronger result, if one uses [31] or the elliptization conjecture as above.

19. Appendix A : Choosing ball covers

Let M be a complete Riemannian manifold and let V be a bounded subset of M . Given $p \in V$ and $r > 0$, we write $B(p, r)$ for the metric ball in M around p of radius r . Let $\mathcal{R} : V \rightarrow \mathbb{R}$ be a (not necessarily continuous) function with range in some compact positive interval. For $p \in V$, we denote $\mathcal{R}(p)$ by \mathcal{R}_p . Put $S_1 = V$, $\rho_1 = \sup_{p \in V} \mathcal{R}_p$ and $\rho_\infty = \inf_{p \in V} \mathcal{R}_p$. Choose a point $p_1 \in V$ so that $\mathcal{R}_{p_1} \geq \frac{1}{2}\rho_1$. Inductively, for $i \geq 1$, let S_{i+1} be the subset of V consisting of points p such that $B(p, \mathcal{R}_p)$ is disjoint from $B(p_1, \mathcal{R}_{p_1}) \cup \dots \cup B(p_i, \mathcal{R}_{p_i})$. If $S_{i+1} = \emptyset$ then stop. If $S_{i+1} \neq \emptyset$, put

$\rho_{i+1} = \sup_{p \in S_{i+1}} \mathcal{R}_p$ and choose a point $p_{i+1} \in S_{i+1}$ so that $\mathcal{R}_{p_{i+1}} \geq \frac{1}{2}\rho_{i+1}$. This process must terminate after a finite number of steps, as the ρ_1 -neighborhood of V cannot contain an infinite number of disjoint balls with radius at least $\frac{1}{2}\rho_\infty$.

Lemma 19.1. — $\{B(p_i, \mathcal{R}_{p_i})\}$ is a finite disjoint collection of balls such that $V \subset \bigcup_i B(p_i, 3\mathcal{R}_{p_i})$. Furthermore, given $q \in V$, there is some N so that $q \in B(p_N, 3\mathcal{R}_{p_N})$ and $\mathcal{R}_q \leq 2\mathcal{R}_{p_N}$.

Proof. — Given $q \in V$, we know that $B(q, \mathcal{R}_q)$ intersects $\bigcup_i B(p_i, \mathcal{R}_{p_i})$. Let N be the smallest number i such that $B(q, \mathcal{R}_q)$ intersects $B(p_i, \mathcal{R}_{p_i})$. Then $q \in S_N$ and so $\rho_N \geq \mathcal{R}_q$. Thus $\mathcal{R}_{p_N} \geq \frac{1}{2}\rho_N \geq \frac{1}{2}\mathcal{R}_q$. As $B(q, \mathcal{R}_q)$ intersects $B(p_N, \mathcal{R}_{p_N})$, we have $d(q, p_N) < \mathcal{R}_q + \mathcal{R}_{p_N} \leq 3\mathcal{R}_{p_N}$. \square

20. Appendix B : Cloudy submanifolds

In this section we define the notion of a cloudy k -manifold. This is a subset of a Euclidean space with the property that near each point, it looks coarsely close to an affine subspace of the Euclidean space. The result of this appendix is that any cloudy k -manifold can be well interpolated by a smooth k -dimensional submanifold of the Euclidean space.

If H is a Euclidean space, let $\text{Gr}(k, H)$ denote the Grassmannian of codimension- k subspaces of H . It is metrized by saying that for $P_1, P_2 \in \text{Gr}(k, H)$, if $\pi_1, \pi_2 \in \text{End}(H)$ are orthogonal projection onto P_1 and P_2 , respectively, then $d(P_1, P_2)$ is the operator norm of $\pi_1 - \pi_2$. If H' is another Euclidean space then there is an isometric embedding $\text{Gr}(k, H) \rightarrow \text{Gr}(k, H \oplus H')$. If X is a k -dimensional submanifold of H then the normal map of X is the map $X \rightarrow \text{Gr}(k, H)$ which assigns to $p \in X$ the normal space of X at p .

Definition 20.1. — Suppose $C, \delta \in (0, \infty)$, $k \in \mathbb{N}$, and H is a Euclidean space. A (C, δ) cloudy k -manifold in H is a triple (\tilde{S}, S, r) , where $S \subset \tilde{S} \subset H$ is a pair of subsets, and $r : \tilde{S} \rightarrow (0, \infty)$ is a (possibly discontinuous) function such that:

- (1) For all $x, y \in \tilde{S}$, $|r(y) - r(x)| \leq C(|x - y| + r(x))$.
- (2) For all $x \in S$, the rescaled pointed subset $(\frac{1}{r(x)}\tilde{S}, x)$ is δ -close in the pointed Hausdorff distance to $(\frac{1}{r(x)}A_x, x)$, where A_x is a k -dimensional affine subspace of H . Here, as usual, $\frac{1}{r(x)}\tilde{S}$ means the subset \tilde{S} equipped with the distance function of H rescaled by $\frac{1}{r(x)}$.

We will sometimes say informally that a pair (\tilde{S}, S) is a cloudy k -manifold if it can be completed to a triple (\tilde{S}, S, r) which is a (C, δ) cloudy k -manifold for some (C, δ) . We will write $A_x^0 \subset H$ for the k -dimensional linear subspace parallel to A_x and we will write $\pi_{A_x^0}$ for orthogonal projection onto A_x^0 . Let $P_{A_x} : H \rightarrow H$ be the nearest point projection to A_x , given by $P_{A_x}(y) = x + \pi_{A_x^0}(y - x)$.

Lemma 20.2. — For all $k, K \in \mathbb{Z}^+$, $\epsilon \in (0, \infty)$ and $C < \infty$, there is a $\delta = \delta(k, K, \epsilon, C) > 0$ with the following property. Suppose (\tilde{S}, S, r) is a (C, δ) cloudy k -manifold in a Euclidean space H , and for every $x \in S$ we denote by A_x an affine subspace as in Definition 20.1. Then there is a k -dimensional smooth submanifold $W \subset H$ such that

- (1) For all $x \in S$, the pointed Hausdorff distance from $(\frac{1}{r(x)}\tilde{S}, x)$ to $(\frac{1}{r(x)}W, x)$ is at most ϵ .
- (2) $W \subset N_{\epsilon r}(\tilde{S})$.
- (3) For all $x \in S$, the restriction of the normal map of W to $B(x, r(x)) \cap W$ has image contained in an ϵ -ball of A_x^\perp in $\text{Gr}(k, H)$.
- (4) If I is a multi-index with $|I| \leq K$ then the I^{th} covariant derivative of the second fundamental form of W at w is bounded in norm by $\epsilon r(x)^{-(|I|+1)}$.
- (5) $W \cap N_r(S)$ is properly embedded in $N_r(S)$.
- (6) The nearest point map $P : N_r(S) \rightarrow W$ is a well-defined smooth submersion.
- (7) If I is a multi-index with $1 \leq |I| \leq K$ then for all $x \in S$, the restriction of $P - P_{A_x}$ to $B(x, r(x))$ has I^{th} derivative bounded in norm by $\epsilon r(x)^{-(|I|-1)}$.




FIGURE 5

Proof. — With the notation of Lemma 19.1, put $V = S$ and $\mathcal{R} = r$. Let T be the finite collection of points $\{p_i\}$ from the conclusion of Lemma 19.1. Then $\{B(\hat{x}, r_{\hat{x}})\}_{\hat{x} \in T}$ is a disjoint collection of balls such that for any $x \in S$, there is some $\hat{x} \in T$ with $x \in B(\hat{x}, 3r(\hat{x}))$ and $r(x) \leq 2r(\hat{x})$. Hence $B(x, r(x)) \subset B(\hat{x}, r(x) + 3r(\hat{x})) \subset B(\hat{x}, 5r(\hat{x}))$. This shows that $\bigcup_{x \in S} B(x, r(x)) \subset \bigcup_{\hat{x} \in T} B(\hat{x}, 5r(\hat{x}))$.

For each $\hat{x} \in T$, let $A_{\hat{x}} \subset H$ be the k -dimensional affine subspace from Definition 20.1, so that $(\frac{1}{r(\hat{x})}\tilde{S}, \hat{x})$ is δ -close in the pointed Hausdorff topology to $(\frac{1}{r(\hat{x})}A_{\hat{x}}, \hat{x})$. Here δ is a parameter which will eventually be made small enough so the proof works. Let $A_{\hat{x}}^0 \subset H$ be the k -dimensional linear subspace which is parallel to $A_{\hat{x}}$. Let $p_{\hat{x}} : H \rightarrow H$ be orthogonal projection onto the orthogonal complement of $A_{\hat{x}}^0$.

In view of the assumptions of the lemma, a packing argument shows that for any $l < \infty$, for sufficiently small δ there is a number $m = m(k, C, l)$ so that for all $\hat{x} \in T$, there are at most m elements of T in $B(\hat{x}, lr(\hat{x}))$. Fix a nonnegative function $\phi \in C^\infty(\mathbb{R})$ which is identically one on $[0, 1]$ and vanishes on $[2, \infty)$. For $\hat{x} \in T$, define $\phi_{\hat{x}} : H \rightarrow \mathbb{R}$ by $\phi_{\hat{x}}(v) = \phi(\frac{|v - \hat{x}|}{10r(\hat{x})})$. Let $E(k, H)$ be the set of pairs consisting of a codimension- k plane in H and a point in that plane. That is, $E(k, H)$ is the total space of the universal bundle over $\text{Gr}(k, H)$. Given $v \in \bigcup_{\hat{x} \in T} B(\hat{x}, 5r(\hat{x}))$, put $O_v = \frac{\sum_{\hat{x} \in T} \phi_{\hat{x}}(v) p_{\hat{x}}}{\sum_{\hat{x} \in T} \phi_{\hat{x}}(v)}$. Note that for small δ , there is a uniform upper bound on the number of nonzero terms in the summation, in terms of k and C ; hence the rank of O_v is also bounded in terms of k and C .

If δ is sufficiently small then since the projection operators $p_{\hat{x}}$ that occur with a nonzero coefficient in the summation are uniformly norm-close to each other, the self-adjoint operator O_v will have k eigenvalues near 0, with the rest of the spectrum being near 1. Let $\nu(v)$ be the orthogonal complement of the span of the eigenvectors corresponding to the k eigenvalues of O_v near 0. Let Q_v be orthogonal projection onto $\nu(v)$.

Recall that $\bigcup_{x \in S} B(x, r(x)) \subset \bigcup_{\hat{x} \in T} B(\hat{x}, 5r(\hat{x}))$. Define $\eta : \bigcup_{\hat{x} \in T} B(\hat{x}, 5r(\hat{x})) \rightarrow H$ by

$$(20.3) \quad \eta(v) = \frac{\sum_{\hat{x} \in T} \phi_{\hat{x}}(v) Q_v(v - \hat{x})}{\sum_{\hat{x} \in T} \phi_{\hat{x}}(v)}.$$

Define $\pi : \bigcup_{x \in S} B(x, r(x)) \rightarrow E(k, H)$ by $\pi(v) = (\nu(v), \eta(v))$.

If δ is sufficiently small then π is uniformly transverse to the zero-section of $E(k, H)$. Hence the inverse image under π of the zero section will be a k -dimensional submanifold W . The map P is defined as in the statement of the lemma.

The conclusions of the lemma follow from a convergence argument. For example, for conclusion (3), suppose that there is a sequence $\delta_j \rightarrow 0$ and a collection of counterexamples to conclusion (3). Let $x_j \in S_j$ be the relevant point. In view of the multiplicity bounds, we can assume without loss of generality that the dimension of the Euclidean space is uniformly bounded above. Hence after passing to a subsequence, we can pass to the case when $\dim(H_j)$ is constant in j . Then $\lim_{j \rightarrow \infty} (\frac{1}{r(x_j)}S_j, x_j)$ exists in the pointed Hausdorff topology and is a k -dimensional plane (S_∞, x_∞) . The

map ν_∞ is a constant map and η is an orthogonal projection. Then W_∞ is a flat k -dimensional manifold, which gives a contradiction. The verifications of the other conclusions of the lemma are similar. \square

21. Appendix C : An isotopy lemma

Lemma 21.1. — *Suppose that $F : Y \times [0, 1] \rightarrow N$ is a smooth map between manifolds, with slices $\{f^t : Y \rightarrow N\}_{t \in [0, 1]}$, and let $X \subset N$ be a submanifold with boundary ∂X . If*

- f^t is transverse to both X and ∂X for every $t \in [0, 1]$ and
- $F^{-1}(X)$ is compact

then $(f^0)^{-1}(X)$ is isotopic in Y to $(f^1)^{-1}(X)$.

Proof. — Suppose first that $\partial X = \emptyset$. Now $F(y, t) = f^t(y)$. For $v \in T_y Y$ and $c \in \mathbb{R}$, we can write $DF(v + c \frac{\partial}{\partial t}) = Df^t(v) + c \frac{\partial f^t}{\partial t}(y)$. We know that if $F(y, t) = x \in X$ then $T_{(y,t)}(F^{-1}(X)) = (DF_{y,t})^{-1}(T_x X)$.

By assumption, for each $t \in [0, 1]$, if $f_t(y) = x \in X$ then we have

$$(21.2) \quad \text{Im}(Df_t)_y + T_x X = T_x N.$$

We want to show that projection onto the $[0, 1]$ -factor gives a submersion $F^{-1}(X) \rightarrow [0, 1]$. Suppose not. Then for some $(y, t) \in F^{-1}(X)$, we have $T_{(y,t)}F^{-1}(X) \subset T_y Y$. That is, putting $F(y, t) = x$, whenever $v \in T_y Y$ and $c \in \mathbb{R}$ satisfy $Df^t(v) + c \frac{\partial f^t}{\partial t}(y) \in T_x X$ then we must have $c = 0$. However, for any $c \in \mathbb{R}$, equation (21.2) implies that we can solve $Df_t(-v) + w = c \frac{\partial f^t}{\partial t}(y)$ for some $v \in T_y Y$ and $w \in T_x X$. This is a contradiction.

Thus we have a submersion from the compact set $F^{-1}(X)$ to $[0, 1]$. This submersion must have a product structure, from which the lemma follows.

The case when $\partial X \neq \emptyset$ is similar. \square

Lemma 21.3. — *Suppose that Y is a smooth manifold, $(X, \partial X) \subset \mathbb{R}^k$ is a smooth submanifold, $f : Y \rightarrow \mathbb{R}^k$ is transverse to both X and ∂X , and $\widehat{X} = f^{-1}(X)$ is compact. Then for any compact subset $Y' \subset Y$ whose interior contains \widehat{X} , there is an $\epsilon > 0$ such that if $f' : Y \rightarrow \mathbb{R}^k$ and $\|f' - f\|_{C^1(Y')} < \epsilon$ then $f'^{-1}(X)$ is isotopic to $f^{-1}(X)$.*

Proof. — This follows from Lemma 21.1. \square

References

- [1] M. T. ANDERSON – “The L^2 structure of moduli spaces of Einstein metrics on 4-manifolds”, *Geom. Funct. Anal.* **2** (1992), no. 1, p. 29–89.

- [2] L. BESSIÈRES, G. BESSON, M. BOILEAU, S. MAILLOT & J. PORTI – “Collapsing irreducible 3-manifolds with nontrivial fundamental group”, *Invent. Math.* **179** (2010), no. 2, p. 435–460.
- [3] D. BURAGO, Y. BURAGO & S. IVANOV – *A course in metric geometry*, Grad. Stud. Math., vol. 33, Amer. Math. Soc., Providence, RI, 2001.
- [4] Y. BURAGO, M. GROMOV & G. PEREL'MAN – “A. D. Aleksandrov spaces with curvatures bounded below”, *Uspekhi Mat. Nauk* **47** (1992), no. 2(284), p. 3–51, 222.
- [5] J. CAO & J. GE – “A simple proof of Perelman’s collapsing theorem for 3-manifolds”, *J. Geom. Anal.* **21** (2011), no. 4, p. 807–869.
- [6] J. CHEEGER – “Critical points of distance functions and applications to geometry”, in *Geometric topology: recent developments (Montecatini Terme, 1990)*, Lecture Notes in Math., vol. 1504, Springer, Berlin, 1991, p. 1–38.
- [7] J. CHEEGER & D. GROMOLL – “On the structure of complete manifolds of nonnegative curvature”, *Ann. of Math. (2)* **96** (1972), p. 413–443.
- [8] J. CHEEGER & M. GROMOV – “Collapsing Riemannian manifolds while keeping their curvature bounded. I”, *J. Differential Geom.* **23** (1986), no. 3, p. 309–346.
- [9] ———, “Collapsing Riemannian manifolds while keeping their curvature bounded. II”, *J. Differential Geom.* **32** (1990), no. 1, p. 269–298.
- [10] J. CHEEGER & G. TIAN – “Curvature and injectivity radius estimates for Einstein 4-manifolds”, *J. Amer. Math. Soc.* **19** (2006), no. 2, p. 487–525 (electronic).
- [11] F. H. CLARKE – *Optimization and nonsmooth analysis*, second ed., Classics Appl. Math., vol. 5, Soc. Ind. Appl. Math. (SIAM), Philadelphia, PA, 1990.
- [12] K. FUKAYA – “Collapsing Riemannian manifolds to ones of lower dimensions”, *J. Differential Geom.* **25** (1987), no. 1, p. 139–156.
- [13] M. GROMOV – “Almost flat manifolds”, *J. Differential Geom.* **13** (1978), no. 2, p. 231–241.
- [14] ———, *Metric structures for Riemannian and non-Riemannian spaces*, Progr. Math., vol. 152, Birkhäuser Boston, Inc., Boston, MA, 1999.
- [15] K. GROVE & P. PETERSEN, V – “Bounding homotopy types by geometry”, *Ann. of Math. (2)* **128** (1988), no. 1, p. 195–206.
- [16] K. GROVE & K. SHIOHAMA – “A generalized sphere theorem”, *Ann. of Math. (2)* **106** (1977), no. 2, p. 201–211.
- [17] R. S. HAMILTON – “Three-manifolds with positive Ricci curvature”, *J. Differential Geom.* **17** (1982), no. 2, p. 255–306.
- [18] ———, “Four-manifolds with positive curvature operator”, *J. Differential Geom.* **24** (1986), no. 2, p. 153–179.
- [19] V. KAPOVITCH – “Perelman’s stability theorem”, in *Surveys in differential geometry. Vol. XI*, Surv. Differ. Geom., vol. 11, Int. Press, Somerville, MA, 2007, p. 103–136.
- [20] A. KATSUDA – “Gromov’s convergence theorem and its application”, *Nagoya Math. J.* **100** (1985), p. 11–48.
- [21] B. KLEINER & J. LOTT – “Notes on Perelman’s papers”, *Geom. Topol.* **12** (2008), no. 5, p. 2587–2855.
- [22] S. MATVEEV – *Algorithmic topology and classification of 3-manifolds*, Algorithms Comput. Math., vol. 9, Springer-Verlag, Berlin, 2003.

- [23] J. MORGAN & G. TIAN – *The geometrization conjecture*, Clay Math. Monogr., vol. 5, Amer. Math. Soc./Clay Math. Inst., Providence, RI/Cambridge, MA, 2014.
- [24] G. PERELMAN – “Ricci Flow with Surgery on Three-Manifolds”, <http://arxiv.org/abs/math.DG/0303109>.
- [25] ———, “Alexandrov’s Spaces with Curvature Bounded from Below II”, preprint, 1991.
- [26] P. PETERSEN – *Riemannian geometry*, second ed., Grad. Texts in Math., vol. 171, Springer, New York, 2006.
- [27] C. C. PUGH – “Smoothing a topological manifold”, *Topology Appl.* **124** (2002), no. 3, p. 487–503.
- [28] E. R. REIFENBERG – “Solution of the Plateau Problem for m -dimensional surfaces of varying topological type”, *Acta Math.* **104** (1960), p. 1–92.
- [29] T. SHIOYA & T. YAMAGUCHI – “Collapsing three-manifolds under a lower curvature bound”, *J. Differential Geom.* **56** (2000), no. 1, p. 1–66.
- [30] ———, “Volume collapsed three-manifolds with a lower curvature bound”, *Math. Ann.* **333** (2005), no. 1, p. 131–155.
- [31] M. SIMON – “Ricci flow of almost non-negatively curved three manifolds”, *J. Reine Angew. Math.* **630** (2009), p. 177–217.
- [32] J. STALLINGS – “On fibering certain 3-manifolds”, in *Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute, 1961)*, Prentice-Hall, Englewood Cliffs, N.J., 1962, p. 95–100.
- [33] T. YAMAGUCHI – “Collapsing and pinching under a lower curvature bound”, *Ann. of Math. (2)* **133** (1991), no. 2, p. 317–357.

BRUCE KLEINER, Courant Institute of Mathematical Sciences, 251 Mercer St., New York, NY 10012
E-mail : bkleiner@cims.nyu.edu
 JOHN LOTT, Department of Mathematics, University of California at Berkeley, Berkeley, CA 94720
E-mail : lott@math.berkeley.edu

**GEOMETRIZATION OF THREE-DIMENSIONAL ORBIFOLDS
VIA RICCI FLOW**

by

Bruce Kleiner & John Lott

Abstract. —

Résumé. —

1. Introduction

1.1. Orbifolds and geometrization. — Thurston’s geometrization conjecture for 3-manifolds states that every closed orientable 3-manifold has a canonical decomposition into geometric pieces. In the early 1980’s Thurston announced a proof of the conjecture for Haken manifolds [56], with written proofs appearing much later [36, 41, 47, 48]. The conjecture was settled completely a few years ago by Perelman in his spectacular work using Hamilton’s Ricci flow [49, 50].

2000 Mathematics Subject Classification. — ???

Key words and phrases. — ???

Research supported by NSF grants DMS-0903076 and DMS-1007508.

Thurston also formulated a geometrization conjecture for orbifolds. We recall that orbifolds are similar to manifolds, except that they are locally modelled on quotients of the form \mathbb{R}^n/G , where $G \subset O(n)$ is a finite subgroup of the orthogonal group. Although the terminology is relatively recent, orbifolds have a long history in mathematics, going back to the classification of crystallographic groups and Fuchsian groups. In this paper, using Ricci flow, we will give a new proof of the geometrization conjecture for orbifolds:

Theorem 1.1. — *Let \mathcal{O} be a closed connected orientable three-dimensional orbifold which does not contain any bad embedded 2-dimensional suborbifolds. Then \mathcal{O} has a geometric decomposition.*

The existing proof of Theorem 1.1 is based on a canonical splitting of \mathcal{O} along spherical and Euclidean 2-dimensional suborbifolds, which is analogous to the prime and JSJ decomposition of 3-manifolds. This splitting reduces Theorem 1.1 to two separate cases – when \mathcal{O} is a manifold, and when \mathcal{O} has a nonempty singular locus and satisfies an irreducibility condition. The first case is Perelman’s theorem for manifolds. Thurston announced a proof of the latter case in [57] and gave an outline. A detailed proof of the latter case was given by Boileau-Leeb-Porti [4], after work of Boileau-Maillot-Porti [5], Boileau-Porti [6], Cooper-Hodgson-Kerckhoff [19] and Thurston [57]. The monographs [5, 19] give excellent expositions of 3-orbifolds and their geometrization.

1.2. Discussion of the proof. — The main purpose of this paper is to provide a new proof of Theorem 1.1. Our proof is an extension of Perelman’s proof of geometrization for 3-manifolds to orbifolds, bypassing [4–6, 19, 57]. The motivation for this alternate approach is twofold. First, anyone interested in the geometrization of general orbifolds as in Theorem 1.1 will necessarily have to go through Perelman’s Ricci flow proof in the manifold case, and also absorb foundational results about orbifolds. At that point, the additional effort required to deal with general orbifolds is relatively minor in comparison to the proof in [4]. This latter proof involves a number of ingredients, including Thurston’s geometrization of Haken manifolds, the deformation and collapsing theory of hyperbolic cone manifolds, and some Alexandrov space theory. Also, in contrast to the existing proof of Theorem 1.1, the Ricci flow argument gives a unified approach to geometrization for both manifolds and orbifolds.

Many of the steps in Perelman’s proof have evident orbifold generalizations, whereas some do not. It would be unwieldy to rewrite all the details of Perelman’s proof, on the level of [38], while making notational changes from manifolds to orbifolds. Consequently, we focus on the steps in Perelman’s proof where an orbifold extension is not immediate. For a step where the orbifold extension is routine, we make the precise orbifold statement and indicate where the analogous manifold proof occurs in [38].

In the course of proving Theorem 1.1, we needed to develop a number of foundational results about the geometry of orbifolds. Some of these may be of independent interest, or of use for subsequent work in this area, such as the compactness theorem for Riemannian orbifolds, critical point theory, and the soul theorem.

Let us mention one of the steps where the orbifold extension could *a priori* be an issue. This is where one characterizes the topology of the thin part of the large-time orbifold. To do this, one first needs a sufficiently flexible proof in the manifold case. We provided such a proof in [37]. The proof in [37] uses some basic techniques from Alexandrov geometry, combined with smoothness results in appropriate places. It provides a decomposition of the thin part into various pieces which together give an explicit realization of the thin part as a graph manifold. When combined with preliminary results that are proved in this paper, we can extend the techniques of [37] to orbifolds. We get a decomposition of the thin part of the large-time orbifold into various pieces, similar to those in [37]. We show that these pieces give an explicit realization of each component of the thin part as either a graph orbifold or one of a few exceptional cases. This is more involved to prove in the orbifold case than in the manifold case but the basic strategy is the same.

1.3. Organization of the paper. — The structure of this paper is as follows. One of our tasks is to provide a framework for the topology and Riemannian geometry of orbifolds, so that results about Ricci flow on manifolds extend as easily as possible to orbifolds. In Section 2 we recall the relevant notions that we need from orbifold topology. We then introduce Riemannian orbifolds and prove the orbifold versions of some basic results from Riemannian geometry, such as the de Rham decomposition and critical point theory.

Section 3 is concerned with noncompact nonnegatively curved orbifolds. We prove the orbifold version of the Cheeger-Gromoll soul theorem. We list the diffeomorphism types of noncompact nonnegatively curved orbifolds with dimension at most three.

In Section 4 we prove a compactness theorem for Riemannian orbifolds. Section 5 contains some preliminary information about Ricci flow on orbifolds, along with the classification of the diffeomorphism types of compact nonnegatively curved three-dimensional orbifolds. We also show how to extend Perelman's no local collapsing theorem to orbifolds.

Section 6 is devoted to κ -solutions. Starting in Section 7, we specialize to three-dimensional orientable orbifolds with no bad 2-dimensional suborbifolds. We show how to extend Perelman's results in order to construct a Ricci flow with surgery.

In Section 8 we show that the thick part of the large-time geometry approaches a finite-volume orbifold of constant negative curvature. Section 9 contains the topological characterization of the thin part of the large-time geometry.

Section 10 concerns the incompressibility of hyperbolic cross-sections. Rather than using minimal disk techniques as initiated by Hamilton [33], we follow an approach

introduced by Perelman [50, Section 8] that uses a monotonic quantity, as modified in [38, Section 93.4].

The appendix contains topological facts about graph orbifolds. We show that a “weak” graph orbifold is the result of performing 0-surgeries (*i.e.*, connected sums) on a “strong” graph orbifold. This material is probably known to some experts but we were unable to find references in the literature, so we include complete proofs.

After writing this paper we learned that Daniel Faessler independently proved Proposition 9.7, which is the orbifold version of the collapsing theorem [24].

Acknowledgements. — We thank Misha Kapovich and Sylvain Maillot for orbidiscussions. We thank the referee for a careful reading of the paper and for corrections.

2. Orbifold topology and geometry

In this section we first review the differential topology of orbifolds. Subsections 2.1 and 2.2 contain information about orbifolds in any dimension. In some cases we give precise definitions and in other cases we just recall salient properties, referring to the monographs [5, 19] for more detailed information. Subsections 2.3 and 2.4 are concerned with low-dimensional orbifolds.

We then give a short exposition of aspects of the differential geometry of orbifolds, in Subsection 2.5. It is hard to find a comprehensive reference for this material and so we flag the relevant notions; see [8] for further discussion of some points. Subsection 2.6 shows how to do critical point theory on orbifolds. Subsection 2.7 discusses the smoothing of functions on orbifolds.

For notation, B^n is the open unit n -ball, D^n is the closed unit n -ball and $I = [-1, 1]$. We let D_k denote the dihedral group of order $2k$.

2.1. Differential topology of orbifolds. — An *orbivector space* is a triple (V, G, ρ) , where

- V is a vector space,
- G is a finite group and
- $\rho : G \rightarrow \text{Aut}(V)$ is a faithful linear representation.

A (closed/ open/ convex/...) *subset* of (V, G, ρ) is a G -invariant subset of V which is (closed/ open/ convex/...) A *linear map* from (V, G, ρ) to (V', G', ρ') consists of a linear map $T : V \rightarrow V'$ and a homomorphism $h : G \rightarrow G'$ so that for all $g \in G$, $\rho'(h(g)) \circ T = T \circ \rho(g)$. The linear map is *injective* (resp. *surjective*) if T is *injective* (resp. *surjective*) and h is *injective* (resp. *surjective*). An *action* of a group K on (V, G, ρ) is given by a short exact sequence $1 \rightarrow G \rightarrow L \rightarrow K \rightarrow 1$ and a homomorphism $L \rightarrow \text{Aut}(V)$ that extends ρ .

A *local model* is a pair (\widehat{U}, G) , where \widehat{U} is a connected open subset of a Euclidean space and G is a finite group that acts smoothly and effectively on \widehat{U} , on the right.

(Effectiveness means that the homomorphism $G \rightarrow \text{Diff}(\widehat{U})$ is injective.) We will sometimes write U for \widehat{U}/G , endowed with the quotient topology.

A *smooth map* between local models (\widehat{U}_1, G_1) and (\widehat{U}_2, G_2) is given by a smooth map $\widehat{f} : \widehat{U}_1 \rightarrow \widehat{U}_2$ and a homomorphism $\rho : G_1 \rightarrow G_2$ so that \widehat{f} is ρ -equivariant, i.e., $\widehat{f}(xg_1) = \widehat{f}(x)\rho(g_1)$. We do not assume that ρ is injective or surjective. The map between local models is an *embedding* if \widehat{f} is an embedding; it follows from effectiveness that ρ is injective in this case.

Definition 2.1. — An *atlas* for an n -dimensional orbifold \mathcal{O} consists of

1. A Hausdorff paracompact topological space $|\mathcal{O}|$,
2. An open covering $\{U_\alpha\}$ of $|\mathcal{O}|$,
3. Local models $\{(\widehat{U}_\alpha, G_\alpha)\}$ with each \widehat{U}_α a connected open subset of \mathbb{R}^n and
4. Homeomorphisms $\phi_\alpha : U_\alpha \rightarrow \widehat{U}_\alpha/G_\alpha$ so that
5. If $p \in U_1 \cap U_2$ then there is a local model (\widehat{U}_3, G_3) with $p \in U_3$ along with embeddings $(\widehat{U}_3, G_3) \rightarrow (\widehat{U}_1, G_1)$ and $(\widehat{U}_3, G_3) \rightarrow (\widehat{U}_2, G_2)$.

An *orbifold* \mathcal{O} is an equivalence class of such atlases, where two atlases are equivalent if they are both included in a third atlas. With a given atlas, the orbifold \mathcal{O} is *oriented* if each \widehat{U}_α is oriented, the action of G_α is orientation-preserving, and the embeddings $\widehat{U}_3 \rightarrow \widehat{U}_1$ and $\widehat{U}_3 \rightarrow \widehat{U}_2$ are orientation-preserving. We say that \mathcal{O} is *connected* (resp. *compact*) if $|\mathcal{O}|$ is connected (resp. compact).

An *orbifold-with-boundary* \mathcal{O} is defined similarly, with \widehat{U}_α being a connected open subset of $[0, \infty) \times \mathbb{R}^{n-1}$. The *boundary* $\partial\mathcal{O}$ is a boundaryless $(n - 1)$ -dimensional orbifold, with $|\partial\mathcal{O}|$ consisting of points in $|\mathcal{O}|$ whose local lifts lie in $\{0\} \times \mathbb{R}^{n-1}$. Note that it is possible that $\partial\mathcal{O} = \emptyset$ while $|\mathcal{O}|$ is a topological manifold with a nonempty boundary.

Remark 2.2. — In this paper we only deal with *effective* orbifolds, meaning that in a local model (\widehat{U}, G) , the group G always acts effectively. It would be more natural in some ways to remove this effectiveness assumption. However, doing so would hurt the readability of the paper, so we will stick to effective orbifolds.

Given a point $p \in |\mathcal{O}|$ and a local model (\widehat{U}, G) around p , let $\widehat{p} \in \widehat{U}$ project to p . The *local group* G_p is the stabilizer group $\{g \in G : \widehat{p}g = \widehat{p}\}$. Its isomorphism class is independent of the choices made. We can always find a local model with $G = G_p$.

The *regular part* $|\mathcal{O}|_{reg} \subset |\mathcal{O}|$ consists of the points with $G_p = \{e\}$. It is a smooth manifold that forms an open dense subset of $|\mathcal{O}|$.

Given an open subset $X \subset |\mathcal{O}|$, there is an induced orbifold $\mathcal{O}|_X$ with $|\mathcal{O}|_X = X$. In some cases we will have a subset $X \subset |\mathcal{O}|$, possibly not open, for which $\mathcal{O}|_X$ is an orbifold-with-boundary.

The *ends* of \mathcal{O} are the ends of $|\mathcal{O}|$.

A smooth map $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ between orbifolds is given by a continuous map $|f| : |\mathcal{O}_1| \rightarrow |\mathcal{O}_2|$ with the property that for each $p \in |\mathcal{O}_1|$, there are

- Local models (\widehat{U}_1, G_1) and (\widehat{U}_2, G_2) for p and $f(p)$, respectively, and
- A smooth map $\widehat{f} : (\widehat{U}_1, G_1) \rightarrow (\widehat{U}_2, G_2)$ between local models

so that the diagram

$$(2.3) \quad \begin{array}{ccc} \widehat{U}_1 & \xrightarrow{\widehat{f}} & \widehat{U}_2 \\ \downarrow & & \downarrow \\ U_1 & \xrightarrow{|f|} & U_2 \end{array}$$

commutes.

There is an induced homomorphism from G_p to $G_{f(p)}$. We emphasize that to define a smooth map f between two orbifolds, one must first define a map $|f|$ between their underlying spaces.

We write $C^\infty(\mathcal{O})$ for the space of smooth maps $f : \mathcal{O} \rightarrow \mathbb{R}$.

A smooth map $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is *proper* if $|f| : |\mathcal{O}_1| \rightarrow |\mathcal{O}_2|$ is a proper map.

A *diffeomorphism* $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a smooth map with a smooth inverse. Then G_p is isomorphic to $G_{f(p)}$.

If a discrete group Γ acts properly discontinuously on a manifold M then there is a quotient orbifold, which we denote by $M//\Gamma$. It has $|M//\Gamma| = M/\Gamma$. Hence if \mathcal{O} is an orbifold and (\widehat{U}, G) is a local model for \mathcal{O} then we can say that $\mathcal{O}|_U$ is diffeomorphic to $\widehat{U}//G$. An orbifold \mathcal{O} is *good* if $\mathcal{O} = M//\Gamma$ for some manifold M and some discrete group Γ . It is *very good* if Γ can be taken to be finite. A *bad* orbifold is one that is not good.

Similarly, suppose that a discrete group Γ acts by diffeomorphisms on an orbifold \mathcal{O} . We say that it acts *properly discontinuously* if the action of Γ on $|\mathcal{O}|$ is properly discontinuous. Then there is a quotient orbifold $\mathcal{O}//\Gamma$, with $|\mathcal{O}//\Gamma| = |\mathcal{O}|/\Gamma$; see Remark 2.15.

An *orbifiber bundle* consists of a smooth map $\pi : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ between two orbifolds, along with a third orbifold \mathcal{O}_3 such that

- $|\pi|$ is surjective, and
- For each $p \in |\mathcal{O}_2|$, there is a local model (\widehat{U}, G_p) around p , where G_p is the local group at p , along with an action of G_p on \mathcal{O}_3 and a diffeomorphism $(\mathcal{O}_3 \times \widehat{U})//G_p \rightarrow \mathcal{O}_1|_{|\pi|^{-1}(U)}$ so that the diagram

$$(2.4) \quad \begin{array}{ccc} (\mathcal{O}_3 \times \widehat{U})//G_p & \longrightarrow & \mathcal{O}_1 \\ \downarrow & & \downarrow \\ \widehat{U}//G_p & \longrightarrow & \mathcal{O}_2 \end{array}$$

commutes.

(Note that if \mathcal{O}_2 is a manifold then the orbifold bundle $\pi : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ has a local product structure.) The *fiber* of the orbifold bundle is \mathcal{O}_3 . Note that for $p_1 \in |\mathcal{O}_1|$, the homomorphism $G_{p_1} \rightarrow G_{|\pi|^{-1}(p_1)}$ is surjective.

A *section* of an orbifold bundle $\pi : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a smooth map $s : \mathcal{O}_2 \rightarrow \mathcal{O}_1$ such that $\pi \circ s$ is the identity on \mathcal{O}_2 .

A *covering map* $\pi : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is an orbifold bundle with a zero-dimensional fiber. Given $p_2 \in |\mathcal{O}_2|$ and $p_1 \in |\pi|^{-1}(p_2)$, there are a local model (\widehat{U}, G_2) around p_2 and a subgroup $G_1 \subset G_2$ so that (\widehat{U}, G_1) is a local model around p_1 and the map π is locally $(\widehat{U}, G_1) \rightarrow (\widehat{U}, G_2)$.

A rank- m *orbivector bundle* $\mathcal{V} \rightarrow \mathcal{O}$ over \mathcal{O} is locally isomorphic to $(V \times \widehat{U})/G_p$, where V is an m -dimensional orbivector space on which G_p acts linearly.

The *tangent bundle* $T\mathcal{O}$ of an orbifold \mathcal{O} is an orbivector bundle which is locally diffeomorphic to $T\widehat{U}_\alpha//G_\alpha$. Given $p \in |\mathcal{O}|$, if $\widehat{p} \in \widehat{U}$ covers p then the *tangent space* $T_p\mathcal{O}$ is isomorphic to the orbivector space $(T_{\widehat{p}}\widehat{U}, G_p)$. The *tangent cone* at p is $C_p|\mathcal{O}| \cong T_{\widehat{p}}\widehat{U}/G_p$.

A smooth *vector field* V is a smooth section of $T\mathcal{O}$. In terms of a local model (\widehat{U}, G) , the vector field V restricts to a vector field on \widehat{U} which is G -invariant.

A smooth map $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ gives rise to the *differential*, an orbivector bundle map $df : T\mathcal{O}_1 \rightarrow T\mathcal{O}_2$. At a point $p \in |\mathcal{O}|$, in terms of local models we have a map $\widehat{f} : (\widehat{U}_1, G_1) \rightarrow (\widehat{U}_2, G_2)$ which gives rise to a G_p -equivariant map $d\widehat{f}_p : T_{\widehat{p}}\widehat{U}_1 \rightarrow T_{\widehat{f}(\widehat{p})}\widehat{U}_2$ and hence to a linear map $df_p : T_p\mathcal{O}_1 \rightarrow T_{|f|(p)}\mathcal{O}_2$.

Given a smooth map $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ and a point $p \in |\mathcal{O}_1|$, we say that f is a *submersion at p* (resp. *immersion at p*) if the map $df_p : T_p\mathcal{O}_1 \rightarrow T_{|f|(p)}\mathcal{O}_2$ is surjective (resp. injective).

Lemma 2.5. — *If f is a submersion at p then there is an orbifold \mathcal{O}_3 on which $G_{|f|(p)}$ acts, along with a local model $(\widehat{U}_2, G_{|f|(p)})$ around $|f|(p)$, so that f is equivalent near p to the projection map $(\mathcal{O}_3 \times \widehat{U}_2)//G_{|f|(p)} \rightarrow \widehat{U}_2//G_{|f|(p)}$.*

Proof. — Let $\rho : G_p \rightarrow G_{|f|(p)}$ be the surjective homomorphism associated to df_p . Let $\widehat{f} : (\widehat{U}_1, G_p) \rightarrow (\widehat{U}_2, G_{|f|(p)})$ be a local model for f near p ; it is necessarily ρ -equivariant. Let $\widehat{p} \in \widehat{U}_1$ be a lift of $p \in U_1$. Put $\widehat{W} = \widehat{f}^{-1}(\widehat{f}(\widehat{p}))$. Since \widehat{f} is a submersion at \widehat{p} , after reducing \widehat{U}_1 and \widehat{U}_2 if necessary, there is a ρ -equivariant diffeomorphism $\widehat{W} \times \widehat{U}_2 \rightarrow \widehat{U}_1$ so that the diagram

$$(2.6) \quad \begin{array}{ccc} \widehat{W} \times \widehat{U}_2 & \longrightarrow & \widehat{U}_1 \\ \downarrow & & \downarrow \\ \widehat{U}_2 & \longrightarrow & \widehat{U}_2 \end{array}$$

commutes and is G_p -equivariant. Now $\text{Ker}(\rho)$ acts on \widehat{W} . Put $\mathcal{O}_3 = \widehat{W}//\text{Ker}(\rho)$.

Then there is a commuting diagram of orbifold maps

$$(2.7) \quad \begin{array}{ccc} \mathcal{O}_3 \times \widehat{U}_2 & \longrightarrow & \widehat{U}_1 // \text{Ker}(\rho) \\ \downarrow & & \downarrow \\ \widehat{U}_2 & \longrightarrow & \widehat{U}_2. \end{array}$$

Further quotienting by $G_{|f|(p)}$ gives a commutative diagram

$$(2.8) \quad \begin{array}{ccc} (\mathcal{O}_3 \times \widehat{U}_2) // G_{|f|(p)} & \longrightarrow & \widehat{U}_1 // G_p \\ \downarrow & & \downarrow \\ \widehat{U}_2 // G_{|f|(p)} & \longrightarrow & \widehat{U}_2 // G_{|f|(p)} \end{array}$$

whose top horizontal line is an orbifold diffeomorphism. \square

We say that $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a *submersion* (resp. *immersion*) if it is a submersion (resp. immersion) at p for all $p \in |\mathcal{O}_1|$.

Lemma 2.9. — *A proper surjective submersion $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$, with \mathcal{O}_2 connected, defines an orbifold bundle with compact fibers.*

We will sketch a proof of Lemma 2.9 in Remark 2.17.

In particular, a proper surjective local diffeomorphism to a connected orbifold is a covering map with finite fibers.

An immersion $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ has a *normal bundle* $N\mathcal{O}_1 \rightarrow \mathcal{O}_1$ whose fibers have the following local description. Given $p \in |\mathcal{O}_1|$, let f be described in terms of local models (\widehat{U}_1, G_p) and $(\widehat{U}_2, G_{|f|(p)})$ by a ρ -equivariant immersion $\widehat{f} : \widehat{U}_1 \rightarrow \widehat{U}_2$. Let $F_p \subset G_{|f|(p)}$ be the subgroup which fixes $\text{Im}(d\widehat{f}_p)$. Then the *normal space* $N_p\mathcal{O}_1$ is the orbivector space $(\text{Coker}(d\widehat{f}_p), F_p)$.

A *suborbifold* of \mathcal{O} is given by an orbifold \mathcal{O}' and an immersion $f : \mathcal{O}' \rightarrow \mathcal{O}$ for which $|f|$ maps $|\mathcal{O}'|$ homeomorphically to its image in $|\mathcal{O}|$. From effectiveness, for each $p \in |\mathcal{O}'|$, the homomorphism $\rho_p : G_p \rightarrow G_{|f|(p)}$ is injective. Note that ρ_p need not be an isomorphism. We will identify \mathcal{O}' with its image in \mathcal{O} . There is a neighborhood of \mathcal{O}' which is diffeomorphic to the normal bundle $N\mathcal{O}'$. We say that the suborbifold \mathcal{O}' is *embedded* if $\mathcal{O}|_{|\mathcal{O}'|} = \mathcal{O}'$. Then for each $p \in |\mathcal{O}'|$, the homomorphism ρ_p is an isomorphism.

If \mathcal{O}' is an embedded codimension-1 suborbifold of \mathcal{O} then we say that \mathcal{O}' is *two-sided* if the normal bundle $N\mathcal{O}'$ has a nowhere-zero section. If \mathcal{O} and \mathcal{O}' are both orientable then \mathcal{O}' is two-sided. We say that \mathcal{O}' is *separating* if $|\mathcal{O}'|$ is separating in $|\mathcal{O}|$.

We can talk about two suborbifolds meeting *transversely*, as defined using local models.

Let \mathcal{O} be an oriented orbifold (possibly disconnected). Let D_1 and D_2 be disjoint codimension-zero embedded suborbifolds-with-boundary, both oriented-diffeomorphic to $D^n//\Gamma$. Then the operation of performing *0-surgery along D_1, D_2* produces the new oriented orbifold $\mathcal{O}' = (\mathcal{O} - \text{int}(D_1) - \text{int}(D_2)) \cup_{\partial D_1 \sqcup \partial D_2} (I \times (D^n//\Gamma))$. In the manifold case, a connected sum is the same thing as a 0-surgery along a pair $\{D_1, D_2\}$ which lie in different connected components of \mathcal{O} . Note that unlike in the manifold case, \mathcal{O}' is generally not uniquely determined up to diffeomorphism by knowing the connected components containing D_1 and D_2 . For example, even if \mathcal{O} is connected, D_1 and D_2 may or may not lie on the same connected component of the singular set.

If \mathcal{O}_1 and \mathcal{O}_2 are oriented orbifolds, with $D_1 \subset \mathcal{O}_1$ and $D_2 \subset \mathcal{O}_2$ both oriented diffeomorphic to $D^n//\Gamma$, then we may write $\mathcal{O}_1 \#_{S^{n-1} // \Gamma} \mathcal{O}_2$ for the connected sum. This notation is slightly ambiguous since the location of D_1 and D_2 is implicit. We will write $\mathcal{O} \#_{S^{n-1} // \Gamma}$ to denote a 0-surgery on a single orbifold \mathcal{O} . Again the notation is slightly ambiguous, since the location of $D_1, D_2 \subset \mathcal{O}$ is implicit.

An *involutive distribution* on \mathcal{O} is a subbundle $E \subset T\mathcal{O}$ with the property that for any two sections V_1, V_2 of E , the Lie bracket $[V_1, V_2]$ is also a section of E .

Lemma 2.10. — *Given an involutive distribution E on \mathcal{O} , for any $p \in |\mathcal{O}|$ there is a unique maximal suborbifold passing through p which is tangent to E .*

Orbifolds have partitions of unity.

Lemma 2.11. — *Given an open cover $\{U_\alpha\}_{\alpha \in A}$ of $|\mathcal{O}|$, there is a collection of functions $\rho_\alpha \in C^\infty(\mathcal{O})$ such that*

- $0 \leq \rho_\alpha \leq 1$.
- $\text{supp}(\rho_\alpha) \subset U_{\alpha'}$ for some $\alpha' = \alpha'(\alpha) \in A$.
- For all $p \in |\mathcal{O}|$, $\sum_{\alpha \in A} \rho_\alpha(p) = 1$.

Proof. — The proof is similar to the manifold case, using local models (\widehat{U}, G) consisting of coordinate neighborhoods, along with compactly supported G -invariant smooth functions on \widehat{U} . □

A *curve* in an orbifold is a smooth map $\gamma : I \rightarrow \mathcal{O}$ defined on an interval $I \subset \mathbb{R}$. A *loop* is a curve γ with $|\gamma|(0) = |\gamma|(1) \in |\mathcal{O}|$.

2.2. Universal cover and fundamental group. — We follow the presentation in [5, Chapter 2.2.1]. Choose a regular point $p \in |\mathcal{O}|$. A *special curve* from p is a curve $\gamma : [0, 1] \rightarrow \mathcal{O}$ such that

- $|\gamma|(0) = p$ and
- $|\gamma|(t)$ lies in $|\mathcal{O}|_{\text{reg}}$ for all but a finite number of t .

Suppose that (\widehat{U}, G) is a local model and that $\widehat{\gamma} : [a, b] \rightarrow \widehat{U}$ is a lifting of $\gamma_{[a,b]}$, for some $[a, b] \subset [0, 1]$. An *elementary homotopy* between two special curves is a smooth homotopy of $\widehat{\gamma}$ in \widehat{U} , relative to $\widehat{\gamma}(a)$ and $\widehat{\gamma}(b)$. A *homotopy* of γ is what's generated by elementary homotopies.

If \mathcal{O} is connected then the *universal cover* $\widetilde{\mathcal{O}}$ of \mathcal{O} can be constructed as the set of special curves starting at p , modulo homotopy. It has a natural orbifold structure. The *fundamental group* $\pi_1(\mathcal{O}, p)$ is given by special loops (*i.e.*, special curves γ with $|\gamma|(1) = p$) modulo homotopy. Up to isomorphism, $\pi_1(\mathcal{O}, p)$ is independent of the choice of p .

If \mathcal{O} is connected and a discrete group Γ acts properly discontinuously on \mathcal{O} then there is a short exact sequence

$$(2.12) \quad 1 \longrightarrow \pi_1(\mathcal{O}, p) \longrightarrow \pi_1(\mathcal{O}/\Gamma, p\Gamma) \longrightarrow \Gamma \longrightarrow 1.$$

Remark 2.13. — A more enlightening way to think of an orbifold is to consider it as a smooth effective proper étale groupoid \mathcal{G} , as explained in [1, 12, 44]. We recall that a *Lie groupoid* \mathcal{G} essentially consists of a smooth manifold $\mathcal{G}^{(0)}$ (the space of units), another smooth manifold $\mathcal{G}^{(1)}$ and submersions $s, r : \mathcal{G}^{(1)} \rightarrow \mathcal{G}^{(0)}$ (the source and range maps), along with a partially defined multiplication $\mathcal{G}^{(1)} \times \mathcal{G}^{(1)} \rightarrow \mathcal{G}^{(1)}$ which satisfies certain compatibility conditions. A Lie groupoid is *étale* if s and r are local diffeomorphisms. It is *proper* if $(s, r) : \mathcal{G}^{(1)} \rightarrow \mathcal{G}^{(0)} \times \mathcal{G}^{(0)}$ is a proper map. There is also a notion of an étale groupoid being *effective*.

To an orbifold one can associate an effective proper étale groupoid as follows. Given an orbifold \mathcal{O} , a local model $(\widehat{U}_\alpha, G_\alpha)$ and some $\widehat{p}_\alpha \in \widehat{U}_\alpha$, let $p \in |\mathcal{O}|$ be the corresponding point. There is a quotient map $A_{\widehat{p}_\alpha} : T_{\widehat{p}_\alpha} \widehat{U}_\alpha \rightarrow C_p|\mathcal{O}|$. The unit space $\mathcal{G}^{(0)}$ is the disjoint union of the \widehat{U}_α 's. And $\mathcal{G}^{(1)}$ consists of the triples $(\widehat{p}_\alpha, \widehat{p}_\beta, B_{\widehat{p}_\alpha, \widehat{p}_\beta})$ where

1. $\widehat{p}_\alpha \in \widehat{U}_\alpha$ and $\widehat{p}_\beta \in \widehat{U}_\beta$,
2. \widehat{p}_α and \widehat{p}_β map to the same point $p \in |\mathcal{O}|$ and
3. $B_{\widehat{p}_\alpha, \widehat{p}_\beta} : T_{\widehat{p}_\alpha} \widehat{U}_\alpha \rightarrow T_{\widehat{p}_\beta} \widehat{U}_\beta$ is an invertible linear map so that $A_{\widehat{p}_\alpha} = A_{\widehat{p}_\beta} \circ B_{\widehat{p}_\alpha, \widehat{p}_\beta}$.

There is an obvious way to compose triples $(\widehat{p}_\alpha, \widehat{p}_\beta, B_{\widehat{p}_\alpha, \widehat{p}_\beta})$ and $(\widehat{p}_\beta, \widehat{p}_\gamma, B_{\widehat{p}_\beta, \widehat{p}_\gamma})$. One can show that this gives rise to a smooth effective proper étale groupoid.

Conversely, given a smooth effective proper étale groupoid \mathcal{G} , for any $\widehat{p} \in \mathcal{G}^{(0)}$ the isotropy group $\mathcal{G}_{\widehat{p}}^{\widehat{p}}$ is a finite group. To get an orbifold, one can take local models of the form $(\widehat{U}, \mathcal{G}_{\widehat{p}}^{\widehat{p}})$ where \widehat{U} is a $\mathcal{G}_{\widehat{p}}^{\widehat{p}}$ -invariant neighborhood of \widehat{p} .

Speaking hereafter just of smooth effective proper étale groupoids, Morita-equivalent groupoids give equivalent orbifolds.

A groupoid morphism gives rise to an orbifold map. Taking into account Morita equivalence, from the groupoid viewpoint the right notion of an orbifold map would be a Hilsun-Skandalis map between groupoids. These turn out to correspond to *good maps* between orbifolds, as later defined by Chen-Ruan [1]. This is a more restricted class of maps between orbifolds than what we consider. The distinction is that one can

pull back orbivector bundles under good maps, but not always under smooth maps in our sense. Orbifold diffeomorphisms in our sense are automatically good maps. For some purposes it would be preferable to only deal with good maps, but for simplicity we will stick with our orbifold definitions.

A Lie groupoid \mathcal{G} has a classifying space $B\mathcal{G}$. In the orbifold case, if \mathcal{G} is the étale groupoid associated to an orbifold \mathcal{O} then $\pi_1(\mathcal{O}) \cong \pi_1(B\mathcal{G})$. The definition of the latter can be made explicit in terms of paths and homotopies; see [12, 29]. In the case of effective orbifolds, the definition is equivalent to the one of the present paper.

More information is in [1, 44] and references therein.

2.3. Low-dimensional orbifolds. — We list the connected compact boundaryless orbifolds of low dimension. We mostly restrict here to the orientable case. (The nonorientable ones also arise; even if the total space of an orbifiber bundle is orientable, the base may fail to be orientable.)

2.3.1. Zero dimensions. — The only possibility is a point.

2.3.2. One dimension. — There are two possibilities : S^1 and $S^1//\mathbb{Z}_2$. For the latter, the nonzero element of \mathbb{Z}_2 acts by complex conjugation on S^1 , and $|S^1//\mathbb{Z}_2|$ is an interval. Note that $S^1//\mathbb{Z}_2$ is not orientable.

2.3.3. Two dimensions. — For notation, if S is a connected oriented surface then $S(k_1, \dots, k_r)$ denotes the oriented orbifold \mathcal{O} with $|\mathcal{O}| = S$, having singular points of order $k_1, \dots, k_r > 1$. Any connected oriented 2-orbifold can be written in this way. An orbifold of the form $S^2(p, q, r)$ is called a *turnover*.

The *bad* orientable 2-orbifolds are $S^2(k)$ and $S^2(k, k')$, $k \neq k'$. The latter is simply-connected if and only if $\gcd(k, k') = 1$.

The *spherical* 2-orbifolds are of the form $S^2//\Gamma$, where Γ is a finite subgroup of $\text{Isom}^+(S^2)$. The orientable ones are S^2 , $S^2(k, k)$, $S^2(2, 2, k)$, $S^2(2, 3, 3)$, $S^2(2, 3, 4)$, $S^2(2, 3, 5)$. (If $S^2(1, 1)$ arises in this paper then it means S^2 .)

The *Euclidean* 2-orbifolds are of the form $T^2//\Gamma$, where Γ is a finite subgroup of $\text{Isom}^+(T^2)$. The orientable ones are T^2 , $S^2(2, 3, 6)$, $S^2(2, 4, 4)$, $S^2(3, 3, 3)$, $S^2(2, 2, 2, 2)$. The latter is called a *pillowcase* and can be identified with the quotient of $T^2 = \mathbb{C}/\mathbb{Z}^2$ by \mathbb{Z}_2 , where the action of the nontrivial element of \mathbb{Z}_2 comes from the map $z \rightarrow -z$ on \mathbb{C} .

The other closed orientable 2-orbifolds are hyperbolic.

We will also need some 2-orbifolds with boundary, namely

- The *discal* 2-orbifolds $D^2(k) = D^2//\mathbb{Z}_k$.
- The *half-pillowcase* $D^2(2, 2) = I \times_{\mathbb{Z}_2} S^1$. Here the nontrivial element of \mathbb{Z}_2 acts by involution on I and by complex conjugation on S^1 . We can also write $D^2(2, 2)$ as the quotient $\{z \in \mathbb{C} : \frac{1}{2} \leq |z| \leq 2\}/\mathbb{Z}_2$, where the nontrivial element of \mathbb{Z}_2 sends z to z^{-1} .

- $D^2//\mathbb{Z}_2$, where \mathbb{Z}_2 acts by complex conjugation on D^2 . Then $\partial|D^2//\mathbb{Z}_2|$ is a circle with one orbifold boundary component and one reflector component. See Figure 1, where the dark line indicates the reflector component.




FIGURE 1.

- $D^2//D_k = D^2(k)//\mathbb{Z}_2$, for $k > 1$, where D_k is the dihedral group and \mathbb{Z}_2 acts by complex conjugation on $D^2(k)$. Then $\partial|D^2//D_k|$ is a circle with one orbifold boundary component, one corner reflector point of order k and two reflector components. See Figure 2.




FIGURE 2.

2.3.4. Three dimensions. — If \mathcal{O} is an orientable three-dimensional orbifold then $|\mathcal{O}|$ is an orientable topological 3-manifold. If \mathcal{O} is boundaryless then $|\mathcal{O}|$ is boundaryless. Each component of the singular locus in $|\mathcal{O}|$ is either

1. a knot or arc (with endpoints on $\partial|\mathcal{O}|$), labelled by an integer greater than one, or
2. a trivalent graph with each edge labelled by an integer greater than one, under the constraint that if edges with labels p, q, r meet at a vertex then $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} > 1$.

That is, there is a neighborhood of the vertex which is a cone over an orientable spherical 2-orbifold.

Specifying such a topological 3-manifold and such a labelled graph is equivalent to specifying an orientable three-dimensional orbifold.

We write $D^3//\Gamma$ for a *discal* 3-orbifold whose boundary is $S^2//\Gamma$. They are

- D^3 . There is no singular locus.
- $D^3(k, k)$. The singular locus is a line segment through D^3 . See Figure 3.




FIGURE 3.

- $D^3(2, 2, k)$, $D^3(2, 3, 3)$, $D^3(2, 3, 4)$ and $D^3(2, 3, 5)$. The singular locus is a tripod in D^3 . See Figure 4.




FIGURE 4.

The *solid-toric* 3-orbifolds are

- $S^1 \times D^2$. There is no singular locus.
- $S^1 \times D^2(k)$. The singular locus is a core curve in a solid torus. See Figure 5




FIGURE 5.

- $S^1 \times_{\mathbb{Z}_2} D^2$. The singular locus consists of two arcs in a 3-disk, each labelled by 2. The boundary is $S^2(2, 2, 2, 2)$. See Figure 6.




FIGURE 6.

- $S^1 \times_{\mathbb{Z}_2} D^2(k)$. The singular locus consists of two arcs in a 3-disk, each labelled by 2, joined in their middles by an arc labelled by k . The boundary is $S^2(2, 2, 2, 2)$. See Figure 7.

Given $\Gamma \in \text{Isom}^+(S^2)$, we can consider the quotient $S^3//\Gamma$ where Γ acts on S^3 by the suspension of its action on S^2 . That is, we are identifying $\text{Isom}^+(S^2)$ with $\text{SO}(3)$ and using the embedding $\text{SO}(3) \rightarrow \text{SO}(4)$ to let Γ act on S^3 .




FIGURE 7.

An orientable three-dimensional orbifold \mathcal{O} is *irreducible* if it contains no embedded bad 2-dimensional suborbifolds, and any embedded orientable spherical 2-orbifold $S^2//\Gamma$ bounds a discal 3-orbifold $D^3//\Gamma$ in \mathcal{O} . Figure 8 shows an embedded bad 2-dimensional suborbifold Σ . Figure 9 shows an embedded spherical 2-suborbifold $S^2(k, k)$ that does not bound a discal 3-orbifold; the shaded regions are meant to indicate some complicated orbifold regions.




FIGURE 8.




FIGURE 9. An essential spherical suborbifold

If S is an orientable embedded 2-orbifold in \mathcal{O} then S is *compressible* if there is an embedded discal 2-orbifold $D \subset \mathcal{O}$ so that ∂D lies in S , but ∂D does not bound a discal 2-orbifold in S . (We call D a *compressing discal orbifold*.) Otherwise, S is *incompressible*. Note that any embedded copy of a turnover $S^2(p, q, r)$ is automatically incompressible, since any embedded circle in $S^2(p, q, r)$ bounds a discal 2-orbifold in $S^2(p, q, r)$.

If \mathcal{O} is a compact orientable 3-orbifold then there is a compact orientable irreducible 3-orbifold \mathcal{O}' so that \mathcal{O} is the result of performing 0-surgeries on \mathcal{O}' ; see [5, Chapter 3]. The orbifold \mathcal{O}' can be obtained by taking an appropriate *spherical system* on \mathcal{O} , cutting along the spherical 2-orbifolds and adding discal 3-orbifolds to the ensuing boundary components. If we take a minimal such spherical system then \mathcal{O}' is canonical.

Note that if $\mathcal{O} = S^1 \times S^2$ then $\mathcal{O}' = S^3$. This shows that if \mathcal{O} is a 3-manifold then \mathcal{O}' is not just the disjoint components in the prime decomposition. That is, we are not dealing with a direct generalization of the Kneser-Milnor prime decomposition from 3-manifold theory. Because the notion of connected sum is more involved for orbifolds than for manifolds, the notion of a prime decomposition is also more involved; see [35, 53]. It is not needed for the present paper.

We assume now that \mathcal{O} is irreducible. The *geometrization conjecture* says that if $\partial\mathcal{O} = \emptyset$ and \mathcal{O} does not have any embedded bad 2-dimensional suborbifolds then there is a finite collection $\{S_i\}$ of incompressible orientable Euclidean 2-dimensional suborbifolds of \mathcal{O} so that each connected component of $\mathcal{O}' - \bigcup_i S_i$ is diffeomorphic to a quotient of one of the eight Thurston geometries. Taking a minimal such collection of Euclidean 2-dimensional suborbifolds, the ensuing geometric pieces are canonical. References for the statement of the orbifold geometrization conjecture are [5, Chapter 3.7], [19, Chapter 2.13].

Our statement of the orbifold geometrization conjecture is a generalization of the manifold geometrization conjecture, as stated in [54, Section 6] and [56, Conjecture 1.1]. The cutting of the orientable three-manifold is along two-spheres and two-tori. An alternative version of the geometrization conjecture requires the pieces to have finite volume [45, Conjecture 2.2.1]. In this version one must also allow cutting along one-sided Klein bottles. A relevant example to illustrate this point is when the three-manifold is the result of gluing $I \times_{\mathbb{Z}_2} T^2$ to a cuspidal truncation of a one-cusped complete noncompact finite-volume hyperbolic 3-manifold.

2.4. Seifert 3-orbifolds. — A Seifert orbifold is the orbifold version of the total space of a circle bundle. We refer to [5, Chapters 2.4 and 2.5] for information about Seifert 3-orbifolds. We just recall a few relevant facts.

A Seifert 3-orbifold fibers $\pi : \mathcal{O} \rightarrow \mathcal{B}$ over a 2-dimensional orbifold \mathcal{B} , with circle fiber. If (\widehat{U}, G_p) is a local model around $p \in |\mathcal{B}|$ then there is a neighborhood V of $|\pi|^{-1}(p) \subset |\mathcal{O}|$ so that $\mathcal{O}|_V$ is diffeomorphic to $(S^1 \times \widehat{U})//G_p$, where G_p acts on S^1 via a representation $G_p \rightarrow O(2)$. We will only consider orientable Seifert 3-orbifolds, so the elements of G_p that preserve orientation on \widehat{U} will act on S^1 via $SO(2)$, while the elements of G_p that reverse orientation on \widehat{U} will act on S^1 via $O(2) - SO(2)$. In particular, if $p \in |\mathcal{B}|_{reg}$ then $|\pi|^{-1}(p)$ is a circle, while if $p \notin |\mathcal{B}|_{reg}$ then $|\pi|^{-1}(p)$ may be an interval. We may loosely talk about the circle fibration of \mathcal{O} .

As $\partial\mathcal{O}$ is an orientable 2-orbifold which fibers over a 1-dimensional orbifold, with circle fibers, any connected component of $\partial\mathcal{O}$ must be T^2 or $S^2(2, 2, 2, 2)$. In the case of a boundary component $S^2(2, 2, 2, 2)$, the generic fiber is a circle on $|S^2(2, 2, 2, 2)|$ which separates it into two 2-disks, each containing two singular points. That is, the pillowcase is divided into two half-pillowcases.

A solid-toric orbifold $S^1 \times D^2$ or $S^1 \times D^2(k)$ has an obvious Seifert fibering over D^2 or $D^2(k)$. Similarly, a solid-toric orbifold $S^1 \times_{\mathbb{Z}_2} D^2$ or $S^1 \times_{\mathbb{Z}_2} D^2(k)$ fibers over $D^2//\mathbb{Z}_2$ or $D^2(k)//\mathbb{Z}_2$.

2.5. Riemannian geometry of orbifolds

Definition 2.14. — A *Riemannian metric* on an orbifold \mathcal{O} is given by an atlas for \mathcal{O} along with a collection of Riemannian metrics on the \widehat{U}_α 's so that

- G_α acts isometrically on \widehat{U}_α and
- The embeddings $(\widehat{U}_3, G_3) \rightarrow (\widehat{U}_1, G_1)$ and $(\widehat{U}_3, G_3) \rightarrow (\widehat{U}_2, G_2)$ from part 5 of Definition 2.1 are isometric.

We say that the Riemannian orbifold \mathcal{O} has sectional curvature bounded below by $K \in \mathbb{R}$ if the Riemannian metric on each \widehat{U}_α has sectional curvature bounded below by K , and similarly for other curvature bounds.

A Riemannian orbifold has an *orthonormal frame bundle* $F\mathcal{O}$, a smooth manifold with a locally free (left) $O(n)$ -action whose quotient space is homeomorphic to $|\mathcal{O}|$.

Local charts for $F\mathcal{O}$ are given by $O(n) \times_G \widehat{U}$. Fixing a bi-invariant Riemannian metric on $O(n)$, there is a canonical $O(n)$ -invariant Riemannian metric on $F\mathcal{O}$.

Conversely, if Y is a smooth connected manifold with a locally free $O(n)$ -action then the slice theorem [11, Corollary VI.2.4] implies that for each $y \in Y$, the $O(n)$ -action near the orbit $O(n) \cdot y$ is modeled by the left $O(n)$ -action on $O(n) \times_{G_y} \mathbb{R}^N$, where the finite stabilizer group $G_y \subset O(n)$ acts linearly on \mathbb{R}^N . There is a corresponding N -dimensional orbifold \mathcal{O} with local models given by the pairs (\mathbb{R}^N, G_y) . If Y_1 and Y_2 are two such manifolds and $F : Y_1 \rightarrow Y_2$ is an $O(n)$ -equivariant diffeomorphism then there is an induced quotient diffeomorphism $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$, as can be seen by applying the slice theorem.

If Y has an $O(n)$ -invariant Riemannian metric then \mathcal{O} obtains a quotient Riemannian metric.

Remark 2.15. — Suppose that a discrete group Γ acts properly discontinuously on an orbifold \mathcal{O} . Then there is a Γ -invariant Riemannian metric on \mathcal{O} . Furthermore, Γ acts freely on $F\mathcal{O}$, commuting with the $O(n)$ -action. Hence there is a locally free $O(n)$ -action on the manifold $F\mathcal{O}/\Gamma$ and a corresponding orbifold \mathcal{O}/Γ .

There is a horizontal distribution $T^H F\mathcal{O}$ on $F\mathcal{O}$ coming from the Levi-Civita connection on \widehat{U} . If γ is a loop at $p \in |\mathcal{O}|$ then a horizontal lift of γ allows one to define the *holonomy* H_γ , a linear map from $T_p \mathcal{O}$ to itself.

If $\gamma : [a, b] \rightarrow \mathcal{O}$ is a smooth map to a Riemannian orbifold then its *length* is $L(\gamma) = \int_a^b |\gamma'(t)| dt$, where $|\gamma'(t)|$ can be defined by a local lifting of γ to a local model. This induces a length structure on $|\mathcal{O}|$. The *diameter* of \mathcal{O} is the diameter of $|\mathcal{O}|$. We say that \mathcal{O} is *complete* if $|\mathcal{O}|$ is a complete metric space. If \mathcal{O} has sectional curvature bounded below by $K \in \mathbb{R}$ then $|\mathcal{O}|$ has Alexandrov curvature bounded below by K , as can be seen from the fact that the Alexandrov condition is preserved upon quotienting by a finite group acting isometrically [13, Proposition 10.2.4].

It is useful to think of \mathcal{O} as consisting of an Alexandrov space equipped with an additional structure that allows one to make sense of smooth functions.

We write $dvol$ for the n -dimensional Hausdorff measure on $|\mathcal{O}|$. Using the above-mentioned relationship between the sectional curvature of \mathcal{O} and the Alexandrov curvature of $|\mathcal{O}|$, we can use [13, Chapter 10.6.2] to extend the Bishop-Gromov inequality from Riemannian manifolds with a lower sectional curvature bound, to Riemannian orbifolds with a lower sectional curvature bound. We remark that a Bishop-Gromov inequality for an orbifold with a lower Ricci curvature bound appears in [9].

A *geodesic* is a smooth curve γ which, in local charts, satisfies the geodesic equation. Any length-minimizing curve γ between two points is a geodesic, as can be seen by looking in a local model around $\gamma(t)$.

Lemma 2.16. — *If \mathcal{O} is a complete Riemannian orbifold then for any $p \in |\mathcal{O}|$ and any $v \in C_p |\mathcal{O}|$, there is a unique geodesic $\gamma : \mathbb{R} \rightarrow \mathcal{O}$ such that $|\gamma|(0) = p$ and $\gamma'(0) = v$.*

Proof. — The proof is similar to the proof of the corresponding part of the Hopf-Rinow theorem, as in [39, Theorem 4.1]. \square

The *exponential map* of a complete orbifold \mathcal{O} is defined as follows. Given $p \in |\mathcal{O}|$ and $v \in C_p|\mathcal{O}|$, let $\gamma : [0, 1] \rightarrow \mathcal{O}$ be the unique geodesic with $|\gamma|(0) = p$ and $|\gamma'|(0) = v$. Put $|\exp|(p, v) = (p, |\gamma|(1)) \in |\mathcal{O}| \times |\mathcal{O}|$. This has the local lifting property to define a smooth orbifold map $\exp : T\mathcal{O} \rightarrow \mathcal{O} \times \mathcal{O}$.

Given $p \in |\mathcal{O}|$, the restriction of \exp to $T_p\mathcal{O}$ gives an orbifold map $\exp_p : T_p\mathcal{O} \rightarrow \mathcal{O}$ so that $|\exp|(p, v) = (p, |\exp_p|(v))$.

Similarly, if \mathcal{O}' is a suborbifold of \mathcal{O} then there is a *normal exponential map* $\exp : N\mathcal{O}' \rightarrow \mathcal{O}$. If \mathcal{O}' is compact then for small $\epsilon > 0$, the restriction of \exp to the open ϵ -disk bundle in $N\mathcal{O}'$ is a diffeomorphism to $\mathcal{O}|_{N_\epsilon(|\mathcal{O}'|)}$.

Remark 2.17. — To prove Lemma 2.9, we can give the proper surjective submersion $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ a Riemannian submersion metric in the orbifold sense. Given $p \in |\mathcal{O}_2|$, let U be a small ϵ -ball around p and let (\widehat{U}, G_p) be a local model with $\widehat{U}/G_p = U$. Pulling back $f|_{f^{-1}(U)} : f^{-1}(U) \rightarrow U$ to \widehat{U} , we obtain a G_p -equivariant Riemannian submersion \widehat{f} to \widehat{U} . If $\widehat{p} \in \widehat{U}$ covers p then $\widehat{f}^{-1}(\widehat{p})$ is a compact orbifold on which G_p acts. Using the submersion structure, its normal bundle $N\widehat{f}^{-1}(\widehat{p})$ is G_p -diffeomorphic to $\widehat{f}^{-1}(\widehat{p}) \times T_{\widehat{p}}\widehat{U}$. If ϵ is sufficiently small then the normal exponential map on the ϵ -disk bundle in $N\widehat{f}^{-1}(\widehat{p})$ provides a G_p -equivariant product neighborhood $\widehat{f}^{-1}(\widehat{p}) \times \widehat{U}$ of $\widehat{f}^{-1}(\widehat{p})$; cf. [3, Proof of Theorem 9.42]. This passes to a diffeomorphism between $f^{-1}(U)$ and $(\widehat{f}^{-1}(\widehat{p}) \times \widehat{U})//G_p$.

If $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a local diffeomorphism and g_2 is a Riemannian metric on \mathcal{O}_2 then there is a pullback Riemannian metric f^*g_2 on \mathcal{O}_1 , which makes f into a local isometry.

We now give a useful criterion for a local isometry to be a covering map.

Lemma 2.18. — *If $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a local isometry, \mathcal{O}_1 is complete and \mathcal{O}_2 is connected then f is a covering map.*

Proof. — The proof is along the lines of the corresponding manifold statement, as in [39, Theorem 4.6]. \square

There is an orbifold version of the de Rham decomposition theorem.

Lemma 2.19. — *Let \mathcal{O} be connected, simply-connected and complete. Given $p \in |\mathcal{O}|_{reg}$, suppose that there is an orthogonal splitting $T_p\mathcal{O} = E_1 \oplus E_2$ which is invariant under holonomy around loops based at p . Then there is an isometric splitting $\mathcal{O} = \mathcal{O}_1 \times \mathcal{O}_2$ so that if we write $p = (p_1, p_2)$ then $T_{p_1}\mathcal{O}_1 = E_1$ and $T_{p_2}\mathcal{O}_2 = E_2$.*

Proof. — The parallel transport of E_1 and E_2 defines involutive distributions D_1 and D_2 , respectively, on \mathcal{O} . Let \mathcal{O}_1 and \mathcal{O}_2 be maximal integrable suborbifolds through p for D_1 and D_2 , respectively.

Given a smooth curve $\gamma : [a, b] \rightarrow \mathcal{O}$ starting at p , there is a development $C : [a, b] \rightarrow T_p\mathcal{O}$ of γ , as in [39, Section III.4]. Let $C_1 : [a, b] \rightarrow E_1$ and $C_2 : [a, b] \rightarrow E_2$ be the orthogonal projections of C . Then there are undevelopments $\gamma_1 : [a, b] \rightarrow \mathcal{O}_1$ and $\gamma_2 : [a, b] \rightarrow \mathcal{O}_2$ of C_1 and C_2 , respectively.

As in [39, Lemma IV.6.6], one shows that $(|\gamma_1|(b), |\gamma_2|(b))$ only depends on $|\gamma|(b)$. In this way, one defines a map $f : \mathcal{O} \rightarrow \mathcal{O}_1 \times \mathcal{O}_2$. As in [39, p. 192], one shows that f is a local isometry. As in [39, p. 188], one shows that \mathcal{O}_1 and \mathcal{O}_2 are simply-connected. The lemma now follows from Lemma 2.18. \square

The regular part $|\mathcal{O}|_{reg}$ inherits a Riemannian metric. The corresponding volume form equals the n -dimensional Hausdorff measure on $|\mathcal{O}|_{reg}$. We define $\text{vol}(\mathcal{O})$, or $\text{vol}(|\mathcal{O}|)$, to be the volume of the Riemannian manifold $|\mathcal{O}|_{reg}$, which equals the n -dimensional Hausdorff mass of the metric space $|\mathcal{O}|$.

If $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a diffeomorphism between Riemannian orbifolds (\mathcal{O}_1, g_1) and (\mathcal{O}_2, g_2) then we can define the C^K -distance between g_1 and f^*g_2 , using local models for \mathcal{O}_1 .

A *pointed orbifold* (\mathcal{O}, p) consists of an orbifold \mathcal{O} and a basepoint $p \in |\mathcal{O}|$. Given $r > 0$, we can consider the pointed suborbifold $\check{B}(p, r) = \mathcal{O}|_{B(p, r)}$.

Definition 2.20. — Let (\mathcal{O}_1, p_1) and (\mathcal{O}_2, p_2) be pointed connected orbifolds with complete Riemannian metrics g_1 and g_2 that are C^K -smooth. (That is, the orbifold transition maps are C^{K+1} and the metric tensor in a local model is C^K .) Given $\epsilon > 0$, we say that the C^K -distance between (\mathcal{O}_1, p_1) and (\mathcal{O}_2, p_2) is bounded above by ϵ if there is a C^{K+1} -smooth map $f : \check{B}(p_1, \epsilon^{-1}) \rightarrow \mathcal{O}_2$ that is a diffeomorphism onto its image, such that

- The C^K -distance between g_1 and f^*g_2 on $B(p_1, \epsilon^{-1})$ is at most ϵ , and
- $d_{|\mathcal{O}_2|}(|f|(p_1), p_2) \leq \epsilon$.

Taking the infimum of all such possible ϵ 's defines the C^K -distance between (\mathcal{O}_1, p_1) and (\mathcal{O}_2, p_2) .

Remark 2.21. — It may seem more natural to require $|f|$ to be basepoint-preserving. However, this would cause problems. For example, given $k \geq 2$, take $\mathcal{O} = \mathbb{R}^2 // \mathbb{Z}_k$. Let $\pi : \mathbb{R}^2 \rightarrow |\mathcal{O}|$ be the quotient map. We would like to say that if i is large then the pointed orbifold $(\mathcal{O}, \pi(i^{-1}, 0))$ is close to $(\mathcal{O}, \pi(0, 0))$. However, there is no *basepoint-preserving* map $f : \check{B}(\pi(i^{-1}, 0), 1) \rightarrow (\mathcal{O}, \pi(0, 0))$ which is a diffeomorphism onto its image, due to the difference between the local groups at the two basepoints.

2.6. Critical point theory for distance functions. — Let \mathcal{O} be a complete Riemannian orbifold and let Y be a closed subset of $|\mathcal{O}|$. A point $p \in |\mathcal{O}| - Y$ is *noncritical* if there is a nonzero G_p -invariant vector $v \in T_p\mathcal{O} \cong T_p\hat{U}$ making an angle

strictly larger than $\frac{\pi}{2}$ with any lift to $T_{\hat{p}}\widehat{U}$ of the initial velocity of any minimizing geodesic segment from p to Y .

In the next lemma we give an equivalent formulation in terms of noncriticality on $|\mathcal{O}|$.

Lemma 2.22. — *A point $p \in |\mathcal{O}| - Y$ is noncritical if and only if there is some $w \in C_p|\mathcal{O}| \cong T_{\hat{p}}\widehat{U}/G_p$ so that the comparison angle between w and any minimizing geodesic from p to Y is strictly greater than $\frac{\pi}{2}$.*

Proof. — Suppose that p is noncritical. Given v as in the definition of noncriticality, put $w = vG_p$.

Conversely, suppose that $w \in C_p|\mathcal{O}| \cong T_{\hat{p}}\widehat{U}/G_p$ is such that the comparison angle between w and any minimizing geodesic from p to Y is strictly greater than $\frac{\pi}{2}$. Let v_0 be a preimage of w in $T_{\hat{p}}\widehat{U}$. Then v_0 makes an angle greater than $\frac{\pi}{2}$ with any lift to $T_{\hat{p}}\widehat{U}$ of the initial velocity of any minimizing geodesic from p to Y . As the set of such initial velocities is G_p -invariant, for any $g \in G_p$ the vector v_0g also makes an angle greater than $\frac{\pi}{2}$ with any lift to $T_{\hat{p}}\widehat{U}$ of the initial velocity of any minimizing geodesic from p to Y . As $\{v_0g\}_{g \in G_p}$ lies in an open half-plane, we can take v to be the nonzero vector $\frac{1}{|G_p|} \sum_{g \in G_p} v_0g$. \square

We now prove the main topological implications of noncriticality.

Lemma 2.23. — *If Y is compact and there are no critical points in the set $d_Y^{-1}(a, b)$ then there is a smooth vector field ξ on $\mathcal{O}|_{d_Y^{-1}(a, b)}$ so that d_Y has uniformly positive directional derivative in the ξ direction.*

Proof. — The proof is similar to that of [14, Lemma 1.4]. For any $p \in |\mathcal{O}| - Y$, there are a precompact neighborhood U_p of p in $|\mathcal{O}| - Y$ and a smooth vector field V_p on U_p so that d_Y has positive directional derivative in the V_p direction, on U_p . Let $\{U_{p_i}\}$ be a finite collection that covers $d_Y^{-1}(a, b)$. From Lemma 2.11, there is a subordinate partition of unity $\{\rho_i\}$. Put $\xi = \sum_i \rho_i V_i$. \square

Lemma 2.24. — *If Y is compact and there are no critical points in the set $d_Y^{-1}(a, b)$ then $\mathcal{O}|_{d_Y^{-1}(a, b)}$ is diffeomorphic to a product orbifold $\mathbb{R} \times \mathcal{O}'$.*

Proof. — Construct ξ as in Lemma 2.23. Choose $c \in (a, b)$. Then $\mathcal{O}|_{d_Y^{-1}(c)}$ is a Lipschitz-regular suborbifold of \mathcal{O} which is transversal to ξ , as can be seen in local models. Working in local models, inductively from lower-dimensional strata of $|\mathcal{O}|$ to higher-dimensional strata, we can slightly smooth $\mathcal{O}|_{d_Y^{-1}(c)}$ to form a smooth suborbifold \mathcal{O}' of \mathcal{O} which is transverse to ξ . Flowing (which is defined using local models) in the direction of ξ gives an orbifold diffeomorphism between $\mathcal{O}|_{d_Y^{-1}(a, b)}$ and $\mathbb{R} \times \mathcal{O}'$. \square

2.7. Smoothing functions. — Let \mathcal{O} be a Riemannian orbifold. Let F be a Lipschitz function on $|\mathcal{O}|$. Given $p \in |\mathcal{O}|$, we define the generalized gradient $\nabla_p^{gen} F \subset T_p \mathcal{O}$ as follows. Let (\widehat{U}, G) be a local model around p . Let \widehat{F} be the lift of F to \widehat{U} . Choose $\widehat{p} \in \widehat{U}$ covering p . Let $\epsilon > 0$ be small enough so that $\exp_{\widehat{p}} : B(0, \epsilon) \rightarrow \widehat{U}$ is a diffeomorphism onto its image. If $\widehat{x} \in B(\widehat{p}, \epsilon)$ is a point of differentiability of \widehat{F} then compute $\nabla_{\widehat{x}} \widehat{F}$ and parallel transport it along the minimizing geodesic to \widehat{p} . Take the closed convex hull of the vectors so obtained and then take the intersection as $\epsilon \rightarrow 0$. This gives a closed convex G_p -invariant subset of $T_{\widehat{p}} \widehat{U}$, or equivalently a closed convex subset of $T_p \mathcal{O}$; we denote this set by $\nabla_p^{gen} F$. The union $\bigcup_{p \in |\mathcal{O}|} \nabla_p^{gen} F \subset T\mathcal{O}$ will be denoted $\nabla^{gen} F$.

Lemma 2.25. — *Let \mathcal{O} be a complete Riemannian orbifold and let $|\pi| : |T\mathcal{O}| \rightarrow |\mathcal{O}|$ be the projection map. Suppose that $U \subset |\mathcal{O}|$ is an open set, $C \subset U$ is a compact subset and S is an open fiberwise-convex subset of $T\mathcal{O}|_{|\pi|^{-1}(U)}$. (That is, S is an open subset of $|\pi|^{-1}(U)$ and for each $p \in |\mathcal{O}|$, the preimage of $(S \cap |\pi|^{-1}(p)) \subset C_p|\mathcal{O}|$ in $T_p \mathcal{O}$ is convex.)*

Then for any $\epsilon > 0$ and any Lipschitz function $F : |\mathcal{O}| \rightarrow \mathbb{R}$ whose generalized gradient over U lies in S , there is a Lipschitz function $F' : |\mathcal{O}| \rightarrow \mathbb{R}$ such that :

1. *There is an open subset of $|\mathcal{O}|$ containing C on which F' is a smooth orbifold function.*
2. *The generalized gradient of F' , over U , lies in S .*
3. $|F' - F|_{\infty} \leq \epsilon$.
4. $F'|_{|\mathcal{O}| - U} = F|_{|\mathcal{O}| - U}$.

Proof. — The proof proceeds by mollifying the Lipschitz function F as in [28, Section 2]. The mollification there is clearly G -equivariant in a local model (\widehat{U}, G) . \square

Corollary 2.26. — *For all $\epsilon > 0$ there is a $\theta > 0$ with the following property.*

Let \mathcal{O} be a complete Riemannian orbifold, let $Y \subset |\mathcal{O}|$ be a closed subset and let $d_Y : |\mathcal{O}| \rightarrow \mathbb{R}$ be the distance function from Y . Given $p \in |\mathcal{O}| - Y$, let $V_p \subset C_p|\mathcal{O}|$ be the set of initial velocities of minimizing geodesics from p to Y . Suppose that $U \subset |\mathcal{O}| - Y$ is an open subset such that for all $p \in U$, one has $\text{diam}(V_p) < \theta$. Let C be a compact subset of U . Then for every $\epsilon_1 > 0$, there is a Lipschitz function $F' : |\mathcal{O}| \rightarrow \mathbb{R}$ such that

- F' is smooth on a neighborhood of C .
- $\|F' - d_Y\|_{\infty} < \epsilon_1$.
- $F'|_{M-U} = d_Y|_{M-U}$
- For every $p \in C$, the angle between $-\nabla_p F'$ and V_p is at most ϵ .
- $F' - d_Y$ is ϵ -Lipschitz.

3. Noncompact nonnegatively curved orbifolds

In this section we extend the splitting theorem and the soul theorem from Riemannian manifolds to Riemannian orbifolds. We give an argument to rule out tight necks in a noncompact nonnegatively curved orbifold. We give the topological description of noncompact nonnegatively curved orbifolds of dimension two and three.

Assumption 3.1. — *In this section, \mathcal{O} will be a complete nonnegatively curved Riemannian orbifold.*

We may emphasize in some places that \mathcal{O} is nonnegatively curved.

3.1. Splitting theorem

Proposition 3.2. — *If $|\mathcal{O}|$ contains a line then \mathcal{O} is an isometric product $\mathbb{R} \times \mathcal{O}'$ for some complete Riemannian orbifold \mathcal{O}' .*

Proof. — As $|\mathcal{O}|$ contains a line, the splitting theorem for nonnegatively curved Alexandrov spaces [13, Chapter 10.5] implies that $|\mathcal{O}|$ is an isometric product $\mathbb{R} \times Y$ for some complete nonnegatively curved Alexandrov space Y . The isometric splitting lifts to local models, showing that $\mathcal{O}|_Y$ is a Riemannian orbifold \mathcal{O}' and that the isometry $|\mathcal{O}| \rightarrow \mathbb{R} \times Y$ is a smooth orbifold splitting $\mathcal{O} \rightarrow \mathbb{R} \times \mathcal{O}'$. \square

Corollary 3.3. — *If \mathcal{O} has more than one end then it has two ends and \mathcal{O} is an isometric product $\mathbb{R} \times \mathcal{O}'$ for some compact Riemannian orbifold \mathcal{O}' .*

Remark 3.4. — A splitting theorem for orbifolds with nonnegative Ricci curvature appears in [10]. As the present paper deals with lower sectional curvature bounds, the more elementary Proposition 3.2 is sufficient for our purposes.

3.2. Cheeger-Gromoll-type theorem. — A subset $Z \subset |\mathcal{O}|$ is *totally convex* if any geodesic segment (possibly not minimizing) with endpoints in Z lies entirely in Z .

Lemma 3.5. — *Let $Z \subset |\mathcal{O}|$ be totally convex and let (\widehat{U}, G) be a local model. Put $U = \widehat{U}/G$ and let $q : \widehat{U} \rightarrow U$ be the quotient map. If γ is a geodesic segment in \widehat{U} with endpoints in $q^{-1}(U \cap Z)$ then γ lies in $q^{-1}(U \cap Z)$.*

Proof. — Suppose that $\gamma(t) \notin q^{-1}(U \cap Z)$ for some t . Then $q \circ \gamma$ is a geodesic in \mathcal{O} with endpoints in Z , but $q(\gamma(t)) \notin Z$. This is a contradiction. \square

Lemma 3.6. — *Let $Z \subset |\mathcal{O}|$ be a closed totally convex set. Let k be the Hausdorff dimension of Z . Let \mathcal{N} be the union of the k -dimensional suborbifolds \mathcal{S} of \mathcal{O} with $|\mathcal{S}| \subset Z$. Then \mathcal{N} is a totally geodesic k -dimensional suborbifold of $|\mathcal{O}|$ and $Z = \overline{|\mathcal{N}|}$. Furthermore, if Y is a closed subset of $|\mathcal{N}|$ and $p \in Z - |\mathcal{N}|$ then there is a $v \in C_p|\mathcal{O}|$ so that the initial velocity of any minimizing geodesic from p to Y makes an angle greater than $\frac{\pi}{2}$ with v .*

Proof. — Using Lemma 3.5, the proof is along the lines of that in [27, Chapter 3.1]. \square

We put $\partial Z = Z - |\mathcal{N}|$. Note that in the definition of \mathcal{N} we are dealing with orbifolds as opposed to manifolds. For example, if $\mathcal{O}|_Z$ is a boundaryless k -dimensional orbifold then $\partial Z = \emptyset$.

A function $f : |\mathcal{O}| \rightarrow \mathbb{R}$ is *concave* if for any geodesic segment $\gamma : [a, b] \rightarrow \mathcal{O}$, for all $c \in [a, b]$ one has

$$(3.7) \quad f(|\gamma|(c)) \geq \frac{b-c}{b-a} f(|\gamma|(a)) + \frac{c-a}{b-a} f(|\gamma|(b)).$$

Lemma 3.8. — *It is equivalent to require (3.7) for all geodesic segments or just for minimizing geodesic segments.*

Proof. — Suppose that (3.7) holds for all minimizing geodesic segments. Let $\gamma : [a, b] \rightarrow \mathcal{O}$ be a geodesic segment, maybe not minimizing. For any $t \in [a, b]$, we can find a neighborhood I_t of t in $[a, b]$ so that the restriction of γ to I_t is minimizing. Then (3.7) holds on I_t . It follows that (3.7) holds on $[a, b]$. \square

Any superlevel set $f^{-1}[c, \infty)$ of a concave function is closed and totally convex.

Let f be a proper concave function on $|\mathcal{O}|$ which is bounded above. Then there is a maximal $c \in \mathbb{R}$ so that the superlevel set $f^{-1}[c, \infty)$ is nonempty, and so $f^{-1}[c, \infty) = f^{-1}\{c\}$ is a closed totally convex set.

Suppose for the rest of this subsection that \mathcal{O} is noncompact.

Lemma 3.9. — *Let $Z \subset |\mathcal{O}|$ be a closed totally convex set with $\partial Z \neq \emptyset$. Then $d_{\partial Z}$ is a concave function on Z . Furthermore, suppose that for a minimizing geodesic $\gamma : [a, b] \rightarrow Z$ in Z , the restriction of $d_{\partial Z} \circ |\gamma|$ is a constant positive function on $[a, b]$. Let $t \rightarrow \exp_{\gamma(a)} tX(a)$ be a minimizing unit-speed geodesic from $|\gamma|(a)$ to ∂Z , defined for $t \in [0, d]$. Let $\{X(s)\}_{s \in [a, b]}$ be the parallel transport of $X(a)$ along γ . Then for any $s \in [a, b]$, the curve $t \rightarrow \exp_{\gamma(s)} tX(s)$ is a minimal geodesic from $|\gamma|(s)$ to ∂Z , of length d . Also, the rectangle $V : [a, b] \times [0, d] \rightarrow Z$ given by $V(s, t) = \exp_{\gamma(s)} tX(s)$ is flat and totally geodesic.*

Proof. — The proof is similar to that of [27, Theorem 3.2.5]. \square

Fix a basepoint $\star \in |\mathcal{O}|$. Let η be a unit-speed ray in $|\mathcal{O}|$ starting from \star ; note that η is automatically a geodesic. Let $b_\eta : |\mathcal{O}| \rightarrow \mathbb{R}$ be the Busemann function;

$$(3.10) \quad b_\eta(p) = \lim_{t \rightarrow \infty} (d(p, \eta(t)) - t).$$

Lemma 3.11. — *The Busemann function b_η is concave.*

Proof. — The proof is similar to that of [27, Theorem 3.2.4]. \square

Lemma 3.12. — *Putting $f = \inf_\eta b_\eta$, where η runs over unit speed rays starting at \star , gives a proper concave function on $|\mathcal{O}|$ which is bounded above.*

Proof. — The proof is similar to that of [27, Proposition 3.2.1]. \square

We now construct the soul of \mathcal{O} , following Cheeger-Gromoll [17]. Let C_0 be the minimal nonempty superlevel set of f . For $i \geq 0$, if $\partial C_i \neq \emptyset$ then let C_{i+1} be the minimal nonempty superlevel set of $d_{\partial C_i}$ on C_i . Let S be the nonempty C_i so that $\partial C_i = \emptyset$. Define the *soul* to be $\mathcal{S} = \mathcal{O}|_S$. Then \mathcal{S} is a totally geodesic suborbifold of \mathcal{O} .

Proposition 3.13. — \mathcal{O} is diffeomorphic to the normal bundle $\mathcal{N}\mathcal{S}$ of \mathcal{S} .

Proof. — Following [27, Lemma 3.3.1], we claim that d_S has no critical points on $|\mathcal{O}| - S$. To see this, choose $p \in |\mathcal{O}| - S$. There is a totally convex set $Z \subset |\mathcal{O}|$ for which $p \in \partial Z$; either a superlevel set of f or one of the sets C_i . Defining \mathcal{N} as in Lemma 3.6, we also know that $S \subset |\mathcal{N}|$. By Lemma 3.6, p is noncritical for d_S .

From Lemma 2.24, for small $\epsilon > 0$, we know that \mathcal{O} is diffeomorphic to $\mathcal{O}|_{N_\epsilon(S)}$. However, if ϵ is small then the normal exponential map gives a diffeomorphism between $\mathcal{N}\mathcal{S}$ and $\mathcal{O}|_{N_\epsilon(S)}$. □

Remark 3.14. — One can define a soul for a general complete nonnegatively curved Alexandrov space X . The soul will be homotopy equivalent to X . However, X need not be homeomorphic to a fiber bundle over the soul, as shown by an example of Perelman [13, Example 10.10.9].

We include a result that we will need later about orbifolds with locally convex boundary.

Lemma 3.15. — Let \mathcal{O} be a compact connected orbifold-with-boundary with nonnegative sectional curvature. Suppose that $\partial\mathcal{O}$ is nonempty and has positive-definite second fundamental form. Then there is some $p \in |\mathcal{O}|$ so that $\partial\mathcal{O}$ is diffeomorphic to the unit distance sphere from the vertex in $T_p\mathcal{O}$.

Proof. — Let $p \in |\mathcal{O}|$ be a point of maximal distance from $|\partial\mathcal{O}|$. We claim that p is unique. If not, let p' be another such point and let γ be a minimizing geodesic between them. Applying Lemma 3.9 with $Z = |\mathcal{O}|$, there is a nontrivial geodesic $s \rightarrow V(s, d)$ of \mathcal{O} that lies in $|\partial\mathcal{O}|$. This contradicts the assumption on $\partial\mathcal{O}$. Thus p is unique. The lemma now follows from the proof of Lemma 3.13, as we are effectively in a situation where the soul is a point. □

3.3. Ruling out tight necks in nonnegatively curved orbifolds

Lemma 3.16. — Suppose that \mathcal{O} is a complete connected Riemannian orbifold with nonnegative sectional curvature. If X is a compact connected 2-sided codimension-1 suborbifold of \mathcal{O} then precisely one of the following occurs :

- X is the boundary of a compact suborbifold of \mathcal{O} .
- X is nonseparating, \mathcal{O} is compact and X lifts to a \mathbb{Z} -cover $\mathcal{O}' \rightarrow \mathcal{O}$, where $\mathcal{O}' = \mathbb{R} \times \mathcal{O}''$ with \mathcal{O}'' compact.
- X separates \mathcal{O} into two unbounded connected components and $\mathcal{O} = \mathbb{R} \times \mathcal{O}'$ with \mathcal{O}' compact.

Proof. — Suppose that X separates \mathcal{O} . If both components of $|\mathcal{O}| - |X|$ are unbounded then \mathcal{O} contains a line. From Proposition 3.2, $\mathcal{O} = \mathbb{R} \times \mathcal{O}'$ for some \mathcal{O}' . As X is compact, \mathcal{O}' must be compact.

The remaining case is when X does not separate \mathcal{O} . If γ is a smooth closed curve in \mathcal{O} which is transversal to X (as defined in local models) then there is a well-defined intersection number $\gamma \cdot X \in \mathbb{Z}$. This gives a homomorphism $\rho : \pi_1(\mathcal{O}, p) \rightarrow \mathbb{Z}$. Since X is nonseparating, there is a γ so that $\gamma \cdot X \neq 0$; hence the image of ρ is an infinite cyclic group. Put $\mathcal{O}' = \tilde{\mathcal{O}}/\text{Ker}(\rho)$; it is an infinite cyclic cover of \mathcal{O} . As \mathcal{O}' contains a line, the lemma follows from Proposition 3.2. \square

Lemma 3.17. — *Suppose that $\mathbb{R}^n//G$ is a Euclidean orbifold with G a finite subgroup of $O(n)$. If $X \subset \mathbb{R}^n//G$ is a connected compact 2-sided codimension-1 suborbifold, then X bounds some $D \subset \mathbb{R}^n//G$ with $\text{diam}_{\mathcal{O}}(D) < 4|G|\text{diam}_X(X)$, where $\text{diam}_{\mathcal{O}}(D)$ denote the extrinsic diameter of D in $|\mathcal{O}|$ while $\text{diam}_X(X)$ denotes the intrinsic diameter of X .*

Proof. — Let \hat{X} be the preimage of X in \mathbb{R}^n . Let Δ be any number greater than $\text{diam}_X(X)$. Let x be a point in $|X|$. Let $\{\hat{x}_i\}_{i \in I}$ be the preimages of x in \hat{X} . Here the cardinality of I is bounded above by $|G|$. We claim that $\hat{X} = \bigcup_{i \in I} B(\hat{x}_i, \Delta)$, where $B(\hat{x}_i, \Delta)$ denotes a distance ball in \hat{X} with respect to its intrinsic metric. To see this, let \hat{y} be an arbitrary point in \hat{X} . Let y be its image in X . Join y to x by a minimizing geodesic γ in X , which is necessarily of length at most Δ . Then a horizontal lift of γ , starting at \hat{y} , joins \hat{y} to some \hat{x}_i and also has length at most Δ .

Let \hat{C} be a connected component of \hat{X} . Since \hat{C} is connected, it has a covering by a subset of $\{B(\hat{x}_i, 2\text{diam}_X(X))\}_{i \in I}$ with connected nerve, and so \hat{C} has diameter at most $4|G|\text{diam}_X(X)$. Furthermore, from the Jordan separation theorem, \hat{C} is the boundary of a domain $\hat{D} \in \mathbb{R}^n$ with extrinsic diameter at most $4|G|\text{diam}_X(X)$. Letting $D \in \mathcal{O}$ be the projection of \hat{D} , the lemma follows. \square

Proposition 3.18. — *Suppose that \mathcal{O} is a complete connected noncompact Riemannian n -orbifold with nonnegative sectional curvature. Then there is a number $\delta > 0$ (depending on \mathcal{O}) so that the following holds. Let X be a connected compact 2-sided codimension-1 suborbifold of \mathcal{O} . Then either*

- X bounds a connected suborbifold D of \mathcal{O} with $\text{diam}_{\mathcal{O}}(D) < 8(\sup_{p \in |\mathcal{O}|} |G_p|) \cdot \text{diam}(X)$, or
- $\text{diam}(X) > \delta$.

Proof. — Suppose that the proposition is not true. Then there is a sequence $\{X_i\}_{i=1}^{\infty}$ of connected compact 2-sided codimension-1 suborbifolds of \mathcal{O} so that $\lim_{i \rightarrow \infty} \text{diam}(X_i) = 0$ but each X_i fails to bound a connected suborbifold whose extrinsic diameter is at most $8 \sup_{p \in |\mathcal{O}|} |G_p|$ times as much.

If all of the $|X_i|$'s lie in a compact subset of $|\mathcal{O}|$ then a subsequence converges in the Hausdorff topology to a point $p \in |\mathcal{O}|$. As a sufficiently small neighborhood

of p can be well approximated metrically by a neighborhood of $0 \in |\mathbb{R}^n//G_p|$ after rescaling, Lemma 3.17 implies that for large i we can find $D_i \subset \mathcal{O}$ with $X_i = \partial D_i$ and $\text{diam}_{\mathcal{O}_i}(D_i) < 8(\sup_{p \in |\mathcal{O}|} |G_p|) \cdot \text{diam}(X_i)$. This is a contradiction. Hence we can assume that the sets $|X_i|$ tend to infinity.

If some X_i does not bound a compact suborbifold of \mathcal{O} then by Lemma 3.16, there is an isometric splitting $\mathcal{O} = \mathbb{R} \times \mathcal{O}'$ with \mathcal{O}' compact. This contradicts the assumed existence of the sequence $\{X_i\}_{i=1}^\infty$ with $\lim_{i \rightarrow \infty} \text{diam}(X_i) = 0$. Thus we can assume that $X_i = \partial D_i$ for some compact suborbifold D_i of \mathcal{O} . If \mathcal{O} had more than one end then it would split off an \mathbb{R} -factor and as before, the sequence $\{X_i\}_{i=1}^\infty$ would not exist. Hence \mathcal{O} is one-ended and after passing to a subsequence, we can assume that $D_1 \subset D_2 \subset \dots$. Fix a basepoint $\star \in |D_1|$. Let η be a unit-speed ray in $|\mathcal{O}|$ starting from \star and let b_η be the Busemann function from (3.10).

Suppose that $p, p' \in |\mathcal{O}|$ are such that $b_\eta(p) = b_\eta(p')$. For t large, consider a geodesic triangle with vertices $p, p', \eta(t)$. Given X_i with i large, if t is sufficiently large then $\overline{p\eta(t)}$ and $\overline{p'\eta(t)}$ pass through X_i . Taking $t \rightarrow \infty$, triangle comparison implies that $d(p, p') \leq \text{diam}(X_i)$. Taking $i \rightarrow \infty$ gives $p = p'$. Thus b_η is injective. This is a contradiction. \square

3.4. Nonnegatively curved 2-orbifolds

Lemma 3.19. — *Let \mathcal{O} be a complete connected orientable 2-dimensional orbifold with nonnegative sectional curvature which is C^K -smooth, $K \geq 3$. We have the following classification of the diffeomorphism type, based on the number of ends. For notation, Γ denotes a finite subgroup of the oriented isometry group of the relevant orbifold and Σ^2 denotes a simply-connected bad 2-orbifold with some Riemannian metric.*

- 0 ends : $S^2//\Gamma, T^2//\Gamma, \Sigma^2//\Gamma$.
- 1 end : $\mathbb{R}^2//\Gamma, S^1 \times_{\mathbb{Z}_2} \mathbb{R}$.
- 2 ends : $\mathbb{R} \times S^1$.

Proof. — If \mathcal{O} has zero ends then it is compact and the classification follows from the orbifold Gauss-Bonnet theorem [5, Proposition 2.9]. If \mathcal{O} has more than one end then Proposition 3.2 implies that \mathcal{O} has two ends and isometrically splits off an \mathbb{R} -factor. Hence it must be diffeomorphic to $\mathbb{R} \times S^1$. Suppose that \mathcal{O} has one end. The soul \mathcal{S} has dimension 0 or 1. If \mathcal{S} has dimension zero then \mathcal{S} is a point and \mathcal{O} is diffeomorphic to the normal bundle of \mathcal{S} , which is $\mathbb{R}^2//\Gamma$. If \mathcal{S} has dimension one then it is S^1 or $S^1//\mathbb{Z}_2$ and \mathcal{O} is diffeomorphic to the normal bundle of \mathcal{S} . As $S^1 \times \mathbb{R}$ has two ends, the only possibility is $S^1 \times_{\mathbb{Z}_2} \mathbb{R}$. \square

3.5. Noncompact nonnegatively curved 3-orbifolds

Lemma 3.20. — *Let \mathcal{O} be a complete connected noncompact orientable 3-dimensional orbifold with nonnegative sectional curvature which is C^K -smooth, $K \geq 3$. We have the following classification of the diffeomorphism type, based on the number of ends.*

For notation, Γ denotes a finite subgroup of the oriented isometry group of the relevant orbifold and Σ^2 denotes a simply-connected bad 2-orbifold with some Riemannian metric.

- 1 end : $\mathbb{R}^3//\Gamma$, $S^1 \times \mathbb{R}^2$, $S^1 \times \mathbb{R}^2(k)$, $S^1 \times_{\mathbb{Z}_2} \mathbb{R}^2$, $S^1 \times_{\mathbb{Z}_2} \mathbb{R}^2(k)$, $\mathbb{R} \times_{\mathbb{Z}_2} (S^2//\Gamma)$, $\mathbb{R} \times_{\mathbb{Z}_2} (T^2//\Gamma)$ or $\mathbb{R} \times_{\mathbb{Z}_2} (\Sigma^2//\Gamma)$.
- 2 ends : $\mathbb{R} \times (S^2//\Gamma)$, $\mathbb{R} \times (T^2//\Gamma)$ or $\mathbb{R} \times (\Sigma^2//\Gamma)$.

Proof. — Because \mathcal{O} is noncompact, it has at least one end. If it has more than one end then Proposition 3.2 implies that \mathcal{O} has two ends and isometrically splits off an \mathbb{R} -factor. This gives rise to the possibilities listed for two ends.

Suppose that \mathcal{O} has one end. The soul \mathcal{S} has dimension 0, 1 or 2. If \mathcal{S} has dimension zero then \mathcal{S} is a point and \mathcal{O} is diffeomorphic to the normal bundle of \mathcal{S} , which is $\mathbb{R}^3//\Gamma$. If \mathcal{S} has dimension one then it is S^1 or $S^1//\mathbb{Z}_2$ and \mathcal{O} is diffeomorphic to the normal bundle of \mathcal{S} , which is $S^1 \times \mathbb{R}^2$, $S^1 \times \mathbb{R}^2(k)$, $S^1 \times_{\mathbb{Z}_2} \mathbb{R}^2$ or $S^1 \times_{\mathbb{Z}_2} \mathbb{R}^2(k)$. If \mathcal{S} has dimension two then since it has nonnegative curvature, it is diffeomorphic to a quotient of S^2 , T^2 or Σ^2 . Then \mathcal{O} is diffeomorphic to the normal bundle of \mathcal{S} , which is $\mathbb{R} \times_{\mathbb{Z}_2} (S^2//\Gamma)$, $\mathbb{R} \times_{\mathbb{Z}_2} (T^2//\Gamma)$ or $\mathbb{R} \times_{\mathbb{Z}_2} (\Sigma^2//\Gamma)$, since \mathcal{O} has one end. \square

3.6. 2-dimensional nonnegatively curved orbifolds that are pointed Gromov-Hausdorff close to an interval. — We include a result that we will need later about 2-dimensional nonnegatively curved orbifolds that are pointed Gromov-Hausdorff close to an interval.

Lemma 3.21. — *There is some $\beta > 0$ so that the following holds. Suppose that \mathcal{O} is a pointed nonnegatively curved complete orientable Riemannian 2-orbifold which is C^K -smooth for some $K \geq 3$. Let $\star \in |\mathcal{O}|$ be a basepoint and suppose that the pointed ball $(B(\star, 10), \star) \subset |\mathcal{O}|$ has pointed Gromov-Hausdorff distance at most β from the pointed interval $([0, 10], 0)$. Then for every $r \in [1, 9]$, the orbifold $\mathcal{O}|_{\overline{B(\star, r)}}$ is a discal 2-orbifold or is diffeomorphic to $D^2(2, 2)$.*

Proof. — As in [37, Pf. of Lemma 3.12], the distance function $d_\star : A(\star, 1, 9) \rightarrow [1, 9]$ defines a fibration with a circle fiber.

The possible diffeomorphism types of \mathcal{O} are listed in Lemma 3.19. Looking at them, if $\overline{B(\star, 1)}$ is not a topological disk then \mathcal{O} must be T^2 and we obtain a contradiction as in [37, Pf. of Lemma 3.12]. Hence $\overline{B(\star, 1)}$ is a topological disk. If $\mathcal{O}|_{\overline{B(\star, 1)}}$ is not a discal 2-orbifold then it has at least two singular points, say $p_1, p_2 \in |\mathcal{O}|$. Choose $q \in |\mathcal{O}|$ with $d(\star, q) = 2$. By triangle comparison, the comparison angles satisfy $\tilde{\angle}_{p_1}(p_2, q) \leq \frac{2\pi}{|G_{p_1}|}$ and $\tilde{\angle}_{p_2}(p_1, q) \leq \frac{2\pi}{|G_{p_2}|}$. If β is small then $\tilde{\angle}_{p_1}(p_2, q) + \tilde{\angle}_{p_2}(p_1, q)$ is close to π . It follows that $|G_{p_1}| = |G_{p_2}| = 2$.

Suppose that there are three distinct singular points $p_1, p_2, p_3 \in |\mathcal{O}|$. We know that they lie in $\overline{B(\star, 1)}$. Let $\overline{p_i q}$ and $\overline{p_k p_j}$ denote minimal geodesics. If β is small then the angle at p_1 between $\overline{p_1 q}$ and $\overline{p_1 p_2}$ is close to $\frac{\pi}{2}$, and similarly for the angle at p_1 between $\overline{p_1 q}$ and $\overline{p_1 p_3}$. As $\dim(\mathcal{O}) = 2$, and p_1 has total cone angle π , it follows that

if β is small then the angle at p_1 between $\overline{p_1 p_2}$ and $\overline{p_1 p_2}$ is small. The same reasoning applies at p_2 and p_3 , so we have a geodesic triangle in $|\mathcal{O}|$ with small total interior angle, which violates the fact that $|\mathcal{O}|$ has nonnegative Alexandrov curvature.

Thus $\mathcal{O}|_{\overline{B(\star,1)}}$ is diffeomorphic to $D^2(2,2)$. □

4. Riemannian compactness theorem for orbifolds

In this section we prove a compactness result for Riemannian orbifolds.

The statement of the compactness result is slightly different from the usual statement for Riemannian manifolds, which involves a lower injectivity radius bound. The standard notion of injectivity radius is not a useful notion for orbifolds. For example, if \mathcal{O} is an orientable 2-orbifold with a singular point p then a geodesic from a regular point q in $|\mathcal{O}|$ to p cannot minimize beyond p . As q could be arbitrarily close to p , we conclude that the injectivity radius of \mathcal{O} would vanish. (We note, however, that there is a modified version of the injectivity radius that does makes sense for constant-curvature cone manifolds [5, Section 9.2.3], [19, Section 6.4].)

Instead, our compactness result is phrased in terms of local volumes. This fits well with Perelman’s work on Ricci flow, where local volume estimates arise naturally.

If one tried to prove a compactness result for Riemannian orbifolds directly, following the proofs in the case of Riemannian manifolds, then one would have to show that orbifold singularities do not coalesce when taking limits. We avoid this issue by passing to orbifold frame bundles, which are manifolds, and using equivariant compactness results there.

Compactness theorems for Riemannian metrics and Ricci flows for orbifolds with isolated singularities were proved in [40]. Compactness results for general orbifolds were stated in [18, Chapter 3.3] with a short sketch of a proof.

Proposition 4.1. — *Fix $K \in \mathbb{Z}^+ \cup \{\infty\}$. Let $\{(\mathcal{O}_i, p_i)\}_{i=1}^\infty$ be a sequence of pointed complete connected C^{K+3} -smooth Riemannian n -dimensional orbifolds. Suppose that for each $j \in \mathbb{Z}^{\geq 0}$ with $j \leq K$, there is a function $A_j : (0, \infty) \rightarrow \infty$ so that for all i , $|\nabla^j \text{Rm}| \leq A_j(r)$ on $B(p_i, r) \subset |\mathcal{O}_i|$. Suppose that for some $r_0 > 0$, there is a $v_0 > 0$ so that for all i , $\text{vol}(B(p_i, r_0)) \geq v_0$. Then there is a subsequence of $\{(\mathcal{O}_i, p_i)\}_{i=1}^\infty$ that converges in the pointed C^{K-1} -topology to a pointed complete connected Riemannian n -dimensional orbifold $(\mathcal{O}_\infty, p_\infty)$.*

Proof. — Let $F\mathcal{O}_i$ be the orthonormal frame bundle of \mathcal{O}_i . Pick a basepoint $\widehat{p}_i \in F\mathcal{O}_i$ that projects to $p_i \in |\mathcal{O}_i|$. As in [26, Section 6], after taking a subsequence we may assume that the frame bundles $\{(F\mathcal{O}_i, \widehat{p}_i)\}_{i=1}^\infty$ converge in the pointed $O(n)$ -equivariant Gromov-Hausdorff topology to a C^{K-1} -smooth Riemannian manifold X with an isometric $O(n)$ -action and a basepoint \widehat{p}_∞ . (We lose one derivative because we are working on the frame bundle.) Furthermore, we may assume that the convergence is

realized as follows : Given any $O(n)$ -invariant compact codimension-zero submanifold-with-boundary $K \subset X$, for large i there is an $O(n)$ -invariant compact codimension-zero submanifold-with-boundary $\widehat{K}_i \subset F\mathcal{O}_i$ and a smooth $O(n)$ -equivariant fiber bundle $\widehat{K}_i \rightarrow K$ with nilmanifold fiber whose diameter goes to zero as $i \rightarrow \infty$ [15, Section 3], [26, Section 9].

Quotienting by $O(n)$, the underlying spaces $\{(|\mathcal{O}_i|, p_i)\}_{i=1}^\infty$ converge in the pointed Gromov-Hausdorff topology to $(O(n)\backslash X, p_\infty)$. Because of the lower volume bound $\text{vol}(B(p_i, r_0)) \geq v_0$, a pointed Gromov-Hausdorff limit of the Alexandrov spaces $\{(|\mathcal{O}_i|, p_i)\}_{i=1}^\infty$ is an n -dimensional Alexandrov space [13, Corollary 10.10.11]. Thus there is no collapsing and so for large i the submersion $\widehat{K}_i \rightarrow K$ is an $O(n)$ -equivariant C^{K-1} -smooth diffeomorphism. In particular, the $O(n)$ -action on X is locally free. There is a corresponding quotient orbifold \mathcal{O}_∞ with $|\mathcal{O}_\infty| = O(n)\backslash X$. As the manifolds $\{(F\mathcal{O}_i, \widehat{p}_i)\}_{i=1}^\infty$ converge in a C^{K-1} -smooth pointed equivariant sense to (X, \widehat{p}_∞) we can take $O(n)$ -quotients to conclude that the orbifolds $\{(\mathcal{O}_i, p_i)\}_{i=1}^\infty$ converge in the pointed C^{K-1} -smooth topology to $(\mathcal{O}_\infty, p_\infty)$. \square

Remark 4.2. — As a consequence of Proposition 4.1, if there is a number N so $|G_{q_i}| \leq N$ for all $q_i \in |\mathcal{O}_i|$ and all i then $|G_{q_\infty}| \leq N$ for all $q_\infty \in |\mathcal{O}_\infty|$. That is, under the hypotheses of Proposition 4.1, the orders of the isotropy groups cannot increase in the limit.

Remark 4.3. — In the proof of Proposition 4.1, the submersions $\widehat{K}_i \rightarrow K$ may not be basepoint-preserving. This is where one has to leave the world of basepoint-preserving maps.

5. Ricci flow on orbifolds

In this section we first make some preliminary remarks about Ricci flow on orbifolds and we give the orbifold version of Hamilton's compactness theorem. We then give the topological classification of compact nonnegatively curved 3-orbifolds. Finally, we extend Perelman's no local collapsing theorem to orbifolds.

5.1. Function spaces on orbifolds. — Let $\rho : O(n) \rightarrow \mathbb{R}^N$ be a representation. Given a local model $(\widehat{U}_\alpha, G_\alpha)$ and a G_α -invariant Riemannian metric on \widehat{U}_α , let $\widehat{V}_\alpha = \mathbb{R}^N \times_{O(n)} F\widehat{U}_\alpha$ be the associated vector bundle. If \mathcal{O} is a n -dimensional Riemannian orbifold then there is an associated orbivector bundle V with local models $(\widehat{V}_\alpha, G_\alpha)$. Its underlying space is $|V| = \mathbb{R}^N \times_{O(n)} F\mathcal{O}$. By construction, V has an inner product coming from the standard inner product on \mathbb{R}^N . A section s of V is given by an $O(n)$ -equivariant map $s : F\mathcal{O} \rightarrow \mathbb{R}^N$. In terms of local models, s is described by G_α -invariant sections s_α of \widehat{V}_α that satisfy compatibility conditions with respect to part 5 of Definition 2.1.

The C^K -norm of s is defined to be the supremum of the C^K -norms of the s_α 's. Similarly, the square of the H^K -norm of s is defined to be the integral over $|\mathcal{O}|_{reg}$ of the local square H^K -norm, the latter being defined using local models. (Note that $|\mathcal{O}|_{reg}$ has full Hausdorff n -measure in $|\mathcal{O}|$.) Then H^{-K} can be defined by duality. One has the rough Laplacian mapping H^K -sections of V to H^{K-2} -sections of V .

One can define differential operators and pseudodifferential operators acting on H^K -sections of V . Standard elliptic and parabolic regularity theory extends to the orbifold setting, as can be seen by working equivariantly in local models.

5.2. Short-time existence for Ricci flow on orbifolds. — Suppose that $\{g(t)\}_{t \in [A, B]}$ is a smooth 1-parameter family of Riemannian metrics on \mathcal{O} . We will call g a *flow of metrics* on \mathcal{O} . The Ricci flow equation $\frac{\partial g}{\partial t} = -2 \text{Ric}$ makes sense in terms of local models. Using the DeTurck trick [20], which is based on local differential analysis, one can reduce the short-time existence problem for the Ricci flow to the short-time existence problem for a parabolic PDE. Then any short-time existence proof for parabolic PDEs on compact manifolds, such as that of [55, Proposition 15.8.2], will extend from the manifold setting to the orbifold setting.

Remark 5.1. — Even in the manifold case, one needs a slight additional argument to reduce the short-time existence of the Ricci-DeTurck equation to that of a standard quasilinear parabolic PDE. In local coordinates the Ricci-DeTurck equation takes the form

$$(5.2) \quad \frac{\partial g_{ij}}{\partial t} = \sum_{kl} g^{kl} \partial_k \partial_l g_{ij} + \dots$$

There is a slight issue since (5.2) is not uniformly parabolic, in that g^{kl} could degenerate with respect to, say, the initial metric g_0 . This issue does not seem to have been addressed in the literature. However, it is easily circumvented. Let \mathcal{M} be the space of smooth Riemannian metrics on a compact manifold M . Let $F : \mathcal{M} \rightarrow \mathcal{M}$ be a smooth map so that for some $\epsilon > 0$, we have $F(g) = g$ if $\|g - g_0\|_{g_0} < \epsilon$, and in addition $\epsilon g_0 \leq F(g) \leq \epsilon^{-1} g_0$ for all g . (Such a map F is easily constructed using the fact that the inner products on $T_p M$, relative to $g_0(p)$, can be identified with $\text{GL}(n, \mathbb{R})/O(n)$, along with the fact that $\text{GL}(n, \mathbb{R})/O(n)$ deformation retracts onto a small ball around its basepoint.) By [55, Proposition 15.8.2], there is a short-time solution to

$$(5.3) \quad \frac{\partial g_{ij}}{\partial t} = \sum_{kl} F(g)^{kl} \partial_k \partial_l g_{ij} + \dots$$

with $g(0) = g_0$. Given this solution, there is some $\delta > 0$ so that $\|g(t) - g_0\|_{g_0} < \epsilon$ whenever $t \in [0, \delta]$. Then $\{g(t)\}_{t \in [0, \delta]}$ also solves the Ricci-DeTurck equation (5.2).

We remark that any Ricci flow results based on the maximum principle will have evident extensions from manifolds to orbifolds. Such results include

- The lower bound on scalar curvature
- The Hamilton-Ivey pinching results for three-dimensional scalar curvature
- Hamilton’s differential Harnack inequality for Ricci flow solutions with nonnegative curvature operator
- Perelman’s differential Harnack inequality.

5.3. Ricci flow compactness theorem for orbifolds. — Let \mathcal{O}_1 and \mathcal{O}_2 be two connected pointed n -dimensional orbifolds, with flows of metrics g_1 and g_2 . If $f : \mathcal{O}_1 \rightarrow \mathcal{O}_2$ is a (time-independent) diffeomorphism then we can construct the pullback flow f^*g_2 and define the C^K -distance between g_1 and f^*g_2 , using local models for \mathcal{O}_1 .

Definition 5.4. — Let \mathcal{O}_1 and \mathcal{O}_2 be connected pointed n -dimensional orbifolds. Given numbers A, B with $-\infty \leq A < 0 \leq B \leq \infty$, suppose that g_i is a flow of metrics on \mathcal{O}_i that exists for the time interval $[A, B]$. Suppose that $g_i(t)$ is complete for each t . Given $\epsilon > 0$, suppose that $f : \check{B}(p_1, \epsilon^{-1}) \rightarrow \mathcal{O}_2$ is a smooth map from the time-zero ball that is a diffeomorphism onto its image. Let $|f| : B(p_1, \epsilon^{-1}) \rightarrow |\mathcal{O}_2|$ be the underlying map. We say that the C^K -distance between the flows $(\mathcal{O}_1, p_1, g_1)$ and $(\mathcal{O}_2, p_2, g_2)$ is bounded above by ϵ if

1. The C^K -distance between g_1 and f^*g_2 on $([A, B] \cap (-\epsilon^{-1}, \epsilon^{-1})) \times \check{B}(p_1, \epsilon^{-1})$ is at most ϵ and
2. The time-zero distance $d_{|\mathcal{O}_2|}(|f|(p_1), p_2)$ is at most ϵ .

Taking the infimum of all such possible ϵ ’s defines the C^K -distance between the flows $(\mathcal{O}_1, p_1, g_1)$ and $(\mathcal{O}_2, p_2, g_2)$.

Note that time derivatives appear in the definition of the C^K -distance between g_1 and f^*g_2 .

Proposition 5.5. — Let $\{g_i\}_{i=1}^\infty$ be a sequence of Ricci flow solutions on pointed connected n -dimensional orbifolds $\{(\mathcal{O}_i, p_i)\}_{i=1}^\infty$, defined for $t \in (A, B)$ and complete for each t , with $-\infty \leq A < 0 \leq B \leq \infty$. Suppose that the following two conditions are satisfied :

1. For every compact interval $I \subset (A, B)$, there is some $K_I < \infty$ so that for all i , we have $\sup_{|\mathcal{O}_i| \times I} |\text{Rm}_{g_i}(p, t)| \leq K_I$, and
2. For some $r_0, v_0 > 0$ and all i , the time-zero volume $\text{vol}(B(p_i, r_0))$ is bounded below by v_0 .

Then a subsequence of the solutions converges in the sense of Definition 5.4 to a Ricci flow solution $g_\infty(t)$ on a pointed connected n -dimensional orbifold $(\mathcal{O}_\infty, p_\infty)$, defined for all $t \in (A, B)$.

Proof. — Using Proposition 4.1, the proof is essentially the same as that in [32, p. 548-551] and [40, p. 1116-1117]. \square

Remark 5.6. — There are variants of Proposition 5.5 that hold, for example, if one just assumes a uniform curvature bound on r -balls, for each $r > 0$. These variants are orbifold versions of the results in [38, Appendix E], to which we refer for details. The proofs of these orbifold extensions use, among other things, the orbifold version of the Shi estimates; the proof of the latter goes through to the orbifold setting with no real change.

5.4. Compact nonnegatively curved 3-orbifolds

Proposition 5.7. — *Any compact nonnegatively curved 3-orbifold \mathcal{O} is diffeomorphic to one of*

1. $S^3//\Gamma$ for some finite group $\Gamma \subset \text{Isom}^+(S^3)$.
2. $T^3//\Gamma$ for some finite group $\Gamma \subset \text{Isom}^+(T^3)$.
3. $S^1 \times (S^2//\Gamma)$ or $S^1 \times_{\mathbb{Z}_2} (S^2//\Gamma)$ for some finite group $\Gamma \subset \text{Isom}(S^2)$.
4. $S^1 \times (\Sigma^2//\Gamma)$ or $S^1 \times_{\mathbb{Z}_2} (\Sigma^2//\Gamma)$ for some finite group $\Gamma \subset \text{Isom}(\Sigma^2)$, where Σ^2 is a simply-connected bad 2-orbifold equipped with its unique (up to diffeomorphism) Ricci soliton metric [58, Theorem 4.1].

Proof. — Let k be the largest number so that the universal cover $\tilde{\mathcal{O}}$ isometrically splits off an \mathbb{R}^k -factor. Write $\tilde{\mathcal{O}} = \mathbb{R}^k \times \mathcal{O}'$.

If \mathcal{O}' is noncompact then by the Cheeger-Gromoll argument [16, Pf. of Theorem 3], $|\mathcal{O}'|$ contains a line. Proposition 3.2 implies that \mathcal{O}' splits off an \mathbb{R} -factor, which is a contradiction. Thus \mathcal{O}' is simply-connected and compact with nonnegative sectional curvature.

If $k = 3$ then $\tilde{\mathcal{O}} = \mathbb{R}^3$ and \mathcal{O} is a quotient of T^3 .

If $k = 2$ then there is a contradiction, as there is no simply-connected compact 1-orbifold.

If $k = 1$ then \mathcal{O}' is diffeomorphic to S^2 or Σ^2 . The Ricci flow on $\tilde{\mathcal{O}} = \mathbb{R} \times \mathcal{O}'$ splits isometrically. After rescaling, the Ricci flow on \mathcal{O}' converges to a constant curvature metric on S^2 or to the unique Ricci soliton metric on Σ^2 [58]. Hence $\pi_1(\mathcal{O})$ is a subgroup of $\text{Isom}(\mathbb{R} \times S^2)$ or $\text{Isom}(\mathbb{R} \times \Sigma^2)$, where the isometry groups are in terms of standard metrics. As $\pi_1(\mathcal{O})$ acts properly discontinuously and cocompactly on $\tilde{\mathcal{O}}$, there is a short exact sequence

$$(5.8) \quad 1 \longrightarrow \Gamma_1 \longrightarrow \pi_1(\mathcal{O}) \longrightarrow \Gamma_2 \longrightarrow 1,$$

where $\Gamma_1 \subset \text{Isom}(\mathcal{O}')$ and Γ_2 is an infinite cyclic group or an infinite dihedral group. It follows that \mathcal{O} is finitely covered by $S^1 \times S^2$ or $S^1 \times \Sigma^2$.

Suppose that $k = 0$. If \mathcal{O} is positively curved then any proof of Hamilton’s theorem about 3-manifolds with positive Ricci curvature [30] extends to the orbifold case, to

show that \mathcal{O} admits a metric of constant positive curvature; c.f. [34]. Hence we can reduce to the case when \mathcal{O} does not have positive curvature and the Ricci flow does not immediately give it positive curvature. From the strong maximum principle as in [31, Section 8], for any $p \in |\mathcal{O}|_{reg}$ there is a nontrivial orthogonal splitting $T_p\mathcal{O} = E_1 \oplus E_2$ which is invariant under holonomy around loops based at p . The same will be true on $\tilde{\mathcal{O}}$. Lemma 2.19 implies that $\tilde{\mathcal{O}}$ splits off an \mathbb{R} -factor, which is a contradiction. \square

5.5. \mathcal{L} -geodesics and noncollapsing. — Let \mathcal{O} be an n -dimensional orbifold and let $\{g(t)\}_{t \in [0, T]}$ be a Ricci flow solution on \mathcal{O} so that

- The time slices $(\mathcal{O}, g(t))$ are complete.
- There is bounded curvature on compact subintervals of $[0, T]$.

Given $t_0 \in [0, T]$ and $p \in |\mathcal{O}|$, put $\tau = t_0 - t$. Let $\gamma : [0, \bar{\tau}] \rightarrow \mathcal{O}$ be a piecewise smooth curve with $|\gamma|(0) = p$ and $\bar{\tau} \leq t_0$. Put

$$(5.9) \quad \mathcal{L}(\gamma) = \int_0^{\bar{\tau}} \sqrt{\tau} (R(\gamma(\tau)) + |\dot{\gamma}(\tau)|^2) d\tau,$$

where the scalar curvature R and the norm $|\dot{\gamma}(\tau)|$ are evaluated using the metric at time $t_0 - \tau$. With $X = \frac{d\gamma}{d\tau}$, the \mathcal{L} -geodesic equation is

$$(5.10) \quad \nabla_X X - \frac{1}{2} \nabla R + \frac{1}{2\tau} X + 2 \operatorname{Ric}(X, \cdot) = 0.$$

Given an \mathcal{L} -geodesic γ , its *initial velocity* is defined to be $v = \lim_{\tau \rightarrow 0} \sqrt{\tau} \frac{d\gamma}{d\tau} \in C_p|\mathcal{O}|$.

Given $q \in |\mathcal{O}|$, put

$$(5.11) \quad L(q, \bar{\tau}) = \inf\{\mathcal{L}(\gamma) : |\gamma|(\bar{\tau}) = q\},$$

where the infimum runs over piecewise smooth curves γ with $|\gamma|(0) = p$ and $|\gamma|(\bar{\tau}) = q$. Then any piecewise smooth curve γ which is a minimizer for L is a smooth \mathcal{L} -geodesic.

Lemma 5.12. — *There is a minimizer γ for L .*

Proof. — The proof is similar to that in [38, p. 2631]. We outline the steps. Given p and q , one considers piecewise smooth curves γ as above. Fixing $\epsilon > 0$, one shows that the curves γ with $\mathcal{L}(\gamma) < L(q, \bar{\tau}) + \epsilon$ are uniformly continuous. In particular, there is an $R < \infty$ so that any such γ lies in $B(p, R)$. Next, one shows that there is some $\rho \in (0, R)$ so that for any $x \in B(p, R)$, there is a local model (\widehat{U}, G_x) with $\widehat{U}/G_x = B(x, \rho)$ such that for any $p', q' \in B(x, \rho)$ and any subinterval $[\bar{\tau}_1, \bar{\tau}_2] \subset [0, \bar{\tau}]$,

- There is a unique minimizer for the functional $\int_{\bar{\tau}_1}^{\bar{\tau}_2} \sqrt{\tau} (R(\gamma(\tau)) + |\dot{\gamma}(\tau)|^2) d\tau$ among piecewise smooth curves $\gamma : [\bar{\tau}_1, \bar{\tau}_2] \rightarrow \mathcal{O}$ with $|\gamma|(\bar{\tau}_1) = p'$ and $|\gamma|(\bar{\tau}_2) = q'$.
- The minimizing γ is smooth and the image of $|\gamma|$ lies in $B(x, \rho)$.

This is shown by working in the local models. Now cover $\overline{B(p, R)}$ by a finite number of ρ -balls $\{B(x_i, \rho)\}_{i=1}^N$. Using the uniform continuity, let $A \in \mathbb{Z}^+$ be such that for any $\gamma : [0, \bar{\tau}] \rightarrow \mathcal{O}$ with $|\gamma|(0) = p$, $|\gamma|(\bar{\tau}) = q$ and $\mathcal{L}(\gamma) < L(q, \bar{\tau}) + \epsilon$, and any $[\bar{\tau}_1, \bar{\tau}_2] \subset [0, \bar{\tau}]$ of length at most $\frac{\bar{\tau}}{A}$, the distance between $|\gamma|(\bar{\tau}_1)$ and $|\gamma|(\bar{\tau}_2)$ is less than the Lebesgue number of the covering. We can effectively reduce the problem of finding a minimizer for L to the problem of minimizing a continuous function defined on tuples $(p_0, \dots, p_A) \in \overline{B(p, R)}^{A+1}$ with $p_0 = p$ and $p_A = q$. This shows that the minimizer exists. \square

Define the \mathcal{L} -exponential map $: T_p\mathcal{O} \rightarrow \mathcal{O}$ by saying that for $v \in C_p|\mathcal{O}|$, we put $\mathcal{L} \exp_{\bar{\tau}}(v) = |\gamma|(\bar{\tau})$, where γ is the unique \mathcal{L} -geodesic from p whose initial velocity is v . Then $\mathcal{L} \exp_{\bar{\tau}}$ is a smooth orbifold map.

Let $\mathcal{B}_{\bar{\tau}} \subset |\mathcal{O}|$ be the set of points q which are either endpoints of more than one minimizing \mathcal{L} -geodesic $\gamma : [0, \bar{\tau}] \rightarrow \mathcal{O}$, or are the endpoint of a minimizing geodesic $\gamma_v : [0, \bar{\tau}] \rightarrow \mathcal{O}$ where $v \in C_p|\mathcal{O}|$ is a critical point of $\mathcal{L} \exp_{\bar{\tau}}$. We call $\mathcal{B}_{\bar{\tau}}$ the *time- $\bar{\tau}$ \mathcal{L} -cut locus* of p . It is a closed subset of $|\mathcal{O}|$. Let $\mathcal{G}_{\bar{\tau}} \subset |\mathcal{O}|$ be the complement of $\mathcal{B}_{\bar{\tau}}$ and let $\Omega_{\bar{\tau}} \subset C_p|\mathcal{O}|$ be the corresponding set of initial conditions for minimizing \mathcal{L} -geodesics. Then $\Omega_{\bar{\tau}}$ is an open set, and the restriction of $\mathcal{L} \exp_{\bar{\tau}}$ to $T_p\mathcal{O}|_{\Omega_{\bar{\tau}}}$ is an orbifold diffeomorphism to $\mathcal{O}|_{\mathcal{G}_{\bar{\tau}}}$.

Lemma 5.13. — $\mathcal{B}_{\bar{\tau}}$ has measure zero in $|\mathcal{O}|$.

Proof. — The proof is similar to that in [38, p. 2632]. By Sard’s theorem, it suffices to show that the subset $\mathcal{B}'_{\bar{\tau}} \subset \mathcal{B}_{\bar{\tau}}$, consisting of regular values of $\mathcal{L} \exp_{\bar{\tau}}$, has measure zero in $|\mathcal{O}|$. One shows that $\mathcal{B}'_{\bar{\tau}}$ is contained in the underlying spaces of a countable union of codimension-1 suborbifolds of \mathcal{O} , which implies the lemma. \square

Therefore one may compute the integral of any integrable function on $|\mathcal{O}|$ by pulling it back to $\Omega_{\bar{\tau}} \subset C_p|\mathcal{O}|$ and using the change of variable formula.

For $q \in |\mathcal{O}|$, put $l(q, \bar{\tau}) = \frac{L(q, \bar{\tau})}{2\sqrt{\bar{\tau}}}$. Define the *reduced volume* by

$$(5.14) \quad \tilde{V}(\bar{\tau}) = \bar{\tau}^{-\frac{n}{2}} \int_{|\mathcal{O}|} e^{-l(q, \bar{\tau})} \, \text{dvol}(q).$$

Lemma 5.15. — *The reduced volume is monotonically nonincreasing in $\bar{\tau}$.*

Proof. — The proof is similar to that in [38, Section 23]. In the proof, one pulls back the integrand to $C_p|\mathcal{O}|$. \square

Lemma 5.16. — *For each $\bar{\tau} > 0$, there is some $q \in |\mathcal{O}|$ so that $l(q, \bar{\tau}) \leq \frac{n}{2}$.*

Proof. — The proof is similar to that in [38, Section 24]. It uses the maximum principle, which is valid for orbifolds. \square

Definition 5.17. — Given $\kappa, \rho > 0$, a Ricci flow solution $g(\cdot)$ defined on a time interval $[0, T)$ is κ -noncollapsed on the scale ρ if for each $r < \rho$ and all $(x_0, t_0) \in |\mathcal{O}| \times [0, T)$ with $t_0 \geq r^2$, whenever it is true that $|\text{Rm}(x, t)| \leq r^{-2}$ for every $x \in B_{t_0}(x_0, r)$ and $t \in [t_0 - r^2, t_0]$, then we also have $\text{vol}(B_{t_0}(x_0, r)) \geq \kappa r^n$.

Lemma 5.18. — If a Ricci flow solution is κ -noncollapsed on some scale then there is a uniform upper bound $|G_p| \leq N(n, \kappa)$ on the orders of the isotropy groups at points $p \in |\mathcal{O}|$.

Proof. — Given $p \in |\mathcal{O}|$, let $B_{t_0}(p, r)$ be a ball such that $|\text{Rm}(x, t_0)| \leq r^{-2}$ for all $x \in B_{t_0}(p, r)$. By assumption $r^{-n} \text{vol}(B_{t_0}(x_0, r)) \geq \kappa$. Let c_n denote the area of the unit $(n - 1)$ -sphere in \mathbb{R}^n . Applying the Bishop-Gromov inequality to $B_{t_0}(p, r)$ gives

$$(5.19) \quad \frac{1}{|G_p|} \geq \frac{r^{-n} \text{vol}(B_{t_0}(x_0, r))}{c_n \int_0^1 \sinh^{n-1}(s) ds} \geq \frac{\kappa}{c_n \int_0^1 \sinh^{n-1}(s) ds}.$$

The lemma follows. □

Proposition 5.20. — Given numbers $n \in \mathbb{Z}^+$, $T < \infty$ and $\rho, K, c > 0$, there is a number $\kappa = \kappa(n, K, c, \rho, T) > 0$ with the following property. Let $(\mathcal{O}^n, g(\cdot))$ be a Ricci flow solution defined on the time interval $[0, T)$, with complete time slices, such that the curvature $|\text{Rm}|$ is bounded on every compact subinterval $[0, T'] \subset [0, T)$. Suppose that $(\mathcal{O}, g(0))$ has $|\text{Rm}| \leq K$ and $\text{vol}(B(p, 1)) \geq c > 0$ for every $p \in |\mathcal{O}|$. Then the Ricci flow solution is κ -noncollapsed on the scale ρ .

Proof. — The proof is similar to that in [38, Section 26]. As in the proof there, we use the fact that the initial conditions give uniformly bounded geometry in a small time interval $[0, \bar{t}/2]$, as follows from Proposition 5.5 and derivative estimates. □

Proposition 5.21. — For any $A \in (0, \infty)$, there is some $\kappa = \kappa(A) > 0$ with the following property. Let $(\mathcal{O}, g(\cdot))$ be an n -dimensional Ricci flow solution defined for $t \in [0, r_0^2]$ having complete time slices and uniformly bounded curvature. Suppose that $\text{vol}(B_0(p_0, r_0)) \geq A^{-1}r_0^n$ and that $|\text{Rm}|(q, t) \leq \frac{1}{nr_0^2}$ for all $(q, t) \in B_0(p_0, r_0) \times [0, r_0^2]$. Then the solution cannot be κ -collapsed on a scale less than r_0 at any point (q, r_0^2) with $q \in B_{r_0^2}(p_0, Ar_0)$.

Proof. — The proof is similar to that in [38, Section 28]. □

6. κ -solutions

In this section we extend results about κ -solutions from manifolds to orbifolds.

Definition 6.1. — Given $\kappa > 0$, a κ -solution is a Ricci flow solution $(\mathcal{O}, g(t))$ that is defined on a time interval of the form $(-\infty, C)$ (or $(-\infty, C]$) such that :

1. The curvature $|\text{Rm}|$ is bounded on each compact time interval $[t_1, t_2] \subset (-\infty, C)$ (or $(-\infty, C]$), and each time slice $(\mathcal{O}, g(t))$ is complete.

2. The curvature operator is nonnegative and the scalar curvature is everywhere positive.
3. The Ricci flow is κ -noncollapsed at all scales.

Lemma 5.18 gives an upper bound on the orders of the isotropy groups. In the rest of this section we will use this upper bound without explicitly restating it.

6.1. Asymptotic solitons. — Let (p, t_0) be a point in a κ -solution $(\mathcal{O}, g(\cdot))$ so that G_p has maximal order. Define the reduced volume $\tilde{V}(\bar{\tau})$ and the reduced length $l(q, \bar{\tau})$ as in Subsection 5.5, by means of curves starting from (p, t_0) , with $\tau = t_0 - t$. From Lemma 5.16, for each $\bar{\tau} > 0$ there is some $q(\bar{\tau}) \in |\mathcal{O}|$ such that $l(q(\bar{\tau}), \bar{\tau}) \leq \frac{n}{2}$. (Note that $l \geq 0$ from the curvature assumption.)

Proposition 6.2. — *There is a sequence $\bar{\tau}_i \rightarrow \infty$ so that if we consider the solution $g(\cdot)$ on the time interval $[t_0 - \bar{\tau}_i, t_0 - \frac{1}{2}\bar{\tau}_i]$ and parabolically rescale it at the point $(q(\bar{\tau}_i), t_0 - \bar{\tau}_i)$ by the factor $\bar{\tau}_i^{-1}$ then as $i \rightarrow \infty$, the rescaled solutions converge to a nonflat gradient shrinking soliton (restricted to $[-1, -\frac{1}{2}]$).*

Proof. — The proof is similar to that in [38, Section 39]. Using estimates on the reduced length as defined with the basepoint (p, t_0) , one constructs a limit Ricci flow solution $(\mathcal{O}_\infty, g_\infty(\cdot))$ defined for $t \in [-1, -\frac{1}{2}]$, which is a gradient shrinking soliton. The only new issue is to show that it is nonflat.

As in [38, Section 39], there is a limiting reduced length function $l_\infty(\cdot, \tau) \in C^\infty(\mathcal{O}_\infty)$, and a reduced volume which is a constant c , strictly less than the $t \rightarrow t_0$ limit of the reduced volume of $(\mathcal{O}, g(\cdot))$. The latter equals $\frac{(4\pi)^{\frac{n}{2}}}{|G_p|}$. If the limit solution were flat then $l_\infty(\cdot, \tau)$ would have a constant positive-definite Hessian. It would then have a unique critical point q . Using the gradient flow of $l_\infty(\cdot, \tau)$, one deduces that \mathcal{O}_∞ is diffeomorphic to $T_q\mathcal{O}_\infty$. As in [38, Section 39], one concludes that

$$(6.3) \quad c = \int_{C_q|\mathcal{O}_\infty| \cong \mathbb{R}^n/G_q} \tau^{-\frac{n}{2}} e^{-\frac{|x|^2}{4\tau}} \, d\text{vol} = \frac{(4\pi)^{\frac{n}{2}}}{|G_q|}.$$

As $|G_q| \leq |G_p|$, we obtain a contradiction. □

6.2. Two-dimensional κ -solutions

Lemma 6.4. — *Any two-dimensional κ -solution $(\mathcal{O}, g(\cdot))$ is an isometric quotient of the round shrinking 2-sphere or is a Ricci soliton metric on a bad 2-orbifold.*

Proof. — The proof is similar to that in [59, Theorem 4.1]. One considers the asymptotic soliton and shows that it has strictly positive scalar curvature outside of a compact region (as in [50, Lemma 1.2]). Using standard Jacobi field estimates, the asymptotic soliton must be compact. The lemma then follows from convergence results for 2-dimensional compact Ricci flow (using [58] in the case of bad 2-orbifolds). □

Remark 6.5. — One can alternatively prove Lemma 6.4 using the fact that if $(\mathcal{O}, g(\cdot))$ is a κ -solution then so is the pullback solution $(\tilde{\mathcal{O}}, \tilde{g}(\cdot))$ on the universal cover. If \mathcal{O} is a bad 2-orbifold then \mathcal{O} is compact and the result follows from [58]. If \mathcal{O} is a good 2-orbifold then $(\tilde{\mathcal{O}}, \tilde{g}(\cdot))$ is a round shrinking S^2 from [38, Section 40].

6.3. Asymptotic scalar curvature and asymptotic volume ratio

Definition 6.6. — If \mathcal{O} is a complete connected Riemannian orbifold then its *asymptotic scalar curvature ratio* is $\mathcal{R} = \limsup_{q \rightarrow \infty} R(q)d(x, p)^2$. It is independent of the basepoint $p \in |\mathcal{O}|$.

Lemma 6.7. — *Let $(\mathcal{O}, g(\cdot))$ be a noncompact κ -solution. Then the asymptotic scalar curvature ratio is infinite for each time slice.*

Proof. — The proof is similar to that in [38, Section 41]. Choose a time t_0 . If $\mathcal{R} \in (0, \infty)$ then after rescaling $(\mathcal{O}, g(t_0))$, one obtains convergence to a smooth annular region in the Tits cone $C_T \mathcal{O}$ at time t_0 . (Here $C_T \mathcal{O}$ denotes a smooth orbifold structure on the complement of the vertex in the Tits cone $C_T |\mathcal{O}|$.) Working on the regular part of the annular region, one obtains a contradiction from the curvature evolution equation.

If $\mathcal{R} = 0$ then the rescaling limit is a smooth flat metric on $C_T \mathcal{O}$, away from the vertex. The unit sphere S_∞ in $C_T \mathcal{O}$ has principal curvatures one. It can be approximated by a sequence of codimension-one compact suborbifolds S_k in \mathcal{O} with rescaled principal curvatures approaching one, which bound compact suborbifolds $\mathcal{O}_k \subset \mathcal{O}$.

Suppose first that $n \geq 3$. By Lemma 3.15, for large k there is some $p_k \in |\mathcal{O}|$ so that the suborbifold S_k is diffeomorphic to the unit sphere in $T_{p_k} \mathcal{O}$. As S_k is diffeomorphic to S_∞ for large k , we conclude that S_∞ is isometric to S^{n-1}/Γ for some finite group $\Gamma \subset \text{Isom}^+(S^{n-1})$. Let $p \in |\mathcal{O}|$ be a point with $G_p \cong \Gamma$. As $C_T |\mathcal{O}|$ is isometric to \mathbb{R}^n/Γ , $\lim_{r \rightarrow \infty} r^{-n} \text{vol}(B(p, r))$ exists and equals the $\frac{1}{|\Gamma|}$ times the volume of the unit ball in \mathbb{R}^n . On the other hand, this equals $\lim_{r \rightarrow 0} r^{-n} \text{vol}(B(p, r))$. As we have equality in the Bishop-Gromov inequality, we conclude that \mathcal{O} is flat, which is a contradiction.

If $n = 2$ then we can adapt the argument in [38, Section 41] to the orbifold setting. \square

Definition 6.8. — If \mathcal{O} is a complete n -dimensional Riemannian orbifold with nonnegative Ricci curvature then its *asymptotic volume ratio* is $\mathcal{V} = \lim_{r \rightarrow \infty} r^{-n} \text{vol}(B(p, r))$. It is independent of the choice of basepoint $p \in |\mathcal{O}|$.

Lemma 6.9. — *Let $(\mathcal{O}, g(\cdot))$ be a noncompact κ -solution. Then the asymptotic volume ratio \mathcal{V} vanishes for each time slice $(\mathcal{O}, g(t_0))$. Moreover, there is a sequence of points $p_k \in |\mathcal{O}|$ going to infinity such that the pointed sequence $\{(\mathcal{O}, (p_k, t_0), g(\cdot))\}_{k=1}^\infty$*

converges, modulo rescaling by $R(p_k, t_0)$, to a κ -solution which isometrically splits off an \mathbb{R} -factor.

Proof. — The proof is similar to that in [38, Section 41]. □

6.4. In a κ -solution, the curvature and the normalized volume control each other

Lemma 6.10. — *Given $n \in \mathbb{Z}^+$, we consider n -dimensional κ -solutions.*

1. *If $B(p_0, r_0)$ is a ball in a time slice of a κ -solution then the normalized volume $r^{-n} \text{vol}(B(p_0, r_0))$ is controlled (i.e., bounded away from zero) \Leftrightarrow the normalized scalar curvature $r_0^2 R(p_0)$ is controlled (i.e., bounded above)*
2. *(Precompactness) If $\{(\mathcal{O}_k, (p_k, t_k), g_k(\cdot))\}_{k=1}^\infty$ is a sequence of pointed κ -solutions and for some $r > 0$, the r -balls $B(p_k, r) \subset (\mathcal{O}_k, g_k(t_k))$ have controlled normalized volume, then a subsequence converges to an ancient solution $(\mathcal{O}_\infty, (p_\infty, 0), g_\infty(\cdot))$ which has nonnegative curvature operator, and is κ -noncollapsed (though a priori the curvature may be unbounded on a given time slice).*
3. *There is a constant $\eta = \eta(n, \kappa)$ so that for all $p \in |\mathcal{O}|$, we have $|\nabla R|(p, t) \leq \eta R^{\frac{3}{2}}(p, t)$ and $|R_t|(p, t) \leq \eta R^2(p, t)$. More generally, there are scale invariant bounds on all derivatives of the curvature tensor, that only depend on n and κ .*
4. *There is a function $\alpha : [0, \infty) \rightarrow [0, \infty)$ depending only on n and κ such that $\lim_{s \rightarrow \infty} \alpha(s) = \infty$, and for every $p, p' \in |\mathcal{O}|$, we have $R(p')d^2(p, p') \leq \alpha(R(p)d^2(p, p'))$.*

Proof. — The proof is similar to that in [38, Section 42]. In the proof by contradiction of the implication \Leftarrow of part (1), after passing to a subsequence we can assume that $|G_{p_k}|$ is a constant C . Then we use the argument in [38, Section 42] with c_n equal to $\frac{1}{C}$ times the volume of the unit Euclidean n -ball. □

6.5. A volume bound

Lemma 6.11. — *For every $\epsilon > 0$, there is an $A < \infty$ with the following property. Suppose that we have a sequence of (not necessarily complete) Ricci flow solutions $g_k(\cdot)$ with nonnegative curvature operator, defined on $\mathcal{O}_k \times [t_k, 0]$, such that:*

- *For each k , the time-zero ball $B(p_k, r_k)$ has compact closure in $|\mathcal{O}_k|$.*
- *For all $(p, t) \in B(p_k, r_k) \times [t_k, 0]$, we have $\frac{1}{2}R(p, t) \leq R(p_k, 0) = Q_k$.*
- $\lim_{k \rightarrow \infty} t_k Q_k = -\infty$.
- $\lim_{k \rightarrow \infty} r_k^2 Q_k = \infty$.

Then for large k , we have $\text{vol}(B(p_k, AQ_k^{-\frac{1}{2}})) \leq \epsilon(AQ_k^{-\frac{1}{2}})^n$ at time zero.

Proof. — The proof is similar to that in [38, Section 44]. □

6.6. Curvature bounds for Ricci flow solutions with nonnegative curvature operator, assuming a lower volume bound

Lemma 6.12. — For every $w > 0$, there are $B = B(w) < \infty$, $C = C(w) < \infty$ and $\tau_0 = \tau_0(w) > 0$ with the following properties.

(a) Take $t_0 \in [-r_0^2, 0)$. Suppose that we have a (not necessarily complete) Ricci flow solution $(\mathcal{O}, g(\cdot))$, defined for $t \in [t_0, 0]$, so that at time zero the metric ball $B(p_0, r_0)$ has compact closure. Suppose that for each $t \in [t_0, 0]$, $g(t)$ has nonnegative curvature operator and $\text{vol}(B_t(p_0, r_0)) \geq wr_0^n$. Then

$$(6.13) \quad R(p, t) \leq Cr_0^{-2} + B(t - t_0)^{-1}$$

whenever $\text{dist}_t(p, p_0) \leq \frac{1}{4}r_0$.

(b) Suppose that we have a (not necessarily complete) Ricci flow solution $(\mathcal{O}, g(\cdot))$, defined for $t \in [-\tau_0 r_0^2, 0]$, so that at time zero the metric ball $B(p_0, r_0)$ has compact closure. Suppose that for each $t \in [-\tau_0 r_0^2, 0]$, $g(t)$ has nonnegative curvature operator. If we assume a time-zero volume bound $\text{vol}(B_0(p_0, r_0)) \geq wr_0^n$ then

$$(6.14) \quad R(p, t) \leq Cr_0^{-2} + B(t + \tau_0 r_0^2)^{-1}$$

whenever $t \in [-\tau_0 r_0^2, 0]$ and $\text{dist}_t(p, p_0) \leq \frac{1}{4}r_0$.

Proof. — The proof is similar to that in [38, Section 45]. □

Corollary 6.15. — For every $w > 0$, there are $B = B(w) < \infty$, $C = C(w) < \infty$ and $\tau_0 = \tau_0(w) > 0$ with the following properties. Suppose that we have a (not necessarily complete) Ricci flow solution $(\mathcal{O}, g(\cdot))$, defined for $t \in [-\tau_0 r_0^2, 0]$, so that at time zero the metric ball $B(p_0, r_0)$ has compact closure. Suppose that for each $t \in [-\tau_0 r_0^2, 0]$, the curvature operator in the time- t ball $B(p_0, r_0)$ is bounded below by $-r_0^{-2}$. If we assume a time-zero volume bound $\text{vol}(B_0(p_0, r_0)) \geq wr_0^n$ then

$$(6.16) \quad R(p, t) \leq Cr_0^{-2} + B(t + \tau_0 r_0^2)^{-1}$$

whenever $t \in [-\tau_0 r_0^2, 0]$ and $\text{dist}_t(p, p_0) \leq \frac{1}{4}r_0$.

Proof. — The proof is similar to that in [38, Section 45]. □

6.7. Compactness of the space of three-dimensional κ -solutions

Proposition 6.17. — Given $\kappa > 0$, the set of oriented three-dimensional κ -solutions $(\mathcal{O}, g(\cdot))$ is compact modulo scaling.

Proof. — If $\{(\mathcal{O}_k, (p_k, 0), g_k(\cdot))\}_{k=1}^\infty$ is a sequence of such κ -solutions with $R(p_k, 0) = 1$ then parts (1) and (2) of Lemma 6.10 imply that there is a subsequence that converges to an ancient solution $(\mathcal{O}_\infty, (p_\infty, 0), g_\infty(\cdot))$ which has nonnegative curvature operator and is κ -noncollapsed. The remaining issue is to show that it has bounded curvature. Since $R_t \geq 0$, it is enough to show that $(\mathcal{O}_\infty, g_\infty(0))$ has bounded scalar curvature.

If not then there is a sequence of points q_i going to infinity in $|\mathcal{O}_\infty|$ such that $R(q_i, 0) \rightarrow \infty$ and $R(q, 0) \leq 2R(q_i, 0)$ for $q \in B(q_i, A_i R(q_i, 0)^{-\frac{1}{2}})$, where $A_i \rightarrow \infty$. Using the κ -noncollapsing, a subsequence of the rescalings $(\mathcal{O}_\infty, q_i, R(q_i, 0)g_\infty)$ will converge to a limit orbifold N_∞ that isometrically splits off an \mathbb{R} -factor. By Lemma 6.4, N_∞ must be a standard solution on $\mathbb{R} \times (S^2//\Gamma)$ or $\mathbb{R} \times (\Sigma^2//\Gamma)$. Thus $(\mathcal{O}_\infty, g_\infty)$ contains a sequence X_i of neck regions, with their cross-sectional radii tending to zero as $i \rightarrow \infty$. This contradicts Proposition 3.18. \square

6.8. Necklike behavior at infinity of a three-dimensional κ -solution

Definition 6.18. — Fix $\epsilon > 0$. Let $(\mathcal{O}, g(\cdot))$ be an oriented three-dimensional κ -solution. We say that a point $p_0 \in |\mathcal{O}|$ is the *center of an ϵ -neck* if the solution $g(\cdot)$ in the set $\{(p, t) : -(\epsilon Q)^{-1} < t \leq 0, \text{dist}_0(p, p_0)^2 < (\epsilon Q)^{-1}\}$, where $Q = R(p_0, 0)$, is, after scaling with the factor Q , ϵ -close in some fixed smooth topology to the corresponding subset of a κ -solution $\mathbb{R} \times \mathcal{O}'$ that splits off an \mathbb{R} -factor. That is, \mathcal{O}' is the standard evolving $S^2//\Gamma$ or $\Sigma^2//\Gamma$ with extinction time 1. Here Σ^2 is a simply-connected bad 2-orbifold with a Ricci soliton metric.

We let $|\mathcal{O}|_\epsilon$ denote the points in $|\mathcal{O}|$ which are not centers of ϵ -necks.

Proposition 6.19. — For all $\kappa > 0$, there exists an $\epsilon_0 > 0$ such that for all $0 < \epsilon < \epsilon_0$ there exists an $\alpha = \alpha(\epsilon, \kappa)$ with the property that for any oriented three dimensional κ -solution $(\mathcal{O}, g(\cdot))$, and at any time t , precisely one of the following holds :

- $(\mathcal{O}, g(\cdot))$ splits off an \mathbb{R} -factor and so every point at every time is the center of an ϵ -neck for all $\epsilon > 0$.
- \mathcal{O} is noncompact, $|\mathcal{O}|_\epsilon \neq \emptyset$, and for all $x, y \in |\mathcal{O}|_\epsilon$, we have $R(x) d^2(x, y) < \alpha$.
- \mathcal{O} is compact, and there is a pair of points $x, y \in |\mathcal{O}|_\epsilon$ such that $R(x)d^2(x, y) > \alpha$,

$$(6.20) \quad |\mathcal{O}|_\epsilon \subset B\left(x, \alpha R(x)^{-\frac{1}{2}}\right) \cup B\left(y, \alpha R(y)^{-\frac{1}{2}}\right),$$

and there is a minimizing geodesic \overline{xy} such that every $z \in |\mathcal{O}| - |\mathcal{O}|_\epsilon$ satisfies $R(z) d^2(z, \overline{xy}) < \alpha$.

- \mathcal{O} is compact and there exists a point $x \in |\mathcal{O}|_\epsilon$ such that $R(x) d^2(x, z) < \alpha$ for all $z \in |\mathcal{O}|$.

Proof. — The proof is similar to that in [38, Section 48]. \square

6.9. Three-dimensional gradient shrinking κ -solutions

Lemma 6.21. — Any three-dimensional gradient shrinking κ -solution \mathcal{O} is one of the following:

- A finite isometric quotient of the round shrinking S^3 .
- $\mathbb{R} \times (S^2//\Gamma)$ or $\mathbb{R} \times_{\mathbb{Z}_2} (S^2//\Gamma)$ for some finite group $\Gamma \subset \text{Isom}(S^2)$.
- $\mathbb{R} \times (\Sigma^2//\Gamma)$ or $\mathbb{R} \times_{\mathbb{Z}_2} (\Sigma^2//\Gamma)$ for some finite group $\Gamma \subset \text{Isom}(\Sigma^2)$.

Proof. — As \mathcal{O} is a κ -solution, we know that \mathcal{O} has nonnegative sectional curvature. If \mathcal{O} has positive sectional curvature then the proofs of [46, Theorem 3.1] or [52, Theorem 1.2] show that \mathcal{O} is a finite isometric quotient of the round shrinking S^3 .

Suppose that \mathcal{O} does not have positive sectional curvature. Let $f \in C^\infty(\mathcal{O})$ denote the soliton potential function. Let $\tilde{\mathcal{O}}$ be the universal cover of \mathcal{O} and let $\tilde{f} \in C^\infty(\tilde{\mathcal{O}})$ be the pullback of f to $\tilde{\mathcal{O}}$. The strong maximum principle, as in [31, Section 8], implies that if $p \in |\mathcal{O}|_{reg}$ then there is an orthogonal splitting $T_p\mathcal{O} = E_1 \oplus E_2$ which is invariant under holonomy around loops based at p . The same will be true on $\tilde{\mathcal{O}}$. Lemma 2.19 implies that $\tilde{\mathcal{O}} = \mathbb{R} \times \mathcal{O}'$ for some two-dimensional simply-connected gradient shrinking κ -solution \mathcal{O}' . From Lemma 6.4, \mathcal{O}' is the round shrinking 2-sphere or the Ricci soliton metric on a bad 2-orbifold Σ^2 . Now \tilde{f} must be $-\frac{s^2}{4} + f'$, where s is a coordinate on the \mathbb{R} -factor and f' is the soliton potential function on \mathcal{O}' . As $\pi_1(\mathcal{O})$ preserves \tilde{f} , and acts properly discontinuously and isometrically on $\mathbb{R} \times \mathcal{O}'$, it follows that $\pi_1(\mathcal{O})$ is a finite subgroup of $\text{Isom}^+(\mathbb{R} \times \mathcal{O}')$. \square

Remark 6.22. — In the manifold case, the nonexistence of noncompact positively-curved three-dimensional κ -noncollapsed gradient shrinkers was first proved by Perelman [50, Lemma 1.2]. Perelman's argument applied the Gauss-Bonnet theorem to level sets of the soliton function. This argument could be extended to orbifolds if one assumes that there are no bad 2-suborbifolds, as in Theorem 1.1. However, it is not so clear how it would extend without this assumption. Instead we use the arguments of [46, Theorem 3.1] or [52, Theorem 1.2], which do extend to the general orbifold setting.

6.10. Getting a uniform value of κ

Lemma 6.23. — *Given $N \in \mathbb{Z}^+$, there is a $\kappa_0 = \kappa_0(N) > 0$ so that if $(\mathcal{O}, g(\cdot))$ is an oriented three-dimensional κ -solution for some $\kappa > 0$, with $|G_p| \leq N$ for all $p \in |\mathcal{O}|$, then it is a κ_0 -solution or it is a quotient of the round shrinking S^3 .*

Proof. — The proof is similar to that in [38, Section 50]. The bound on $|G_p|$ gives a finite number of possible noncompact asymptotic solitons from Lemma 6.21, since a given closed two-dimensional orbifold has a unique Ricci soliton metric up to scaling, and the topological type of $S^2//\Gamma$ (or $\Sigma//\Gamma$) is determined by the number of singular points (which is at most three) and the isotropy groups of those points. \square

Lemma 6.24. — *Given $N \in \mathbb{Z}^+$, there is a universal constant $\eta = \eta(N) > 0$ such that at each point of every three-dimensional ancient solution $(\mathcal{O}, g(\cdot))$ that is a κ -solution for some $\kappa > 0$, and has $|G_p| \leq N$ for all $p \in |\mathcal{O}|$, we have estimates*

$$(6.25) \quad |\nabla R| < \eta R^{\frac{3}{2}}, \quad |R_t| \leq \eta R^2.$$

Proof. — The proof is similar to that in [38, Section 59]. \square

7. Ricci flow with surgery for orbifolds

In this section we construct the Ricci-flow-with-surgery for three-dimensional orbifolds.

Starting in Subsection 7.2, we will assume that there are no bad 2-dimensional suborbifolds. Starting in Subsection 7.5, we will assume that the Ricci flows have normalized initial conditions, as defined there.

7.1. Canonical neighborhood theorem

Definition 7.1. — Let $\Phi \in C^\infty(\mathbb{R})$ be a positive nondecreasing function such that for positive s , $\frac{\Phi(s)}{s}$ is a decreasing function which tends to zero as $s \rightarrow \infty$. A Ricci flow solution is said to have Φ -almost nonnegative curvature if for all (p, t) , we have

$$(7.2) \quad \text{Rm}(p, t) \geq -\Phi(R(p, t)).$$

Our example of Φ -almost nonnegative curvature comes from the Hamilton-Ivey pinching condition [38, Appendix B], which is valid for any three-dimensional orbifold Ricci flow solution which has complete time slices, bounded curvature on compact time intervals, and initial curvature operator bounded below by $-I$.

Proposition 7.3. — Given $\epsilon, \kappa, \sigma > 0$ and a function Φ as above, one can find $r_0 > 0$ with the following property. Let $(\mathcal{O}, g(\cdot))$ be a Ricci flow solution on a three-dimensional orbifold \mathcal{O} , defined for $0 \leq t \leq T$ with $T \geq 1$. We suppose that for each t , $g(t)$ is complete, and the sectional curvature is bounded on compact time intervals. Suppose that the Ricci flow has Φ -almost nonnegative curvature and is κ -noncollapsed on scales less than σ . Then for any point (p_0, t_0) with $t_0 \geq 1$ and $Q = R(p_0, t_0) \geq r_0^{-2}$, the solution in $\{(p, t) : \text{dist}_{t_0}^2(p, p_0) < (\epsilon Q)^{-1}, t_0 - (\epsilon Q)^{-1} \leq t \leq t_0\}$ is, after scaling by the factor Q , ϵ -close to the corresponding subset of a κ -solution.

Proof. — The proof is similar to that in [38, Section 52]. We have to allow for the possibility of neck-like regions approximated by $\mathbb{R} \times (S^2//\Gamma)$ or $\mathbb{R} \times (\Sigma^2//\Gamma)$. In the proof of [38, Lemma 52.12], the “injectivity radius” can be replaced by the “local volume”. □

7.2. Necks and horns

Assumption 7.4. — Hereafter, we only consider three-dimensional orbifolds that do not contain embedded bad 2-dimensional suborbifolds.

In particular, neck regions will be modeled on $\mathbb{R} \times (S^2//\Gamma)$, where $S^2//\Gamma$ is a quotient of the round shrinking S^2 .

We let $B(p, t, r)$ denote the open metric ball of radius r , with respect to the metric at time t , centered at $p \in |\mathcal{O}|$.

We let $P(p, t, r, \Delta t)$ denote a parabolic neighborhood, that is the set of all points (p', t') with $p' \in B(p, t, r)$ and $t' \in [t, t + \Delta t]$ or $t' \in [t + \Delta t, t]$, depending on the sign of Δt .

Definition 7.5. — An open set $U \subset |\mathcal{O}|$ in a Riemannian 3-orbifold \mathcal{O} is an ϵ -neck if modulo rescaling, it has distance less than ϵ , in the $C^{[1/\epsilon]+1}$ -topology, to a product $(-L, L) \times (S^2//\Gamma)$, where $S^2//\Gamma$ has constant scalar curvature 1 and $L > \epsilon^{-1}$. If a point $p \in |\mathcal{O}|$ and a neighborhood U of p are specified then we will understand that “distance” refers to the pointed topology. With an ϵ -approximation $f : (-L, L) \rightarrow (S^2//\Gamma) \rightarrow U$ being understood, a *cross-section* of the neck is the image of $\{\lambda\} \times (S^2//\Gamma)$ for some $\lambda \in (-L, L)$.

Definition 7.6. — A subset of the form $\mathcal{O}|_U \times [a, b] \subset \mathcal{O} \times [a, b]$ sitting in the spacetime of a Ricci flow, where $U \subset |\mathcal{O}|$ is open, is a *strong ϵ -neck* if after parabolic rescaling and time shifting, it has distance less than ϵ to the product Ricci flow defined on the time interval $[-1, 0]$ which, at its final time, is isometric to $(-L, L) \times (S^2//\Gamma)$, where $S^2//\Gamma$ has constant scalar curvature 1 and $L > \epsilon^{-1}$.

Definition 7.7. — A metric on $(-1, 1) \times (S^2//\Gamma)$ such that each point is contained in an ϵ -neck is called an ϵ -tube, an ϵ -horn or a *double ϵ -horn* if the scalar curvature stays bounded on both ends, stays bounded on one end and tends to infinity on the other, or tends to infinity on both ends, respectively.

A metric on $B^3//\Gamma$ or $(-1, 1) \times_{\mathbb{Z}_2} (S^2//\Gamma)$, such that each point outside some compact subset is contained in an ϵ -neck, is called an ϵ -cap or a *capped ϵ -horn*, if the scalar curvature stays bounded or tends to infinity on the end, respectively.

Lemma 7.8. — Let U be an ϵ -neck in an ϵ -tube (or horn) and let $S = S^2//\Gamma$ be a cross-sectional 2-sphere quotient in U . Then S separates the two ends of the tube (or horn).

Proof. — The proof is similar to that in [38, Section 58]. □

7.3. Structure of three-dimensional κ -solutions. — Recall the definition of $|\mathcal{O}|_\epsilon$ from Subsection 6.8.

Lemma 7.9. — If $(\mathcal{O}, g(t))$ is a time slice of a noncompact three-dimensional κ -solution and $|\mathcal{O}|_\epsilon \neq \emptyset$ then there is a compact suborbifold-with-boundary $X \subset \mathcal{O}$ so that $|\mathcal{O}|_\epsilon \subset X$, X is diffeomorphic to $D^3//\Gamma$ or $I \times_{\mathbb{Z}_2} (S^2//\Gamma)$, and $\mathcal{O} - \text{int}(X)$ is diffeomorphic to $[0, \infty) \times (S^2//\Gamma)$.

Proof. — The proof is similar to that in [38, Section 59]. □

Lemma 7.10. — If $(\mathcal{O}, g(t))$ is a time slice of a three-dimensional κ -solution with $|\mathcal{O}|_\epsilon = \emptyset$ then the Ricci flow is the evolving round cylinder $\mathbb{R} \times (S^2//\Gamma)$.

Proof. — The proof is similar to that in [38, Section 59]. □

Lemma 7.11. — *If a three-dimensional κ -solution $(\mathcal{O}, g(\cdot))$ is compact and has a non-compact asymptotic soliton then \mathcal{O} is diffeomorphic to $S^3//\mathbb{Z}_k$ or $S^3//D_k$ for some $k \geq 1$.*

Proof. — The proof is similar to that in [38, Section 59]. □

Lemma 7.12. — *For every sufficiently small $\epsilon > 0$ one can find $C_1 = C_1(\epsilon)$ and $C_2 = C_2(\epsilon)$ such that for each point (p, t) in every κ -solution, there is a radius $r \in [R(p, t)^{-\frac{1}{2}}, C_1 R(p, t)^{-\frac{1}{2}}]$ and a neighborhood $B, \overline{B(p, t, r)} \subset B \subset B(p, t, 2r)$, which falls into one of the four categories :*

- (a) *B is a strong ϵ -neck, or*
- (b) *B is an ϵ -cap, or*
- (c) *B is a closed orbifold diffeomorphic to $S^3//\mathbb{Z}_k$ or $S^3//D_k$ for some $k \geq 1$.*
- (d) *B is a closed orbifold of constant positive sectional curvature.*

Furthermore:

- *The scalar curvature in B at time t is between $C_2^{-1}R(p, t)$ and $C_2R(p, t)$.*
- *The volume of B in cases (a), (b) and (c) is greater than $C_2^{-1}R(p, t)^{-\frac{3}{2}}$.*
- *In case (b), there is an ϵ -neck $U \subset B$ with compact complement in B such that the distance from p to U is at least $10000R(p, t)^{-\frac{1}{2}}$.*
- *In case (c) the sectional curvature in B is greater than $C_2^{-1}R(p, t)$.*

Proof. — The proof is similar to that in [38, Section 59]. □

7.4. Standard solutions. — Put $\mathcal{O} = \mathbb{R}^3//\Gamma$, where Γ is a finite subgroup of $SO(3)$. We fix a smooth $SO(3)$ -invariant metric g_0 on \mathbb{R}^3 which is the result of gluing a hemispherical-type cap to a half-infinite cylinder $[0, \infty) \times S^2$ of scalar curvature 1. We also use g_0 to denote the quotient metric on \mathcal{O} . Among other properties, g_0 is complete and has nonnegative curvature operator. We also assume that g_0 has scalar curvature bounded below by 1.

Definition 7.13. — A Ricci flow $(\mathbb{R}^3//\Gamma, g(\cdot))$ defined on a time interval $[0, a)$ is a *standard solution* if it has complete time slices, it has initial condition g_0 , the curvature $|\text{Rm}|$ is bounded on compact time intervals $[0, a'] \subset [0, a)$, and it cannot be extended to a Ricci flow with the same properties on a strictly longer time interval.

Lemma 7.14. — *Let $(\mathbb{R}^3//\Gamma, g(\cdot))$ be a standard solution. Then:*

1. *The curvature operator of g is nonnegative.*
2. *All derivatives of curvature are bounded for small time, independent of the standard solution.*
3. *The blowup time is 1 and the infimal scalar curvature on the time-t slice tends to infinity as $t \rightarrow 1^-$ uniformly for all standard solutions.*

4. $(\mathbb{R}^3//\Gamma, g(\cdot))$ is κ -noncollapsed at scales below 1 on any time interval contained in $[0, 1)$, where κ depends only on g_0 and $|\Gamma|$.
5. $(\mathbb{R}^3//\Gamma, g(\cdot))$ satisfies the conclusion of Proposition 7.3.
6. $R_{\min}(t) \geq \text{const} \cdot (1-t)^{-1}$, where the constant does not depend on the standard solution.
7. The family \mathcal{ST} of pointed standard solutions $\{(\mathcal{M}, (p, 0))\}$ is compact with respect to pointed smooth convergence.

Proof. — Working equivariantly, the proof is the same as that in [38, Sections 60-64]. \square

7.5. Structure at the first singularity time

Definition 7.15. — Given $v_0 > 0$, a compact Riemannian three-dimensional orbifold \mathcal{O} is *normalized* if $|\text{Rm}| \leq 1$ everywhere and for every $p \in |\mathcal{O}|$, we have $\text{vol}(B(p, 1)) \geq v_0$.

Here v_0 is a global parameter in the sense that it will be fixed throughout the rest of the paper. If \mathcal{O} is normalized then the Bishop-Gromov inequality implies that there is a uniform upper bound $N = N(v_0) < \infty$ on the order of the isotropy groups; cf. the proof of Lemma 5.18. The next lemma says that by rescaling we can always achieve a normalized metric.

Lemma 7.16. — Given $N \in \mathbb{Z}^+$, there is a $v_0 = v_0(N) > 0$ with the following property. Let \mathcal{O} be a compact orientable Riemannian three-dimensional orbifold, whose isotropy groups have order at most N . Then a rescaling of \mathcal{O} will have a normalized metric.

Proof. — Let c_3 be the volume of the unit ball in \mathbb{R}^3 . Consider a ball B_r of radius $r > 0$ with arbitrary center in a Euclidean orbifold $\mathbb{R}^3//G$, where G is a finite subgroup of $O(3)$ with order at most N . Applying the Bishop-Gromov inequality to compare the volume of B_r with the volume of a very large ball having the same center, we see that $\text{vol}(B_r) \geq \frac{c_3}{N} r^3$. Put $v_0 = \frac{c_3}{2N}$. We claim that this value of v_0 satisfies the lemma.

To prove this by contradiction, suppose that there is an orbifold \mathcal{O} which satisfies the hypotheses of the lemma but for which the conclusion fails. Then there is a sequence $\{r_i\}_{i=1}^\infty$ of positive numbers with $\lim_{i \rightarrow \infty} r_i = 0$ along with points $\{p_i\}_{i=1}^\infty$ in $|\mathcal{O}|$ so that for each i , we have $\text{vol}(B(p_i, r_i)) < v_0 r_i^3$. After passing to a subsequence, we can assume that $\lim_{i \rightarrow \infty} p_i = p'$ for some $p' \in |\mathcal{O}|$. Using the inverse exponential map, for large i the ball $B(p_i, r_i)$ will, up to small distortion, correspond to a ball of radius r_i in the tangent space $T_{p'}\mathcal{O}$. In view of our choice of v_0 , this is a contradiction. \square

Assumption 7.17. — Hereafter we assume that our Ricci flows have normalized initial condition.

Consider the labels on the edges in the singular part of the orbifold. They clearly do not change under a smooth Ricci flow. If some components of the orbifold are discarded at a singularity time then the set of edge labels can only change by deletion of some labels. Otherwise, the surgery procedure will be such that the set of edge labels does not change, although the singular graphs will change. Hence the normalized initial condition implies a uniform upper bound on the orders of the isotropy groups for all time.

Let \mathcal{O} be a connected closed oriented 3-dimensional orbifold. Let $g(\cdot)$ be a Ricci flow on \mathcal{O} defined on a maximal time interval $[0, T)$ with $T < \infty$. For any $\epsilon > 0$, we know that there are numbers $r = r(\epsilon) > 0$ and $\kappa = \kappa(\epsilon) > 0$ so that for any point (p, t) with $Q = R(p, t) \geq r^{-2}$, the solution in $P(p, t, (\epsilon Q)^{-\frac{1}{2}}, (\epsilon Q)^{-1})$ is (after rescaling by the factor Q) ϵ -close to the corresponding subset of a κ -solution.

Definition 7.18. — Define a subset Ω of $|\mathcal{O}|$ by

$$(7.19) \quad \Omega = \left\{ p \in |\mathcal{O}| : \sup_{t \in [0, T)} |\text{Rm}|(p, t) < \infty \right\}.$$

Lemma 7.20. — We have

- Ω is open in $|\mathcal{O}|$.
- Any connected component of Ω is noncompact.
- If $\Omega = \emptyset$ then \mathcal{O} is diffeomorphic to $S^3//\Gamma$ or $(S^1 \times S^2)//\Gamma$.

Proof. — The proof is similar to that in [38, Section 67]. □

Definition 7.21. — Put $\bar{g} = \lim_{t \rightarrow T^-} g(t)|_{\Omega}$, a smooth Riemannian metric on $\mathcal{O}|_{\Omega}$. Let \bar{R} denote its scalar curvature.

Lemma 7.22. — (Ω, \bar{g}) has finite volume.

Proof. — The proof is similar to that in [38, Section 67]. □

Definition 7.23. — For $\rho < \frac{r}{2}$, put $\Omega_{\rho} = \{p \in |\Omega| : \bar{R}(p) \leq \rho^{-2}\}$.

Lemma 7.24. — We have

- Ω_{ρ} is a compact subset of $|\mathcal{O}|$.
- If C is a connected component of Ω which does not intersect Ω_{ρ} then C is a double ϵ -horn or a capped ϵ -horn.
- There is a finite number of connected components of Ω that intersect Ω_{ρ} , each such component having a finite number of ends, each of which is an ϵ -horn.

Proof. — The proof is similar to that in [38, Section 67]. □

7.6. δ -necks in ϵ -horns. — We define a *Ricci flow with surgery* \mathcal{M} to be the obvious orbifold extension of [38, Section 68]. The objects defined there have evident analogs in the orbifold setting.

The *r-canonical neighborhood assumption* is the obvious orbifold extension of what's in [38, Section 69], with condition (c) replaced by “ \mathcal{O} is a closed orbifold diffeomorphic to an isometric quotient of S^3 ”.

The Φ -*pinching assumption* is the same as in [38, Section 69].

The *a priori assumptions* consist of the Φ -pinching assumption and the *r-canonical neighborhood assumption*.

Lemma 7.25. — *Given the pinching function Φ , a number $\widehat{T} \in (0, \infty)$, a positive nonincreasing function $r : [0, \widehat{T}] \rightarrow \mathbb{R}$ and a number $\delta \in (0, \frac{1}{2})$, there is a nonincreasing function $h : [0, \widehat{T}] \rightarrow \mathbb{R}$ with $0 < h(t) < \delta^2 r(t)$ so that the following property is satisfied. Let \mathcal{M} be a Ricci flow with surgery defined on $[0, T)$, with $T < \widehat{T}$, which satisfies the a priori assumptions and which goes singular at time T . Let (Ω, \bar{g}) denote the time- T limit. Put $\rho = \delta r(T)$ and*

$$(7.26) \quad \Omega_\rho = \{(p, T) \in \Omega : \bar{R}(p, T) \leq \rho^{-2}\}.$$

Suppose that (p, T) lies in an ϵ -horn $\mathcal{H} \subset \Omega$ whose boundary is contained in Ω_ρ . Suppose also that $\bar{R}(p, T) \geq h^{-2}(T)$. Then the parabolic region $P(p, T, \delta^{-1}\bar{R}(p, T)^{-\frac{1}{2}}, -\bar{R}(p, T)^{-1})$ is contained in a strong δ -neck.

Proof. — The proof is similar to that in [38, Section 71]. □

7.7. Surgery and the pinching condition

Lemma 7.27. — *There exists $\delta' = \delta'(\delta) > 0$ with $\lim_{\delta \rightarrow 0} \delta'(\delta) = 0$ and a constant $\delta_0 > 0$ such that the following holds. Suppose that $\delta < \delta_0$, $p \in \{0\} \times (S^2//\Gamma)$ and h_0 is a Riemannian metric on $(-A, \frac{1}{\delta}) \times (S^2//\Gamma)$ with $A > 0$ and $R(p) > 0$ such that:*

- h_0 satisfies the time- t Hamilton-Ivey pinching condition.
- $R(p)h_0$ is δ -close to g_{cyl} in the $C^{[\frac{1}{\delta}] + 1}$ -topology.

Then there are a $B = B(A) > 0$ and a smooth metric h on $\mathbb{R}^3//\Gamma = (D^3//\Gamma) \cup ((-B, \frac{1}{\delta}) \times (S^2//\Gamma))$ such that

- h satisfies the time- t pinching condition.
- The restriction of h to $[0, \frac{1}{\delta}) \times (S^2//\Gamma)$ is h_0 .
- The restriction of $R(p)h$ to $(-B, -A) \times (S^2//\Gamma)$ is g_0 , the initial metric of a standard solution.

Proof. — The proof is the same as that in [38, Section 72], working equivariantly. □

We define a *Ricci flow with (r, δ) -cutoff* by the obvious orbifold extension of the definition in [38, Section 73].

In the surgery procedure, one first throws away all connected components of Ω which do not intersect Ω_ρ . For each connected component Ω_j of Ω that intersects

Ω_ρ and for each ϵ -horn of Ω_j , take a cross-sectional S^2 -quotient that lies far in the ϵ -horn. Let X be what's left after cutting the ϵ -horns at the 2-sphere quotients and removing the tips. The (possibly disconnected) postsurgery orbifold \mathcal{O}' is the result of capping off ∂X by discal 3-orbifolds.

Lemma 7.28. — *The presurgery orbifold can be obtained from the postsurgery orbifold by applying the following operations finitely many times:*

- Taking the disjoint union with a finite isometric quotient of $S^1 \times S^2$ or S^3 .
- Performing a 0-surgery.

Proof. — The proof is similar to that in [38, Section 73]. □

7.8. Evolution of a surgery cap

Lemma 7.29. — *For any $A < \infty$, $\theta \in (0, 1)$ and $\widehat{r} > 0$, one can find $\widehat{\delta} = \widehat{\delta}(A, \theta, \widehat{r}) > 0$ with the following property. Suppose that we have a Ricci flow with (r, δ) -cutoff defined on a time interval $[a, b]$ with $\min r = r(b) \geq \widehat{r}$. Suppose that there is a surgery time $T_0 \in (a, b)$ with $\delta(T_0) \leq \widehat{\delta}$. Consider a given surgery at the surgery time and let $(p, T_0) \in \mathcal{M}_{T_0}^+$ be the center of the surgery cap. Let $\widehat{h} = h(\delta(T_0), \epsilon, r(T_0), \Phi)$ be the surgery scale given by Lemma 7.25 and put $T_1 = \min(b, T_0 + \theta \widehat{h}^2)$. Then one of the two following possibilities occurs:*

1. *The solution is unscathed on $P(p, T_0, A\widehat{h}, T_1 - T_0)$. The pointed solution there is, modulo parabolic rescaling, A^{-1} -close to the pointed flow on $U_0 \times [0, (T_1 - T_0)\widehat{h}^{-2}]$, where U_0 is an open subset of the initial time slice $|\mathcal{S}_0|$ of a standard solution \mathcal{S} and the basepoint is the center of the cap in $|\mathcal{S}_0|$.*
2. *Assertion 1 holds with T_1 replaced by some $t^+ \in [T_0, T_1)$, where t^+ is a surgery time. Moreover, the entire ball $B(p, T_0, A\widehat{h})$ becomes extinct at time t^+ .*

Proof. — The proof is similar to that in [38, Section 74]. □

7.9. Existence of Ricci flow with surgery

Proposition 7.30. — *There exist decreasing sequences $0 < r_j < \epsilon^2$, $\kappa_j > 0$, $0 < \overline{\delta}_j < \epsilon^2$ for $1 \leq j \leq \infty$, such that for any normalized initial data on an orbifold \mathcal{O} and any nonincreasing function $\delta : [0, \infty) \rightarrow (0, \infty)$ such that $\delta \leq \overline{\delta}_j$ on $[2^{j-1}\epsilon, 2^j\epsilon]$, the Ricci flow with (r, δ) -cutoff is defined for all time and is κ -noncollapsed at scales below ϵ . Here r and κ are the functions on $[0, \infty)$ so that $r|_{[2^{j-1}\epsilon, 2^j\epsilon]} = r_j$ and $\kappa|_{[2^{j-1}\epsilon, 2^j\epsilon]} = \kappa_j$, and $\epsilon > 0$ is a global constant.*

Proof. — The proof is similar to that in [38, Sections 77-80]. □

Remark 7.31. — We restrict to 3-orbifolds without bad 2-suborbifolds in order to perform surgery. Without this assumption, there could be a neckpinch whose cross-section is a bad 2-orbifold Σ . In the case of a nondegenerate neckpinch, the blowup limit would be the product of \mathbb{R} with an evolving Ricci soliton metric on Σ . The

problem in performing surgery is that after slicing at a bad cross-section, there is no evident way to cap off the ensuing pieces with 3-dimensional orbifolds so as to preserve the Hamilton-Ivey pinching condition.

8. Hyperbolic regions

In this section we show that the w -thick part of the evolving orbifold approaches a finite-volume Riemannian orbifold with constant curvature $-\frac{1}{4}$.

As a standing assumption in this section, we suppose that we have a solution to the Ricci flow with (r, δ) -cutoff and with normalized initial data.

8.1. Double sided curvature bounds in the thick part

Proposition 8.1. — *Given $w > 0$, one can find $\tau = \tau(w) > 0$, $K = K(w) < \infty$, $\bar{r} = \bar{r}(w) > 0$ and $\theta = \theta(w) > 0$ with the following property. Let $h_{max}(t_0)$ be the maximal surgery radius on $[t_0/2, t_0]$. Let r_0 satisfy*

1. $\theta^{-1}h_{max}(t_0) \leq r_0 \leq \bar{r}\sqrt{t_0}$.
2. The ball $B(p_0, t_0, r_0)$ has sectional curvatures at least $-r_0^{-2}$ at each point.
3. $\text{vol}(B(p_0, t_0, r_0)) \geq wr_0^3$.

Then the solution is unscathed in $P(p_0, t_0, r_0/4, -\tau r_0^2)$ and satisfies $R < Kr_0^{-2}$ there.

Proof. — The proof is similar to that in [38, Sections 81-86]. In particular, it uses Proposition 5.21. \square

8.2. Noncollapsed regions with a lower curvature bound are almost hyperbolic on a large scale

Proposition 8.2

(a) *Given $w, r, \xi > 0$, one can find $T = T(w, r, \xi) < \infty$ so that the following holds. If the ball $B(p_0, t_0, r\sqrt{t_0}) \subset \mathcal{M}_{t_0}^+$ at some $t_0 \geq T$ has volume at least $wr^3r_0^{\frac{3}{2}}$ and sectional curvatures at least $-r^{-2}t_0^{-1}$ then the curvature at (p_0, t_0) satisfies*

$$(8.3) \quad |2tR_{ij}(p_0, t_0) + g_{ij}|^2 \leq \xi^2.$$

(b) *Given in addition $A < \infty$ and allowing T to depend on A , we can ensure (8.3) for all points in $B(p_0, t_0, Ar\sqrt{t_0})$.*

(c) *The same is true for $P(p_0, t_0, Ar\sqrt{t_0}, Ar^2t_0)$.*

Proof. — The proof is similar to that in [38, Sections 87 and 88]. \square

8.3. Hyperbolic rigidity and stabilization of the thick part

Definition 8.4. — Let \mathcal{O} be a complete Riemannian orbifold. Define the curvature scale as follows. Given $p \in |\mathcal{O}|$, if the connected component of \mathcal{O} containing p has nonnegative sectional curvature then put $R_p = \infty$. Otherwise, let R_p be the unique number $r \in (0, \infty)$ such that $\inf_{B(p,r)} \text{Rm} = -r^{-2}$.

Definition 8.5. — Let \mathcal{O} be a complete Riemannian orbifold. Given $w > 0$, the w -thin part $\mathcal{O}^-(w) \subset |\mathcal{O}|$ is the set of points $p \in \mathcal{O}$ so that either $R_p = \infty$ or

$$(8.6) \quad \text{vol}(B(p, R_p)) < wR_p^3.$$

The w -thick part is $\mathcal{O}^+(w) = |\mathcal{O}| - \mathcal{O}^-(w)$.

In what follows, we take “hyperbolic” to mean “constant curvature $-\frac{1}{4}$ ”. When applied to a hyperbolic orbifold, the definitions of the thick and thin parts are essentially equivalent to those in [5, Chapter 6.2], to which we refer for more information about hyperbolic 3-orbifolds.

Recall that a hyperbolic 3-orbifold can be written as $H^3//\Gamma$ for some discrete group $\Gamma \subset \text{Isom}^+(H^3)$ [19, Theorem 2.26].

Definition 8.7. — A *Margulis tube* is a compact quotient of a normal neighborhood of a geodesic in H^3 by an elementary Kleinian group.

A *rank-2 cusp neighborhood* is the quotient of a horoball in H^3 by an elementary rank-2 parabolic group.

In either case, the boundary is a compact Euclidean 2-orbifold.

There is a Margulis constant $\mu_0 > 0$ so that for any finite-volume hyperbolic 3-orbifold \mathcal{O} , if $\mu \leq \mu_0$ then the connected components of the μ -thin part of \mathcal{O} are Margulis tubes or rank-2 cusp neighborhoods.

Furthermore, given a finite-volume hyperbolic 3-orbifold \mathcal{O} , if $\mu > 0$ is sufficiently small then the connected components of the μ -thin part are rank-2 cusp neighborhoods.

Mostow-Prasad rigidity works just as well for finite-volume hyperbolic orbifolds as for finite-volume hyperbolic manifolds. Indeed, the rigidity statements are statements about lattices in $\text{Isom}(H^n)$.

Lemma 8.8. — *Let (\mathcal{O}, p) be a pointed complete connected finite-volume three-dimensional hyperbolic orbifold. Then for each $\zeta > 0$, there exists $\xi > 0$ such that if \mathcal{O}' is a complete connected finite-volume three-dimensional hyperbolic orbifold with at least as many cusps as \mathcal{O} , and $f : (\mathcal{O}, p) \rightarrow \mathcal{O}'$ is a ξ -approximation in the pointed smooth topology as in [38, Definition 90.6], then there is an isometry $f' : (\mathcal{O}, p) \rightarrow \mathcal{O}'$ which is ζ -close to f in the pointed smooth topology.*

Proof. — The proof is similar to that in [38, Section 90], replacing “injectivity radius” by “local volume”. □

If \mathcal{M} is a Ricci flow with surgery then we let $\mathcal{O}^-(w, t) \subset |\mathcal{M}_t^+|$ denote the w -thin part of the orbifold at time t (postsurgery if t is a surgery time), and similarly for the w -thick part $\mathcal{O}^+(w, t)$.

Proposition 8.9. — *Given a Ricci flow with surgery \mathcal{M} , there exist a number $T_0 < \infty$, a nonincreasing function $\alpha : [T_0, \infty) \rightarrow (0, \infty)$ with $\lim_{t \rightarrow \infty} \alpha(t) = 0$, a (possibly empty) collection $\{(H_1, x_1), \dots, (H_N, x_N)\}$ of complete connected pointed finite-volume three-dimensional hyperbolic orbifolds and a family of smooth maps*

$$(8.10) \quad f(t) : B_t = \bigcup_{i=1}^N H_i \Big|_{B(x_i, 1/\alpha(t))} \rightarrow \mathcal{M}_t,$$

defined for $t \in [T_0, \infty)$, such that

1. $f(t)$ is close to an isometry:

$$(8.11) \quad \|t^{-1}f(t)^*g_{\mathcal{M}_t} - g_{B_t}\|_{C^{1/\alpha(t)}} < \alpha(t).$$

2. $f(t)$ defines a smooth family of maps which changes smoothly with time:

$$(8.12) \quad |\dot{f}(p, t)| < \alpha(t)t^{-\frac{1}{2}}$$

for all $p \in |B_t|$, and

3. $f(t)$ parametrizes more and more of the thick part: $\mathcal{O}^+(\alpha(t), t) \subset \text{Im}(|f(t)|)$ for all $t \geq T_0$.

Proof. — The proof is similar to that in [38, Section 90]. □

9. Locally collapsed 3-orbifolds

In this section we consider compact Riemannian 3-orbifolds \mathcal{O} that are locally collapsed with respect to a local lower curvature bound. Under certain assumptions about smoothness and boundary behavior, we show that \mathcal{O} is either the result of performing 0-surgery on a strong graph orbifold or is one of a few special types. We refer to Definition 11.8 for the definition of a strong graph orbifold.

We first consider the boundaryless case.

Proposition 9.1. — *Let c_3 be the volume of the unit ball in \mathbb{R}^3 , let $K \geq 10$ be a fixed integer and let N be a positive integer. Fix a function $A : (0, \infty) \rightarrow (0, \infty)$. Then there is a $w_0 \in (0, c_3/N)$ such that the following holds.*

Suppose that (\mathcal{O}, g) is a connected closed orientable Riemannian 3-orbifold. Assume in addition that for all $p \in |\mathcal{O}|$,

1. $|G_p| \leq N$.
2. $\text{vol}(B(p, R_p)) \leq w_0 R_p^3$, where R_p is the curvature scale at p , Definition 8.4.

3. For every $w' \in [w_0, c_3/N)$, $k \in [0, K]$ and $r \leq R_p$ such that $\text{vol}(B(p, r)) \geq w'r^3$, the inequality

$$(9.2) \quad |\nabla^k \text{Rm}| \leq A(w')r^{-(k+2)}$$

holds in the ball $B(p, r)$.

Then \mathcal{O} is the result of performing 0-surgeries on a strong graph orbifold or is diffeomorphic to an isometric quotient of S^3 or T^3 .

Remark 9.3. — We recall that a strong graph orbifold is allowed to be disconnected. By Proposition 11.12, a weak graph orbifold is the result of performing 0-surgeries on a strong graph orbifold. Because of this, to prove Proposition 9.1 it is enough to show that \mathcal{O} is the result of performing 0-surgeries on a weak graph orbifold or is diffeomorphic to an isometric quotient of S^3 or T^3 .

Remark 9.4. — A 3-manifold which is an isometric quotient of S^3 or T^3 is a Seifert 3-manifold [54, Section 4]. The analogous statement for orbifolds is false [23].

Proof. — We follow the method of proof of [37]. The basic strategy is to construct a partition of \mathcal{O} into pieces whose topology can be recognized. Many of the arguments in [37], such as the stratification, are based on the underlying Alexandrov space structure. Such arguments will extend without change to the orbifold setting. Other arguments involve smoothness, which also makes sense in the orbifold setting. We now mention the relevant places in [37] where manifold smoothness needs to be replaced by orbifold smoothness.

- The critical point theory in [37, Section 3.4] can be extended to the orbifold setting using the results in Subsection 2.6.
- The results about the topology of nonnegatively curved manifolds in [37, Lemma 3.11] can be extended to the orbifold setting using Lemma 3.20 and Proposition 5.7.
- The smoothing results of [37, Section 3.6] can be extended to the orbifold setting using Lemma 2.25 and Corollary 2.26.
- The C^K -precompactness result of [37, Lemma 6.10] can be proved in the orbifold setting using Proposition 4.1.
- The C^K -splitting result of [37, Lemma 6.16] can be proved in the orbifold setting using Proposition 3.2.
- The result about the topology of the edge region in [37, Lemma 9.21] can be extended to the orbifold setting using Lemma 3.21.
- The result about the topology of the slim stratum in [37, Lemma 10.3] can be extended to the orbifold setting using Lemma 3.19.
- The results about the topology and geometry of the 0-ball regions in [37, Sections 11.1 and 11.2] can be extended to the orbifold setting using Lemma 2.24 and Proposition 3.13.

- The adapted coordinates in [37, Lemmas 8.2, 9.12, 9.17, 10.1 and 11.3] and their use in [37, Sections 12-14] extend without change to the orbifold setting.

The upshot is that we can extend the results of [37, Sections 1-14] to the orbifold setting. This gives a partition of \mathcal{O} into codimension-zero suborbifolds-with-boundary $\mathcal{O}^{0\text{-stratum}}$, \mathcal{O}^{slim} , \mathcal{O}^{edge} and $\mathcal{O}^{2\text{-stratum}}$, with the following properties.

- Each connected component of $\mathcal{O}^{0\text{-stratum}}$ is diffeomorphic either to a closed nonnegatively curved 3-dimensional orbifold, or to the unit disk bundle in the normal bundle of a soul in a complete connected noncompact nonnegatively curved 3-dimensional orbifold.
- Each connected component of \mathcal{O}^{slim} is the total space of an orbundle whose base is S^1 or I , and whose fiber is a spherical or Euclidean orientable compact 2-orbifold.
- Each connected component of \mathcal{O}^{edge} is the total space of an orbundle whose base is S^1 or I , and whose fiber is $D^2(k)$ or $D^2(2, 2)$.
- Each connected component of $\mathcal{O}^{2\text{-stratum}}$ is the total space of a circle bundle over a smooth compact 2-manifold.
- Intersections of $\mathcal{O}^{0\text{-stratum}}$, \mathcal{O}^{slim} , \mathcal{O}^{edge} and $\mathcal{O}^{2\text{-stratum}}$ are 2-dimensional orbifolds, possibly with boundary. The fibration structures coming from two intersecting strata are compatible on intersections.

In order to prove the proposition, we now follow the method of proof of [37, Section 15].

Each connected component of $\mathcal{O}^{0\text{-stratum}}$ has boundary which is empty, a spherical 2-orbifold or a Euclidean 2-orbifold. By Proposition 5.7, if the boundary is empty then the component is diffeomorphic to a finite isometric quotient of $S^1 \times S^2$, S^3 or T^3 . In the $S^1 \times S^2$ case, \mathcal{O} is a Seifert orbifold [22, p. 70-71]. Hence we can assume that the boundary is nonempty. By Lemma 3.20, if the boundary is a spherical 2-orbifold then the component is diffeomorphic to $D^3//\Gamma$ or $I \times_{\mathbb{Z}_2} (S^2//\Gamma)$. We group together such components as $\mathcal{O}_{Sph}^{0\text{-stratum}}$. By Lemma 3.20 again, if the boundary is a Euclidean 2-orbifold then the component is diffeomorphic to $S^1 \times D^2$, $S^1 \times D^2(k)$, $S^1 \times_{\mathbb{Z}_2} D^2$, $S^1 \times_{\mathbb{Z}_2} D^2(k)$ or $I \times_{\mathbb{Z}_2} (T^2//\Gamma)$. We group together such components as $\mathcal{O}_{Euc}^{0\text{-stratum}}$.

If a connected component of \mathcal{O}^{slim} fibers over S^1 then \mathcal{O} is closed and has a geometric structure based on \mathbb{R}^3 , $\mathbb{R} \times S^2$, Nil or Sol [22, p. 72]. If the structure is $\mathbb{R} \times S^2$ or Nil then \mathcal{O} is a Seifert orbifold [22, Theorem 1]. If the structure is Sol then \mathcal{O} can be cut along a fiber to see that it is a weak graph orbifold. Hence we can assume that each component of \mathcal{O}^{slim} fibers over I . We group these components into \mathcal{O}_{Sph}^{slim} and \mathcal{O}_{Euc}^{slim} , where the distinction is whether the fiber is a spherical 2-orbifold or a Euclidean 2-orbifold.

Lemma 9.5. — *Let $\mathcal{O}_i^{0\text{-stratum}}$ be a connected component of $\mathcal{O}^{0\text{-stratum}}$. If $\mathcal{O}_i^{0\text{-stratum}} \cap \mathcal{O}^{slim} \neq \emptyset$ then $\partial\mathcal{O}_i^{0\text{-stratum}}$ is a boundary component of a connected component of \mathcal{O}^{slim} .*

If $\mathcal{O}_i^{0\text{-stratum}} \cap \mathcal{O}^{slim} = \emptyset$ then we can write $\partial\mathcal{O}_i^{0\text{-stratum}} = A_i \cup B_i$ where

1. $A_i = \mathcal{O}_i^{0\text{-stratum}} \cap \mathcal{O}^{edge}$ is a disjoint union of discal 2-orbifolds and $D^2(2, 2)$'s.
2. $B_i = \mathcal{O}_i^{0\text{-stratum}} \cap \mathcal{O}^{2\text{-stratum}}$ is the total space of a circle bundle and
3. $A_i \cap B_i = \partial A_i \cap \partial B_i$ is a union of circle fibers.

Furthermore, if $\partial\mathcal{O}_i^{0\text{-stratum}}$ is Euclidean then $A_i = \emptyset$ unless $\partial\mathcal{O}_i^{0\text{-stratum}} = S^2(2, 2, 2, 2)$, in which case A_i consists of two $D^2(2, 2)$'s. If $\partial\mathcal{O}_i^{0\text{-stratum}}$ is spherical then the possibilities are

1. $\partial\mathcal{O}_i^{0\text{-stratum}} = S^2$ and A_i consists of two disks D^2 .
2. $\partial\mathcal{O}_i^{0\text{-stratum}} = S^2(k, k)$ and A_i consists of two $D^2(k)$'s.
3. $\partial\mathcal{O}_i^{0\text{-stratum}} = S^2(2, 2, k)$ and A_i consists of $D^2(2, 2)$ and $D^2(k)$.

Proof. — The proof is similar to that of [37, Lemma 15.1]. □

Lemma 9.6. — Let \mathcal{O}_i^{slim} be a connected component of \mathcal{O}^{slim} . Let Y_i be one of the connected components of $\partial\mathcal{O}_i^{slim}$. If $Y_i \cap \mathcal{O}^{0\text{-stratum}} \neq \emptyset$ then $Y_i = \partial\mathcal{O}_i^{0\text{-stratum}}$ for some connected component $\mathcal{O}_i^{0\text{-stratum}}$ of $\mathcal{O}^{0\text{-stratum}}$.

If $Y_i \cap \mathcal{O}^{0\text{-stratum}} = \emptyset$ then we can write $\partial Y_i = A_i \cup B_i$ where

1. $A_i = Y_i \cap \mathcal{O}^{edge}$ is a disjoint union of discal 2-orbifolds and $D^2(2, 2)$'s,
2. $B_i = Y_i \cap \mathcal{O}^{2\text{-stratum}}$ is the total space of a circle bundle and
3. $A_i \cap B_i = \partial A_i \cap \partial B_i$ is a union of circle fibers.

Furthermore, if Y_i is Euclidean then $A_i = \emptyset$ unless $Y_i = S^2(2, 2, 2, 2)$, in which case A_i consists of two $D^2(2, 2)$'s. If Y_i is spherical then the possibilities are

1. $Y_i = S^2$ and A_i consists of two disks D^2 .
2. $Y_i = S^2(k, k)$ and A_i consists of two $D^2(k)$'s.
3. $Y_i = S^2(2, 2, k)$ and A_i consists of $D^2(2, 2)$ and $D^2(k)$.

Proof. — The proof is similar to that of [37, Lemma 15.2]. □

Let \mathcal{O}'_{Sph} be the union of the connected components of $\mathcal{O}^{0\text{-stratum}}_{Sph} \cup \mathcal{O}^{slim}_{Sph}$ that do not intersect \mathcal{O}^{edge} . Then \mathcal{O}'_{Sph} is either empty or is all of \mathcal{O} , in which case \mathcal{O} is diffeomorphic to the gluing of two connected components of $\mathcal{O}^{0\text{-stratum}}_{Sph}$ along a spherical 2-orbifold. As each connected component is diffeomorphic to some $D^3//\Gamma$ or $I \times_{\mathbb{Z}_2} (S^2//\Gamma)$, it then follows that \mathcal{O} is diffeomorphic to $S^3//\Gamma$, $(S^3//\Gamma)//\mathbb{Z}_2$ or $S^1 \times_{\mathbb{Z}_2} (S^2//\Gamma)$, the latter of which is a Seifert 3-orbifold. Hence we can assume that each connected component of $\mathcal{O}^{0\text{-stratum}}_{Sph} \cup \mathcal{O}^{slim}_{Sph}$ intersects \mathcal{O}^{edge} . A component of \mathcal{O}^{slim}_{Sph} which intersects $\mathcal{O}^{0\text{-stratum}}_{Sph}$ can now only do so on one side, so we can collapse such a component of \mathcal{O}^{slim}_{Sph} without changing the diffeomorphism type. Thus we can assume that each connected component of $\mathcal{O}^{0\text{-stratum}}_{Sph}$ and each connected component of \mathcal{O}^{slim}_{Sph} intersects \mathcal{O}^{edge} , and that $\mathcal{O}^{0\text{-stratum}}_{Sph} \cap \mathcal{O}^{slim}_{Sph} = \emptyset$. By Lemmas 9.5 and 9.6, each of their boundary components is one of S^2 , $S^2(k, k)$ and $S^2(2, 2, k)$.

Consider the connected components of $\mathcal{O}_{Euc}^{0\text{-stratum}} \cup \mathcal{O}_{Euc}^{slim}$ whose boundary components are $S^2(2, 3, 6)$, $S^2(2, 4, 4)$ or $S^2(3, 3, 3)$. They cannot intersect any other strata, so if there is one such connected component then \mathcal{O} is formed entirely of such components. In this case \mathcal{O} is diffeomorphic to the result of gluing together two copies of $I \times_{\mathbb{Z}_2} (T^2 // \Gamma)$. Hence \mathcal{O} fibers over $S^1 // \mathbb{Z}_2$ and has a geometric structure based on \mathbb{R}^3 , Nil or Sol [22, p. 72]. If the structure is Nil then \mathcal{O} is a Seifert orbifold [22, Theorem 1]. If the structure is Sol then we can cut \mathcal{O} along a generic fiber to see that it is a weak graph orbifold. Hence we can assume that there are no connected components of $\mathcal{O}_{Euc}^{0\text{-stratum}} \cup \mathcal{O}_{Euc}^{slim}$ whose boundary components are $S^2(2, 3, 6)$, $S^2(2, 4, 4)$ or $S^2(3, 3, 3)$. Next, consider the connected components of $\mathcal{O}_{Euc}^{0\text{-stratum}} \cup \mathcal{O}_{Euc}^{slim}$ with T^2 -boundary components. They are weak graph orbifolds that do not intersect any strata other than $\mathcal{O}^{2\text{-stratum}}$. If X_1 is their complement in \mathcal{O} then in order to show that \mathcal{O} is a weak graph orbifold, it suffices to show that X_1 is a weak graph orbifold. Hence we can assume that each connected component of $\mathcal{O}_{Euc}^{0\text{-stratum}} \cup \mathcal{O}_{Euc}^{slim}$ has $S^2(2, 2, 2, 2)$ -boundary components, in which case it necessarily intersects \mathcal{O}^{edge} . As above, after collapsing some components of \mathcal{O}_{Euc}^{slim} , we can assume that each connected component of $\mathcal{O}_{Euc}^{0\text{-stratum}}$ and each connected component of \mathcal{O}_{Euc}^{slim} intersects \mathcal{O}^{edge} , and that $\mathcal{O}_{Euc}^{0\text{-stratum}} \cap \mathcal{O}_{Euc}^{slim} = \emptyset$.

A connected component of \mathcal{O}_{Sph}^{slim} is now diffeomorphic to $I \times \mathcal{O}'$, where \mathcal{O}' is diffeomorphic to S^2 , $S^2(k, k)$ or $S^2(2, k, k)$. We cut each such component along $\{\frac{1}{2}\} \times \mathcal{O}'$ and glue on two discal caps. If X_2 is the ensuing orbifold then X_1 is the result of performing a 0-surgery on X_2 , so it suffices to prove that X_2 satisfies the conclusion of the proposition. Therefore we assume henceforth that $\mathcal{O}_{Sph}^{slim} = \emptyset$.

A remaining connected component of \mathcal{O}_{Euc}^{slim} is diffeomorphic to $I \times \mathcal{O}'$, where $\mathcal{O}' = S^2(2, 2, 2, 2)$. It intersects \mathcal{O}^{edge} in four copies of $D^2(2, 2)$. We cut the connected component of \mathcal{O}_{Euc}^{slim} along $\{\frac{1}{2}\} \times \mathcal{O}'$. The result is two copies of $I \times \mathcal{O}'$, each with one free boundary component and another boundary component which intersects \mathcal{O}^{edge} in two copies of $D^2(2, 2)$. If the result X_3 of all such cuttings satisfies the conclusion of the proposition then so does X_2 , it being the result of gluing Euclidean boundary components of X_3 together.

A connected component C of \mathcal{O}^{edge} fibers over I or S^1 . Suppose that it fibers over S^1 . Then it is diffeomorphic to $S^1 \times D^2(k)$ or $S^1 \times D^2(2, 2)$, or else is the total space of a bundle over S^1 with holonomy that interchanges the two singular points in a fiber $D^2(2, 2)$; this is because the mapping class group of $D^2(2, 2)$ is a copy of \mathbb{Z}_2 , as follows from [25, Proposition 2.3]. If C is diffeomorphic to $S^1 \times D^2(k)$ or $S^1 \times D^2(2, 2)$ then it is clearly a weak graph orbifold. In the third case, $|C|$ is a solid torus and the singular locus consists of a circle labelled by 2 that wraps twice around the solid torus. See Figure 10. We can decompose C as $C = (S^1 \times_{\mathbb{Z}_2} D^2) \cup_{S^2(2,2,2,2)} C_1$, where $C_1 = S^1 \times_{\mathbb{Z}_2} (S^2 - 3B^2)$ with one B^2 being sent to itself by the \mathbb{Z}_2 -action and the other two B^2 's being switched. See Figure 11. As C_1 is a Seifert orbifold, in any case C is a weak graph orbifold. Put $X_4 = X_3 - \text{int}(C)$. If we can show that X_4 is a weak




FIGURE 10.




FIGURE 11. C_1

graph orbifold then it follows that X_3 is a weak graph orbifold. Hence we can assume that each connected component of \mathcal{O}^{edge} fibers over I .

A connected component Z of $X_4 - \text{int}(\mathcal{O}^{2\text{-stratum}})$ can be described by a graph, *i.e.*, a one-dimensional CW-complex, of degree 2. Its vertices correspond to copies of

- A connected component of $\mathcal{O}_{Sph}^{0\text{-stratum}}$ with boundary S^2 or $S^2(k, k)$,
- A connected component of $\mathcal{O}_{Euc}^{0\text{-stratum}}$ with boundary $S^2(2, 2, 2, 2)$, or
- $I \times S^2(2, 2, 2, 2)$.

Each edge corresponds to a copy of

- $I \times D^2$,
- $I \times D^2(k)$ or
- $I \times D^2(2, 2)$.

If a vertex is of type $I \times S^2(2, 2, 2, 2)$ then the edge orbifolds only intersect the vertex orbifold on a single one of its two boundary components. Note that $|Z|$ is a solid torus with a certain number of balls removed.

A connected component of $\mathcal{O}_{Sph}^{0-stratum}$ is diffeomorphic to D^3 , $D^3(k, k)$, $D^3(2, 2, k)$, $I \times_{\mathbb{Z}_2} S^2$, or $I \times_{\mathbb{Z}_2} S^2(2, 2, k)$. Now $I \times_{\mathbb{Z}_2} S^2$ is diffeomorphic to $\mathbb{R}P^3 \# D^3$, $I \times_{\mathbb{Z}_2} S^2(k, k)$ is diffeomorphic to $(S^3(k, k) // \mathbb{Z}_2) \#_{S^2(k, k)} D^3(k, k)$ and $I \times_{\mathbb{Z}_2} S^2(2, 2, k)$ is diffeomorphic to $(S^3(2, 2, k) // \mathbb{Z}_2) \#_{S^2(2, 2, k)} D^3(2, 2, k)$, where \mathbb{Z}_2 acts by the antipodal action. Hence we can reduce to the case when each connected component of $\mathcal{O}_{Sph}^{0-stratum}$ is diffeomorphic to D^3 , $D^3(k, k)$ or $D^3(2, 2, k)$, modulo performing connected sums with the Seifert orbifolds $\mathbb{R}P^3$, $S^3(k, k) // \mathbb{Z}_2$ and $S^3(2, 2, k) // \mathbb{Z}_2$.

Any connected component of $\mathcal{O}_{Euc}^{0-stratum}$ with boundary $S^2(2, 2, 2, 2)$ can be written as the gluing of a weak graph orbifold with $I \times S^2(2, 2, 2, 2)$. Hence we may assume that there are no vertices corresponding to connected components of $\mathcal{O}_{Euc}^{0-stratum}$ with boundary $S^2(2, 2, 2, 2)$.

Suppose that there are no edges of type $I \times D^2(2, 2)$. Then Z is $I \times D^2$ or $I \times D^2(k)$, which is a weak graph orbifold.

Now suppose that there is an edge of type $I \times D^2(2, 2)$. We build up a skeleton for Z . First, the orbifold corresponding to a graph with a single vertex of type $I \times S^2(2, 2, 2, 2)$, and a single edge of type $I \times D^2(2, 2)$, can be identified as the Seifert orbifold $C_1 = S^1 \times_{\mathbb{Z}_2} (S^2 - 3B^2)$ of before. Let C_m be the orbifold corresponding to a graph with m vertices of type $I \times S^2(2, 2, 2, 2)$ and m edges of type $I \times D^2(2, 2)$. See Figure 12. Then C_m is an m -fold cover of C_1 and is also a Seifert orbifold.




FIGURE 12. C_m , $m = 3$

Returning to the orbifold Z , there is some m so that Z is diffeomorphic to the result of starting with C_m and gluing some $S^1 \times_{\mathbb{Z}_2} D^2(k_i)$'s onto some of the boundary $S^2(2, 2, 2, 2)$'s, where $k_i \geq 1$. See Figure 13 for an illustrated example.

Thus Z is a weak graph orbifold.

As X_3 is the result of gluing Z to a circle bundle over a surface, X_3 is a weak graph orbifold. Along with Proposition 11.12, this proves the proposition. \square

Proposition 9.7. — *Let c_3 be the volume of the unit ball in \mathbb{R}^3 , let $K \geq 10$ be a fixed integer and let N be a positive integer. Fix a function $A : (0, \infty) \rightarrow (0, \infty)$. Then there is a $w_0 \in (0, c_3/N)$ such that the following holds.*




FIGURE 13.

Suppose that (\mathcal{O}, g) is a compact connected orientable Riemannian 3-orbifold with boundary. Assume in addition that

1. $|G_p| \leq N$.
2. The diameters of the connected components of $\partial\mathcal{O}$ are bounded above by w_0 .
3. For each component X of $\partial\mathcal{O}$, there is a hyperbolic orbifold cusp \mathcal{H}_X with boundary $\partial\mathcal{H}_X$, along with a C^{K+1} -embedding of pairs $e : (N_{100}(\partial\mathcal{H}_X), \partial\mathcal{H}_X) \rightarrow (\mathcal{O}, X)$ which is w_0 -close to an isometry.
4. For every $p \in |\mathcal{O}|$ with $d(p, \partial\mathcal{O}) \geq 10$, we have, $\text{vol}(B(p, R_p)) \leq w_0 R_p^3$.
5. For every $p \in |\mathcal{O}|$, $w' \in [w_0, c_3/N]$, $k \in [0, K]$ and $r \leq R_p$ such that $\text{vol}(B(p, r)) \geq w' r^3$, the inequality

$$(9.8) \quad |\nabla^k \text{Rm}| \leq A(w') r^{-(k+2)}$$

holds in the ball $B(p, r)$.

Then \mathcal{O} is diffeomorphic to

- The result of performing 0-surgeries on a strong graph orbifold,
- A closed isometric quotient of S^3 or T^3 ,
- $I \times S^2(2, 3, 6)$, $I \times S^2(2, 4, 4)$ or $I \times S^2(3, 3, 3)$, or
- $I \times_{\mathbb{Z}_2} S^2(2, 3, 6)$, $I \times_{\mathbb{Z}_2} S^2(2, 4, 4)$ or $I \times_{\mathbb{Z}_2} S^2(3, 3, 3)$.

Proof. — We follow the method of proof of [38, Section 16]. The effective difference from the proof of Proposition 9.1 is that we have additional components of $\mathcal{O}^{0\text{-stratum}}$, which are diffeomorphic to $I \times (T^2/\Gamma)$. If such a component is diffeomorphic to $I \times T^2$ or $I \times S^2(2, 2, 2, 2)$ then we can incorporate it into the weak graph orbifold structure. The other cases give rise to the additional possibilities listed in the conclusion of the proposition. \square

10. Incompressibility of cuspidal cross-sections and proof of Theorem 1.1

In this section we complete the proof of Theorem 1.1.

With reference to Proposition 8.9, given a sequence $t^\alpha \rightarrow \infty$, let Y^α be the truncation of $\coprod_{i=1}^N H_i$ obtained by removing horoballs at distance approximately $\frac{1}{2\beta(t^\alpha)}$ from the basepoints x_i . Put $\mathcal{O}^\alpha = \mathcal{O}_{t^\alpha} - f_{t^\alpha}(Y^\alpha)$.

Proposition 10.1. — *For large α , the orbifold \mathcal{O}^α satisfies the hypotheses of Proposition 9.7.*

Proof. — The proof is similar to that of [37, Theorem 17.3]. □

So far we know that if α is large then the 3-orbifold \mathcal{O}_{t^α} has a (possibly empty) hyperbolic piece whose complement satisfies the conclusion of Proposition 9.7. In this section we show that there is such a decomposition of \mathcal{O}_{t^α} so that the hyperbolic cusps, if any, are incompressible in \mathcal{O}_{t^α} .

The corresponding manifold result was proved by Hamilton in [33] using minimal disks. He used results of Meeks-Yau [43] to find embedded minimal disks with boundary on an appropriate cross-section of the cusp. The Meeks-Yau proof in turn used a tower construction [42] similar to that used in the proof of Dehn's Lemma in 3-manifold topology. It is not clear to us whether this line of proof extends to three-dimensional orbifolds, or whether there are other methods using minimal disks which do extend. To circumvent these issues, we use an alternative incompressibility argument due to Perelman [50, Section 8.2] that exploits certain quantities which change monotonically under the Ricci flow. Perelman's monotonic quantity involved the smallest eigenvalue of a certain Schrödinger-type operator. We will instead use a variation of Perelman's argument involving the minimal scalar curvature, following [38, Section 93.4].

Before proceeding, we need two lemmas:

Lemma 10.2. — *Suppose $\epsilon > 0$, and \mathcal{O}' is a Riemannian 3-orbifold with scalar curvature $\geq -\frac{3}{2}$. Then any orbifold \mathcal{O} obtained from \mathcal{O}' by 0-surgeries admits a Riemannian metric with scalar curvature $\geq -\frac{3}{2}$, such that $\text{vol}(\mathcal{O}) < \text{vol}(\mathcal{O}') + \epsilon$.*

Proof. — If a 0-surgery adds a neck $(S^2//\Gamma) \times I$ then we can put a metric on the neck which is an isometric quotient of a slight perturbation of the doubled Schwarzschild metric [2, (1.23)] on $S^2 \times I$. Hence we can perform the 0-surgery so that the scalar curvature is bounded below by $-\frac{3}{2} + \frac{\epsilon}{10}$ and the volume increases by at most $\frac{\epsilon}{10}$; see [2, p. 155] and [51] for the analogous result in the manifold case. The lemma now follows from an overall rescaling to make $R \geq -\frac{3}{2}$. □

Lemma 10.3. — *Suppose that \mathcal{O} is a strong graph orbifold with boundary components C_1, \dots, C_k . Let H_1, \dots, H_k be truncated hyperbolic cusps, where ∂H_i is diffeomorphic to C_i for all $i \in \{1, \dots, k\}$. Then for all $\epsilon > 0$, there is a metric on \mathcal{O} with scalar*

curvature $\geq -\frac{3}{2}$ such that $\text{vol}(\mathcal{O}) < \epsilon$, and C_i has a collar which is isometric to one side of a collar neighborhood of a cuspidal 2-orbifold in H_i .

Proof. — We first prove the case when \mathcal{O} is a closed strong graph manifold. The strong graph manifold structure gives a graph whose vertices $\{v_a\}$ correspond to the Seifert blocks and whose edges $\{e_b\}$ correspond to 2-tori. For each vertex v_a , let M_a be the corresponding Seifert block. We give it a Riemannian metric g_a which is invariant under the local S^1 -actions and with the property that the quotient metric on the orbifold base is a product near its boundary. Then g_a has a product structure near ∂M_a . Given $\delta > 0$, we uniformly shrink the Riemannian metric on g_a by δ in the fiber directions. As $\delta \rightarrow 0$, the volume of M_a goes to zero while the curvature stays bounded.

Let T_b^2 be the torus corresponding to the edge e_b . There are associated toral boundary components $\{B_1, B_2\}$ of Seifert blocks. Given $\delta > 0$ and $i \in \{1, 2\}$, consider the warped product metric $ds^2 + e^{-2s}g_{B_i}$ on a product manifold $P_{\delta,i} = [0, L_{\delta,i}] \times B_i$. We attach this at B_i to obtain a C^0 -metric, which we will smooth later. The sectional curvatures of $P_{\delta,i}$ are -1 and the volume of $P_{\delta,i}$ is bounded above by the area of B_i . We choose $L_{\delta,i}$ so that the areas of the cross-sections $\{L_{\delta,1}\} \times B_1$ and $\{L_{\delta,2}\} \times B_2$ are both equal to some number A . Finally, consider \mathbb{R}^3 with the Sol-invariant metric $e^{-2z}dx^2 + e^{2z}dy^2 + dz^2$. Let Γ be a \mathbb{Z}^2 -subgroup of the normal \mathbb{R}^2 -subgroup of Sol. Note that the curvature of \mathbb{R}^3/Γ is independent of Γ . The z -coordinate gives a fibering $z : \mathbb{R}^3/\Gamma \rightarrow \mathbb{R}$ with T^2 -fibers. We can choose $\Gamma = \Gamma_\delta$ and an interval $[c_1, c_2] \subset \mathbb{R}$ so that $z^{-1}(c_1)$ is isometric to $\{L_{\delta,1}\} \times B_1$ and $z^{-1}(c_2)$ is isometric to $\{L_{\delta,2}\} \times B_2$. Note that $[c_1, c_2]$ can be taken independent of A . We attach $z^{-1}([c_1, c_2])$ to the previously described truncated cusps, at the boundary components $\{L_{\delta,1}\} \times B_1$ and $\{L_{\delta,2}\} \times B_2$. See Figure 14.




FIGURE 14.

Taking A sufficiently small we can ensure that

$$(10.4) \quad \text{vol}(P_{\delta,1}) + \text{vol}(P_{\delta,2}) + \text{vol}(z^{-1}([c_1, c_2])) < \text{area}(B_1) + \text{area}(B_2) + \delta.$$

We repeat this process for all of the tori $\{T_b^2\}$, to obtain a piecewise-smooth C^0 -metric g_δ on \mathcal{O} .

As $\delta \rightarrow 0$, the sectional curvature stays uniformly bounded on the smooth pieces. Furthermore, the volume of (\mathcal{O}, g_δ) goes to zero. By slightly smoothing g_δ and performing an overall rescaling to ensure that the scalar curvature is bounded below by $-\frac{3}{2}$, if δ is sufficiently small then we can ensure that $\text{vol}(\mathcal{O}, g_\delta) < \epsilon$. This proves the lemma when \mathcal{O} is a closed strong graph manifold.

If \mathcal{O} is a strong graph manifold but has nonempty boundary components, as in the hypotheses of the lemma, then we treat each boundary component C_i analogously to a factor B_1 in the preceding construction. That is, given parameters $0 < c_{1,C_i} < c_{2,C_i}$, we start by putting a truncated hyperbolic metric $ds^2 + e^{-2s}g_{\partial H_i}$ on $[c_{1,C_i}, c_{2,C_i}] \times C_i$. This will be the metric on the collar neighborhood of C_i , where $\{c_{1,C_i}\} \times C_i$ will end up becoming a boundary component of \mathcal{O} . We take c_{2,C_i} so that the area of $\{c_{2,C_i}\} \times C_i$ matches the area of a relevant cross-section of the truncated cusp extending from a boundary component $B_{2,i}$ of a Seifert block. We then construct a metric g_δ on \mathcal{O} as before. If we additionally take the parameters $\{c_{1,C_i}\}$ sufficiently large then we can ensure that $\text{vol}(\mathcal{O}, g_\delta) < \epsilon$.

Finally, if \mathcal{O} is a strong graph orbifold then we can go through the same steps. The only additional point is to show that elements of the (orientation-preserving) mapping class group of an oriented Euclidean 2-orbifold $T^2//\Gamma$ are represented by affine diffeomorphisms, in order to apply the preceding construction using the Sol geometry. To see this fact, if Γ is trivial then the mapping class group of T^2 is isomorphic to $\text{SL}(2, \mathbb{Z})$ and the claim is clear. To handle the case when $T^2//\Gamma$ is a sphere with three singular points, we use the fact that the mapping class group of a sphere with three marked points is isomorphic to the permutation group of the three points [25, Proposition 2.3]. The mapping class group of the orbifold $T^2//\Gamma$ will then be the subgroup of the permutation group that preserves the labels. If $T^2//\Gamma$ is $S^2(2, 3, 6)$ then its mapping class group is trivial. If $T^2//\Gamma$ is $S^2(2, 4, 4)$ then its mapping class group is isomorphic to \mathbb{Z}_2 . Picturing $S^2(2, 4, 4)$ as two right triangles glued together, the nontrivial mapping class group element is represented by the affine diffeomorphism which is a flip around the “2” vertex that interchanges the two triangles. If $T^2//\Gamma$ is $S^2(3, 3, 3)$ then its mapping class group is isomorphic to S_3 . Picturing $S^2(3, 3, 3)$ as two equilateral triangles glued together, the nontrivial mapping class group elements are represented by affine diffeomorphisms as rotations and flips. Finally, if $T^2//\Gamma$ is $S^2(2, 2, 2, 2)$ then its mapping class group is isomorphic to $\text{PSL}(2, \mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z})$ [25, Proposition 2.7]. These all lift to \mathbb{Z}_2 -equivariant affine diffeomorphisms of T^2 . Elements of $\text{PSL}(2, \mathbb{Z})$ are represented by linear actions of $\text{SL}(2, \mathbb{Z})$ on T^2 . Generators of $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ are represented by rotations of the S^1 -factors in $T^2 = S^1 \times S^1$ by π . \square

Let \mathcal{O} be a closed connected orientable three-dimensional orbifold. If \mathcal{O} admits a metric of positive scalar curvature then by finite extinction time, \mathcal{O} is diffeomorphic to the result of performing 0-surgeries on a disjoint collection of isometric quotients of S^3 and $S^1 \times S^2$.

Suppose that \mathcal{O} does not admit a metric of positive scalar curvature. Put

$$(10.5) \quad \sigma(\mathcal{O}) = \sup_g R_{\min}(g)V(g)^{\frac{2}{3}}.$$

Then $\sigma(\mathcal{O}) \leq 0$.

Suppose that we have a given representation of \mathcal{O} as the result of performing 0-surgeries on the disjoint union of an orbifold \mathcal{O}' and isometric quotients of S^3 and $S^1 \times S^2$, and that there exists a (possibly empty, possibly disconnected) finite-volume complete hyperbolic orbifold N which can be embedded in \mathcal{O}' so that the connected components of the complement (if nonempty) satisfy the conclusion of Proposition 9.7. Let V_{hyp} denote the hyperbolic volume of N . We do not assume that the cusps of N are incompressible in \mathcal{O}' .

Let \widehat{V} denote the minimum of V_{hyp} over all such decompositions of \mathcal{O} . (As the set of volumes of complete finite-volume three-dimensional hyperbolic orbifolds is well-ordered, there is a minimum. If there is a decomposition with $N = \emptyset$ then $V_{hyp} = 0$.)

Lemma 10.6

$$(10.7) \quad \sigma(\mathcal{O}) = -\frac{3}{2}\widehat{V}^{\frac{2}{3}}.$$

Proof. — Using Lemmas 7.28, 10.2 and 10.3, the proof is similar to that of [38, Proposition 93.10]. □

Proposition 10.8. — *Let N be a hyperbolic orbifold as above for which $\text{vol}(N) = \widehat{V}$. Then the cuspidal cross-sections of N are incompressible in \mathcal{O}' .*

Proof. — As in [38, Section 93], it suffices to show that if a cuspidal cross-section of N is compressible in \mathcal{O}' then there is a metric g on \mathcal{O} with $R(g) \geq -\frac{3}{2}$ and $\text{vol}(\mathcal{O}, g) < \text{vol}(N)$.

Put $Y = \mathcal{O}' - N$. Suppose that some connected component C_0 of ∂Y is compressible, with compressing discal 2-orbifold $Z \subset \mathcal{O}'$. We can make Z transverse to ∂Y and then count the number of connected components of the intersection $Z \cap \partial Y$. Minimizing this number among all such compressing disks for all compressible components of ∂Y , we may assume – after possibly replacing C_0 with a different component of ∂Y – that Z intersects ∂Y only along ∂Z .

By assumption, the components of Y satisfy the conclusion of Proposition 9.7. Hence Y has a decomposition into connected components $Y = Y_0 \sqcup \dots \sqcup Y_n$, where Y_0 is the component containing C_0 , and Y_0 arises from a strong graph orbifold by 0-surgeries, as otherwise there would not be a compressing discal orbifold. By Lemma 11.16, Y_0 comes from a disjoint union $A \sqcup B$ via 0-surgeries, where A is one

of the four solid-toric possibilities of that Lemma, and B is a strong graph orbifold. By Lemmas 10.2 and 10.3, we may assume without loss of generality that $B = \emptyset$.

To construct the desired metric on \mathcal{O}' , we proceed as follows. Let H_0, \dots, H_n be the cusps of the hyperbolic orbifold N , where H_0 corresponds to the component C_0 of Y . We first truncate N along totally umbilic cuspidal 2-orbifolds C_0, \dots, C_n . Pick $\epsilon > 0$. For each $i \geq 1$ such that the component Y_i comes from 0-surgeries on a strong graph orbifold, we use Lemmas 10.2 and 10.3 to find a metric with $R \geq -\frac{3}{2}$ on Y_i , which glues isometrically along the corresponding cusps in $C_1 \sqcup \dots \sqcup C_n$, and which can be arranged to have volume $< \epsilon$ by taking the C_i 's to be deep in their respective cusps. For the components Y_i , $i \geq 1$, which do not come from a strong graph orbifold *via* 0-surgery, we may also find metrics with $R \geq -\frac{3}{2}$ and arbitrarily small volume, which glue isometrically onto the corresponding truncated cusps of N (when they have nonempty boundary). Our final step will be to find a metric on $Y_0 = A$ with $R \geq -\frac{3}{2}$ which glues isometrically to C_0 , and has volume strictly smaller than the portion of the cusp H_0 cut off by C_0 . Since ϵ is arbitrary, this will yield a contradiction.

Suppose first that A is $S^1 \times D^2$ or $S^1 \times D^2(k)$. In the $S^1 \times D^2$ case, after going far enough down the cusp, the desired metric g on $S^1 \times D^2$ is constructed in [2, Pf. of Theorem 2.9]. (The condition $f_2(0) = a > 0$ in [2, (2.47)] should be changed to $f_2(0) > 0$.) In the $S^1 \times D^2(k)$ -case, [2, (2.46)] gets changed to $f_1'(0)(1-a^2)^{1/2} = 1/k$. One can then make the appropriate modifications to [2, (2.54)-(2.56)] to construct the desired metric g on $S^1 \times D^2(k)$.

If A is $S^1 \times_{\mathbb{Z}_2} D^2$ or $S^1 \times_{\mathbb{Z}_2} D^2(k)$ we can perform the construction of the previous paragraph equivariantly with respect to the \mathbb{Z}_2 -action, to form the desired metric on $S^1 \times_{\mathbb{Z}_2} D^2$ (or $S^1 \times_{\mathbb{Z}_2} D^2(k)$). \square

10.1. Proof of Theorem 1.1. — As mentioned before, if \mathcal{O} admits a metric of positive scalar curvature then \mathcal{O} is diffeomorphic to the result of performing 0-surgeries on a disjoint collection of isometric quotients of S^3 and $S^1 \times S^2$, so the theorem is true in that case. If \mathcal{O} does not admit a metric of positive scalar curvature then by Proposition 10.8,

1. \mathcal{O} is the result of performing 0-surgeries on an orbifold \mathcal{O}' and a disjoint collection of isometric quotients of S^3 and $S^1 \times S^2$, such that
2. There is a finite-volume complete hyperbolic orbifold N which can be embedded in \mathcal{O}' so that each connected component \mathcal{P} of the complement (if nonempty) has a metric completion $\overline{\mathcal{P}}$ which satisfies the conclusion of Proposition 9.7, and
3. The cuspidal cross-sections of N are incompressible in \mathcal{O}' .

Referring to Proposition 9.7, if $\overline{\mathcal{P}}$ is an isometric quotient of S^3 or T^3 then it already has a geometric structure. If $\overline{\mathcal{P}}$ is $I \times S^2(p, q, r)$ with $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$ then we can remove it without losing any information. If $\overline{\mathcal{P}}$ is $I \times_{\mathbb{Z}_2} S^2(p, q, r)$ with $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$ then \mathcal{P} has a Euclidean structure.

Finally, suppose that $\overline{\mathcal{P}}$ is the result of performing 0-surgeries on a collection of strong graph orbifolds in the sense of Definition 11.8. A Seifert-fibered 3-orbifold with no bad 2-dimensional suborbifolds is geometric in the sense of Thurston [5, Proposition 2.13]. This completes the proof of Theorem 1.1. \square

Remark 10.9. — The geometric decomposition of \mathcal{O} that we have produced, using strong graph orbifolds, will not be minimal if \mathcal{O} has Sol geometry. In such a case, \mathcal{O} fibers over a 1-dimensional orbifold. Cutting along a fiber and taking the metric completion gives a product orbifold, which is a graph orbifold. Of course, the minimal geometric decomposition of \mathcal{O} would leave it with its Sol structure.

Remark 10.10. — Theorem 1.1 implies that \mathcal{O} is very good, *i.e.*, the quotient of a manifold by a finite group action [4, Corollary 1.3]. Hence one could obtain the geometric decomposition of \mathcal{O} by running Perelman’s proof equivariantly, as is done in detail for elliptic and hyperbolic manifolds in [21]. However, one cannot prove the geometrization of orbifolds this way, as the reasoning would be circular; one only knows that \mathcal{O} is very good after proving Theorem 1.1.

11. Appendix A : Weak and strong graph orbifolds

In this appendix we provide proofs of some needed facts about graph orbifolds. We show that a weak graph orbifold is the result of performing 0-surgeries on a strong graph orbifold. (Since we don’t require strong graph orbifolds to be connected, we need only one.) A similar result appears in [24, Section 2.4].

In order to clarify the arguments, we prove the corresponding manifold results before proving the orbifold results.

Definition 11.1. — A *weak graph manifold* is a compact orientable 3-manifold M for which there is a collection $\{T_i\}$ of disjoint embedded tori in $\text{int}(M)$ so that after splitting M open along $\{T_i\}$, the result has connected components that are Seifert-fibered 3-manifolds (possibly with boundary).

We do not assume that M is connected. Here “splitting M open along $\{T_i\}$ ” means taking the metric completion of $M - \bigcup_i T_i$ with respect to an arbitrary Riemannian metric on M .

Remark 11.2. — In the definition of a weak graph manifold, we could have instead required that the connected components of the metric completion of $M - \bigcup_i T_i$ are circle bundles over surfaces. This would give an equivalent notion, since any Seifert-fibered 3-manifold can be cut along tori into circle bundles over surfaces.

For notation, we will write $S^2 - kB^2$ for the complement of k disjoint separated open 2-balls in S^2 .

Definition 11.3. — A *strong graph manifold* is a compact orientable 3-manifold M for which there is a collection $\{T_i\}$ of disjoint embedded tori in $\text{int}(M)$ such that

1. After splitting M open along $\{T_i\}$, the result has connected components that are Seifert manifolds (possibly with boundary).
2. For any T_i , the two circle fibrations on T_i coming from the adjacent Seifert bundles are not isotopic.
3. Each T_i is incompressible in M .

11.1. Appendix A.1 : Weak graph manifolds are connected sums of strong graph manifolds. —

The next lemma states if we glue two solid tori (respecting orientations) then the result is a Seifert manifold. The lemma itself is trivial, since we know that the manifold is $S^1 \times S^2$, S^3 or a lens space, each of which is a Seifert manifold. However, we give a proof of the lemma which will be useful in the orbifold case.

Lemma 11.4. — *Let U and V be two oriented solid tori. Let $\phi : \partial U \rightarrow \partial V$ be an orientation-reversing diffeomorphism. Then $U \cup_\phi V$ admits a Seifert fibration.*

Proof. — We first note that the circle fiberings of T^2 are classified, up to isotopy, by the image of the fiber in $(H^1(T^2; \mathbb{Z}) - \{0\})/\{\pm 1\} \simeq (\mathbb{Z}^2 - \{0\})/\{\pm 1\}$. There is one circle fibering of ∂U (up to isotopy) whose fibers bound compressing disks in U . Any other circle fibering of ∂U is the boundary fibration of a Seifert fibration of U . Hence we can choose a circle fibering \mathcal{F} of ∂U so that \mathcal{F} is the boundary fibration of a Seifert fibration of U , and $\phi_*\mathcal{F}$ is the boundary fibration of a Seifert fibration of V . The ensuing Seifert fibrations of U and V join together to give a Seifert fibration of $U \cup_\phi V$. \square

Proposition 11.5. — *If a connected strong graph manifold contains an essential embedded 2-sphere then it is diffeomorphic to $S^1 \times S^2$ or $\mathbb{R}P^3 \# \mathbb{R}P^3$.*

Proof. — Suppose that a connected strong graph manifold M contains an essential embedded 2-sphere S . We can assume that S is transverse to $\bigcup_i T_i$. We choose S among all such essential embedded 2-spheres so that the number of connected components of $S \cap \bigcup_i T_i$ is as small as possible.

If $S \cap \bigcup_i T_i = \emptyset$ then S is an essential 2-sphere in one of the Seifert components.

If $S \cap \bigcup_i T_i \neq \emptyset$, let C be an innermost circle in $S \cap \bigcup_i T_i$. Then $C \subset T_k$ for some k and $C = \partial D$ for some 2-disk D embedded in a Seifert component U with $T_k \subset \partial U$. As T_k is incompressible, $C = \partial D'$ for some 2-disk $D' \subset T_k$. If $D \cup D'$ bounds a 3-ball in U then we can isotope S to remove the intersection with T_k , which contradicts the choice of S . Thus $D \cup D'$ is an essential 2-sphere in U .

In any case, we found an essential 2-sphere in one of the Seifert pieces. It follows that the Seifert piece, and hence all of M , is diffeomorphic to $S^1 \times S^2$ or $\mathbb{R}P^3 \# \mathbb{R}P^3$ [54, p. 432]. \square

Proposition 11.6. — *A weak graph manifold is the result of performing 0-surgeries on a strong graph manifold.*

Proof. — Suppose that Proposition 11.6 fails. Let n be the minimal number of decomposing tori among weak graph manifolds which are counterexamples, and let M be a counterexample with decomposing tori $\{T_i\}_{i=1}^n$.

We first look for a torus T_j for which the two induced circle fibrations (coming from the adjacent Seifert bundles) are isotopic. If there is one then we extend the Seifert fibration over T_j . In this case, by removing T_j from $\{T_i\}_{i=1}^n$, we get a weak graph decomposition of M with $(n - 1)$ tori, contradicting the definition of n .

Therefore there is no such torus. Since M is a counterexample to Proposition 11.6, there must be a torus in $\{T_i\}_{i=1}^n$ which is compressible. Let D be a compressing disk, which we can assume to be transversal to $\bigcup_{i=1}^n T_i$. We choose such a compressing disk so that $D \cap \bigcup_{i=1}^n T_i$ has the smallest possible number of connected components. Let C be an innermost circle in $D \cap \bigcup_{i=1}^n T_i$, say lying in T_k . Then C bounds a disk D' in a Seifert bundle V which has T_k as a boundary component.

If C also bounds a disk $D'' \subset T_k$ then $D' \cup D''$ is an embedded 2-sphere S in V . If S is not essential in V then we can isotope D so that it does not intersect T_k , which contradicts the choice of D . So S is essential in V . Then V is diffeomorphic to $S^1 \times S^2$ or $\mathbb{R}P^3 \# \mathbb{R}P^3$, which contradicts the assumption that it has T_k as a boundary component.

Thus we can assume that D' is a compressing disk for V , which is necessarily a solid torus [54, Corollary 3.3].

Let U be the Seifert bundle on the other side of T_k from V . Let B be the orbifold base of U , with projection $\pi : U \rightarrow B$. There is a circle boundary component $R \subset \partial B$ so that $T_k = \pi^{-1}(R)$. That is, V is glued to U along $\pi^{-1}(R)$. Choose a D^2 -fibration $\sigma : V \rightarrow R$ that extends $\pi : T_k \rightarrow R$.

If $C = \partial D' \subset T_k$ is not isotopic to a fiber of $\pi|_{T_k}$, let $u > 0$ be their algebraic intersection number in T_k . Then $U \cup_{T_k} V$ has a Seifert fibration over $B \cup_R D^2(u)$. Removing T_k from $\{T_i\}_{i=1}^n$, we again have a weak graph decomposition of M , now with $(n - 1)$ tori, which is a contradiction.

Therefore $C = \partial D' \subset T_k$ is isotopic to a fiber of $\pi|_{T_k}$.

Step 1: If B is diffeomorphic to D^2 , $D^2(r)$ or $S^1 \times I$ then put $M' = M$ and $B' = B$, and go to Step 2. Otherwise, let $\{\gamma_j\}_{j=1}^J$ be a maximal disjoint collection of smooth embedded arcs $\gamma_j : [0, 1] \rightarrow B_{reg}$, with $\{\gamma_j(0), \gamma_j(1)\} \subset R$, which determine distinct nontrivial homotopy classes for the pair (B_{reg}, R) . (Note that $\partial B \subset B_{reg}$.) If B' is the result of splitting B open along $\{\gamma_j\}_{j=1}^J$, then the connected components of B' are diffeomorphic to D^2 , $D^2(r)$ for some $r > 1$, or $S^1 \times I$. See Figure 15. Let R' be the result of splitting the 1-manifold R along the finite subset $\bigcup_{j=1}^J \{\gamma_j(0), \gamma_j(1)\}$.

Define a 2-sphere $S_j^2 \subset M$ by $S_j^2 = \sigma^{-1}(\gamma_j(0)) \cup_{\pi^{-1}(\gamma_j(0))} \pi^{-1}(\gamma_j) \cup_{\pi^{-1}(\gamma_j(1))} \sigma^{-1}(\gamma_j(1))$. Let Y be the result of splitting M open along $\{S_j^2\}_{j=1}^J$. It has $2J$

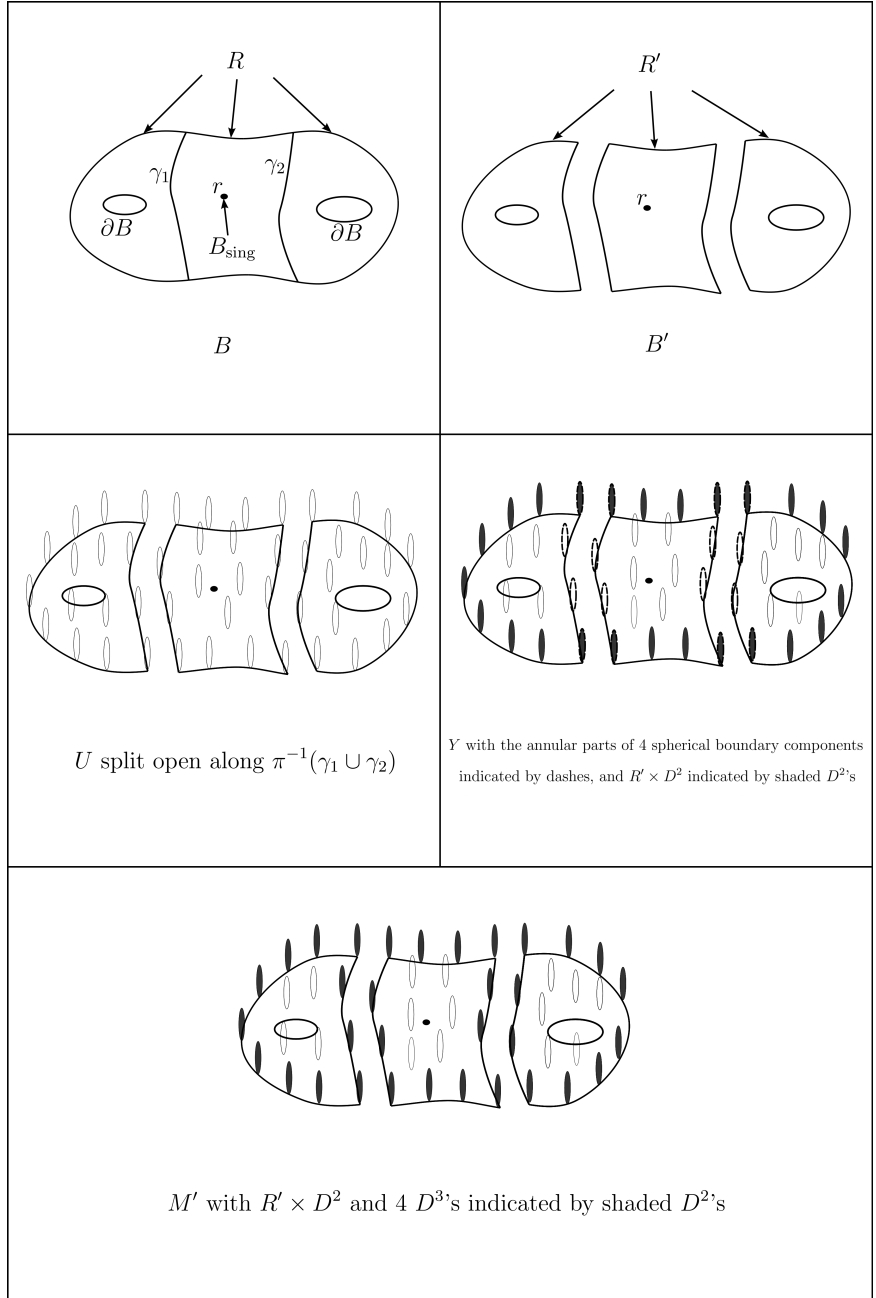


FIGURE 15.

spherical boundary components corresponding to the spherical cuts. We glue on $2J$ 3-disks there, to obtain M' . By construction, M is the result of performing J 0-surgeries on M' .

We claim that M' is a weak graph manifold. To see this, note that the union W of the D^2 -bundle over R' and the $2J$ 3-disks is a disjoint union of solid tori in M' ; see Figure 15. The metric completion of $M' - W$ inherits a weak graph structure from M . This shows that M' is a weak graph manifold.

Step 2: For each component P of B' that is diffeomorphic to D^2 or $D^2(r)$, the corresponding component of M' is the result of gluing two solid tori: one being $\pi^{-1}(P)$ and the other one being a connected component of W . By Lemma 11.4, this component of M' is Seifert-fibered and hence is a strong graph manifold. We discard all such components of M' and let \widehat{M} denote what's left.

A component P of B' diffeomorphic to $S^1 \times I$ has a boundary consisting of two circles C_1 and C_2 , of which exactly one, say C_1 , does not intersect R . In \widehat{M} , the preimage $\pi^{-1}(C_1)$ is attached to the union of $\pi^{-1}(P)$ with a solid torus. This union is itself a solid torus.

In this way, we see that \widehat{M} has a weak graph decomposition with $(n - 1)$ tori, since T_k has disappeared. Since M was a counterexample to Proposition 11.6, it follows that \widehat{M} is also a counterexample. This contradicts the definition of n and so proves the proposition. \square

11.2. Appendix A.2 : Weak graph orbifolds are connected sums of strong graph orbifolds. — In this section we only consider 3-dimensional orbifolds that do not admit embedded bad 2-dimensional suborbifolds.

Definition 11.7. — A *weak graph orbifold* is a compact orientable 3-orbifold \mathcal{O} for which there is a collection $\{E_i\}$ of disjoint embedded orientable Euclidean 2-orbifolds in $\text{int}(\mathcal{O})$ so that after splitting \mathcal{O} open along $\{E_i\}$, the result has connected components that are Seifert-fibered orbifolds (possibly with boundary).

Definition 11.8. — A *strong graph orbifold* is a compact orientable 3-orbifold \mathcal{O} for which there is a collection $\{E_i\}$ of disjoint embedded orientable Euclidean 2-orbifolds in $\text{int}(\mathcal{O})$ such that

1. After splitting \mathcal{O} open along $\{E_i\}$, the result has connected components that are Seifert orbifolds (possibly with boundary).
2. For any E_i , the two circle fibrations on E_i coming from the adjacent Seifert bundles are not isotopic.
3. Each E_i is incompressible in \mathcal{O} .

From Subsection 2.4, each E_i is diffeomorphic to T^2 or $S^2(2, 2, 2, 2)$.

Lemma 11.9. — *Let U and V be two oriented solid-toric 3-orbifolds with diffeomorphic boundaries. Let $\phi : \partial U \rightarrow \partial V$ be an orientation-reversing diffeomorphism. Then $U \cup_{\phi} V$ admits a Seifert orbifold structure.*

Proof. — Suppose first that ∂U is a 2-torus. Then U is diffeomorphic to $S^1 \times D^2$ or $S^1 \times D^2(k)$. The Seifert orbifold structures on U are in one-to-one correspondence with the Seifert manifold structures on $|U|$ [7, p. 36-37]. There is one circle fibering of ∂U (up to isotopy) whose fibers bound compressing discal 2-orbifolds in U . Any other circle fibering of ∂U is the boundary fibration of a Seifert fibration of U . As in the proof of Lemma 11.4, we can choose a circle fibering \mathcal{F} of ∂U so that \mathcal{F} is the boundary fibration of a Seifert fibration of U , and $\phi_*\mathcal{F}$ is the boundary fibration of a Seifert fibration of V . The ensuing Seifert fibrations of U and V join together to give a Seifert fibration of $U \cup_{\phi} V$.

Now suppose that ∂U is diffeomorphic to $S^2(2, 2, 2, 2)$. The orbifiberings of $S^2(2, 2, 2, 2)$ with one-dimensional fiber are the \mathbb{Z}_2 -quotients of \mathbb{Z}_2 -invariant circle fiberings of T^2 . In particular, there is an infinite number of such orbifiberings up to isotopy. (More concretely, given an orbifiberings, there are two disjoint arc fibers connecting pairs of singular points. The complement of the two arcs in $|S^2(2, 2, 2, 2)|$ is an open cylinder with an induced circle fibering. The isotopy class of the orbifiberings is specified by the isotopy class of the two disjoint arcs.)

From [7, p. 38-39], the Seifert fibrations of U are the \mathbb{Z}_2 -quotients of \mathbb{Z}_2 -invariant Seifert fibrations of its solid-toric double cover. It follows that there is one orbifiberings of ∂U (up to isotopy) whose fibers bound compressing discal 2-orbifolds in U . Any other orbifiberings of ∂U is the boundary fibration of a Seifert fibration of U . Hence we can choose an orbifiberings \mathcal{F} of ∂U so that \mathcal{F} is the boundary fibration of a Seifert fibration of U , and $\phi_*\mathcal{F}$ is the boundary fibration of a Seifert fibration of V . The ensuing Seifert fibrations of U and V join together to give a Seifert fibration of $U \cup_{\phi} V$. \square

Proposition 11.10. — *If a connected strong graph orbifold contains an essential embedded spherical 2-orbifold then it is diffeomorphic to a finite isometric quotient of $S^1 \times S^2$.*

Proof. — Suppose that a connected strong graph orbifold \mathcal{O} contains an essential embedded spherical 2-orbifold S .

Lemma 11.11. — *After an isotopy of S , we can assume that $S \cap \bigcup_i E_i$ is a disjoint collection of closed curves in the regular part of S .*

Proof. — If E_i is diffeomorphic to T^2 then a neighborhood of E_i lies in $|\mathcal{O}|_{reg}$ and after isotopy, $S \cap E_i$ is a disjoint collection of closed curves in the regular part of S . Suppose that E_i is diffeomorphic to $S^2(2, 2, 2, 2)$. A neighborhood of E_i is diffeomorphic to $I \times E_i$. Suppose that $p \in S$ is a singular point of E_i . Then the local group of p in S must be \mathbb{Z}_2 . After pushing a neighborhood of $p \in S$ slightly in the I -direction

of $I \times E_i$, we can remove the intersection of S with that particular singular point of E_i . In this way, we can arrange so that S intersects $\bigcup_i E_i$ transversely, with the intersection lying in the regular part of S . \square

We choose S among all such essential embedded spherical 2-orbifolds so that the number of connected components of $|S \cap \bigcup_i E_i|$ is as small as possible.

If $S \cap \bigcup_i E_i = \emptyset$ then S is an essential embedded spherical 2-orbifold in one of the Seifert pieces.

If $S \cap \bigcup_i E_i \neq \emptyset$, let $C \subset |S|$ be an innermost circle in $|S \cap \bigcup_i E_i|$. Then $C \subset |E_k|$ for some k , and $C = \partial D$ for some discal 2-orbifold D embedded in a Seifert component U with $E_k \subset \partial U$. As E_k is incompressible, $C = \partial D'$ for some discal 2-orbifold $D' \subset E_k$. Then $D \cup D'$ is an embedded 2-orbifold with underlying space S^2 and at most two singular points. As \mathcal{O} has no bad 2-suborbifolds, $D \cup D'$ must be diffeomorphic to $S^2(r, r)$ for some $r \geq 1$. If $D \cup D'$ bounds some $D^3(r, r)$ in U then we can isotope S to remove the intersection with E_k , which contradicts the choice of S . Thus $D \cup D'$ is an essential embedded spherical 2-orbifold in U .

In any case, we found an essential embedded spherical 2-orbifold in one of the Seifert pieces. Then the universal cover of the Seifert piece contains an essential embedded S^2 . It follows that the universal cover of the Seifert piece is $\mathbb{R} \times S^2$ [5, Proposition 2.13]. The Seifert piece, and hence all of \mathcal{O} , must then be diffeomorphic to a finite isometric quotient of $S^1 \times S^2$. \square

Proposition 11.12. — *A weak graph orbifold is the result of performing 0-surgeries on a strong graph orbifold.*

Proof. — Suppose that Proposition 11.12 fails. Let n be the minimal number of decomposing Euclidean 2-orbifolds among weak graph orbifolds which are counterexamples, and let \mathcal{O} be a counterexample with decomposing Euclidean 2-orbifolds $\{E_i\}_{i=1}^n$.

We first look for a 2-orbifold E_j for which the two induced circle fibrations (coming from the adjacent Seifert bundles) are isotopic, in the sense of [5, Chapter 2.5]. If there is one then we extend the Seifert fibration over E_j . In this case, by removing E_j from $\{E_i\}$, we get a weak graph decomposition of \mathcal{O} with $(n-1)$ Euclidean 2-orbifolds, contradicting the definition of n .

Therefore there is no such Euclidean 2-orbifold. Since \mathcal{O} is a counterexample to Proposition 11.12, there must be a Euclidean 2-orbifold in $\{E_i\}$ which is compressible. Let D be a compressing discal 2-orbifold. As in Lemma 11.11, we can assume that D intersects $\bigcup_i E_i$ transversally, with the intersection lying in the regular part of D . We choose such a compressing discal 2-orbifold so that $D \cap \bigcup_i E_i$ has the smallest possible number of connected components. Let C be an innermost circle in $D \cap \bigcup_i E_i$, say lying in $|E_k|$. Then C bounds a discal 2-orbifold D' lying in a Seifert bundle V which has E_k as a boundary component.

If C also bounds a discal 2-orbifold $D'' \subset E_k$ then $D' \cup D''$ is an embedded 2-orbifold S in the Seifert bundle. As there are no bad 2-orbifolds in \mathcal{O} , the suborbifold

S must be diffeomorphic to $S^2(r, r)$ for some $r \geq 1$. If S is not essential in V then it bounds a $D^3(r, r)$ in V and we can isotope D so that it does not intersect E_k , which contradicts the choice of D . So S is essential in V . From Proposition 11.10, the Seifert bundle V is diffeomorphic to a finite isometric quotient of $S^1 \times S^2$, which contradicts the assumption that it has E_k as a boundary component.

Thus we can assume that C bounds a compressing discal 2-orbifold for V , which is necessarily a solid-toric orbifold diffeomorphic to $S^1 \times D^2(r)$ or $S^1 \times_{\mathbb{Z}_2} D^2(r)$ for some $r \geq 1$ [19, Lemma 2.47].

Let U be the Seifert bundle on the other side of E_k from V . Let B be the orbifold base of U , with projection $\pi : U \rightarrow B$. There is a 1-orbifold boundary component $R \subset \partial B$, diffeomorphic to S^1 or $S^1 // \mathbb{Z}_2$, so that $E_k = \pi^{-1}(R)$. That is, V is glued to U along $\pi^{-1}(R)$. Choose a discal orbifibration $\sigma : V \rightarrow R$ that extends $\pi : E_k \rightarrow R$.

We refer to [5, Chapter 2.5] for a discussion of Dehn fillings, *i.e.*, gluings of V to $\pi^{-1}(R)$. If the meridian curve of V is not isotopic to a fiber of $\pi|_{E_k}$, let $u > 0$ be the algebraic intersection number (computed using the maximal abelian subgroup of $\pi_1(E_k)$). Then the gluing of V to U , along $\pi^{-1}(R)$, has a Seifert fibration. Removing E_k from $\{E_i\}$, we again have a weak graph orbifold decomposition of \mathcal{O} , now with $(n - 1)$ Euclidean 2-orbifolds, which is a contradiction.

Therefore, the meridian curve of V is isotopic to a fiber of $\pi|_{E_k}$.

Step 1: If one of the following possibilities holds then put $\mathcal{O}' = \mathcal{O}$ and $B' = B$, and go to Step 2:

1. $B = D^2$.
2. $B = D^2(s)$ for some $s > 1$.
3. $B = D^2 // \mathbb{Z}_2$.
4. $B = D^2(s) // \mathbb{Z}_2$ for some $s > 1$.
5. $B = S^1 \times I$.
6. $B = (S^1 // \mathbb{Z}_2) \times I$.

Otherwise, we split B open along a disjoint collection of smooth embedded arcs $\{\gamma_j\}_{j=1}^J \cup \{\gamma'_{j'}\}_{j'=1}^{J'}$ of the following type. A curve $\gamma_j : [0, 1] \rightarrow B$ lies in B_{reg} and has $|\gamma_j|(0), |\gamma_j|(1) \in \text{int}(|R|)$. A curve $\gamma'_{j'} : [0, 1] \rightarrow B$ has $|\gamma'_{j'}|(0) \in \text{int}(|R|)$ and lies in B_{reg} , except for its endpoint $|\gamma'_{j'}|(1)$ which is in the interior of a reflector component of $\partial|B|$ but is not a corner reflector point. We can find a collection of such curves so that if B' is the result of splitting B open along them, then each connected component of B' is of type (1)-(6) above. Put

$$(11.13) \quad R' = R - \bigcup_{j=1}^J \{|\gamma_j|(0), |\gamma_j|(1)\} - \bigcup_{j'=1}^{J'} \{|\gamma'_{j'}|(0)\}.$$

Associated to γ_j is a spherical 2-orbifold X_j , diffeomorphic to $S^2(r, r)$, given by

$$(11.14) \quad X_j = \sigma^{-1}(\gamma_j(0)) \cup_{\pi^{-1}(\gamma_j(0))} \pi^{-1}(\gamma_j) \cup_{\pi^{-1}(\gamma_j(1))} \sigma^{-1}(\gamma_j(1)).$$

Associated to $\gamma'_{j'}$ is a spherical 2-orbifold $X'_{j'}$, diffeomorphic to $S^2(2, 2, r)$, given by

$$(11.15) \quad X'_{j'} = \sigma^{-1}(\gamma'_{j'}(0)) \cup_{\pi^{-1}(\gamma'_{j'}(0))} \pi^{-1}(\gamma'_{j'}).$$

Let Y be the result of splitting \mathcal{O} open along $\{X_j\}_{j=1}^J \cup \{X'_{j'}\}_{j'=1}^{J'}$. It has $2(J + J')$ spherical boundary components corresponding to the spherical cuts. We glue on $2J$ copies of $D^3(r, r)$ and $2J'$ copies of $D^3(2, 2, r)$, to obtain \mathcal{O}' . By construction, \mathcal{O} is the result of performing 0-surgeries on \mathcal{O}' .

We claim that \mathcal{O}' is a weak graph orbifold. To see this, note that the union W of $\sigma^{-1}(R')$ and the $2(J + J')$ discal 3-orbifolds is a disjoint union of solid-toric 3-orbifolds in \mathcal{O}' . The metric completion of $|\mathcal{O}'| - |W|$ in $|\mathcal{O}'|$ inherits a weak graph orbifold structure from \mathcal{O} . This shows that \mathcal{O}' is a weak graph orbifold.

Step 2: For each connected component of B' of type (1)-(4) above, the corresponding component of \mathcal{O}' is the result of gluing two solid-toric orbifolds: one being the Seifert orbifold over that component of B' , and the other one being a connected component of W . By Lemma 11.9, this component of \mathcal{O}' is Seifert-fibered and hence is a strong graph orbifold. We discard all such components of \mathcal{O}' and let $\widehat{\mathcal{O}}$ denote what's left.

Turning to the remaining possibilities, an annular component P of B' has a boundary consisting of two circles C_1 and C_2 , of which exactly one, say C_1 , does not intersect R . In $\widehat{\mathcal{O}}$, the preimage $\pi^{-1}(C_1)$ is attached to the union of $\pi^{-1}(P)$ with a solid-toric orbifold diffeomorphic to $S^1 \times D^2(r)$. This union is itself diffeomorphic to $S^1 \times D^2(r)$, since $\pi^{-1}(P)$ is diffeomorphic to $S^1 \times S^1 \times I$.

Finally, if a component P of B' is diffeomorphic to $(S^1 // \mathbb{Z}_2) \times I$ then $\partial|P|$ consists of a circle with two reflector components and two nonreflector components. Exactly one of the nonreflector components, say C_1 , does not intersect R . In $\widehat{\mathcal{O}}$, the preimage $\pi^{-1}(C_1)$ is attached to the union of $\pi^{-1}(P)$ with a solid-toric orbifold diffeomorphic to $S^1 \times_{\mathbb{Z}_2} D^2(r)$. This union is itself diffeomorphic to $S^1 \times_{\mathbb{Z}_2} D^2(r)$, since $\pi^{-1}(P)$ is diffeomorphic to $(S^1 \times_{\mathbb{Z}_2} S^1) \times I$.

In this way, we see that $\widehat{\mathcal{O}}$ has a weak graph orbifold decomposition with $(n - 1)$ Euclidean 2-orbifolds, since E_k has disappeared. Since \mathcal{O} was a counterexample to Proposition 11.12, it follows that $\widehat{\mathcal{O}}$ is also a counterexample. This contradicts the definition of n and so proves the proposition. \square

11.3. Appendix A.3 : Weak graph orbifolds with a compressible boundary component

Lemma 11.16. — *Suppose that \mathcal{O} is a weak graph orbifold, and that $C \subset \partial\mathcal{O}$ is a compressible boundary component. Then \mathcal{O} arises from 0-surgery on a disjoint collection $\mathcal{O}_0 \sqcup \dots \sqcup \mathcal{O}_n$, where:*

- \mathcal{O}_i is a strong graph manifold for all i .
- $\partial\mathcal{O}_0 = C$.
- \mathcal{O}_0 is a solid-toric 3-orbifold.

Proof. — Let Z be a compressing disc orbifold for C .

By Proposition 11.12 we know that \mathcal{O} comes from 0-surgery on a collection $\mathcal{O}_0, \dots, \mathcal{O}_n$ of strong graph orbifolds, where $\partial\mathcal{O}_0$ contains C . Consider a collection $\mathcal{S} = \{S_1, \dots, S_k\} \subset \mathcal{O}$ of spherical 2-suborbifolds associated with such a 0-surgery description of \mathcal{O} . We may assume that Z is transverse to \mathcal{S} , and that the number of connected components in the intersection $Z \cap \mathcal{S}$ is minimal among such compressing disc orbifolds. Reasoning as in the proof of Lemma 11.11, we conclude that Z is disjoint from \mathcal{S} . Therefore after splitting \mathcal{O} open along \mathcal{S} and filling in the boundary components to undo the 0-surgeries, we get that Z lies in \mathcal{O}_0 . Similar reasoning shows that Z must lie in a single Seifert component U of \mathcal{O}_0 . An orientable Seifert 3-orbifold with a compressible boundary component must be a solid-toric 3-orbifold [19, Lemma 2.47]. The lemma follows. \square

References

- [1] A. ADEM, J. LEIDA & Y. RUAN – *Orbifolds and stringy topology*, Cambridge Tracts in Math. vol. 171, Cambridge Univ. Press, Cambridge, 2007.
- [2] M. T. ANDERSON – “Scalar curvature and the existence of geometric structures on 3-manifolds. I”, *J. Reine Angew. Math.* **553** (2002), p. 125–182.
- [3] A. L. BESSE – *Einstein manifolds*, *Ergeb. Math. Grenzgeb. (3)*, vol. 10, Springer-Verlag, Berlin, 1987.
- [4] M. BOILEAU, B. LEEB & J. PORTI – “Geometrization of 3-dimensional orbifolds”, *Ann. of Math. (2)* **162** (2005), no. 1, p. 195–290.
- [5] M. BOILEAU, S. MAILLOT & J. PORTI – *Three-dimensional orbifolds and their geometric structures*, Panoramas & Synthèses, vol. 15, Soc. Math. France, Paris, 2003.
- [6] M. BOILEAU & J. PORTI – “Geometrization of 3-orbifolds of cyclic type”, *Astérisque* (2001), no. 272, p. 208.
- [7] F. BONAHOON & L. SIEBENMANN – “The classification of Seifert fibred 3-orbifolds”, in *Low-dimensional topology (Chelwood Gate, 1982)*, London Math. Soc. Lecture Note Ser., vol. 95, Cambridge Univ. Press, Cambridge, 1985, p. 19–85.
- [8] J. E. BORZELLINO – “Riemannian geometry of orbifolds”, Ph.D. Thesis, University of California, Los Angeles, 1992, http://www.calpoly.edu/~jborzell/Publications/Publication%20PDFs/phd_thesis.pdf.
- [9] ———, “Orbifolds of maximal diameter”, *Indiana Univ. Math. J.* **42** (1993), no. 1, p. 37–53.
- [10] J. E. BORZELLINO & S.-H. ZHU – “The splitting theorem for orbifolds”, *Illinois J. Math.* **38** (1994), no. 4, p. 679–691.
- [11] G. E. BREDON – *Introduction to compact transformation groups*, Pure Applied Math., vol. 46, Academic Press, New York-London, 1972.
- [12] M. R. BRIDSON & A. HAEFLIGER – *Metric spaces of non-positive curvature*, Grundlehren Math. Wiss., vol. 319, Springer-Verlag, Berlin, 1999.
- [13] D. BURAGO, Y. BURAGO & S. IVANOV – *A course in metric geometry*, Grad. Stud. Math., vol. 33, Amer. Math. Soc., Providence, RI, 2001.

- [14] J. CHEEGER – “Critical points of distance functions and applications to geometry”, in *Geometric topology: recent developments (Montecatini Terme, 1990)*, Lecture Notes in Math., vol. 1504, Springer, Berlin, 1991, p. 1–38.
- [15] J. CHEEGER, K. FUKAYA & M. GROMOV – “Nilpotent structures and invariant metrics on collapsed manifolds”, *J. Amer. Math. Soc.* **5** (1992), no. 2, p. 327–372.
- [16] J. CHEEGER & D. GROMOLL – “The splitting theorem for manifolds of nonnegative Ricci curvature”, *J. Differential Geom.* **6** (1971/72), p. 119–128.
- [17] ———, “On the structure of complete manifolds of nonnegative curvature”, *Ann. of Math. (2)* **96** (1972), p. 413–443.
- [18] B. CHOW, S.-C. CHU, D. GLICKENSTEIN, C. GUENTHER, J. ISENBERG, T. IVEY, D. KNOPF, P. LU, F. LUO & L. NI – *The Ricci flow: techniques and applications. Part I: Geometric aspects*, Math. Surveys Monogr., vol. 135, Amer. Math. Soc., Providence, RI, 2007.
- [19] D. COOPER, C. D. HODGSON & S. P. KERCKHOFF – *Three-dimensional orbifolds and cone-manifolds*, MSJ Mem., vol. 5, Math. Soc. Japan, Tokyo, 2000.
- [20] D. M. DETURCK – “Deforming metrics in the direction of their Ricci tensors”, *J. Differential Geom.* **18** (1983), no. 1, p. 157–162.
- [21] J. DINKELBACH & B. LEEB – “Equivariant Ricci flow with surgery and applications to finite group actions on geometric 3-manifolds”, *Geom. Topol.* **13** (2009), no. 2, p. 1129–1173.
- [22] W. D. DUNBAR – “Geometric orbifolds”, *Rev. Mat. Univ. Complut. Madrid* **1** (1988), no. 1-3, p. 67–99.
- [23] ———, “Nonfibering spherical 3-orbifolds”, *Trans. Amer. Math. Soc.* **341** (1994), no. 1, p. 121–142.
- [24] D. FAESSLER – “On the Topology of Locally Volume Collapsed Riemannian 3-Orbifolds”, preprint, <http://arxiv.org/abs/1101.3644>, 2011.
- [25] B. FARB & D. MARGALIT – *A primer on mapping class groups*, Princeton Math. Ser., vol. 49, Princeton Univ. Press, Princeton, NJ, 2012.
- [26] K. FUKAYA – “A boundary of the set of the Riemannian manifolds with bounded curvatures and diameters”, *J. Differential Geom.* **28** (1988), no. 1, p. 1–21.
- [27] D. GROMOLL & G. WALSHAP – *Metric foliations and curvature*, Progr. Math., vol. 268, Birkhäuser Verlag, Basel, 2009.
- [28] K. GROVE & K. SHIOHAMA – “A generalized sphere theorem”, *Ann. of Math. (2)* **106** (1977), no. 2, p. 201–211.
- [29] A. HAEFLIGER – “Orbi-espaces”, in *Sur les groupes hyperboliques d’après Mikhael Gromov* (E. Ghys & P. de la Harpe, eds.), Progr. Math., vol. 83, Birkhäuser Boston, Inc., Boston, MA, 1990, p. 203–213.
- [30] R. S. HAMILTON – “Three-manifolds with positive Ricci curvature”, *J. Differential Geom.* **17** (1982), no. 2, p. 255–306.
- [31] ———, “Four-manifolds with positive curvature operator”, *J. Differential Geom.* **24** (1986), no. 2, p. 153–179.
- [32] ———, “A compactness property for solutions of the Ricci flow”, *Amer. J. Math.* **117** (1995), no. 3, p. 545–572.
- [33] ———, “Non-singular solutions of the Ricci flow on three-manifolds”, *Comm. Anal. Geom.* **7** (1999), no. 4, p. 695–729.

- [34] ———, “Three-orbifolds with positive Ricci curvature”, in *Collected papers on Ricci flow*, Ser. Geom. Topol., vol. 37, Int. Press, Somerville, MA, 2003, p. 521–524.
- [35] C. HOG-ANGELONI & S. MATVEEV – “Roots in 3-manifold topology”, in *The Zieschang Gedenkschrift*, Geom. Topol. Monogr., vol. 14, Geom. Topol. Publ., Coventry, 2008, p. 295–319.
- [36] M. KAPOVICH – *Hyperbolic manifolds and discrete groups*, Progress in Mathematics, vol. 183, Birkhäuser Boston, Inc., Boston, MA, 2001.
- [37] B. KLEINER & J. LOTT – “Locally collapsed 3-manifolds”, in this volume.
- [38] ———, “Notes on Perelman’s papers”, *Geom. Topol.* **12** (2008), no. 5, p. 2587–2855.
- [39] S. KOBAYASHI & K. NOMIZU – *Foundations of differential geometry. Vol I*, Interscience Publishers, New York, 1963.
- [40] P. LU – “A compactness property for solutions of the Ricci flow on orbifolds”, *Amer. J. Math.* **123** (2001), no. 6, p. 1103–1134.
- [41] C. MCMULLEN – “Iteration on Teichmüller space”, *Invent. Math.* **99** (1990), no. 2, p. 425–454.
- [42] W. H. MEEKS, III & S. T. YAU – “The classical Plateau problem and the topology of three-dimensional manifolds. The embedding of the solution given by Douglas-Morrey and an analytic proof of Dehn’s lemma”, *Topology* **21** (1982), no. 4, p. 409–442.
- [43] ———, “The existence of embedded minimal surfaces and the problem of uniqueness”, *Math. Z.* **179** (1982), no. 2, p. 151–168.
- [44] I. MOERDIJK – “Orbifolds as groupoids: an introduction”, in *Orbifolds in mathematics and physics (Madison, WI, 2001)*, Contemp. Math., vol. 310, Amer. Math. Soc., Providence, RI, 2002, p. 205–222.
- [45] J. W. MORGAN – “Recent progress on the Poincaré conjecture and the classification of 3-manifolds”, *Bull. Amer. Math. Soc. (N.S.)* **42** (2005), no. 1, p. 57–78 (electronic).
- [46] L. NI & N. WALLACH – “On a classification of gradient shrinking solitons”, *Math. Res. Lett.* **15** (2008), no. 5, p. 941–955.
- [47] J.-P. OTAL – “Le théorème d’hyperbolisation pour les variétés fibrées de dimension 3”, *Astérisque* (1996), no. 235, p. x+159.
- [48] ———, “Thurston’s hyperbolization of Haken manifolds”, in *Surveys in differential geometry. Vol. III (Cambridge, MA, 1996)*, Int. Press, Boston, MA, 1998, p. 77–194.
- [49] G. PERELMAN – “The Entropy Formula for the Ricci Flow and its Geometric Applications”, preprint, <http://arXiv.org/abs/math.DG/0211159>, 2002.
- [50] ———, “Ricci Flow with Surgery on Three-Manifolds”, preprint, <http://arxiv.org/abs/math.DG/0303109>, 2003.
- [51] J. PETEAN & G. YUN – “Surgery and the Yamabe invariant”, *Geom. Funct. Anal.* **9** (1999), no. 6, p. 1189–1199.
- [52] P. PETERSEN & W. WYLIE – “On the classification of gradient Ricci solitons”, *Geom. Topol.* **14** (2010), no. 4, p. 2277–2300.
- [53] C. PETRONIO – “Spherical splitting of 3-orbifolds”, *Math. Proc. Cambridge Philos. Soc.* **142** (2007), no. 2, p. 269–287.
- [54] P. SCOTT – “The geometries of 3-manifolds”, *Bull. London Math. Soc.* **15** (1983), no. 5, p. 401–487.

- [55] M. E. TAYLOR – *Partial differential equations. III*, Appl. Math. Sciences, vol. 117, Springer-Verlag, New York, 1997.
- [56] W. P. THURSTON – “Three-dimensional manifolds, Kleinian groups and hyperbolic geometry”, *Bull. Amer. Math. Soc. (N.S.)* **6** (1982), no. 3, p. 357–381.
- [57] ———, “Three-Manifolds with Symmetry”, preprint, 1982.
- [58] L.-F. WU – “The Ricci flow on 2-orbifolds with positive curvature”, *J. Differential Geom.* **33** (1991), no. 2, p. 575–596.
- [59] R. YE – “Notes on 2-Dimensional κ -Solutions”, <http://www.math.ucsb.edu/~yer/kappa-solutions.pdf>.

BRUCE KLEINER, Courant Institute of Mathematical Sciences, 251 Mercer St., New York, NY 10012
E-mail : bkleiner@cims.nyu.edu

JOHN LOTT, Department of Mathematics, University of California at Berkeley, Berkeley, CA 94720
E-mail : lott@math.berkeley.edu

ASTÉRISQUE

2014

361. SÉMINAIRE BOURBAKI, volume 2012/2013, exposés 1059–1073
360. J.I. BURGOS GIL, P. PHILIPPON, M. SOMBRA – *Géométrie arithmétique des variétés toriques. Métriques, mesures et hauteurs*
359. M. BROUÉ, G. MALLE, J. MICHEL – *Split Spetses for primitive reflection group*

2013

358. A. GETMANENKO, D. TAMARKIN – *Microlocal Properties of Sheaves and Complex WKB*
357. A. AVILA, J. SANTAMARIA, M. VIANA, A. WILKINSON – *Cocycles over partially hyperbolic maps*
356. D. SCHÄPPI – *The Formal Theory of Tannaka Duality*
355. J.-P. RAMIS, J. SAULOY, C. ZHANG – *Local analytic classification of q -difference equations*
354. S. CROVISIER – *Perturbation de la dynamique de difféomorphismes en topologie C^1*
353. N.-G. KANG, N. G. MAKAROV – *Gaussian free field and conformal field theory*
352. SÉMINAIRE BOURBAKI, volume 2011/2012, exposés 1043–1058
351. R. MELROSE, A. VASY, J. WUNSCH – *Diffraction of Singularities for the Wave Equation on Manifolds with Corners*
350. F. LE ROUX – *L'ensemble de rotation autour d'un point fixe*
349. J. T. COX, R. DURRETT, E. A. PERKINS – *Voter Model Perturbations and Reaction Diffusion Equations*

2012

348. SÉMINAIRE BOURBAKI, volume 2010/2011, exposés 1027–1042
347. C. MOEGLIN, J.-L. WALDSPURGER – *Sur les conjectures de Gross et Prasad (II)*
346. W. T. GAN, B. H. GROSS, D. PRASAD, J.-L. WALDSPURGER – *Sur les conjectures de Gross et Prasad (I)*
345. M. KASHIWARA, P. SCHAPIRA – *Deformation quantization modules*
344. M. MITREA, M. WRIGHT – *Boundary value problems for the Stokes system in arbitrary Lipschitz domains*
343. K. BEHREND, G. GINOT, B. NOOHI, P. XU – *String topology for stacks*
342. H. BAHOURI, C. FERMANIAN-KAMMERER, I. GALLAGHER – *Phase-space analysis and pseudodifferential calculus on the Heisenberg group*
341. J.-M. DELORT – *A quasi-linear Birkhoff normal forms method. Application to the quasi-linear Klein-Gordon equation on S^1*

2011

340. T. MOCHIZUKI – *Wild harmonic bundles and wild pure twistor D -modules*
339. SÉMINAIRE BOURBAKI, volume 2009/2010, exposés 1012–1026
338. G. ARONE, M. CHING – *Operads and chain rules for the calculus of functors*
337. U. BUNKE, T. SCHICK, M. SPITZWECK – *Periodic twisted cohomology and T -duality*
336. P. GYRYA, L. SALOFF-COSTE – *Neumann and Dirichlet Heat Kernels in Inner Uniform Domains*
335. P. PELAEZ – *Multiplicative Properties of the Slice Filtration*

2010

334. J. POINEAU – *La droite de Berkovich sur \mathbf{Z}*
333. K. PONTO – *Fixed point theory and trace for bicategories*
332. SÉMINAIRE BOURBAKI, volume 2008/2009, exposés 997–1011
331. Représentations p -adiques de groupes p -adiques III : méthodes globales et géométriques, L. BERGER, C. BREUIL, P. COLMEZ, éditeurs
330. Représentations p -adiques de groupes p -adiques II : représentations de $\mathrm{GL}_2(\mathbf{Q}_p)$ et (φ, Γ) -modules, L. BERGER, C. BREUIL, P. COLMEZ, éditeurs
329. T. LÉVY – *Two-dimensional Markovian holonomy fields*

2009

- 328. From probability to geometry (II), Volume in honor of the 60th birthday of Jean-Michel Bismut, X. DAI, R. LÉANDRE, X. MA, W. ZHANG, editors
- 327. From probability to geometry (I), Volume in honor of the 60th birthday of Jean-Michel Bismut, X. DAI, R. LÉANDRE, X. MA, W. ZHANG, editors
- 326. SÉMINAIRE BOURBAKI, volume 2007/2008, exposés 982–996
- 325. P. HAÏSSINSKY, K. M. PILGRIM – *Coarse expanding conformal dynamics*
- 324. J. BELLAÏCHE, G. CHENEVIER – *Families of Galois representations and Selmer groups*
- 323. Équations différentielles et singularités en l'honneur de J. M. Aroca, F. CANO, F. LORAY, J. J. MORALES-RUIZ, P. SAD, M. SPIVAKOVSKY, éditeurs

2008

- 322. Géométrie différentielle, Physique mathématique, Mathématiques et société (II). Volume en l'honneur de Jean Pierre Bourguignon, O. HIJAZI, éditeur
- 321. Géométrie différentielle, Physique mathématique, Mathématiques et société (I). Volume en l'honneur de Jean Pierre Bourguignon, O. HIJAZI, éditeur
- 320. J.-L. LODAY – *Generalized bialgebras and triples of operads*
- 319. Représentations p -adiques de groupes p -adiques I : représentations galoisiennes et (φ, Γ) -modules, L. BERGER, C. BREUIL, P. COLMEZ, éditeurs
- 318. X. MA, W. ZHANG – *Bergman kernels and symplectic reduction*
- 317. SÉMINAIRE BOURBAKI, volume 2006/2007, exposés 967-981

2007

- 316. M. C. OLSSON – *Crystalline cohomology of algebraic stacks and Hyodo-Kato cohomology*
- 315. J. AYOUB – *Les six opérations de Grothendieck et le formalisme des cycles évanescents dans le monde motivique (II)*
- 314. J. AYOUB – *Les six opérations de Grothendieck et le formalisme des cycles évanescents dans le monde motivique (I)*
- 313. T. NGO DAC – *Compactification des champs de chtoucas et théorie géométrique des invariants*
- 312. ARGOS seminar on intersections of modular correspondences
- 311. SÉMINAIRE BOURBAKI, volume 2005/2006, exposés 952–966

2006

- 310. J. NEKOVÁŘ – *Selmer Complexes*
- 309. T. MOCHIZUKI – *Kobayashi-Hitchin correspondence for tame harmonic bundles and an application*
- 308. D.-C. CISINSKI – *Les préfaisceaux comme modèles des types d'homotopie*
- 307. SÉMINAIRE BOURBAKI, volume 2004/2005, exposés 938–951
- 306. C. BONNAFÉ – *Sur les caractères des groupes réductifs finis à centre non connexe : applications aux groupes spéciaux linéaires et unitaires*
- 305. M. JUNGE, C. LE MERDY, Q. XU – *H^∞ functional calculus and square functions on noncommutative L^p -spaces*

Instructions aux auteurs

Revue internationale de haut niveau, *Astérisque* publie en français et en anglais des monographies de qualité, des séminaires prestigieux, ou des comptes-rendus de grands colloques internationaux. Les textes sont choisis pour leur contenu original ou pour la nouveauté de la présentation qu'ils donnent d'un domaine de recherche. Chaque volume est consacré à un seul sujet, et tout le spectre des mathématiques est en principe couvert.

Le manuscrit doit être envoyé en *double* exemplaire au secrétariat des publications en précisant le nom de la revue.

Le fichier *source* T_EX (un seul fichier par article ou monographie) peut aussi être envoyé par courrier électronique ou ftp, *sous réserve* que sa compilation par le secrétariat SMF soit possible. Contacter le secrétariat à l'adresse électronique revues@smf.ens.fr pour obtenir des précisions.

La SMF recommande *vivement* l'utilisation d'*AMS-L^AT_EX* avec les classes *smfart* ou *smfbook*,

disponibles ainsi que leur documentation sur le serveur <http://smf.emath.fr/> ou sur demande au secrétariat des publications SMF.

Les fichiers *AMS-L^AT_EX* au format *amsart* ou *amsbook*, ainsi que les fichiers *L^AT_EX* au format *article* ou *book* sont aussi les bienvenus. Ils seront saisis suivant les normes suivantes:

- taille des caractères égale à 10 points (option 10pt);
- largeur du texte (`textwidth`) de 13 cm;
- hauteur du texte (`textheight`) égale à 21.5 cm;
- le texte étant en outre centré sur une feuille A4 (option `a4paper`).

Les autres formats T_EX et les autres types de traitement de texte ne sont pas utilisables par le secrétariat et sont *fortement* déconseillés.

Avant de saisir leur texte, les auteurs sont invités à prendre connaissance du document *Recommandations aux auteurs* disponible au secrétariat des publications de la SMF ou sur le serveur de la SMF.

Instructions to Authors

Astérisque is a high level international journal which publishes excellent research monographs in French or in English, and proceedings of prestigious seminars or of outstanding international meetings. The texts are selected for the originality of their contents or the new presentation they give of some area of research. Each volume is devoted to a single topic, chosen, in principle, from the whole spectrum of mathematics.

Two copies of the original manuscript should be sent to the editorial board of the SMF, indicating to which publication the paper is being submitted.

The T_EX *source* file (a single file for each article or monograph) may also be sent by electronic mail or ftp, in a format suitable for typesetting by the secretary. Please, send an email to revues@smf.ens.fr for precise information.

The SMF has a *strong* preference for *AMS-L^AT_EX* together with the documentclasses *smfart*

or *smfbook*, available with their user's guide at <http://smf.emath.fr/> (Internet) or on request from the editorial board of the SMF.

The *AMS-L^AT_EX* files using the documentclasses *amsart* or *amsbook*, or the *L^AT_EX* files using the documentclasses *article* or *book* are also encouraged. They will be prepared following the rules below:

- font size equal to 10 points (10pt option);
- text width (`textwidth`): 13 cm;
- text height (`textheight`): 21.5 cm;
- the text being centered on a A4 page (`a4paper` option).

Files prepared with other T_EX dialects or other word processors cannot be used by the editorial board and are *not* encouraged.

Before preparing their electronic manuscript, the authors should read the *Advice to authors*, available on request from the editorial board of the SMF or from the web site of the SMF.

This volume has two papers, which can be read separately. The first paper concerns local collapsing in Riemannian geometry. We prove that a three-dimensional compact Riemannian manifold which is locally collapsed, with respect to a lower curvature bound, is a graph manifold. This theorem was stated by Perelman without proof and was used in his proof of the geometrization conjecture. The second paper is about the geometrization of orbifolds. A three-dimensional closed orientable orbifold, which has no bad suborbifolds, is known to have a geometric decomposition from work of Perelman in the manifold case, along with earlier work of Boileau-Leeb-Porti, Boileau-Maillot-Porti, Boileau-Porti, Cooper-Hodgson-Kerckhoff and Thurston. We give a new, logically independent, unified proof of the geometrization of orbifolds, using Ricci flow.