The Topology of locally volume collapsed 3-Orbifolds

Daniel Faessler

Dissertation
an der Fakultät für Mathematik, Informatik und Statistik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Daniel Faessler
aus Steinebach (Wörthsee)

München, den 5. April 2011

Erstgutachter: Prof. Bernhard Leeb, Ph.D.

Zweitgutachter: Prof. John Lott, Ph.D.

Tag der mündlichen Prüfung: 30. Juni 2011

Zusammenfassung

In dieser Arbeit untersuchen wir Geometrie und Topologie von Riemannschen 3-Orbifolds, die bezüglich einer Krümmungsskala lokal volumenkollabiert sind. Unser Hauptergebnis ist, dass eine hinreichend kollabierte geschlossene 3-Orbifold ohne schlechte 2-Unterorbifolds Thurstons Geometrisierungsvermutung genügt. Wir beweisen auch eine Version dieses Ergebnisses mit Rand. Kleiner und Lott haben unabhängig und zeitgleich ähnliche Ergebnisse bewiesen ([KL11]).
Hauptschritt unseres Beweises ist die Konstruktion einer Graphenzerlegung von hinreichend kollabierten (geschlossenen) 3-Orbifolds. Wir beschreiben eine grobe Stratifizierung von ungefähr 2-dimensionalen Alexandrov-Räumen, die wir dann für kollabierte 3-Orbifolds zu einer Zerlegung verfeinern; diese Zerlegung kann dann zu einer Graphenzerlegung vereinfacht werden. Wir schließen unseren Beweis ab, indem wir zeigen, dass Graphenorbifolds ohne schlechte 2-Unterorbifolds der Geometrisierungsvermutung genügen.

Abstract

In this thesis we study the geometry and topology of Riemannian 3-orbifolds which are locally volume collapsed with respect to a curvature scale. Our main result is that a sufficiently collapsed closed 3-orbifold without bad 2-suborbifolds satisfies Thurston's Geometrization Conjecture. We also prove a version of this result with boundary. Kleiner and Lott indepedently and simultanously proved similar results ([KL11]).

The main step of our proof is to construct a graph decomposition of sufficiently collapsed (closed) 3-orbifolds. We describe a coarse stratification of roughly 2-dimensional Alexandrov spaces which we then promote to a decomposition into suborbifolds for collapsed 3-orbifolds; this decomposition can then be reduced to a graph decomposition. We complete our proof by showing that graph orbifolds without bad 2-suborbifolds satisfy the Geometrization Conjecture.

Contents

1. Introduction

Like manifolds, orbifolds are defined as topological spaces admitting certain local models. However, whereas manifolds are locally modelled on n-dimensional Euclidean space $\mathbb{R}^n$, orbifolds generalize this definition by admitting as local models all quotients of $\mathbb{R}^n$ by finite groups of diffeomorphisms.

Although I. Satake [Sa56] introduced a very similar definition in 1956 (so-called *V-manifolds*), the term *orbifold* was first introduced in 1976/77 by W. Thurston ([Th78]). Thurston's main interest in orbifolds seems to have been that they can be used in studying the topology and geometry of 3-manifolds; for instance, the basis of a Seifert fibered 3-manifold has a natural structure as a 2-orbifold (cf. [Sc83]). However, orbifolds also naturally occur in other fields of geometrical studies such as knot theory (cf. [BS87]). Beginning with Thurston, orbifolds have also been studied as geometrical objects sui generis. For instance, they can be assigned a *Riemannian metric* and one can investigate the geometric properties of the resulting length space structure (cf. [Bo92]). Orbifolds with lower sectional curvature bounds are Alexandrov spaces.

In 1981, Thurston also extended his Geometrization Conjecture from manifolds to orbifolds. The (original) Geometrization Conjecture for manifolds stated that every closed 3-manifold admitted a (unique) decomposition as follows: In a first step, the manifold was reduced by *spherical* surgery to irreducible (connected sum) components by cutting it open along embedded 2-spheres and glueing 3-balls into the resulting boundary components. The summands would then further be cut up along incompressible *tori* into *geometric* components, i.e. compact 3-orbifolds with toric boundary whose interior admitted a Riemannian metric modelled on one of the eight (homogeneous) model geometries. Thurston conjectured that all closed 3-orbifolds had a similar decomposition into geometric pieces unless they contained *bad* 2-suborbifolds, i.e. closed 2-suborbifolds not globally covered by a manifold. This extended conjecture was motivated by the following observation: For an orbifold which is the quotient of a manifold by a group of diffeomorphisms, a geometrization of the orbifold is equivalent to an invariant geometrization of the covering manifold, equivalently, to a geometrization respecting some symmetry properties of the manifold.

In the years following the formulation of the Geometrization Conjecture, Thurston was able to prove it for Haken manifolds (cf. [Th86], [Ot88]). However, the full conjecture still remained unproved by the beginning of the 21st century. For orbifolds, the Geometrization Conjecture was confirmed for (locally) orientable orbifolds with non-empty singular locus by M. Boileau, B. Leeb and J. Porti ([BLP05]). Note that an orbifold is not necessarily locally orientable: it may locally be modelled on a quotient of $\mathbb{R}^n$ by a group of diffeomorphisms which need not be orientation-preserving.

The complete Geometrization Conjecture was proved in 2003 by G. Perelman's seminal series of papers on the Ricci flow, based on previous work by R. Hamilton. The central idea of the proof is to endow a closed 3-manifold M_0 with an arbitrary Riemannian metric g_0 and let this metric develop with the Ricci flow equation $\dot{g}(t) = -2 \operatorname{Ric} g(t)$. If the initial metric g_0 has positive curvature, the Ricci flow exists for all times and the metrics $g(t)$ converge after rescaling to a spherical limit metric ([Ha82]). For arbitrary initial metrics, the Ricci flow may become singular ([Ha95]). However, Perelman ([Pe03]) showed that in this case it is still possible to construct a *Ricci flow with surgery* for all times by performing spherical surgery whenever singularities occur while at the same time maintaining certain geometric and analytic controls on the Ricci flow. (In particular, the surgery times do not accumulate.) In a Ricci flow with surgery $(M(t), g(t))$, the diffeomorphism type of $M(t)$ may change at surgery times. Following Perelman's work, several detailed treatments of the Ricci flow with surgery have been given ([MT07], [Ba07], [KL08]).

As suggested by the corresponding decomposition of hyperbolic 3-manifolds, Perelman subdivided the Riemannian manifolds $(M(t), g(t))$ into a thick part and a thin part: For small $v > 0$, he defined the thin part $M_-(t)$ of $(M(t), g(t))$ to be the set of all points $x \in M(t)$ with the property that $\operatorname{vol} B_r(x) \leq \rho_{-1}(x)^3 v$ where $\rho_{-1}(x)$ is the smallest real number r such that $\sec \geq -r^2$ on $B_r(x)$ (the so-called curvature scale). The thick part $M_+(t)$ is defined as the closure of the complement of the thin part, $M_+(t) = \overline{M(t) - M_-(t)}$. One obtains the following properties for this thick-thin decomposition of the Ricci flow ([Pe03, Sec. 7.3], cf. also [KL10, Sec. 17]):

For arbitrarily large $s_0 > 0$ there is a function $K : (0, \omega_3) \to (0, \infty)$ (where $\omega_3 = \operatorname{vol} B_1^{\mathbb{R}^3}(0)$) such that for sufficiently small v and large t there following hold: There is a (possibly non-connected) compact hyperbolic 3-manifold with toric boundary H with constant sectional curvature $-\frac{1}{4}$ and an almost-isometry $\phi : H \hookrightarrow M(t)$ such that $M_+(t) \subset \phi(H)$ and that $\phi(\partial H)$ is a family of incompressible tori in $M(t)$. Moreover, on the complement $\tilde{M}(t) = M(t) - \phi(H)$ we have that

(i) for all points $x \in \tilde{M}(t)$ we have $\operatorname{vol}(B_{g(t)}(x, \rho_{-1}(x))) < \rho_{-1}(x)^3 v_0$ (with ρ_{-1} as above),

(ii) for every boundary component C of $\tilde{M}(t)$ there is a hyperbolic manifold cusp X_C (a quotient of a horoball with sectional curvature equal $-\frac{1}{4}$) such that the pairs $(N_{100}(C), C)$

and $(N_{100}(\partial X_C), \partial X_C)$ habe distance $\leq v$ in the $\mathcal{C}^{s_0}$-topology, and

(iii) if $\operatorname{vol} B(x,r) \geq r^3 v'$ for $x \in \tilde{M}(t)$, $v' \in [v, \omega_3)$ and $r \in (0, \rho_{-1}(x)]$, then $\|\nabla^s R\| \leq K(v) r^{-2-s}$ on $B(x,r)$ for $s = 0, \ldots, s_0$.

We sum up these properties by saying that the Riemannian manifold $\tilde{M}(t)$ is v-*collapsed* at the scale ρ_{-1} and has (v, s_0)-*almost cuspidal ends* and (v, s_0, K)-*curvature control* below scale ρ.

Geometrization of the manifold $M(t)$ (and hence of the original manifold M_0) now follows from the following

Theorem 1.0.1 (cf. [Pe03, Thm. 7.4]). *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds: If a 3-manifold (M, g) is closed or compact with (v_0, s_0)-almost cuspidal ends, is v_0-collapsed at the scale ρ_{-1} and has (v_0, s_0, K)-curvature control below the scale ρ_{-1}, then M satisfies Thurston's Geometrization Conjecture.*

Theorem 1.0.1 was stated in [Pe03] without proof. In the meantime several proofs have been published ([SY05], [MT08], [BBBMP10], [KL10]). They are all disjoint from the analytic study of evolution equations; instead, they are essentially geometric in nature.

The aim of this thesis is to generalize Theorem 1.0.1 to orbifolds. More precisely, our main result is

Theorem 1.0.2 (cf. Theorem 7.0.2). *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds: If a 3-orbifold (O, g) is closed or compact with (v_0, s_0)-almost cuspidal ends, is v_0-collapsed at the scale ρ_{-1}, has (v_0, s_0, K)-curvature control below the scale ρ_{-1} and contains no bad 2-suborbifolds, then O is either closed and admits a $\mathcal{C}^5$ Riemannian metric with* $\sec \geq 0$ *or satisfies Thurston's Geometrization Conjecture.*

We expect 3-orbifolds satisfying the conditions of the theorem to arise when generalizing Perelman's construction of the Ricci flow with surgery to orbifolds (which we do not propose to do in this thesis). Similarly, we expect closed 3-orbifolds with nonnegative sectional curvature to be geometric by an orbifold version of Hamilton's corresponding result for 3-manifolds [Ha82]. Thus, we consider Theorem 1.0.2 as the final step in a possible proof of the Geometrization Conjecture for manifolds. It should be noted that if we only consider closed 3-orbifolds O which are *very good*, i.e. quotients of a manifold by a finite group of diffeomorphisms, there is an equivariant version of the Ricci flow with surgery by [DL09].

After finishing a preprint of our main results, we learned that B. Kleiner and J. Lott independently and simultanously proved results similar to our main result, cf. [KL11, Prop. 9.7]. Their method is an extension of their work [KL10] in the manifold case to the orbifold case, whereas our approach is closer to an extension of the approach in [MT08]. This is most noticable in the conical approximation argument we use to obtain a coarse stratification of sufficiently collapsed 3-orbifolds, and in our construction of the (graph) decomposition of the collapsed 3-orbifolds (which is quite different from the construction in [KL11, Sec. 12]). In order to formulate a more concise version of our main result, we refer to one result of [KL11] in this thesis, namely to their generalization of Hamilton's application of the Ricci flow to orbifolds of non-negative curvature (Theorem 6.1.13). The proof of this result is entirely independent of the methods used in this thesis.

Corollary 1.0.3 (cf. Corollary 7.0.3). *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds: If a 3-orbifold (O, g) is closed or compact with (v_0, s_0)-almost cuspidal ends, is v_0-collapsed at the scale ρ_{-1}, has (v_0, s_0, K)-curvature control below the scale ρ_{-1} and contains no bad 2-suborbifolds, then O satisfies Thurston's Geometrization Conjecture.*

Structure of the text.

In section 2 or this thesis, we first review basic facts on orbifolds in low dimensions. We then discuss decompositions of 3-orbifolds along spherical and toric 2-suborbifolds and prove that *graph* orbifolds in the sense of Waldhausen (cf. section 2.2.3) satisfy Thurston's Geometrization Conjecture (Corollary 2.3.3). A standard reference for much of the material covered in this section is contained in [BMP03].

In the third section, we turn to the study of Riemannian orbifolds. As a preliminary to the injectivity radius bound proved in section 4.1, we discuss the geometry of the cut locus of a Riemannian manifold in section 3.3. We also discuss comparison geometry properties of Riemannian orbifolds with lower curvature bounds: Following [Bo92], we discuss an orbifold version for Toponogov's theorem (i.e. we prove that Riemannian orbifolds with lower curvature bounds are Alexandrov spaces). We also discuss a version of the Soul Theorem for orbifolds and its implications on the topology of 3-dimensional orbifolds of non-negative curvature. In this second part of section 3, we use both material from the corresponding results for Riemannian manifolds (cf. [CE75], [Ka89]) and from the much more general study of Alexandrov spaces (cf. [BGP92], [BBI01]).

In section 4, we study limits of Riemannian orbifolds. The main results of this section (Theorems 4.2.7 and 4.2.21) state that after passing to a subsequence, a sequence of complete Riemannian orbifolds with thick base points and uniform bounds on the covariant derivatives of the curvature operator converges to a Riemannian orbifold satisfying the same properties, both in the Gromov-Hausdorff and in the smooth sense. In other words, the space of thick orbifolds with uniform curvature bounds is compact both in the Gromov-

Hausdorff and the smooth topology. This result has already been proved for orbifolds with isolated singularities by P. Lu in [Lu01] and for good orbifolds by K. Fukaya in [Fu86]. It has been stated without proof in [CC07].

We prove our convergence results via a bound on the injectivity radius (Proposition 4.1.4). More precisely, we show that on a thick orbifold with sectional curvature bounds, the injectivity radius can be uniformly bounded below unless one approaches the boundary of a singular stratum (equivalently, a lower-dimensional stratum). For manifolds, this injectivity radius bound is a well-known result (cf. [CGT82]); it is proved in [Lu01] for orbifolds with isolated singularities. Using our injectivity radius bound, we deduce our convergence results by the suitable adaptation of a standard argument (cf. [Ba07], [Lu01], [BLP05]). Part of the results of section 4.1 concerning the local geometry of thick orbifolds have been obtained jointly with Bernhard Leeb ([FL]).

In the fifth section, we turn to the study of collapsed 3-orbifolds by discussing a coarse stratification of roughly 2-dimensional Alexandrov spaces. More precisely, we use a conical approximation argument to show that the points in such a space which do not admit 1-strainers of a certain length and quality accumulate in isolated regions. Outside these regions, the Alexandrov space is 1-strained which allows us to perform a (coarse) dimension reduction by considering cross sections to these strainers. We further distinguish points according to whether they lie in coarse necks, edges or the interior of the Alexandrov space and study their geometric properties. These considerations are similar in spirit to considerations in [MT08] and [KL10].

In section 6, we restrict our attention to closed volume collapsed 3-orbifolds. We consider them as Alexandrov spaces which are roughly of dimension ≤ 2 and promote their coarse stratification to a certain decomposition into 3-suborbifolds. To determine the local topology of the components in this decomposition, we use a variation (and extension to additional situations) of the blow-up arguments in [SY00]. We derive a graph decomposition of the collapsed 3-orbifolds. Combined with the results of section 2, the main result follows for closed orbifolds (Theorem 6.1.12).

We note that we use the condition of (v, s_0, K)-curvature control below scale ρ at only one step of our argument, namely when determining the local topology of the components of our decomposition. Using our compactness results from section 4, we find that a sequence of blow-ups eventually become diffeomorphic to the limit, thus establishing their topology. Probably, this part of the argument vould be replaced by an orbifold version of Perelman's Stability Theorem. We require our curvature control condition precisely to avoid using the Stability Theorem. In the manifold case, there are proofs of Theorem 1.0.1 using the Stability Theorem (which do therefore not require curvature conditions below a certain scale, cf. [KL10, Sec. 18]).

In the last section, we discuss the case with boundary to obtain the more general Theorem 7.0.2. We show how almost cuspidal neighbourhoods of boundary components can be integrated into the decomposition according to the rough stratification from section 6. Since we control the topology (indeed the geometry) of these almost cuspidal neighbourhoods, our main result follows. The results of the last three chapters are contained in the preprint [Fa11].

Acknowledgments.

I would like to thank everybody who helped and supported me during my work on this thesis. In particular, I am most grateful to Professor Bernhard Leeb for calling my attention to 3-orbifolds and their geometrization. His continuing support and the great number of helpful discussions we had have been exceedingly valuable. I would also like to thank Richard Bamler and Hartmut Weiß for their helpful discussions, in particular on the Ricci flow. I am grateful for funding by Studienstiftung des Deutschen Volkes which I recieved during my work on this thesis. Finally, my thanks go out to everybody who supported me personally over the last years, in particular to my wife Valerie.

2. Decomposition of 3-orbifolds along 2-orbifolds

2.1 Smooth orbifolds

2.1.1 Definitions

Intuitively speaking, (smooth) manifolds are spaces which locally look like the Euclidean space $\mathbb{R}^n$. Orbifolds generalize this concept by allowing as local models not only Euclidean space itself, but also its quotients by linear actions of finite groups. The following definition formalizes this idea (cf. [Th78, Sec. 13.2], [BMP03, Sec. 2.1.1]).

Definition 2.1.1 (Smooth orbifolds). An n-dimensional (smooth) orbifold O^n is a metrizable topological space together with a maximal atlas of orbifold charts.

An orbifold chart $(U, \tilde{U}, \Gamma_U, \pi_U)$ consists of an open subset $U \subseteq O$, a connected smooth n-manifold $\tilde{U}$ (e.g. a smooth ball), a finite subgroup $\Gamma_U \subset \mathrm{Diff}(\tilde{U})$, and a continuous map $\pi_U : \tilde{U} \to U$ inducing a homeomorphism $\tilde{U}/\Gamma_U \xrightarrow{\cong} U$.

Any two charts $(U_i, \tilde{U}_i, \Gamma_{U_i}, \pi_{U_i})$, $i = 1, 2$, must be *compatible* in the following sense: If $\tilde{x}_i \in \tilde{U}_i$ are points with $\pi_{U_1}(\tilde{x}_1) = \pi_{U_2}(\tilde{x}_2)$, then there exists a diffeomorphism $\tilde{\phi} : \tilde{V}_1 \to \tilde{V}_2$ of open neighbourhoods $\tilde{V}_i$ of the $\tilde{x}_i$ with $\pi_{U_2} \circ \tilde{\phi} = \pi_{U_1}$.

Finally, the charts must *cover* O.

It follows from the definition that every manifold is in particular an orbifold of the same dimension. All orbifolds are locally compact and locally path connected topological spaces. In particular, they are connected if and only if they are path connected.

Definition 2.1.2 (Orbifold maps). A continuous map $f : O \to O'$ of smooth orbifolds is called *smooth* (in the orbifold sense) if it lifts locally to a smooth map (in the manifold sense) of charts, i.e. if for any point $x \in O$ exist charts $(U, \tilde{U}, \Gamma_U, \pi_U)$ around x and $(U', \tilde{U}', \Gamma_{U'}, \pi_{U'})$ around $x' = f(x)$ and a smooth map $\tilde{f}_U : \tilde{U} \to \tilde{U}'$ with $f \circ \pi_U = \pi_{U'} \circ \tilde{f}_U$.

A smooth map f is an *immersion* (*submersion*, respectively) if it lifts locally to an immersion (submersion, respectively) of charts. A smooth map $f : O^n \to O'^n$ between orbifolds of equal dimensions is a *local diffeomorphism* if it can be inverted locally by a smooth map. It is called a (global) *diffeomorphism* if it is both a local diffeomorphism and a homeomorphism.

A sequence of smooth orbifold maps $f_n : O \to O'$ is said to *converge smoothly* to a smooth orbifold map $f_\infty : O \to O'$, $f_n \overset{\mathcal{C}^\infty}{\to} f_\infty$, if for any point $x \in O$ exist charts $(U, \tilde{U}, \Gamma_U, \pi_U)$ around x and $(U', \tilde{U}', \Gamma_{U'}, \pi_{U'})$ around $x' = f_\infty(x)$, and smooth lifts $\tilde{f}_{n,U}, \tilde{f}_{\infty,U} : \tilde{U} \to \tilde{U}'$ such that $\tilde{f}_{n,U} \overset{\mathcal{C}^\infty}{\to} \tilde{f}_{\infty,U}$.

It is important to note that an orbifold may be homeomorphic, but not diffeomorphic to a manifold. In order to emphasize this distinction, the *underlying space* of an orbifold O, i.e. the orbifold O considered as a topological space without additional structure, is occasionally denoted by $|O|$.

We define smooth n-dimensional orbifolds O *with boundary* by allowing the chart domains $\tilde{U}$ to be smooth n-manifolds with boundary.

Since the local coordinate changes $\tilde{\phi}$ in definition 2.1.1 are smooth, they preserve boundaries. Thus, the set ∂O consisting of those points whose preimages in the chart domains are boundary points is a well defined closed subset of O. We call ∂O the *boundary* and $O \setminus \partial O$ the *interior* of O.

The boundary ∂O of an orbifold inherits from O a structure as a smooth $(n-1)$-dimensional orbifold without boundary. It has an open collar in O, i.e. an open neighbourhood in O diffeomorphic to the product orbifold $\partial O \times [0, 1)$ where $[0, 1)$ is to be understood as a 1-dimensional manifold with boundary.

We define an m-dimensional smooth *suborbifold* of a smooth orbifold to be a subset whose preimages in local charts are smooth m-dimensional submanifolds, cf. [BS87, sec. B]. (This is the more restrictive one of two definitions used in the literature. For the other one, see e.g. [BMP03, 2.1.3].) Analogously, a subset of an orbifold with boundary is called a *proper* suborbifold if its preimages in local charts are proper (smooth) submanifolds.

Orbifolds arise naturally from the following construction (cf. [Th78, Prop. 13.2.1]):

Proposition 2.1.3. *Let M^n be an n-dimensional manifold and $\Gamma \subset \mathrm{Diff}(M)$ a group of diffeomorphisms acting properly discontinuously (but not necessarily freely) on M. Then the quotient space M/Γ is an n-dimensional orbifold.*

Proof. Consider the maximal atlas on the metrizable topological space M/Γ containing the orbifold chart $(M/\Gamma, M, \Gamma, \pi_\Gamma)$ where $\pi_\Gamma : M \to M/\Gamma$ is the natural projection map. □

In particular, it follows from the proof of the proposition that the quotient space M/Γ is again a manifold if and only if the group Γ acts freely on M.

Proposition 2.1.3 motivates the following

Definition 2.1.4. An n-dimensional orbifold O^n is called *good* if there is an n-dimensional manifold M^n and a group of diffeomorphisms $\Gamma \subset \mathrm{Diff}(M)$ acting properly discontinuously on M such that O is diffeomorphic to the orbifold M/Γ. It is called is called very good if moreover the group of diffeomorphisms Γ is finite. An orbifold is called *bad* if it is not good.

That bad orbifolds do exist will become clear in section 2.1.3 where we will encounter examples of 2-dimensional bad orbifolds. A very good orbifold is called *spherical* (*discal*, *toric*, *solid toric*) if it is diffeomorphic to the quotient of a round sphere S^n (a closed unit disc D^n, a flat torus T^n, the compact 3-dimensional solid torus $D^2 \times S^1$) by a finite isometric group action.

We mention without proof that there is another, equivalent way of stating definition 2.1.4 which is often used in the literature: A *covering map* of orbifolds is a continuous map $p : \hat{O} \to O$ such that every point $x \in O$ has a neighbourhood U with the following property: For every component V of $p^{-1}(U) \subset \hat{O}$ there is an orbifold chart $(V, \tilde{V}, \Gamma_V, \pi_V)$ and a (possibly larger) finite group of diffeomorphisms $\Gamma_V \subset \Gamma'_V \subset \mathrm{Diff}(\tilde{V})$ such that $(U, \tilde{V}, \Gamma'_V, p \circ \pi_V)$ is an orbifold chart for O. In this sense, an orbifold is good if it is covered by a manifold, and very good if it is finitely covered by a manifold, i.e. covered by a manifold such that every point has finitely many preimages.

The proof that the two definitions for good orbifolds given above uses that covering maps for orbifolds share many of the properties of manifold coverings. In particular, one can define the deck transformation group of an orbifold covering, and orbifold equivalents to the fundamental group and the universal covering. (Note that these are not the same as the fundamental group and the universal covering of the underlying space!) Since we will not require any of these notions for the purposes of this work, we refer the reader to the discussion in [Th78, Sec. 13.2].

Throughout this section, we will only discuss smooth ($\mathcal{C}^\infty$) orbifolds. However, it is equally possible to define differentiable orbifolds of lower regularity. Thus, in the charts $(U, \tilde{U}, \Gamma_U, \pi_U)$ for a $\mathcal{C}^k$ orbifold the open sets $\tilde{U}$ are subsets of $\mathcal{C}^k$ manifolds, and the finite groups Γ_U consist of $\mathcal{C}^k$-diffeomorphisms of $\tilde{U}$. Similarly, we only require the coordinate changes to be $\mathcal{C}^k$-diffeomorphisms. Provided that k is not too small, say $k \geq 4$, the following discussions also apply to orbifolds of regularity $\mathcal{C}^k$. We will use (Riemannian) orbifolds of lower regularity in section 6 where they arise as limits of smooth orbifolds.

2.1.2 Local groups and stratifications of orbifolds

We begin this section by collecting some basic facts about lifting maps of ball quotients to maps of balls. We restrict our attention to maps between ball quotients of equal dimensions which lift locally to smooth embeddings and are interested in the existence and uniqueness of global lifts.

Let $\widetilde{B}_1, \widetilde{B}_2 \subset \mathbb{R}^n$ be open metric balls and let $\Gamma_i \subset \mathrm{Isom}(\widetilde{B}_i)$ be finite groups. We denote by $p_i : \widetilde{B}_i \to B_i := \widetilde{B}_i/\Gamma_i$ the quotient projections.

Lemma 2.1.5 (Ambiguity of lifts). *(i) Let $\tilde{\phi}, \tilde{\phi}' : \widetilde{B}_1 \to \widetilde{B}_2$ be smooth immersions with $p_2 \circ \tilde{\phi} = p_2 \circ \tilde{\phi}'$. Then there exists a unique transformation $\gamma_2 \in \Gamma_2$ such that $\tilde{\phi}' = \gamma_2 \circ \tilde{\phi}$. In particular, if $\tilde{\phi}$ and $\tilde{\phi}'$ coincide on an open set, then $\tilde{\phi}' = \tilde{\phi}$.*

(ii) Let $\widetilde{\psi} : \widetilde{B}_1 \to \widetilde{B}_2$ be a smooth embedding which maps Γ_1-orbits into Γ_2-orbits. Then $\widetilde{\psi}$ is equivariant with respect to a unique monomorphism $\alpha : \Gamma_1 \to \Gamma_2$, i.e. for every $\gamma_1 \in \Gamma_1$ holds $\widetilde{\psi} \circ \gamma_1 = \alpha(\gamma_1) \circ \widetilde{\psi}$.

Proof. (i) Suppose first that $\tilde{\phi}$ and $\tilde{\phi}'$ are embeddings. Then $\tilde{\phi}' \circ \tilde{\phi}^{-1}$ is a diffeomorphism $U \to U'$ between connected open subsets of $\widetilde{B}_2$ which moves points inside their Γ_2-orbits, $p_2 \circ (\tilde{\phi}' \circ \tilde{\phi}^{-1}) = p_2$.

The subset $W \subseteq U$ of points with trivial Γ_2-stabilizer is open and dense in U because it is locally the complement of finitely many proper submanifolds. For any point $\tilde{x}_2 \in W$ exists a unique transformation $\gamma_2(\tilde{x}_2) \in \Gamma_2$ such that $\tilde{\phi}' \circ \tilde{\phi}^{-1} = \gamma_2(\tilde{x}_2)$ near $\tilde{x}_2$. The assignment $\tilde{x}_2 \mapsto \gamma_2(\tilde{x}_2)$ is locally constant.

To see that it is globally constant, we fix a Γ_2-invariant Riemannian metric $\tilde{g}_2$ on $\tilde{B}_2$ and observe that $\tilde{\phi}' \circ \tilde{\phi}^{-1}|_W$ becomes a Riemannian isometry. By continuity, $\tilde{\phi}' \circ \tilde{\phi}^{-1}$ itself is isometric. It follows that the identities $\tilde{\phi}' \circ \tilde{\phi}^{-1} = \gamma_2(\tilde{x}_2)$ for $\tilde{x}_2 \in W$ hold on all of U, because both sides are Riemannian isometries and they coincide on an open subset. In particular, $\gamma_2(\tilde{x}_2)$ is independent of $\tilde{x}_2$. This proves the assertion in the case of embeddings.

In the general case when $\tilde{\phi}$ and $\tilde{\phi}'$ are immersions, the previous argument yields that for every $\tilde{x}_1 \in \tilde{B}_1$ exists a unique $\gamma_2(\tilde{x}_1) \in \Gamma_2$ such that $\tilde{\phi}' = \gamma_2(\tilde{x}_1) \circ \tilde{\phi}$ near $\tilde{x}_1$. Since $\tilde{B}_1$ is connected, $\gamma_2(\tilde{x}_1)$ does not depend on $\tilde{x}_1$.

(ii) This follows from part (i) because $p_2 \circ \tilde{\psi} \circ \gamma_1 = p_2 \circ \tilde{\psi}$ for all $\gamma_1 \in \Gamma_1$. □

Lemma 2.1.6 (Existence of lifts). *(i) Suppose that $\phi : \widetilde{B}_1 \to B_2$ is a continuous map such that for any point $\tilde{x}_1 \in \widetilde{B}_1$ exists an open neighbourhood $\widetilde{V}_{\tilde{x}_1}$ and a smooth embedding $\widetilde{\phi}_{\tilde{x}_1} : \widetilde{V}_{\tilde{x}_1} \to \widetilde{B}_2$ with $p_2 \circ \widetilde{\phi}_{\tilde{x}_1} = \phi|_{\widetilde{V}_{\tilde{x}_1}}$. Then there exists a smooth immersion $\widetilde{\phi} : \widetilde{B}_1 \to \widetilde{B}_2$ with $p_2 \circ \widetilde{\phi} = \phi$.*

(ii) Suppose that $\psi : B_1 \to B_2$ is a continuous map such that for any point $x_1 \in B_1$

exists an open subset $\widetilde{V}_{x_1} \subset \widetilde{B}_1$ *with* $x_1 \in p_1(\widetilde{V}_{x_1})$ *and a smooth embedding* $\tilde{\psi}_{x_1} : \widetilde{V}_{x_1} \to \widetilde{B}_2$ *with* $p_2 \circ \tilde{\psi}_{x_1} = \psi \circ p_1|_{\widetilde{V}_{x_1}}$. *Then there exists a smooth immersion* $\widetilde{\psi} : \widetilde{B}_1 \to \widetilde{B}_2$ *with* $p_2 \circ \widetilde{\psi} = \psi \circ p_1$.

Proof. We fix a connected open subset $\widetilde{V}_0$ and a smooth embedding $\widetilde{\phi}_0 : \widetilde{V}_0 \to \widetilde{B}_2$ with $p_2 \circ \widetilde{\phi}_0 = \phi|_{\widetilde{V}_0}$, and consider finite sequences of connected open subsets $\widetilde{V}_i \subset \widetilde{B}_1$, points $\tilde{y}_i \in \widetilde{V}_i \cap \widetilde{V}_{i-1}$ and smooth embeddings $\widetilde{\phi}_i : \widetilde{V}_i \to \widetilde{B}_2$ with $p_2 \circ \widetilde{\phi}_i = \phi|_{\widetilde{V}_i}$, $1 \leq i \leq k$. For any such chain, we can match the local lifts $\widetilde{\phi}_i$ by postcomposing them with suitable transformations in Γ_2 such that $\widetilde{\phi}_i = \widetilde{\phi}_{i-1}$ near $\tilde{y}_i$, cf. Lemma 2.1.5. This implies a lifting property for paths: For any continuous path $\tilde{c}_1 : [0,1] \to \widetilde{B}_1$ exists a unique path $\tilde{c}_2 : [0,1] \to \widetilde{B}_2$ such that $\tilde{c}_2 = \widetilde{\phi}_0 \circ \tilde{c}_1$ near 0 and $\phi \circ \tilde{c}_1 = p_2 \circ \tilde{c}_2$. Moreover, smooth families of smooth paths lift to smooth families of smooth paths. It follows that $\widetilde{\phi}_0$ extends to a smooth lift $\widetilde{\phi}$ as desired, e.g. by extending it along segments emanating from a basepoint in $\widetilde{V}_0$.

(ii) The map $\psi \circ p_1$ satisfies the assumptions of part (i), since we can precompose the local lifts $\tilde{\psi}_{x_1}$ with transformations in Γ_1. The claim follows. □

By Lemma 2.1.5 the coordinate changes of an orbifold are equivariant. More precisely, consider two charts $(U_i, \tilde{U}_i, \Gamma_{U_i}, \pi_{U_i})$ and corresponding points $\tilde{x}_i \subset \tilde{U}_i$ with $\pi_{U_1}(\tilde{x}_1) = \pi_{U_2}(\tilde{x}_2)$ we have by Definition 2.1.1 a coordinate change $\tilde{\phi} : \tilde{V}_1 \to \tilde{V}_2$ defined on open sets $\tilde{x}_i \in \tilde{V}_i \subset U_i$ such that $\pi_{U_2} \circ \tilde{\phi} = \pi_{U_1}$.

We fix a Γ_{U_1}-invariant metric $\tilde{g}$ on $\tilde{U}_1$. Then a sufficiently small metric ball $\tilde{V}_1' \subset \tilde{V}_1$ centered at $\tilde{x}_1$ is invariant under the group $\mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1)$. Consider a radial segment c in $\tilde{V}_1'$, i.e. a segment starting at $\tilde{x}_1$ By construction, the lifts of $\pi_{U_1} \circ c$ to $\tilde{U}_1$ starting at $\tilde{x}_1$ are precisely the radial segments $\gamma \circ c$ for $\gamma \in \mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1)$, and these segments are contained in $\tilde{V}_1'$.

Now, for any $\gamma' \in \mathrm{Stab}_{\Gamma_{U_2}}(\tilde{x}_2)$, the curve $c_{\gamma'} := \gamma' \tilde{\phi} \circ c$ is a curve in $\tilde{U}_2$ starting at $\tilde{x}_2$ and projecting to $\pi_{U_1} \circ c$. It follows that $\tilde{\phi}^{-1} \circ c_{\gamma'}$ is one of the radial segments $\gamma \circ c$ with $\gamma \in \mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1)$, and hence that $c_{\gamma'}$ is contained in $\tilde{V}_2' = \tilde{\phi}(\tilde{V}_1') \subset \tilde{U}_2$. In other words, the ball $\tilde{V}_2'$ is $\mathrm{Stab}_{\Gamma_{U_2}}(\tilde{x}_2)$-invariant. Moreover, the metric $(\tilde{\phi}^{-1})_* \tilde{g}$ on V_2' is invariant under the operation of the group $\mathrm{Stab}_{\Gamma_{U_2}}(\tilde{x}_2)$. This is clear on the dense set of point where $\mathrm{Stab}_{\Gamma_{U_2}}(\tilde{x}_2)$ acts freely, and hence everywhere by continuity.

Using the exponential map, we can identify $\tilde{V}_1'$ with a metric ball in $T_{\tilde{x}_1}\tilde{U}_1$ on which the linearized group $d\,\mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1)$ operates isometrically; this identification is equivariant with respect to the natural identification $\mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1) \cong d\,\mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1)$. We can perform the same construction for $\tilde{V}_2'$. It now follows from Lemma 2.1.5 that the coordinate change $\tilde{\phi}$ is equivariant on $\tilde{V}_1'$ with respect to a unique group isomorphism $\mathrm{Stab}_{\Gamma_{U_1}}(\tilde{x}_1) \to \mathrm{Stab}_{\Gamma_{U_2}}(\tilde{x}_2)$. In particular, the linear representations $\mathrm{Stab}_{\Gamma_{U_i}}(\tilde{x}_i) \to O(n) \subset \mathrm{Aut}(T_{\tilde{x}_i}\tilde{U}_i)$ are isomorphic, and unique up to conjugacy.

Definition 2.1.7. For every point $x \in O$, we define its *local (isotropy) group* $\Gamma_x \subset O(n)$ as the (conjugacy class of the) image of the linear representation $\mathrm{Stab}_{\Gamma_{U_i}}(\tilde{x}_i) \to O(n)$. A point x is called *regular* if Γ_x is trivial, and *singular* otherwise.

We will often implicitely identify the local isotropy group Γ_x of a point $x \in O$ with the stabilizer group $\mathrm{Stab}_{\Gamma_U}(\tilde{x})$ with respect to some orbifold chart.

An orbifold O has a natural *stratification* into *strata* of points of equal type. In an orbifold chart $(U, \tilde{U}, \Gamma_U, \pi_U)$ around x, the fixed point set of $\mathrm{Stab}_{\Gamma_U}(\tilde{x}) \cong \Gamma_x$ in $\tilde{U}$ is a submanifold and its connected component through $\tilde{x}$ projects to points in O with the same local group as x. Hence the equivalence classes of points with the same local groups inherit natural structures as smooth manifolds, and together form a stratification of O. The dimension of the stratum S_x through a point $x \in O$ equals the dimension of the linear subspace $\mathrm{Fix}(\Gamma_x) \subseteq \mathbb{R}^n$. For $0 \leq d \leq n = \dim O$, we write $O^{(d)}$ for the union of all d-dimensional strata. The boundary of a stratum is a locally finite union of strata of lower dimensions. Note that, when passing from a stratum to its boundary, the local isotropy group strictly increases.

The regular points of an orbifold O form the top-dimensional stratum O^{reg}. It is open, dense and path connected. Its complement $O^{\mathrm{sing}} = O - O^{\mathrm{reg}}$ is called the *singular locus* of O. The singular $(n-1)$-dimensional stratum $O^{(n-1)}$ consists of the points with local group $\cong \mathbb{Z}_2$ generated by a hyperplane reflection. Its closure $\partial_{\mathrm{refl}} O := \overline{O^{(n-1)}}$ is usually referred to as *reflector boundary* or *silvered boundary* of O. It is not usually contained in the boundary ∂O; instead we have the relation $\partial_{\mathrm{refl}} O \cap \partial O = \partial_{\mathrm{refl}} \partial O$. The reflector boundary of O consists of all points whose local group contains a hyperplane reflection. We call $O^{(n-1)}$ the regular part of the reflector boundary.

A component of the zero-dimensional stratum $O^{(0)}$ must be a point, and is often called a singular *vertex*. A component of a one-dimensional stratum $O^{(1)}$ is called a singular *edge* (assuming that $\dim O \geq 2$).

Using local isotropy groups, the notion of a tangent space carries over to orbifolds in a natural way. Thus, the *tangent bundle* $TO \to O$ of a smooth orbifold O can be defined locally, with respect to a chart $(U, \tilde{U}, \Gamma_U, \pi_U)$, as the quotient $T\tilde{U}/\Gamma_U \to U$ of the tangent bundle $T\tilde{U} \to \tilde{U}$. The restriction $TO|_{O^{\mathrm{reg}}} \to O^{\mathrm{reg}}$ of the orbifold tangent bundle to the regular stratum is the tangent bundle in the manifold sense[1]. For a point $x \in O$, the tangent space $T_x O$ can be canonically identified with $\mathbb{R}^n/\Gamma_x$. The tangent space $T_x S_x \subseteq T_x O$ of

[1]In section 2.2.1, we will discuss orbifold fiber bundles; the natural map $TO \to O$ is an orbifold fibration with generic fiber $\mathbb{R}^n$ as defined there. However, we will not discuss or use this "singular bundle structure" of TO away from the regular stratum.

the stratum S_x through x corresponds to the linear subspace $\mathrm{Fix}(\Gamma_x) \subset \mathbb{R}^n$. A (smooth) vector field on an orbifold O is a section of TO which is locally covered by a (smooth) vector field.

2.1.3 Examples: Low-dimensional orbifolds

In this section, we give a brief discussion of the local structure of orbifolds of dimension ≤ 3. We are particularly interested in the closed (i.e. compact without boundary) case.

A point in a 1-orbifold is either regular, a boundary point or a reflector boundary point (with local group $O(1) \cong \mathbb{Z}_2$). In particular, there are precisely two connected closed 1-orbifolds, namely the circle S^1 and the *mirrored interval* I. The latter has as underlying topological space $|I|$ the compact interval $[0, 1]$. However, the boundary points of $|I|$ are reflector boundary points of I, i.e. $\partial_{\mathrm{refl}} I = I^{(0)} = \partial|I|$.

Next, let O^2 be a 2-orbifold, possibly with boundary. A point in the 1-stratum is a reflection boundary point with local group $\cong D_1 \cong \mathbb{Z}_2$ acting by a reflection (on the disc or half-disc). It may be a boundary point, namely one of the points in $\partial O \cap \partial_{\mathrm{refl}} O$. If it not a boundary point, it has a neighbourhood diffeomorphic to $V^2(1) := D^2/D_1$.

There are two possibilities for the local structure at a singular point in $O^{(0)}$. It its local group is $\cong \mathbb{Z}_p$, $p \geq 2$, acting by rotations, it is called a *cone point* of order p. (In the Riemannian case which we discuss in section 3, this corresponds to a *cone angle* of $\frac{2\pi}{p}$.) It then is an isolated singular point in the interior of $|O|$ and has a neighbourhood diffeomorphic to the disc $D^2(p) := D^2/\mathbb{Z}_p$ with cone point of order p. Alternatively, the local group $\Gamma_x \subset O(2)$ of a singular point x may be a dihedral group D_q, $q \geq 2$, in which case x is called a *corner vertex* of order q.[2] Corner vertices are are reflector boundary points but not boundary points and have neighbourhood diffeomorphic to the sector $V^2(q) := D^2/D_q$. Note that $\partial_{\mathrm{refl}} V^2(q) = \partial|V^2(q)|$. Sometimes it will more conventient to also admit cone points and corner vertices of order 1 which are nothing else than regular interior, respectively, regular reflector boundary points.

The singular locus O^{sing} consists of the reflector boundary $\partial_{refl} O = \overline{O^{(1)}}$, which contains the set of all corner vertices, and of the subset of $O^{(0)}$ of (isolated) cone points. We call a connected component of $O^{(1)}$ a *reflector edge*. In a corner vertex, locally two reflector edges meet.

A connected component of $\partial|O|$ can be a connected component of ∂O or of $\partial_{\mathrm{refl}} O$, or it can be a chain of consecutive boundary arcs and reflector edges. In the latter case, any

[2] As a subgroup of $O(2)$, the dihedral group D_q is defined as the isometry group of a regular q-gon. It is generated by the reflections at two lines through the origin with angle $\frac{\pi}{q}$. As an abstract group, it has the presentation $\langle s_1, s_2 | s_1^2 = s_2^2 = (s_1 s_2)^q = 1 \rangle$.

two of the boundary arcs are disjoint, but there may be sequences of consecutive reflector edges meeting at corner vertices.

Following Thurston [Th78], we will use the following notation for *closed* 2-orbifolds. Let Σ be a 2-manifold, possibly with boundary. Then we denote by $\Sigma(p_1, \ldots, p_k; q_1, \ldots, q_l)$ the closed 2-orbifold with underlying space Σ, reflector boundary $\partial\Sigma$, k cone points of orders p_i located in the interior of Σ and l corner vertices of orders q_j lying on $\partial\Sigma$. Thus for example $D^2(p;) = D^2(p)$ is a closed 2-orbifold with a reflector boundary circle (and one cone point of order p). If there are no cone points ($k = 0$) we write $\Sigma(; q_1, \ldots, q_l)$. If no corner vertices occur ($l = 0$), in particular if $\partial\Sigma = \emptyset$, we write briefly $\Sigma(p_1, \ldots, p_k;) =: \Sigma(p_1, \ldots, p_k)$, and also $\Sigma(;) =: \Sigma$. We will only apply this notation if the diffeomorphism type of the 2-orbifold is uniquely determined.

For such a 2-orbifold $O = \Sigma(p_1, \ldots, p_k; q_1, \ldots, q_l)$, we can define its *Euler characteristic* as

$$\chi(O) = \chi(\Sigma) - \sum_{i=1}^{k}(1 - \frac{1}{p_i}) - \frac{1}{2}\sum_{j=1}^{l}(1 - \frac{1}{q_j}).$$

The Euler characteristic has a multiplicative property under finite orbifold coverings ([Th78, Prop. 13.3.4]). Thus, for an orbifold covering $O' \to O$ of closed 2-orbifolds where every regular point of O has k (regular) preimages, we have $\chi(O') = k \cdot \chi(O)$.

It is a simple combinatorial exercise to derive the well-known classification of connected closed orbifolds with $\chi(O) \geq 0$ (cf. [Th78, Thm. 13.3.6]). They can be arranged in three classes:

The 2-orbifolds S^2, $\mathbb{R}P^2$, $S^2(p,p)$, $S^2(2,2,p)$, $S^2(2,3,3)$, $S^2(2,3,4)$, $S^2(2,3,5)$, $\mathbb{R}P^2(p)$, D^2, $D^2(p)$, $D^2(;p,p)$, $D^2(;2,2,p)$, $D^2(;2,3,3)$, $D^2(;2,3,4)$, $D^2(;2,3,5)$, $D^2(2;p)$ and $D^2(3;2)$ with $p \geq 2$ are *spherical* 2-orbifolds, i.e. quotients of the round 2-sphere (by a finite group of isometries).

The 2-orbifolds T^2, K^2, $S^2(2,3,6)$, $S^2(2,4,4)$, $S^2(3,3,3)$, $S^2(2,2,2,2)$, $\mathbb{R}P^2(2,2)$, Ann^2, $\mathrm{Möb}^2$, $D^2(;2,3,6)$, $D^2(;2,4,4)$, $D^2(;3,3,3)$, $D^2(;2,2,2,2)$, $D^2(4;2)$, $D^2(3;3)$, $D^2(2;2,2)$ and $D^2(2,2)$ are *toric*, i.e. quotients of the flat 2-torus.

The remaining closed 2-orbifolds with $\chi(O) \geq 0$ are diffeomorphic to $S^2(p)$, $S^2(p,q)$, $D^2(;p)$ or $D^2(;p,q)$ with $2 \leq p < q$. These orbifolds are *bad*, as can be verified by decomposing them into two discal 2-orbifolds (with boundary). A manifold covering such an orbifold would likewise have to consist out of two 2-discs which is impossible.

In fact these are the only bad 2-orbifolds without boundary ([Sc83, Thm. 2.3]). Since spherical and toric 2-orbifolds clearly have nonnegative Euler characteristic, we have

Proposition 2.1.8. *A closed 2-orbifold has nonnegative Euler characteristic if and only if it is spherical, toric or bad.*

Consider now a 3-orbifold O^3, possibly with boundary. If $x \in O^{(1)}$, then the local group $\Gamma_x \subset O(3)$ fixes a line, i.e. $\Gamma_x \cong \mathbb{Z}_p$ or D_p with $p \geq 2$ and $S^2/\Gamma_x \cong S^2(p,p)$ or $\overline{D}^2(;p,p)$. We call the connected component of $O^{(1)}$ containing x a *singular edge*, respectively, *reflector edge* (or circle) of order p. In the reflector case $\Gamma_x \cong D_p$, locally two reflector faces, i.e. components of $O^{(2)}$, meet at the edge. The boundary points of the singular and reflector edges (i.e. of their underlying 1-manifolds) are the cone points and corner vertices of the 2-orbifold ∂O.

If x is a singular vertex, i.e. if $x \in O^{(0)}$, then $\Gamma_x \subset O(3)$ has no nontrivial fixed vector. Thus, the spherical 2-orbifold S^2/Γ_x must be diffeomorphic to one of $\mathbb{R}P^2$, $S^2(2,2,p)$, $S^2(2,3,3)$, $S^2(2,3,4)$, $S^2(2,3,5)$, $\mathbb{R}P^2(p)$, $D^2(p)$, $D^2(2;p)$ $D^2(3;2)$, $D^2(;2,2,p)$, $D^2(;2,3,3)$, $D^2(;2,3,4)$ or $D^2(;2,3,5)$ with $p \geq 2$. We have $x \in \partial_{\text{refl}} O$ if and only if $\partial_{\text{refl}}(S^2/\Gamma_x) \neq \emptyset$. The cone points and corner vertices of S^2/Γ_x correspond to singular edges emanating from the point x.

2.2 Fibrations and decompositions

2.2.1 Fibered 3-orbifolds

Let F be an orbifold without boundary. Following [BMP03, sec. 2.4] we define an *orbifold fiber bundle* or *orbifold fibration* with generic fiber F as a submersion $p : O \to B$ of orbifolds, possibly with boundary, with the following property: For every point $x \in B$ there is a chart $\phi : \tilde{U} \to U$ around x, a smooth operation $\Gamma_x \curvearrowright F$ and a submersion $\sigma : \tilde{U} \times F \to O$ inducing a diffeomorphism between $(\tilde{U} \times F)/\Gamma_x$ (where we divide out the diagonal action) and $p^{-1}(U)$ such that $p \circ \sigma = \phi \circ \pi_{\tilde{U}}$. In the case with boundary we require that $p^{-1}(\partial B) = \partial O$; in this case p restricts over the boundary to the orbifold fiber bundle $p|_{\partial O} : \partial O \to \partial B$. Note that orbifold coverings are (the same as) orbifold fiber bundles with 0-dimensional fiber. We say that a compact orbifold *fibers* if it is the total space of an orbifold fibration whose base and fiber have strictly positive dimension and whose generic fiber is a closed orbifold.

For the remainder of this text, we will restrict ourselves to fibrations of 3-dimensional orbifolds.

An *orbifold Seifert fibration* is an orbifold fibration $p : O^3 \to B^2$ with 3-dimensional total space O, 2-dimensional base B and 1-dimensional closed connected generic fiber F, i.e. F is the circle S^1 or the mirrored interval I. A *Seifert orbifold* is a 3-orbifold admitting a Seifert fibration. In an orbifold Seifert fibration, every fiber has a neighbourhood which is fiber preserving diffeomorphic to a solid toric orbifold equipped with a canonical Seifert fibration. More precisely, suppose that $x \in B$ is a point in the base and let Γ_x be its local group. Then the fiber $p^{-1}(x)$ has a saturated neighbourhood of the form $(D^2 \times F)/\Gamma_x$ with the natural fibration $(D^2 \times F)/\Gamma_x \to D^2/\Gamma_x$. The action $\Gamma_x \curvearrowright D^2$ is effective, whereas the action $\Gamma_x \curvearrowright F$ is in general not. Fibers in the boundary have similar model neighbourhoods. A classification of Seifert orbifolds, both locally and globally, has been given in [BS85].

Seifert fibrations of solid toric 3-orbifolds as well as 1-dimensional fibrations of their toric boudaries are in general not unique. The next result describes which fibrations of the boundary extend to Seifert fibrations. Let $V \cong (D^2 \times S^1)/\Gamma$ be a solid toric 3-orbifold. We call a 1-dimensional fibration of ∂V *horizontal* if it is isotopic to the fibration $\partial V \to S^1/\Gamma$.

Lemma 2.2.1. *A 1-dimensional fibration of the boundary ∂V of a solid toric orbifold V extends to a Seifert fibration of V if and only if it is not horizontal.*

Proof. This is a consequence of the fact that 1-dimensional fibrations of closed flat 2-orbifolds can be isotoped to be geodesic. □

A *toric fibration* of a 3-orbifold is an orbifold fibration whose generic fiber is a toric 2-orbifold. (Fibrations with 2-dimensional fibers of other topological types will play no role in this text.)

2.2.2 2-suborbifolds in 3-orbifolds

We recall from section 2.1.1 that a (proper) 2-suborbifold Σ in a 3-orbifold O is a subset whose preimages in local charts are (proper) 2-submanifolds. It is called *two-sided* if it has a product neighbourhood of the form $\Sigma \times (-1, 1)$. It is called *locally two-sided* at a point x if it has such a product neighbourhood locally near x, i.e. if Γ_x does not switch the sides of the hypersurface $\pi_U^{-1}(\Sigma) \subset \tilde{U}$ locally at $\tilde{x}$, equivalently, if Γ_x acts trivially on the line $T_{\tilde{x}}\tilde{U}/T_{\tilde{x}}\pi_U^{-1}(\Sigma)$ (in some local trivialization $\tilde{U}$ of O near x).

If Σ is not (globally) two-sided then it has a tubular neighbourhood of the form $(\Sigma' \times (-1, 1))/\mathbb{Z}_2$ where $\mathbb{Z}_2$ acts by a reflection on $(-1, 1)$ and by a (possibly trivial) involution on Σ'. (To verify this, one can take Σ' to be the boundary of a tubular neighbourhood of Σ.) When speaking of a 3-suborbifold $O' \subset O$ we usually suppose that the components of $\partial O'$ are either components of ∂O or disjoint from ∂O, i.e. two-sided suborbifolds of $\mathrm{int}(O)$.

A 3-orbifold O is called *irreducible* if it does not contain any bad 2-suborbifold and if every two-sided spherical 2-suborbifold bounds a discal 3-suborbifold. It is called *(topologically) atoroidal* if every incompressible two-sided toric 2-suborbifold $\Sigma \subset O$ is *boundary parallel*, i.e. bounds a collar neighbourhood $\cong \Sigma \times [0, 1]$ of a boundary component $\cong \Sigma$.

Let O be a 3-orbifold and let $\Sigma \subset O$ be a proper 2-suborbifold. A *compressing discal 2-suborbifold* or *compression disc* for Σ is a discal 2-suborbifold $D \subset O$ which intersects Σ

transversally in $\partial D = D \cap \Sigma$ such that ∂D does not bound a discal 2-suborbifold in Σ. (If D is one-sided, we understand this to mean that splitting the connected component of Σ containing ∂D along ∂D does not yield a discal 2-orbifold. Anyway, one-sided compression discs can be replaced by two-sided ones by passing to the boundary of a tubular neighbourhood.) Note that a spherical 2-suborbifold has no compression discs because every closed 1-suborbifold of a spherical 2-orbifold bounds a discal 2-suborbifold. A *compression* of Σ is either a discal 3-suborbifold whose boundary is a component of Σ or a compression disc for Σ. If Σ admits a compression then it is called *compressible*, and otherwise *incompressible*. Thus a 3-orbifold is irreducible if it contains no bad 2-suborbifolds and if all two-sided spherical 2-suborbifolds are compressible. The notion of incompressibility is particularly useful in the irreducible case because then the position e.g. of closed 2-suborbifolds Σ relative to incompressible 2-suborbifolds Σ_{inc} can be simplified by isotopies. Namely, in this case we can be achieve that Σ_{inc} divides Σ into non-discal components.

Discal 3-orbifolds are irreducible. This is formulated but not proved in [BMP03, Thm. 3.1]. A proof can be found in [DL09, 2.4].

More generally, every closed 2-suborbifold of a discal 3-orbifold is compressible. For nonspherical suborbifolds this follows from the Equivariant Loop Theorem [MY80], cf. [BMP03, Thm. 3.6]. We will use it only for toric 2-suborbifolds.

2.2.3 Decompositions of 3-orbifolds along 2-suborbifolds

In the following, let O be a compact 3-orbifold.

Consider a finite family of disjoint two-sided closed 2-suborbifolds $\Sigma_j \subset \text{int}(O)$. The operation of removing from O an open tubular neighbourhood of $\cup_j \Sigma_j$ is called *splitting O* along the Σ_j. We call the splitting spherical (toric, incompressible) if all Σ_j are spherical (toric, incompressible). We will refer to the Σ_j as *splitting 2-suborbifolds*. A *connected sum decomposition* of O or a (spherical) *surgery* on O is performed by first splitting O along a family of spherical 2-suborbifolds and then filling discal 3-orbifolds into the additional spherical boundary components created by the splitting. (These discal 3-orbifolds are uniquely determined by the diffeomorphism types of the splitting 2-suborbifolds.) Conversely, O is called a *connected sum* of the 3-orbifolds resulting from this decomposition. Note that we allow connected sums of connected orbifolds (components) with themselves.

The following result reduces the study of compact 3-orbifolds without bad 2-suborbifolds to the study of irreducible ones. It is due to Kneser [Kn29] in the manifold case, see [BMP03, 3.3] for a proof in the case of orientable orbifolds. The argument given there also extends to the nonorientable case.

Theorem 2.2.2 (Spherical decomposition). *A compact 3-orbifold without bad 2-suborbi-*

folds can be decomposed by surgery into finitely many irreducible compact 3-orbifolds. The non-spherical components of this decomposition are unique up to diffeomorphism.

If a 3-orbifold O contains no bad 2-suborbifolds, then this is also true for its connected summands: If a summand were to contain a bad 2-suborbifold, it would also contain a 2-sided bad 2-suborbifold. By radially pushing it out of the discal 3-orbifold glued into the splitting 2-suborbifold, we could then make it disjoint from the surgery region, thus obtaining a contradiction. (Note that a 2-sided 2-suborbifold of O cannot meet $O^{(0)}$ and must intersect $O^{(1)}$ and $O^{(2)}$ transversally.)

Remark 2.2.3. If a 3-orbifold O contains bad 2-suborbifolds, it is still possible to decompose it into finitely many components in which spherical 2-suborbifolds bound discal 3-orbifolds. However, in this case the splitting is in general not uniquely determined. To see what can go wrong, consider the following simple counter-example:

Let $O_1 = S^2(2,4,4) \times S^1$ (toric) and $O_2 = S^2(2,4) \times S^1$ (bad). In both orbifolds, consider a tubular neighbourhood of a singular edge of order 4, and decompose it into two components $U \cong D^2(4) \times [0,1]$. Similarly, we decompose the complement of this tubular neighbourhood in O_2 into two components $V \cong D^2(2) \times [0,1]$. We now can write $O_1 = X_1 \cup U_1' \cup U_2'$ and $O_2 = U_1 \cup U_2 \cup V_1 \cup V_2$ with $U_1 \cup V_1 \cong U_2 \cup V_2 \cong S^2(2,4) \times [0,1]$.

If we remove U_1' and U_1 from O_1 and O_2 respectively, glue the rests to each other along the resulting boundary component $\cong S^2(4,4)$ and call the result O, we have by construction a decomposition by surgery of O into O_1 and O_2 and one easily verifies that these components are irreducible (and non-spherical) by considering their universal coverings.

On the other hand, consider in O the union $X = U_2' \cup U_2 \cup V_2$; it has boundary $\partial X \cong S^2(2,2)$. Removing X from O has the same effect as removing $U_1' \cup U_2'$ from O_1 and glueing V_1 into the boundary along $S^1 \times [0,1]$. Thus, performing surgery along ∂X yields the components $O_1' = S^2(2,2,4) \times S^1$ and $O_2 = S^2(2,4) \times S^1$ (bad); and O_1' can be further surgered along a spherical cross section into the spherical 3-orbifold obtained by glueing two copies of the spherical cone over $S^2(2,2,4)$ to each other. These two decopositions of O by surgery are clearly non-equivalent.

Let us now turn to toric splittings.

Lemma 2.2.4. *Suppose that O contains no bad 2-suborbifolds and is split along a toric family $\mathcal{T}$ into compact pieces O_i. Then O is irreducible and $\mathcal{T}$ is incompressible (in O) if and only if all pieces O_i are irreducible and for each piece O_i the portion $\partial O_i - \partial O$ of its boundary corresponding to $\mathcal{T}$ is incompressible (in O_i). Moreover, if in this situation all boundary components of the O_i are incompressible, then O has incompressible boundary.*

Proof. The standard proof in the manifold case (cf. [Wa67, 1.8 and 1.9]) carries over. The "only if" direction uses the fact that toric 2-suborbifolds of discal 3-orbifolds are always compressible. □

There is a canonical splitting of irreducible compact 3-orbifolds along incompressible toric suborbifolds. It is due to Jaco, Shalen and Johannson in the manifold case and has be extended to orbifolds by Bonahon and Siebenmann [BS87], see also [BMP03, 3.3 and 3.15].

Theorem 2.2.5 (JSJ-splitting). *An irreducible compact 3-orbifold admits an incompressible toric splitting into components each of which is atoroidal or Seifert fibered (or both). A minimal such splitting is unique up to isotopy.*

We will also consider a class of toric splittings with weaker properties. Following Waldhausen's definition [Wa67] in the manifold case, we define a *graph splitting* of a compact 3-orbifold with toric boundary to be a (not necessarily incompressible) toric splitting into pieces which admit orbifold fibrations with 1- or 2-dimensional closed fibers. Moreover, the 2-dimensional fibers are required to be toric. We will refer to the pieces with 2-dimensional fibrations as pieces *with toric fibrations*. A 3-orbifold admitting a graph splitting is called a *graph orbifold*. Briefly, it is a 3-orbifold which can be "cut up into fibered pieces".

Connected compact 3-orbifolds with toric fibrations and nonempty boundaries are diffeomorphic to $T\times[-1,1]$ or $(T\times[-1,1])/\mathbb{Z}_2$ with a toric 2-orbifold T and, in the latter case, with $\mathbb{Z}_2$ acting by a reflection on $[-1,1]$. Unlike in the manifold case, they are not always Seifert. This is due to the fact that, whereas a 2-torus or a Klein bottle admits (infinitely many, respectively, two) circle fibrations, not all toric 2-orbifolds admit 1-dimensional orbifold fibrations [3]. Hence, in the orbifold case a graph splitting may comprise non-Seifert pieces.

Seifert orbifolds with discal base orbifold are solid toric. Any other connected Seifert orbifold O with nonempty boundary has a base orbifold B of Euler characteristic $\chi \leq 0$. If we consider the boundary ∂B of the basis as reflector boundary, we turn B into a good closed 2-orbifold which by [Th78, Thm. 13.3.6] can be realized as the quotient of a closed surface Σ with constant curvature -1 or 0 by a finite group of isometries. By pulling back the Seifert fibration over B to Σ, we can realize O as the quotient of an S^1-bundle over Σ by a finite fiber-preserving group Γ. It is now possible to construct a homogeneous metric on this S^1-bundle such that Γ operates isometrically (see [BMP03, Prop. 2.13]). Because Γ contains a reflection at the pre-image of ∂B under this quotient, this pre-image must be totally geodesic, and the curvature of the S^1-bundle is 0. It follows that O is a quotient

[3] Equivalently, the rotational part of a 2-dimensional crystallographic group with torsion is in general not reducible (cf. the discussion in [Du88]).

of a closed submanifold of $\mathbb{H}^2 \times \mathbb{R}$ or $\mathbb{R}^3$ with totally geodesic boundary, depending on whether $\chi(B) < 0$ or $\chi(B) = 0$. Equivalently, the interior of such a 3-orbifold O is covered by $\mathbb{H}^2 \times \mathbb{R}$ or $\mathbb{R}^3$.

Similarly, the 3-orbifolds with toric fibrations and nonempty boundary can be realized as quotients of flat manifolds with totally geodesic boundary. The existence of these geometric structures on the non-solid toric pieces of a graph splitting implies that they are irreducible and have incompressible boundaries.[4] Moreover, the pieces with toric fibrations and nonempty boundaries are atoroidal.

If in a non-trivial graph splitting of O no solid toric pieces occur, then all pieces are irreducible atoroidal or Seifert orbifolds with nonempty incompressible toric boundaries. Hence O is irreducible with incompressible toric boundary and the splitting is incompressible, compare Lemma 2.2.4. In particular, a minimal incompressible graph splitting of an irreducible compact connected 3-orbifold with incompressible toric boundary is *canonical* up to isotopy because it coincides with the JSJ-splitting, unless the orbifold admits a toric fibration over a closed 1-orbifold. Indeed, suppose that a nontrivial minimal incompressible graph splitting were not minimal as a splitting into atoroidal and Seifert components. Then for some splitting toric 2-suborbifold T the union of the (one or two) components adjacent to it cannot be Seifert and must therefore be atoroidal. The definition of atoroidality then implies that one of these components is $\cong T \times [0, 1]$, contradicting the minimality of the graph splitting.

2.3 Thurston's Geometrization Conjecture

2.3.1 Geometric 3-orbifolds

A 3-orbifold carries a *geometric structure* modelled on one of the eight 3-dimensional Thurston geometries S^3, $\mathbb{R}^3$, $\mathbb{H}^3$, $S^2 \times \mathbb{R}$, $\mathbb{H}^2 \times \mathbb{R}$, Nil, $\widetilde{PSL(2, \mathbb{R})}$ or Solv if all local uniformizations can be chosen as submanifolds of the respective model geometry such that operations of the local groups and local coordinate changes are Riemannian isomteries. (For a detailed discussion of the eight model geometries, see e.g. [Sc83, §4].)

A compact 3-orbifold is called *geometric* if its interior admits a complete geometric structure. Geometric orbifolds are good by [Th78, 13.3.2] and therefore a quotient of the respective model geometry by a discrete group of isometries. Since every such group has

[4] For the irreducibility one uses the fact that discal 3-orbifolds are irreducible. Namely, consider an embedded 2-sphere $S \subset \tilde{P}$ in the universal cover of such a piece P preserved by a finite group Γ_S of isometries. With the Hadamard-Cartan Theorem it follows that S is contained in a Γ_S-invariant closed ball on which the Γ_S-action is standard. Hence the action on the ball bounded by S is also standard.

a finite index subgroup operating freely, they are in fact very good. (For the discussion discrete subgroups of the eight isometry groups, see [Th97, Sec. 4.3, 4.4 and 4.7]; for the closed caes see also [Du88, Thm. 1].) The uniqueness of geometric structures for 3-manifolds (see [Sc83, Thm. 5.2]) implies that for a closed geometric 3-orbifold the model geometry is unique.

All compact Seifert 3-orbifolds without bad 2-suborbifolds are geometric, see [Th97, ch. 3] and [BMP03, 2.4]. More precisely, a connected closed Seifert orbifold without bad 2-suborbifolds admits a geometric structure modelled on a unique Thurston geometry different from hyperbolic and solvgeometry (i.e. on $S^2 \times \mathbb{R}^1$, S^3, $\mathbb{R}^3$, Nil, $\mathbb{H}^2 \times \mathbb{R}^1$, or $\widetilde{PSL}(2, \mathbb{R})$).

A solid toric 3-orbifold admits, depending on its topological type, geometric structures modelled on some or all of the six contractible model geometries. We have already seen that a non-solid toric connected compact Seifert orbifold with nonempty boundary contains no bad 2-suborbifolds and admits an $\mathbb{H}^2 \times \mathbb{R}$- or $\mathbb{R}^3$-structure. (If it admits an $\mathbb{R}^3$-structure then also an $\mathbb{H}^2 \times \mathbb{R}$-structure.) In fact it can be geometrized in a stronger sense; namely, it admits a Riemannian metric with totally geodesic boundary locally modelled on either $\mathbb{H}^2 \times \mathbb{R}$ or $\mathbb{R}^3$.

Connected 3-orbifolds admitting toric fibrations are geometric with one of the three model geometries $\widetilde{PSL}(2, \mathbb{R})$, Nil or $\mathbb{R}^3$. Those with nonempty boundary admit euclidean metrics with totally geodesic boundary and complete $\mathbb{R}^3$-structures on their interior. They are diffeomorphic to $T \times [-1, 1]$ or $(T \times [-1, 1])/\mathbb{Z}_2$ with a toric 2-orbifold T and, in the latter case, with $\mathbb{Z}_2$ acting by a reflection on $[-1, 1]$.

A *geometric splitting* of an irreducible compact connected 3-orbifold O is an incompressible toric splitting into geometric pieces, i.e. into irreducible compact 3-orbifolds whose interiors admit complete geometric structures. We refer to the components of the splitting as *geometric pieces*. Note that if O itself is not geometric, then the pieces have nonempty boundaries and admit geometric structures modelled on $\mathbb{H}^3$, $\mathbb{H}^2 \times \mathbb{R}$ or $\mathbb{R}^3$. In particular, a nontrivial incompressible graph splitting of O is a geometric splitting into pieces admitting $\mathbb{H}^2 \times \mathbb{R}$- or $\mathbb{R}^3$-structures.

A compact 3-orbifold without bad 2-suborbifolds is said to be *decomposable into geometric pieces* or to *satisfy Thurston's Geometrization Conjecture* if it can be decomposed by surgery into irreducible compact connected 3-orbifolds which are geometric or admit a geometric splitting, cf. [BMP03, 3.7]. (Thurston's Geometrization Conjecture simply states that all closed orbifolds without bad 2-suborbifolds can be decomposed into geometric pieces.)

2.3.2 Graph orbifolds are geometrizable

Graph splittings of compact 3-orbifolds have fairly weak properties; in particular they are far from being unique. In this section we show that a graph splitting can be improved to a geometric decomposition. The manifold case of this discussion is due to Waldhausen [Wa67].

Theorem 2.3.1. *Suppose that O is a compact connected 3-orbifold with toric boundary which contains no bad 2-suborbifolds and admits a graph splitting. If O is irreducible, then it is either solid toric or has incompressible toric boundary. In the latter case, it is geometric with model geometry different from $S^2 \times \mathbb{R}$ and $\mathbb{H}^3$, or it admits an incompressible graph splitting (and hence a JSJ-splitting without hyperbolic components). If O is not irreducible, it can be decomposed by surgery into irreducible orbifolds of this kind.*

Proof. Consider a graph splitting of O along a toric family $\mathcal{T}$. If $\mathcal{T} = \emptyset$, then O is solid toric, or it is closed and admits an $S^2 \times \mathbb{R}$-structure, or it is irreducible with (possibly empty) incompressible boundary and admits a geometric structure modelled on one of the six geometries different from $S^2 \times \mathbb{R}$ and $\mathbb{H}^3$. If O admits an $S^2 \times \mathbb{R}$-structure, it can be decomposed by surgery along a spherical cross section into one or two spherical 3-orbifolds.

If $\mathcal{T} \neq \emptyset$, then the pieces of the splitting have nonempty boundary. If one of the pieces P is of the form $T \times [0, 1]$ with a toric 2-orbifold T, we may reduce $\mathcal{T}$ by removing one of the components of ∂P, unless P is the only piece. In the latter case, O fibers over the circle and is geometric (with model geometry $\mathbb{R}^3$, Nil or Solv).

If no solid toric piece occurs in the graph splitting (and if $\mathcal{T} \neq \emptyset$), then O is irreducible with incompressible boundary and the graph splitting is incompressible, cf. Lemma 2.2.4. If there is a solid toric piece and if adjacent to it there is another solid toric piece or a one-ended piece with toric fibration, i.e. a piece diffeomorphic to $(T \times [-1, 1])/\mathbb{Z}_2$ with T toric and $\mathbb{Z}_2$ reflecting on $[-1, 1]$, then O is closed and geometric with model geometry S^3 or $S^2 \times \mathbb{R}$. Since in all other cases we are done or can reduce the splitting by removing a component from $\mathcal{T}$, we assume that at least one solid toric piece $V_0 \cong (D^2 \times S^1)/\Gamma$ occurs in the graph splitting and that adjacent to V_0 there is a non-solid toric Seifert piece S. We may further assume that all pieces with toric fibrations are one-ended, i.e. diffeomorphic to $(T \times [-1, 1])/\mathbb{Z}_2$ with T toric and $\mathbb{Z}_2$ reflecting on $[-1, 1]$.

We denote by $q : V_0 \cong (D^2 \times S^1)/\Gamma \to S^1/\Gamma$ the fibration of V_0 by discal cross sections. Let $K \triangleleft \Gamma$ be the kernel of the action $\Gamma \curvearrowright S^1$. Via the action $K \curvearrowright D^2$ we may regard K as a subgroup $K \subset O(2)$. The generic discal cross section of V_0 is $\cong D^2/K$. Let $p : S \to B$ denote the Seifert fibration of S. The base B is a non-discal 2-orbifold with nonempty boundary. Let $T_0 = \partial V_0 = S \cap V_0 \in \mathcal{T}$ denote the toric 2-suborbifold separating V_0 and S, and $\partial_0 B = p(T_0)$ the boundary component of B corresponding to T_0. It is either a circle or an arc connecting two points of $\partial_{\text{refl}}B \cap \partial B$.

If the fibrations $p|_{T_0}$ and $q|_{T_0}$ of T_0 by closed 1-orbifolds are not isotopic, then the Seifert fibration p can be extended over V_0, compare Lemma 2.2.1, i.e. $S \cup V_0$ is Seifert and we reduce the graph splitting by removing the component T_0 from $\mathcal{T}$. (The Seifert piece "swallows" the adjacent solid toric piece.)

Otherwise, if $p|_{T_0}$ and $q|_{T_0}$ are isotopic, we may assume that they agree, i.e. that the discal cross sections of V_0 fill in Seifert fibers. We then have the identification $\partial_0 B \cong S^1/\Gamma$. In this situation we find the following class of two-sided spherical 2-suborbifolds adapted to the graph structure. Let $\alpha \subset B - B^{\text{sing}}$ be a properly embedded arc with endpoints in $\partial_0 B - \partial_{\text{refl}}B$. It yields the spherical 2-suborbifold $\Sigma_\alpha \subset S \cup V_0$ obtained by taking $p^{-1}(\alpha) \subset S$ and attaching to it the pair of discal 2-suborbifolds $q^{-1}(\partial\alpha) \subset V_0$. Hence $\Sigma_\alpha \cong S^2/K$ where we extend the action of $K \subset O(2)$ to $\mathbb{R}^3$ using the canonical embedding $O(2) \subset O(3)$. If $\partial_{\text{refl}}B \neq \emptyset$ we get another similar class of spherical 2-suborbifolds by taking embedded arcs $\hat{\alpha}$ connecting a regular boundary point on $\partial_0 B$ to an interior point of a reflector edge. In this case we have $\Sigma_{\hat{\alpha}} \cong S^2/\hat{K}$ where $\hat{K} \subset O(2)$ is an index two extension of K such that the elements in $\hat{K} - K$ switch the poles $(0, 0, \pm 1)$ of S^2.

We note that the discal 3-orbifold D^3/K has a decomposition into a Seifert and a solid toric piece analogous to the decomposition of $S \cup V_0$. Indeed, take an $O(2)$-invariant decomposition of D^3 into a tubular neighbourhood of the equator $S^1 \times \{0\}$ and the complement of this neighbourhood. By dividing out K, one sees that D^3/K is obtained by attaching the cylinder $(D^2/K) \times [-1, 1]$ to a product fibration with fiber S^1/K and base a bigon (topological disc) such that the fibrations of the boundaries match. Thus we may split O along Σ_α, fill in copies of D^3/K into the two spherical boundary components resulting from the splitting to obtain a (possibly disconnected) 3-orbifold O_α, and extend the graph splitting to all of O_α by attaching copies of $(\partial D^2/K) \times [-1, 1]$ to the pieces of T_0. The effect on the base of the Seifert piece is that B is split along α and two bigons are attached along the copies of α. The discal 3-orbifold $D^3/\hat{K}$ is obtained analogously by attaching the cylinder quotient $(D^2 \times [-1, 1])/\hat{K}$, where $\hat{K}/K$ acts on $[-1, 1]$ by a reflection, to a Seifert manifold with base orbifold a triangular disc Δ whose boundary consists of two boundary arcs and one reflector boundary arc. When performing the connected sum decomposition (surgery) of O along $\Sigma_{\hat{\alpha}}$ and extending the graph splitting over $O_{\hat{\alpha}}$, the effect on the base of the Seifert piece is that B is again split along $\hat{\alpha}$, but this time copies of Δ are attached along the copies of $\hat{\alpha}$.

If $\mathcal{A}$ is a finite system of disjoint properly embedded arcs α_i and $\hat{\alpha}_j$ in B as above, we denote the result of performing simultaneous surgeries along the system of spherical 2-suborbifolds Σ_{α_i} and $\Sigma_{\hat{\alpha}_j}$ by $O_{\mathcal{A}}$ and equip it with an induced graph splitting along a toric family $\mathcal{T}_{\mathcal{A}}$ as explained. The components of $\mathcal{T}$ different from T_0 correspond to components of $\mathcal{T}_{\mathcal{A}}$, whereas T_0 may split up into several components. Furthermore, we denote by $S_{\mathcal{A}} \subset O_{\mathcal{A}}$ the Seifert suborbifold corresponding to S and by $B_{\mathcal{A}}$ the base of its Seifert fibration.

Now we choose the system of arcs $\mathcal{A}$ so that they split B into pieces as simple as possible. After making a suitable choice, every connected component B' of $B_{\mathcal{A}}$ is diffeomorphic to one of the compact 2-orbifolds in the following list:

- Ann^2 or the quadrangle Q bounded by two boundary and two reflector edges occuring in alternating order (a quotient of Ann^2 by an involution);
- $D^2(p)$ or $V^2(p)$ with $p \geq 1$.

Lemma 2.3.2. *Suppose that the compact 3-orbifold O' has a decomposition $O' = S' \cup V'$ along a toric 2-suborbifold $T' = S' \cap V'$ into a Seifert piece S' and a solid toric piece V' such that the discal cross sections of V' fill in Seifert fibers of S'. If B' belongs to the above list, then O' is either spherical or solid toric.*

Proof. If B' splits as the product of the compact interval and a connected closed 1-orbifold, i.e. if B' is the annulus or the quadrangle, then $O' = S' \cup V' \cong V'$ is solid toric.

If B' is discal, then S' is also solid toric and has the form $S' \cong (D^2 \times F')/\Gamma'$ with faithful action $\Gamma' \curvearrowright D^2$ and generic Seifert fiber F'. Let Δ' denote the discal cross section of V'. Using an identification $\partial\Delta' \cong F'$, we form the closed 3-orbifold $\hat{O}'$ by gluing $D^2 \times F'$ and $\partial D^2 \times \Delta'$ canonically along their boundaries. We extend the Γ'-action from $D^2 \times F'$ to $\hat{O}'$ by choosing an extension of the Γ'-action on F' to Δ'. There exists a Γ'-invariant spherical structure on $\hat{O}'$. Hence $O' \cong \hat{O}'/\Gamma'$ is spherical. □

Proof of Theorem 2.3.1 continued. As a consequence of Lemma 2.3.2, the splitting of $O_{\mathcal{A}}$ along the subfamily of $\mathcal{T}_{\mathcal{A}}$ consisting of those toric 2-suborbifolds, which correspond to the components of $\mathcal{T}$ different from T_0, is still a graph splitting. (The pieces resulting from splitting V_0 swallow the corresponding adjacent pieces of S.)

Our discussion yields so far: O is irreducible and satisfies the conclusion of the theorem, or O is closed and admits an $S^2 \times \mathbb{R}$-structure, or O can be decomposed by surgery into components which admit graph splittings along strictly fewer toric 2-suborbifolds. By repeating this process finitely many times, it follows that O can be decomposed by surgery into irreducible orbifolds satisfying the conclusion of the theorem. In particular, the assertion holds if O is not irreducible. If O is irreducible, then O is diffeomorphic to one of the components arising from the surgery, and the assertion holds as well. □

Corollary 2.3.3. *A compact connected 3-orbifold with toric boundary which contains no bad 2-suborbifolds and admits a graph splitting is decomposable into geometric pieces (i.e. satisfies Thurston's Geometrization Conjecture).*

Remark 2.3.4. Let O be as in Theorem 2.3.1. Then the boundary of O is incompressible if and only if no solid toric components occur in the surgery decomposition of O. Indeed, suppose that V is a solid toric component. Then there exists a finite family of disjoint embedded discal 3-suborbifolds $B_i \subset V$ such that $V - \cup_i B_i$ embeds into O. There exists a compression disc for ∂V in V avoiding the B_i, and hence a compression disc for $\partial V \subset \partial O$ in O. Conversely, suppose that ∂O is compressible and consider a compression disc Δ. Using the property that O contains no bad 2-suborbifolds, we can make Δ step by step disjoint from the family of spherical 2-suborbifolds along which the surgery is performed until Δ is contained in an irreducible component. This component must be solid toric.

3. Geometric properties of Riemannian orbifolds

Definition 3.0.1. A *Riemannian orbifold* is a smooth orbifold with the additional structure of (smooth) Riemannian metrics on the local uniformizations $\tilde{U}$ of its orbifold charts. These metrics must be compatible in the sense that the operations $\Gamma_U \curvearrowright \tilde{U}$ and the local coordinate changes $\tilde{\phi}$ are Riemannian isometries.

We also define Riemannian orbifolds of lower ($\mathcal{C}^k$) regularity as $\mathcal{C}^k$ orbifolds with $\mathcal{C}^k$-differentiable Riemannian metrics on the local uniformizations $\tilde{U}$ satisfying the compatibility conditions in Definition 3.0.1. Again, we will only discuss smooth Riemannian orbifolds throughout this section, but our results also hold for $\mathcal{C}^k$ Riemannian orbifolds provided that k is sufficiently large.

On a Riemannian orbifold, there is a canonical representative for the local group $\Gamma_x \subset O(n)$, namely the action of $d\,\mathrm{Stab}_{\Gamma_U}(\tilde{x})$ on the metric vector space $T_{\tilde{x}}\tilde{U}$ for an orbifold chart $(U, \tilde{U}, \Gamma_U, \pi_U)$ with $x \in U$ and $\tilde{x} \in \pi_U^{-1}(x)$. For a Riemannian orbifold O of dimension $n \geq 1$, the tangent space T_xO at a point x is equipped with the angle metric. It is the (complete euclidean) cone over the unit tangent "sphere" Σ_xO which is also called the *link* or the *space of directions* at x. There are canonical isometric identifications $T_xO \cong \mathbb{R}^d/\Gamma_x$ and $\Sigma_xO \cong S^{d-1}/\Gamma_x$. We have the following formula for angles: Let $u, v \subset T_xO$ be two directions at $x \in O$, $\tilde{x} \in \tilde{U}$ a lift of x to a local uniformization $\tilde{U}$ and $\tilde{u}, \tilde{v} \in T_{\tilde{x}}\tilde{U}$ lifts of u and v respectively. Then

$$\angle_x(u, v) = \min_{\gamma \in \Gamma_x} \angle_{\tilde{x}}(\tilde{u}, \gamma \cdot \tilde{v}). \qquad (3.0.2)$$

Using a partition of unity, we can give to every smooth orbifold the additional structure of a Riemannian orbifold. This construction proceeds exactly as in the manifold case. If $f : O \to O'$ is a smooth immersion of orbifolds, one can pull back a Riemannian metric g' on O' to a Riemannian metric f^*g' on O. (More precisely, we pull back the metrics on the local uniformizations. Alternatively, we can consider the orbifold Riemannian metric as a singular bundle metric g on the orbifold tangent bundle.)

By definition, spherical, toric, and more generally geometric orbifolds carry a natural structure as Riemannian orbifolds. Similarly, discal and solid toric orbifolds have natural Riemannian structures.

We say that a Riemannian orbifolds satisfies certain *curvature bounds* (e.g. upper or lower bounds on $|\nabla^l R|$, sec, Ric or scal) if these bounds are satisfied on all orbifolds charts. Note that for a Riemannian orbifold, expressions like $R(u,v)w$ for $u,v,w \in TO$ are in general not well-defined.

Given a closed 2-orbifold O, we remove discal neighbourhoods of the singular vertices in $O^{(0)}$ and apply the Gauss-Bonnet formula to the uniformizations of the resulting components. We thus obtain the *Gauss-Bonnet formula for orbifolds* $\chi(O) = \int_O K dA$ (cf. [Th78, Sec. 13.3]). Since one can easily construct Riemannian metrics with $K \geq 0$ on all bad 2-orbifolds, we can rephrase Proposition 2.1.8 as follows: A closed 2-orbifold admits a Riemannian metric with $K \geq 0$ if and only if it is spheric, toric or bad.

3.1 Geodesics and exponential map

Let O be a Riemannian orbifold and consider a smooth curve c in O (i.e. a smooth immersion of orbifolds $c : [0,1] \to O$). By lifting c to smooth curves in local uniformizations of O, we can compute its length $L(c)$. Similarly, we can determine the length of piecewise smooth curves in O. The corresponding *path metric* gives O a natural structure of a metric *length space*. (A standard reference on length spaces is [BBI01]; for a very thorough discussion of Riemannian orbifolds as metric length spaces see [Bo92].)

A curve in a Riemannian orbifold O is called a *geodesic* if it can be locally lifted to geodesics in the orbifold charts of O. For a tangent vector $v \in TO$ there is a unique maximal geodesic $c_v : (\alpha_v, \omega_v) \to O$, $-\infty \leq \alpha_v < 0 < \omega_v \leq \infty$, with initial condition $\dot{c}_v(0) = v$.

By the convexity of sufficiently small balls in Riemannian manifolds, a geodesic c in a Riemannian orbifold is always (distance) *minimizing* on $[0,\epsilon)$ and $(-\epsilon, 0]$ for some $\epsilon > 0$. However, it is in general not locally minimizing, i.e. minimizing on $(-\epsilon,\epsilon)$ for some positive ϵ. Indeed, it is locally minimizing at $t = 0$ if and only if it is either constant or if $\angle_{c(0)}(\dot{c}(t^+), \dot{c}(t^-)) = \pi$, equivalently, if and only if c remains in the same stratum near 0. (Intuitively, if a geodesic crosses a stratum, "short cuts" are possible.) Similarly, geodesics are in general not suborbifolds: They are locally suborbifolds at $t = 0$ if and only if for a local lift $\tilde{c}$ we have $\Gamma_{c(0)} \cdot \dot{\tilde{c}} \subseteq \{\pm \dot{\tilde{c}}\}$.

The strata of a Riemannian orbifold O carry natural structures as Riemannian manifolds; they are totally geodesic suborbifolds of O and hence locally convex. However, their metric structures are in general not complete, even if O itself is complete. Note that the tangent space T_xS_x of the stratum through $x \in O$ is the maximal euclidean factor of T_xO.

A *geodesic variation* of a geodesic $c : [0,1] \to O$ is a smooth map $\alpha : (-\epsilon, \epsilon) \times [0,1] \to O$, $\epsilon > 0$, such that $\alpha(0,\cdot) = c$ and all $\alpha(s,\cdot)$ are geodesics. The vector field $J = \frac{\partial\alpha}{\partial s}(0^+,\cdot) : [0,1] \to TO$ along c is a *Jacobi field*. The relevant constructions and computations can be adopted from Riemannian manifolds by working in chains of orbifold charts along c. For instance, after choosing local lifts $\tilde\alpha$ of α near the endpoints of c, the first variation formula takes the form

$$\frac{d}{ds}|_{s=0^+} L(\alpha(s,\cdot)) = -|J(0)| \cdot \cos \angle_{\tilde c(0)}(\tilde J(0), \dot{\tilde c}(0^+)) - |J(1)| \cdot \cos \angle_{\tilde c(1)}(\tilde J(1), -\dot{\tilde c}(1^-)). \quad (3.1.1)$$

By our formula for angles 3.0.2, we obtain the first variation *inequality*

$$\frac{d}{ds}|_{s=0^+} L(\alpha(s,\cdot)) \geq -|J(0)| \cdot \cos \angle_{c(0)}(J(0), \dot c(0^+)) - |J(1)| \cdot \cos \angle_{c(1)}(J(1), -\dot c(1^-)). \quad (3.1.2)$$

Note that equality always holds if c remains in one stratum near each of its endpoints.

The endpoints $p = c(0)$ and $q = c(1)$ of c are said to be *conjugate* (along c) if there exists a non-zero Jacobi field J along c with boundary values zero. If p and q are not conjugate, then Jacobi fields J exist for any given boundary values $J(0) \in T_pO$ and $J(1) \in T_qO$. Unlike in the manifold case, due to the ambiguity of lifts to charts, there may exist several Jacobi fields with the same boundary values, however at most finitely many. This occurs if p or q is singular. More precisely, suppose that β_i are (not necessarily minimizing) geodesic segments of length $< \pi$ in $\Sigma_{c(i)}O$, starting in $\dot c(0^+)$, respectively, $-\dot c(1^-)$ and that $\lambda_0, \lambda_1 \geq 0$. Then there exists a geodesic variation α of c with the following property: If $\lambda_0 > 0$, then for a local lift $\tilde\alpha$ of α near $c(0)$ holds $|\tilde J(0)| = \lambda_0$, $\angle_{\tilde c(0)}(\tilde J(0), \dot{\tilde c}(0^+)) < \pi$, and the geodesic segment $\tilde\beta_0$ in $\Sigma_{\tilde c(0)}$ connecting $\dot{\tilde c}(0^+)$ to $\lambda_0^{-1}\tilde J(0)$ projects via the differential of the chart projection to β_0. We can proceed analogously if $\lambda_1 > 0$. In particular, given boundary values $J(0)$ and $J(1)$, the Jacobi field J and a corresponding geodesic variation α can be chosen so that equality holds in (3.1.2).

We will mostly study Riemannian orbifolds O which are *complete* with respect to the path metric, equivalently (by Hopf-Rinow), which are *geodesically complete*, i.e. on which the geodesics c_v are defined on all of $\mathbb{R}$. In this case, closed metric balls are compact and the path metric is geodesic.

On a complete Riemannian orbifold, the *exponential map* $\exp : TO \to O, v \mapsto c_v(1)$ is globally defined, and we will use the notation $\exp_x := \exp|_{T_xO} : T_xO \to O$.

The *injectivity radius* $\mathrm{inj}(x)$ in x is defined as the maximal radius $r \in (0,\infty]$ such that $exp_x|_{B_r^{T_xO}(0)}$ is an embedding. The *conjugate radius* $\mathrm{conj}(x)$ in x is defined as the distance to the closest conjugate point, i.e. the minimal norm of a tangent vector $v \in T_xO$ such that the endpoints of $c_x|_{[0,1]}$ are conjugate to each other. One has $0 < \mathrm{inj}(x) \leq \mathrm{conj}(x) \leq \infty$

(due to the Jacobi Lemma). Lower bounds for $\mathrm{conj}(x)$ and local bilipschitz bounds for $\exp_x$ inside the conjugate radius can be given in terms of local upper sectional curvature bounds (Rauch estimates). We call an open metric ball $B_r(x) \subset O$ a *standard ball* if $r \leq \mathrm{inj}(x)$.

Regarding the continuity properties of the injectivity radius on a complete Riemannian orbifold, we have

Proposition 3.1.3. *The funtion* $\mathrm{inj} : O \to (0, \infty]$ *is upper semicontinuous on* O *and continuous on every stratum* $O^{(d)}$.

Proof. Since the limit of minimizing geodesic segments is minimizing, inj is clearly upper semicontinuous on O. The continuity on strata follows from the continuity of the injectivity radius on Riemannian manifolds because the strata are locally convex. □

For a point $x \in O$, we define $D_x \subseteq T_xO$ as the subset of tangent vectors v in x such that $c_v|_{[0,1+\epsilon)}$ is minimizing for some $\epsilon > 0$. As in the smooth case, it is an open star-shaped subset containing the origin, and $exp_x|_{D_x}$ is an embedding. The *tangential cut locus* with respect to x is defined as $\partial D_x \subset T_xO$ and its image $\mathrm{Cut}(x) := \exp_x(\partial D_x) \subset O$ is called the *cut locus* with respect to x. A point y belongs to the *cut locus* of x, $y \in \mathrm{Cut}(x)$ if there exists no minimizing geodesic from x *through* y, i.e. if there exists no $v \in T_xO$ such that $y = \exp_x(v)$ and $c_v|_{[0,1+\epsilon)}$ is minimizing for some $\epsilon > 0$.

If $\mathrm{inj}(x) < \infty$, then $\exp_x$ is not injective on $\partial B_{\mathrm{inj}(x)}(0)$ or $\exp_x$ is not locally injective near some $v \in \partial B_{\mathrm{inj}(x)}(0)$. In either case there exist cut points in $\partial B_{\mathrm{inj}(x)}(x)$ and hence $\mathrm{inj}(x) = d(x, \mathrm{Cut}(x))$. We will discuss this in more detail in section 3.3 below.

We observe that the well-known proof of Bishop-Gromov *volume comparison* for manifolds (see e.g. [Ka89, 1.9.4]) can also be applied to Riemannian orbifolds with lower Ricci curvature bounds. Thus, we have (cf. [Lu01, sec. 1.2])

Proposition 3.1.4 (Bishop-Gromov for orbifolds). *Let* O^n *be an* n*-dimensional Riemannian orbifold with* $\mathrm{Ric} \geq (n-1)k$. *Then for each point* $x \in O$, *the quotient function* $vol\, B_r^O(x)/v_k^n(r)$ *is nondecreasing in* r *with* $\lim_{r\to 0^+} vol\, B_r^O(x)/v_k^n(r) = |\Gamma_x|^{-1}$. *Here,* $v_k^n(r)$ *is the volume of the* r*-ball in the simply-connected* n*-dimensional space form of constant curvature* k.

3.2 Spherical orbifolds

We will derive the properties of complete spherical orbifolds which are relevant to our purposes from the point of view of isometric group actions on unit spheres. Let $\Gamma \subset O(d)$, $n \geq 1$, be a finite subgroup and consider the complete spherical orbifold $L^{n-1} = S^{n-1}/\Gamma$. We will denote the path metrics on L and S^{n-1} by $\angle(\cdot,\cdot)$. Points $v \in L$ correspond to

Γ-orbits $\Gamma \cdot \tilde{v} \subset S^{n-1}$. The *Dirichlet fundamental domain* $D = D(\Gamma, \tilde{v}) \subseteq S^{n-1}$ around $\tilde{v}$ is the closure of the set of points which are closer to $\tilde{v}$ than to the other points of the orbit $\Gamma \cdot \tilde{v}$. Either $D = S^{n-1}$ or D is a finite intersection of hemispheres, i.e. a convex spherical polyhedron. In the latter case, ∂D projects onto $\text{Cut}(v)$ and the interior of D is identified with $L - \text{Cut}(v)$. It follows that the cut locus is polyhedral and L is obtained from D by gluing the boundary to itself according to isometric face pairings. Clearly, $\text{inj}(v)$ equals the inradius of D with respect to $\tilde{v}$ (i.e. the radius of the largest ball around $\tilde{v}$ contained in D) and $\text{rad}(L, v)$ equals the circumradius. Note that $\text{conj} \equiv \pi$ on L.

The fixed point set $\text{Fix}(\Gamma)$ is a great sphere or empty, and so is its orthogonal complement $\text{Fix}(\Gamma)^\perp = \{v \in S^{n-1} | \angle(v, w) = \frac{\pi}{2} \ \forall w \in \text{Fix}(\Gamma)\}$. The spherical join splitting

$$S^{n-1} = \text{Fix}(\Gamma) \circ \text{Fix}(\Gamma)^\perp$$

is Γ-invariant and descends to the spherical join splitting

$$L \cong \text{Fix}(\Gamma) \circ L' \tag{3.2.1}$$

with $L' \cong \text{Fix}(\Gamma)^\perp / \Gamma$. The unit sphere $\text{Fix}(\Gamma)$ is the *spherical de Rham factor* of L. (If one of the join factors is empty then L coincides with the other one.)

Since the action of Γ on the unit sphere $\text{Fix}(\Gamma)^\perp$ is fixed point free, no orbit can be contained in a ball of radius $< \frac{\pi}{2}$ because otherwise it would have a well-defined center fixed by Γ. Thus every orbit has Hausdorff distance $\leq \frac{\pi}{2}$ from $\text{Fix}(\Gamma)^\perp$. It follows that

$$\text{diam}(L') \leq \frac{\pi}{2}. \tag{3.2.2}$$

In particular, if $\text{diam}(L) > \frac{\pi}{2}$ then $\text{diam}(L) = \pi$. Furthermore, $\text{rad}(L, v) = \pi$ if and only if $v \in \text{Fix}(\Gamma)$, and in this case $\text{inj}(v) = \pi$.

If $\text{rad}(L, v) < \pi$, then $\text{inj}(v) \leq \frac{\pi}{2}$ with equality if and only if $D(\Gamma, \tilde{v})$ is the hemisphere with center $\tilde{v}$. In the equality case $\text{inj}(v) = \frac{\pi}{2}$ we also have $\text{rad}(L, v) = \frac{\pi}{2}$ and every geodesic starting from v remains distance minimizing precisely up to distance $\frac{\pi}{2}$. Moreover, $\Gamma \cdot \tilde{v} = \{\tilde{v}, -\tilde{v}\}$ and $v \in L'$.

If O is a Riemannian orbifold and $L = \Sigma_x O$ the space of directions in a point x, then the splitting (3.2.1) becomes

$$\Sigma_x O \cong \Sigma_x S_x \circ N_x \tag{3.2.3}$$

where N_x denotes the *normal space* to the stratum S_x of O through x, i.e. the space of directions at x orthogonal to S_x. Note that $N_x = \emptyset$ if and only if the point x is regular, and $\Sigma_x S_x = \emptyset$, equivalently, $\text{diam}(\Sigma_x O) \leq \frac{\pi}{2}$ if and only if x is a singular vertex.

For $v \in \Sigma_x O$ with $\text{inj}(v) = \text{rad}(\Sigma_x O, v) = \frac{\pi}{2}$ our discussion above shows that $v \in N_x$ and the geodesic c_v is *reflected* at x, $c_v(-t) = c_v(t)$.

Remark 3.2.4. A finite group has finitely many irreducible orthogonal representations. As a consequence, there are finitely many isometry classes of complete spherical d-orbifolds with a given lower volume bound.

Thus, given a bound on the order of local isotropy groups, there are only finitely many possible types of strata in the stratification of O.

The following technical result will be needed later.

Lemma 3.2.5. *For $a > 0$ there exists $\epsilon = \epsilon(n, a) > 0$ such that: Let $v_1, v_2 \in L$ and suppose that $L - B_{\frac{\pi}{2}-\epsilon}(v_1) - B_{\frac{\pi}{2}-\epsilon}(v_2)$ has volume $\geq a$. Then $\angle(v_i, L') \geq \epsilon$ and the projections v_i'' of the v_i to* $\mathrm{Fix}(\Gamma)$ *satisfy $\angle(v_1'', v_2'') \leq \pi - 2\epsilon$.*

In particular, the spherical de Rham factor $\mathrm{Fix}(\Gamma)$ *of L is nontrivial and* $\mathrm{diam}(L) = \pi$.

Proof. According to Proposition 3.1.4, the subset $B_{\frac{\pi}{2}+\epsilon}(v_i) - B_{\frac{\pi}{2}-\epsilon}(v_i)$ has volume $\leq 2C(d) \cdot ((1 - \frac{2\epsilon}{\pi})^{-(n-1)} - 1)$. Hence, if $\epsilon = \epsilon(d, a) > 0$ is chosen sufficiently small, then the hypotheses of the lemma imply that $L \neq B_{\frac{\pi}{2}+\epsilon}(v_1) \cup B_{\frac{\pi}{2}+\epsilon}(v_2)$. By (3.2.2), for every $v' \in L'$ we have $\mathrm{rad}(L, v') = \frac{\pi}{2}$. It follows that $\angle(v_i, L') \geq \epsilon$.

Using again (3.2.2), we note that $B_{\frac{\pi}{2}+\epsilon}(v_i'') \circ L' \subseteq B_{\frac{\pi}{2}+\epsilon}(v_i)$. (An elementary calculation in $S^1 \circ S^1 \cong S^3 \subset \mathbb{C}^2$ shows that the spherical join of two arcs with lengths $\frac{\pi}{2} + \epsilon$ and $\frac{\pi}{2}$ has diameter $\frac{\pi}{2} + \epsilon$.) Therefore $\mathrm{Fix}(\Gamma) \neq B_{\frac{\pi}{2}+\epsilon}(v_1'') \cup B_{\frac{\pi}{2}+\epsilon}(v_2'')$. This in turn implies that $\angle(v_1'', v_2'') \leq \pi - 2\epsilon$. □

3.3 Geometry at closest cut points

Let O be a complete Riemannian orbifold and $x \in O$ a point. In order to better understand how $\exp_x$ can fail to be injective at radius $\mathrm{inj}(x)$, we begin with the following two observations:

Lemma 3.3.1. *Suppose that $d(x, y) = \mathrm{inj}(x) < \mathrm{conj}(x)$ and that $c_{v_i}|_{[0,1]}$, $i = 1, 2$, are two different minimizing geodesics from x to y. Then they fit together to form a geodesic loop of length $2\,\mathrm{inj}(x)$ based at x with midpoint y which is locally minimizing also at y.*

Proof. Let us denote by $u_i := -\dot{c}_{v_i}(1^-) \in \Sigma_y O$ the directions of the segments at y. Since y is not conjugate to x along c_{v_i}, there exist geodesic variations $\alpha_i(s, t)$ of $c_{v_i}|_{[0,1]}$ which fix the endpoint x and for which $\frac{\partial}{\partial s}|_{s=0^+}\alpha_1(\cdot, 1) = \frac{\partial}{\partial s}|_{s=0^+}\alpha_2(\cdot, 1) \in \Sigma_y O$ is equal to a midpoint (angle bisector) w of u_1 and u_2. The α_i may be chosen so that equality holds in the first variation inequality (3.1.2). Therefore, if $\angle_y(u_1, u_2) < \pi$, these variations are shortening and $\exp_x$ is not injective on the open $d(x, y)$-ball around the origin. This implies $\mathrm{inj}(x) < d(x, y)$ which is absurd. □

Lemma 3.3.2. *Suppose that* $0 < d(x,y) = \mathrm{inj}(x) < \mathrm{conj}(x)$ *and that there exists a unique minimizing geodesic* $c_v|_{[0,1]}$ *from* x *to* y. *With* $\overrightarrow{qp} := -\dot{c}_v(1^-)$ *we have either*

(i) $\overrightarrow{yx} \in T_y S_y$ *and* $\exp_x$ *is locally injective at* v, *or*

(ii) $\mathrm{inj}(\overrightarrow{yx}) = \mathrm{rad}(\Sigma_y O, \overrightarrow{yx}) = \frac{\pi}{2}$ *and* c_v *is reflected at* y, *i.e.* $c_v(1-t) = c_v(1+t)$.

Note that in case (ii) the geodesic $c_v|_{[0,2]}$ *is also a geodesic* loop *of length* $2\,\mathrm{inj}(p)$ *based at* p.

Proof. We assume that $\mathrm{inj}(\overrightarrow{yx}) < \frac{\pi}{2}$ in $\Sigma_y O$. Then $\exp_{\overrightarrow{yx}}^{\Sigma_y O}$ is not injective on the open $\frac{\pi}{2}$-ball around the origin, i.e. there exist two different geodesic segments β_0 and β_1 in $\Sigma_y O$ with equal length $l < \frac{\pi}{2}$ from $\overrightarrow{yx}$ to a direction u with $\angle(\overrightarrow{yx}, u) < \frac{\pi}{2}$. (The segments β_i could be chosen minimizing.) Since y is not conjugate to x along c_v, it follows (cf. section 3.1) that there exist two different geodesic variations $(\alpha_i(s,t))$ of $c_v|_{[0,1]}$ fixing the endpoint x with $\alpha_1(,\cdot,1) = \alpha_2(,\cdot,1)$ and $\frac{\partial}{\partial s}|_{s=0^+}\alpha_1(\cdot,1) = \frac{\partial}{\partial s}|_{s=0^+}\alpha_2(\cdot,1) = u$. They may be chosen so that for local lifts $\tilde{\alpha}_i$ of the α_i near y the geodesic segments $\tilde{\beta}_i$ in $\Sigma_{\tilde{y}}$ connecting $\overrightarrow{\tilde{y}\tilde{x}}$ to $\tilde{J}_i(1)$ project to β_i. (This distinguishes them.) By (3.1.2), the variations α_i shorten $c_v|_{[0,1]}$. Hence, $\exp_x$ is not injective on the open $d(x,y)$-ball around the origin, a contradiction. This shows that $\mathrm{inj}(\overrightarrow{yx}) \geq \frac{\pi}{2}$.

If $\mathrm{inj}(\overrightarrow{yx}) = \pi$, then $\overrightarrow{yx} \in T_y S_y$ and c_v is contained in S_y for t close to 1. Since y is not conjugate to x, in this case $\exp_x$ is locally injective in v. On the other hand, if $\frac{\pi}{2} \leq \mathrm{inj}(\overrightarrow{yx}) < \pi$, then our discussion of complete spherical orbifolds in section 3.2 implies that alternative (ii) holds. □

Suppose now that $\mathrm{inj}(x) < \infty$. Then, as already mentioned, $\exp_x$ is not injective on $\partial B_{\mathrm{inj}(x)}(0)$ or $\exp_x$ is not locally injective near some $v \in \partial B_{\mathrm{inj}(x)}(0)$. From the previous two lemmas we conclude

Proposition 3.3.3 (Closest cut points). *If* $\mathrm{inj}(x) < \mathrm{conj}(x)$ *and if* $y \in \mathrm{Cut}(x)$ *with* $d(x,y) = \mathrm{inj}(x)$, *then there exists a geodesic loop* λ *of length* $2\,\mathrm{inj}(x)$ *based at* x. *The midpoint* y *of* λ *divides it into two minimizing segments* σ_1 *and* σ_2 *from* x *to* y. *Either* $\angle_y(\sigma_1,\sigma_2) = \pi$, *or* $\sigma_1 = \sigma_2$ *and it is the unique minimizing segment connecting* x *and* y. *In the latter case holds* $\mathrm{inj}(\overrightarrow{yx}) = \mathrm{rad}(\Sigma_y O, \overrightarrow{yx}) = \frac{\pi}{2}$.

Proof. Let $c_v|_{[0,1]}$ be a minimizing geodesic from y to x. If $\exp_x$ is locally injective in v, then there must exist another minimizing geodesic from x to y because $c_v|_{[0,1+\epsilon]}$ is not minimizing for any $\epsilon > 0$. Lemma 3.3.1 applies and we are in the case when $\angle_y(\sigma_1,\sigma_2) = \pi$. On the other hand, if there is a unique minimizing geodesic $\overline{yx}$, then $\exp_x$ cannot be locally injective in v and we are in case (ii) of Lemma 3.3.2, i.e. $\sigma_1 = \sigma_2$. □

Remark 3.3.4. If λ is a geodesic loop based at x (not necessarily with minimizing halves), then $\mathrm{inj}(x) \leq \frac{1}{2}L(\lambda)$.

3.4 Toponogov's Theorem for orbifolds

Throughout this section, we fix a *comparison curvature value* $k \in \mathbb{R}$. We denote by (M_k^2, d_k) the simply-connected 2-dimensional model space of constant curvature k with its standard metric.

Let (X, d) be a complete metric length space and consider inside X a *triangle* consisting of three points $x_1, x_2, x_3 \in X$ and three minimizing segments $|x_1x_2|$, $|x_2x_3|$ and $|x_1x_3|$. We also require that the circumference of a triangle (the sum of the lengths of its sides) be $\leq 2\pi k^{-\frac{1}{2}}$. (We consider this condition to be trivially satisfied if $k \leq 0$.)

For such a triangle, there is a *comparison triangle* in M_k^2 consisting of three points $\bar{x}_1, \bar{x}_2, \bar{x}_3 \in M_k^2$ such that $d_k(\bar{x}_i, \bar{x}_j) = d(x_i, x_j)$. These points and the minimizing geodesics $|\bar{x}_i\bar{x}_j|$ between them are determined up to isometry unless we have $k > 0$ and (after permuting indices if necessary) $d(x_1, x_2) = d(x_2, x_3) + d(x_1, x_3) = \pi k^{-\frac{1}{2}}$. In this case, we arrange the comparison triangle such that $\bar{x}_3 \in |\bar{x}_1\bar{x}_2|$. The angle at $\bar{x}_1$ between the geodesics $|\bar{x}_1\bar{x}_2|$ and $|\bar{x}_1\bar{x}_3|$ is therefore well-defined; we call it the *comparison angle* (with respect to the comparison curvature value k) at x_1 between x_2 and x_3 and denote it by $\tilde{\angle}_{x_1}(x_2, x_3)$. We define $\tilde{\angle}_{x_2}(x_1, x_3)$ and $\tilde{\angle}_{x_3}(x_1, x_2)$ analogously.

We recall that for a metric length space (X, d) and two minimizing segments c_1, c_2 emanating from a point $x \in X$, the angle at x between c_1 and c_2 is determined by

$$\cos \angle_x(c_1, c_2) = \lim_{s,t \to 0} \frac{s^2 + t^2 - d(c_1(t), c_2(t))^2}{2st}$$

if this limit exists. If c_1 and c_2 end at two points y_1 and y_2 respectively, we sometimes write $\angle_x(y_1, y_2) := \angle_x(c_1, c_2)$ although the value of this expression also depends on c_1 and c_2.

For a complete metric length space (X, d), the following three conditions are equivalent (cf. [BGP92, 2.7], [BBI01, Thm. 4.3.5]). They all express the notion that triangles are "thicker" than in M_k^2.

(i) For every triangle with vertices $x_1, x_2, x_3 \in X$ and $y \in |x_1x_2|$, one has $d(x_3, y) \geq d_k(\bar{x}_3, \bar{y})$ where $\bar{y} \in |\bar{x}_1\bar{x}_2|$ is the point on the comparison triangle determined by $d(x_1, y) = d_k(\bar{x}_1, \bar{y})$.

(ii) For any two minimizing segments c_1, c_2 emanating from some point $x \in X$, the function $(s, t) \mapsto \tilde{\angle}_x(c_1(s), c_2(t))$ is well-defined and non-increasing in s and t (when the other variable is unchanged).

(iii) Angles between segments in X are well-defined and complementary in the sense that for two segments $|x_1x_2|$ and $|y_2y_2|$ with $y_1 \in |x_1x_2|$ we have $\angle_{y_1}(x_1, y_2) + \angle_{y_1}(x_2, y_2) = \pi$. Moreover, for every triangle with vertices $x_1, x_2, x_3 \in X$ we have $\tilde{\angle}_{x_3}(x_1, x_2) \leq \angle_{x_3}(x_1, x_2)$.

Definition 3.4.1. A complete metric length space (X, d) is an *Alexandrov space* of curvature $\geq k$ if it satisfies conditions (i) - (iii) above.

By Perelman's version of Toponogov's Theorem [BGP92, Thm. 3.2], a complete metric length space (X, d) is already Alexandrov (of curvature $\geq k$) if conditions (i) - (iii) hold locally, i.e. if every point $x \in X$ there is a neighbourhood $x \in U_x \in X$ such that conditions (i) - (iii) hold on U_x. A (not necessarily complete) metric length space satisfying this condition is also called a *local Alexandrov space.*

The classical version of Toponogov's Theorem (see e.g. [Ka89, 4.2]) asserts that a (complete) Riemannian manifold with sectional curvature $\geq k$ is an Alexandrov space of curvature $\geq k$. We will now show (cf. [Bo92, Thm. 1]) that the same is also true for Riemannian orbifolds with sectional curvature $\geq k$. Indeed, we will see that such an orbifold is locally "more strongly curved" near its singularities than a Riemannian manifold satisfying the same curvature condition (such as its local uniformization). Intuitively speaking, positive curvature is "concentrated" in the singularities.

Proposition 3.4.2. *A complete Riemannian orbifold with sectional curvature $\geq k$ is an Alexandrov space of curvature $\geq k$.*

Proof. Let O be a Riemannian orbifold with sectional curvature $\geq k$. By our discussion above, it is sufficient to verify that condition (iii) holds locally on O. Angles between segments in O exist and are complementary by our discussion in section 3.1.

For $x \in O$, let U be a standard ball around x with (finite) cover $\tilde{U}$. We choose $U_x \subset U$ sufficiently small such that all minimizing geodesics between points in U_x are entirely contained in U. Consider a triangle with edges $x_1, x_2, x_3 \in U_x$. For some lift $\tilde{x}_3$ of x_3 in $\tilde{U}$, we lift the geodesics $c_1 = |x_3x_1|$ and $c_2 = |x_3x_2|$ to geodesics $\tilde{c}_1$ and $\tilde{c}_2$ in $\tilde{U}$ starting from $\tilde{x}_3$ such that $\angle_{\tilde{x}_3}(\dot{\tilde{c}}_1(0), \dot{\tilde{c}}_2(0)) = \angle_{x_3}(\dot{c}_1(0), \dot{c}_2(0))$, equivalently, such that $\angle_{\tilde{x}_3}(\tilde{x}_1, \tilde{x}_2) = \angle_{x_3}(x_1, x_2)$ where $\tilde{x}_1$ and $\tilde{x}_2$ are the end points of $\tilde{c}_1$ and $\tilde{c}_2$, respectively.

By construction we have $d(x_1, x_2) \leq d(\tilde{x}_1, \tilde{x}_2)$, and hence $\tilde{\angle}_{x_3}(x_1, x_2) \leq \tilde{\angle}_{\tilde{x}_3}(\tilde{x}_1, \tilde{x}_2)$. Applying Toponogov's Theorem to the Riemannian manifold $\tilde{U}$ with sectional curvature $\geq k$, we now conclude that $\tilde{\angle}_{x_3}(x_1, x_2) \leq \tilde{\angle}_{\tilde{x}_3}(\tilde{x}_1, \tilde{x}_2) \leq \angle_{\tilde{x}_3}(\tilde{x}_1, \tilde{x}_2) = \angle_{x_3}(x_1, x_2)$ which completes the proof. □

Remark 3.4.3. The observation in the proof that $d(x_1, x_2) \leq d(\tilde{x}_1, \tilde{x}_2)$ makes clear how an orbifold with sectional curvature $\geq k$ is more strongly curved near its singularities than its local uniformization. If for instance $\tilde{U}$ has constant sectional curvature (so that $\tilde{\angle} = \angle$ on $\tilde{U}$), it is still possible that $\tilde{\angle}_{x_3}(x_1, x_2) < \angle_{x_3}(x_1, x_2)$ for a suitably chosen triangle in U_x.

On a generic Alexandrov space, minimizing segments cannot in general be naturally extended beyonf their end points. Conversely, geodesics on a complete Riemannian orbifold O have natural extensions, and we can also consider triangles with one side a not necessarily minimizing geodesic. Thus, we also allow triangles consisting of three points $x_1, x_2, x_3 \in O$, two minimizing geodesics $c_1 = |x_1x_3|$ and $c_2 = |x_2x_3|$ and a geodesic c connecting x_1 and x_2. However, for a comparison triangle to exist we have to require that $l(c) \leq d(x_1, x_3) + d(x_2, x_3)$ and that the circumference of the triangle be $\leq 2\pi k^{-\frac{1}{2}}$. In this case, we choose the comparison triangle such that $d_k(\bar{x}_1, \bar{x}_2) = l(c)$.

Proposition 3.4.4 (Toponogov for orbifolds). *Let O be a complete Riemannian orbifold with sectional curvature $\geq k$. For a triangle in O consisting of three points $x_1, x_2, x_3 \in O$, two minimizing geodesics $c_1 = |x_1x_3|$ and $c_2 = |x_2x_3|$ and a geodesic c connecting x_1 and x_2, we have $\tilde{\angle}_{x_1}(x_2, x_3) \leq \angle_{x_1}(x_2, x_3)$ and $\tilde{\angle}_{x_2}(x_1, x_3) \leq \angle_{x_2}(x_1, x_3)$.*

Proof. Consider a triangle with corners $x_1, x_2, x_3 \in O$ as in the proposition. We subdivide the (not necessarily minimizing) geodesic c by choosing points $x_1 = y_1, \ldots, y_l = x_2$ such that c is minimizing between y_i and y_{i+1} for all $1 \leq i < l$. (That this is possible follows from our discussion in 3.1.) We also choose minimizing geodesics γ_l from x_3 to the y_l such that $\gamma_1 = c_1$ and $\gamma_l = c_2$.

By applying condition (iii) to the resulting triangles with sides γ_i, γ_{i+1} and a part of c, we obtain $\tilde{\angle}_{y_i}(y_{i+1}, x_3) \leq \angle_{y_i}(y_{i+1}, x_3)$ and $\tilde{\angle}_{y_{i+1}}(y_i, x_3) \leq \angle_{y_{i+1}}(y_i, x_3)$ for all $1 \leq i < l$. In particular, we have

$$\tilde{\angle}_{y_i}(y_{i-1}, x_3) + \tilde{\angle}_{y_i}(y_{i+1}, x_3) \leq \pi \tag{3.4.5}$$

for all $1 < i < l$.

We now arrange the comparison triangles next to each other in M_k^2 such that the geodesics $\bar{\gamma}_i = |\bar{x}_3\bar{y}_i|$ coincide. The resulting gegodesic polygon is convex by (3.4.5); it consists of two geodesics of length $l(c_1)$ and $l(c_2)$, respectively, and a number of geodesics with total length $l(c)$. If we now "straighten out" this polygon to the comparison triangle for our original triangle, we observe that $\tilde{\angle}_{x_1}(x_2, x_3) \leq \tilde{\angle}_{x_1}(y_2, x_3) \leq \angle_{x_1}(\dot{\gamma}(0), \dot{\gamma}_1(0)) = \angle_{x_1}(x_2, x_3)$ and analogously $\tilde{\angle}_{x_2}(x_1, x_3) \leq \angle_{x_2}(x_1, x_3)$. □

3.5 The orbifold Soul Theorem

Throughout this section, we closely follow the well-known proof of the Soul Theorem for manifolds. Let us consider a complete non-compact non-negatively curved orbifold O without boundary. Then for every point $p \in O$ there is a *ray* ρ (an isometric embedding $\rho : [0, \infty) \to O$) emanating from p. For any such ray we define the associated *Busemann function* as $\beta_\rho(x) := \lim_{t\to\infty}(d(\rho(t), x) - t)$ for all $y \in O$. (Because $(d(\rho(t), x) - t)$ is nonincreasing in t and bounded below by $-d(p,x)$, β_ρ is well defined on all of O and 1-Lipschitz.)

By Proposition 3.4.4 the function β_ρ is *concave* along (not necessarily minimizing) geodesics: Suppose there is a geodesic segment $c : [0,1] \to O$ and $s \in (0,1)$ such that $\beta_\rho(c(s)) < s \cdot \beta_\rho(c(1)) + (1-s) \cdot \beta_\rho(c(0))$. For sufficiently large $t > 0$, we choose minimizing segments from $c(0), c(s)$ and $c(t)$ to $\rho(t)$. For these segments we still have $d(c(s), \rho(t)) < s \cdot d(c(1), \rho(t)) + (1-s) \cdot d(c(0), \rho(t))$. On the other hand, we compute for the Euclidean comparison angle that $d_0(\bar{c}(s), \bar{\rho}(t)) \leq s \cdot d_0(\bar{c}(1), \bar{\rho}(t)) + (1-s) \cdot d(\bar{c}(0), \bar{\rho}(t))$. It follows that we have comparison angles (with respect to the comparison curvature value $k = 0$) satisfying $\tilde{\angle}_{c(s)}(c(0), \rho(t)) + \tilde{\angle}_{c(s)}(c(1), \rho(t)) > \pi$, and hence $\angle_{c(s)}(c(0), \rho(t)) + \angle_{c(s)}(c(1), \rho(t)) > \pi$ which is impossible. Thus, the function $\beta := \min_\rho \beta_\rho$ (where the minimum runs over all rays ρ in O with $\rho(0) = p$; this is a compact family of rays) is also well-defined, 1-Lipschitz and concave along geodesics. For every point $x \in O$ there is a ray ρ_x emanating from x such that $\beta(\rho_x(s)) = \beta(x) - s$.

It follows that for all $t \geq 0$ the subsets $C_t := \{\beta \geq t\} \subset O$ are *totally convex* in the sense that for any two points $x_1, x_2 \in C_t$ all geodesic segments between x_1 and x_2 are entirely contained in C_t. Moreover, C_0 is closed and compact (or it would contain another ray emanating from $p \in C_0$), and hence so are all C_t for $t \geq 0$. Hence in any local uniformization $\pi_U : \tilde{U} \to U$ near a point $x \in C_t$, the lift $\pi^{-1}(C_t \cap U) \subset \tilde{U}$ is a totally convex subset of $\tilde{U}$. But locally closed totally convex subsets of Riemannian manifolds are embedded submanifolds with (possibly nonsmooth) boundary locally supported by a cone in the tangential space (see [CE75, 8.6 – 8.8]). Hence the sets C_t are contained in the closure of embedded suborbifolds in O. We set $C^{(1)} := C_{t_{\max}}$ where $t_{\max}$ is the maximal value of t such that C_t is nonempty. If $N^{(1)}$ is the maximal embedded suborbifold in O such that $C^{(1)} = \overline{N^{(1)}}$, we set $\partial C^{(1)} = C^{(1)} - N^{(1)}$. We have $\dim N^{(1)} < \dim O$.

Suppose that $\partial C^{(1)} \neq \emptyset$. Using Rauch comparison, we can proceed as in the manifold case (see [CE75, 8.9 – 8.10]) to show that the sets $C_a := \{x \in C^{(1)} \mid d(x, \partial C^{(1)}) \geq a\}$ are also totally convex closed compact subsets of O. We set $C^{(2)} = C_{a_{\max}}$ where again $a_{\max}$ is the maximal value of a such that C_a is nonempty. If $N^{(2)}$ is the maximal embedded suborbifold in O such that $C^{(2)} = \overline{N^{(2)}}$, we have $\dim N^{(2)} < \dim N^{(1)}$. After repeating this process a finite number of times if necessary, we arrive at a totally convex, totally geodesic compact suborbifold S without boundary such that $\dim S < \dim O$, a so-called *soul* of O. Moreover, the distance function $d(S, \cdot)$ has no *critical points* on $O - S$: Every point $x \in O - S$ is contained in ∂C for some totally convex compact subset of O (either one of the C_t or one of the C_a), and therefore a lift of x to a local uniformization is contained in the boundary of a totally convex subset of a Riemannian manifold. Again from the study of such sets, we can find near x a vector field making angle $> \frac{\pi}{2}$ with all minimizing geodesics to S (cf. [CE75, 8.8]).

The normal bundle $\nu(S)$ of the soul S is diffeomorphic to a small tubular neighbourhood of $B_\epsilon(S) \subset O$. On the other hand, we can construct a smooth *gradient-like* vector field X on $O - B_{\epsilon/2}(S)$ such that $d(S,\cdot)$ decreases along the integral curves of X at least linearly (i.e. has uniform negative directional derivative): By the last observation in the last paragraph, such vector fields exist locally and we can average them using a partition of unity (cf. [Ka89, 5.2.2]). We have now completed an orbifold version of the proof of the Soul Theorem which was originally given by Cheeger and Gromoll [CG72] for non-compact non-negatively curved manifolds:

Proposition 3.5.1 (Soul Theorem for orbifolds). *A complete non-compact Riemannian orbifold with sectional curvature ≥ 0 contains a totally convex and totally geodesic closed suborbifold, a so-called* soul*, and is diffeomorphic to the normal bundle of the soul.*

3.5.1 Classification of non-compact 3-orbifolds with sec ≥ 0

To derive the classification of all non-compact complete 3-orbifolds with sec ≥ 0 from the Soul Theorem, we will also use the *splitting theorem* for general Alexandrov spaces (see [BBI01, Sec. 10.5]) which states that if an Alexandrov space of curvature ≥ 0 contains a line (an isomtric embedding of $\mathbb{R}$), it splits off a line metrically. Thus, if a complete Riemannian orbifold O with sectional curvature ≥ 0 contains a line, it splits as the metrical product $\mathbb{R} \times O'$ with O' a complete Riemannian orbifold with $\dim O' = \dim O - 1$.

Proposition 3.5.2. *Every connected complete non-compact Riemannian 3-orbifold with sectional curvature ≥ 0 is either diffeomorphic to a discal ($\dim S = 0$) or a solid toric ($\dim S = 1$) 3-orbifold, or isometric to $\Sigma^2 \times \mathbb{R}$ or $(\Sigma^2 \times \mathbb{R})/\mathbb{Z}_2$ where Σ^2 is a closed 2-orbifold with sectional curvature ≥ 0, and $\mathbb{Z}_2$ operates isometrically on $\Sigma^2 \times \mathbb{R}$ ($\dim S = 1$).*

Proof. The proposition is clear for $\dim S = 0$. For $\dim S = 2$ it is sufficient to observe that after passing to a double covering if necessary, S separates $\nu(S)$ and hence O. Thus, O contains a line and the splitting theorem applies.

If $\dim S = 1$ and S is a circle, we obtain O by starting with $D^2(p) \times [-1,1]$ or $V^2(p) \times [-1,1]$ and identifying the boundary components. Thus, O is diffeomorphic to a bundle over S^1 with fiber $\mathbb{R}^2/\Gamma$ for some finite subgroup $\Gamma \subset \mathrm{O}(2)$. If S is diffeomorphic to the mirrored interval, O can be obtained from $D^2(p) \times [-1,1]$, $p \geq 1$, by gluing each of the boundary components $D^2(p) \times \{\pm 1\}$ either to itself via a half-rotation or reflection or by making it a reflector boundary component (there are six such orbifolds), or from $V^2(p) \times [-1,1]$, $p \geq 1$, by gluing each of the boundary components $V^2(p) \times \{\pm 1\}$ either to itself via the reflection at its bisector or by making it a reflector boundary component (there are three such orbifolds). In both cases the orbifold O is diffeomorphic to a solid toric 3-orbifold. □

Remark 3.5.3. We also note an alternative construction for the case where the soul is a mirrored interval: In this case, O can also be obtained by starting with two quotients of the 3-ball, each with boundary $\mathbb{R}P^2(p)$, $S^2(2,2,p)$ or $D^2(p)$ for some fixed $p \geq 1$, and glueing them together along a closed pointed disc $D^2(p)$ contained in both boundaries. The (interior of the) resulting 3-orbifolds are then diffeomorphic to the ones obtained from $D^2(p) \times [-1,1]$ discussed in the proof.

Similarly, we could start with two quotients of the 3-ball, each with boundary $D^2(;2,2,p)$ or $D^2(2;p)$ for some fixed $p \geq 1$, and then identify two sectors $V^2(p)$. In this way, we obtain precisely the orbifolds diffeomorphic to the ones obtained from $V^2(p) \times [-1,1]$.

4. Convergence of thick orbifolds

Let $O^{n\geq 2}$ be a complete connected Riemannian orbifold with bounded sectional curvature $|\sec| \leq \kappa$. We will assume a certain *thickness* (non-collapsedness) in the volume sense in a base point $p \in O$, namely that

$$\mathrm{vol}(B_r(p)) \geq v$$

for certain constants $r, v > 0$. Since we are assuming a lower curvature bound, thickness in the base point propagates over the entire orbifold.

Lemma 4.0.1 (Thickness propagates). *For $R, r' > 0$ there is $v' = v'(n, \kappa, r, v, R, r') > 0$ such that the following holds: If for $p, p' \in O$ holds vol$(B_r(p)) \geq v$ and $d(p', p) \leq R$, then vol$(B_{r'}(p')) \geq v'$.*

Proof. Since $\mathrm{vol}(B_{r+R}(p')) \geq v$ by our assumption, Bishop-Gromov volume comparison (3.1.4) yields a lower bound for $\mathrm{vol}(B_{r'}(p'))$ if $r' \leq r + R$. □

If one imposes an upper bound on the order of local isotropy groups, equivalently, a lower bound on the volume of links, then small local volume implies small injectivity radius (again by Bishop-Gromov). The converse is not true, since the injectivity radius tends to zero as one approaches a singular stratum of $\mathcal{S}$. The main result of the next section (Prop. 4.1.4) is that this is the only way for the injectivity radius to become small in a thick (i.e. volume non-collapsed) region, or in other words that close to a point with small injectivity radius there is a more singular point.

4.1 An injectivity radius bound on thick orbifolds

We begin with another application of volume comparison to show that in a thick point a certain amount of directions can be extended up to a certain length as *minimizing* geodesic segments. Let us denote by $L_{r'} \subseteq \Sigma_p O$ the open subset of directions u such that c_u remains distance minimizing up to a distance $> r'$.

Lemma 4.1.1 (Long directions). *There are $r' = r'(n, \kappa, v) \in (0, r)$ and $a = a(n, \kappa, r, v) > 0$ such that the following holds: If vol$(B_r(p)) \geq v$, then vol$(L_{r'}) \geq a$.*

Proof. We denote by $v_{n,-\kappa}(\rho)$ the volume of a ρ-ball in the n-dimensional model space $M^n_{-\kappa}$ of constant sectional curvature $-\kappa$. We choose r' so that $v_{n,-\kappa}(r') < \frac{v}{2}$ (so $r' < r$.) Since

$$\mathrm{vol}(B_r(p)) \leq v_{n,-\kappa}(r') + \frac{\mathrm{vol}(L_{r'})}{\mathrm{vol}(S^{n-1})} v_{n,-\kappa}(r),$$

we obtain a lower bound for vol$(L_{r'})$. □

If the injectivity radius in p is small, then the directions pointing to closest cut points are *short* in the sense that they cannot be extended beyond length inj(p) as minimizing segments. The lower curvature bound forces via triangle comparison that long and short directions have at least almost right angles. We now apply our earlier discussion of closest cut points in section 3.3.

Lemma 4.1.2 (Almost right angles between long and short directions). *For $\epsilon > 0$ there is $i = i(\kappa, r, \epsilon) \in (0, \frac{\pi}{\sqrt{\kappa}})$ such that the following holds: Suppose that* inj$(p) \leq i$ *and that q is a closest cut point for p. If x is a point at distance $d(p, x) \geq r$ from p, then every minimizing segment $\overline{px}$ has angle $\geq \frac{\pi}{2} - \epsilon$ with every minimizing segment $\overline{pq}$.*

Proof. Due to the upper curvature bound κ we have conj$(p) \geq \frac{\pi}{\sqrt{\kappa}}$ and therefore look for a value of i in $(0, \frac{\pi}{\sqrt{\kappa}})$. Let λ be a geodesic loop of length $2\,\mathrm{inj}(p)$ based at p and with midpoint q, as given by Proposition 3.3.3. We consider the geodesic "triangle" with twice the (minimizing) side $\overline{px}$ and third (non-minimizing) side λ. Triangle comparison as in Proposition 3.4.4 yields the assertion. (If $r \geq \frac{\pi}{\sqrt{\kappa}}$, we replace x by a point x' on $\overline{px}$ closer to p.) □

If the injectivity radius in some point is small compared to the local volume, then the positions of long and short directions are restricted because they are almost angle $\frac{\pi}{2}$ apart and the set of long directions has a certain volume. This inhibits the short directions from spreading out too much and one finds directions in which the injectivity radius decays at a definite rate.

Lemma 4.1.3 (Decay of injectivity radius along strata). *There exist $i = i(n, \kappa, r, v) > 0$ and $\delta = \delta(n, \kappa, r, v) > 0$ such that the following holds: If vol$(B_r(p)) \geq v$ and* inj$(p) \leq i$, *then there exists a direction $w \in \Sigma_p O$ tangent to the stratum S_p of O through p such that the injectivity radius has decay $\leq -\delta$ along S_p in the direction of w (in the barrier sense). If in particular p is a singular vertex of O with vol$(B_r(p)) \geq v$, then* inj$(p) > i$.

By "barrier sense" we mean that for a smooth curve $\gamma : [0, \epsilon) \to S_p$ with $\dot{\gamma}(0) = w$ there is a smooth function $f : [0, \epsilon) \to \mathbb{R}$ with $\dot{f}(0) = -\delta$ such that inj$(\gamma(t)) \leq f(t)$ for t near 0 with equality in $t = 0$.

Proof. Let $r', a > 0$ be the constants depending on n, κ, r, v given by Lemma 4.1.1, let $\epsilon = \epsilon(n, a) > 0$ be the constant given by Lemma 3.2.5, and let $i = i(\kappa, r', \epsilon) > 0$ be the constant given by Lemma 4.1.2.

Suppose that $\mathrm{vol}(B_r(p)) \geq v$ and $\mathrm{inj}(p) \leq i$. The subset $L_{r'} \subseteq \Sigma_p O$ of r'-long directions in p has volume $\geq a$ by Lemma 4.1.1. Let λ be a geodesic loop of length $2\,\mathrm{inj}(p)$ based at p as given by Proposition 3.3.3, and let $v_i = \Sigma_p \sigma_i \in \Sigma_p O$ denote the directions of its ends. By Lemma 4.1.2, $\Sigma_p O - B_{\frac{\pi}{2}-\epsilon}(v_1) - B_{\frac{\pi}{2}-\epsilon}(v_2)$ contains $L_{r'}$ and hence has volume $\geq a$. We consider the join decomposition (3.2.3) and denote by v_i'' the projection of v_i to the sphere factor $\Sigma_p S_p$. As a consequence of Lemma 3.2.5, the v_i'' are well-defined and $\angle(v_1'', v_2'') \leq \pi - 2\epsilon$. In particular, $\Sigma_p S_p \neq \emptyset$ and p cannot be a singular vertex.

We choose $w \in \Sigma_p S_p$ as the midpoint (bisector) of the directions v_1'' and v_2''. The points v_i, v_i'', w are the vertices of a spherical triangle embedded in $\Sigma_p O$. It has a right angle at v_i'' and, by Lemma 3.2.5, the two adjacent sides have length $\leq \frac{\pi}{2} - \epsilon$. Therefore the side opposite to v_i'' has length $\leq \frac{\pi}{2} - \epsilon' < \frac{\pi}{2}$, where $\frac{\pi}{2} - \epsilon'$ is defined as the length of the hypothenuse of an isosceles right-angled spherical triangle with two sides of length $\frac{\pi}{2} - \epsilon$.

Let $\gamma : [0, \epsilon'') \to S_p$ be a smooth curve with $\dot\gamma(0) = w$. According to our discussion in section 3.1, there exists a variation $(\lambda_s)_{s\in[0,\epsilon)}$ of λ by geodesic loops λ_s with base points $\gamma(s)$ such that $\dot\gamma(0) = w$ and such that equality holds in the first variation inequality (3.1.2). So,

$$\frac{d}{ds}|_{s=0} L(\lambda_s) = -\cos\angle_p(v_1, w) - \cos\angle_p(v_2, w) \leq -2\sin\epsilon' =: -\delta(d, \kappa, r, v).$$

The assertion follows, cf. Remark 3.3.4. □

By "integrating" Lemma 4.1.3, we obtain the desired lower injectivity radius bound away from the boundary of strata. Since the comparison arguments based on the curvature bounds are local, we obtain a more general *local* version of the result.

Proposition 4.1.4 (Lower injectivity radius bound away from the boundary of strata). *There exist $i = i(n, \kappa, r, v) > 0$ and $c = c(n, \kappa, r, v) > 0$ such that the following holds: Let $O^{n\geq 2}$ be a complete Riemannian orbifold and let $p \in O$. Suppose that vol$(B_r(p)) \geq v$ and that on $B_r(p)$ the sectional curvature is bounded by $|\sec| \leq \kappa$. Then* $\mathrm{inj}(p) \geq \min(i, c \cdot d(p, \partial S_p))$.

Proof. Assume first that $|\sec| \leq \kappa$ everywhere on O and $\mathrm{vol}(B_r(x)) \geq v$ for all $x \in O$. Let i and δ be the constants in Lemma 4.1.3. For some small $\epsilon > 0$, the function $\phi := \mathrm{inj} + (1-\epsilon)\delta \cdot d(p, \cdot)$ is continuous on the stratum S_p, cf. Proposition 3.1.3. Due to Lemma 4.1.3, it has no local minimum on $\{\mathrm{inj} \leq i\} \cap S_p$. Putting $\rho := (1-\epsilon)^{-1}\frac{i}{\delta}$ we observe that for $\lambda \in (0, 1]$ we have $\phi > \lambda i$ on $\partial B_{\lambda\rho}(p) \cap S_p$. Therefore, if $\mathrm{inj}(p) \leq \lambda i$, the closed ball

$\overline{B}_{\lambda\rho}(p) \cap S_p$ cannot be compact. This means that $\overline{B}_{\lambda\rho}(p) \cap \partial S_p \neq \emptyset$, i.e. $d(p, \partial S_p) \leq \lambda\rho$. The assertion follows with $c = \frac{i}{\rho}$.

In the general case, if we assume the curvature bounds $|\sec| \leq \kappa$ only on $B_r(p)$ and the volume bound $\mathrm{vol}(B_r(p)) \geq v$ only in p, we first decrease v using Lemma 4.0.1 so that $\mathrm{vol}(B_{\frac{r}{2}}(x)) \geq v$ for all $x \in B_{\frac{r}{2}}(p)$. Then we proceed as before (with $\frac{r}{2}$ taking the role of r). In order not to leave $B_r(p)$, we decrease ρ if necessary, so that $\rho \leq \frac{r}{2}$. □

Remark 4.1.5. One can strengthen the conclusion of the proposition by replacing $d(\cdot, \partial S_p)$ with the *intrinsic* distance from ∂S_p inside the stratum S_p.

The following result generalizes [BLP05, Prop. 3.19] to the case of variable curvature and arbitrary dimension.

Corollary 4.1.6 (Small ball contained in small standard ball). *For $s > 0$ there exist radii $0 < \rho(n, \kappa, r, v, s) < \rho'(n, \kappa, r, v, s) < s$ such that the following holds: If O and p are as in Proposition 4.1.4, then the metric ball $B_\rho(p)$ is contained in a standard ball with radius $\leq \rho'$.*

Proof. To facilitate notation, let us again first assume that $|\sec| \leq \kappa$ on all of O and $\mathrm{vol}(B_r(x)) \geq v$ for all $x \in O$. Lemma 4.1.3 provides a lower injectivity radius bound $i_0 = i_0(n, \kappa, r, v) > 0$ on the set of singular vertices $O^{(0)}$. We may decrease it artificially so that $i_0 = i_0(n, \kappa, r, v, s) < s$. Proceeding by induction, one obtains from Proposition 4.1.4 lower injectivity radius bounds $i_l = i_l(n, \kappa, r, v, s) > 0$ on $O^{(l)} - N_{i_{l-1}/4}(O^{(l-1)})$ for $l = 1, \ldots, n$ and one can arrange that $s > i_0 >> i_1 >> \cdots >> i_n$. It follows that for every point $x \in N_{i_l/2}(O^{(l)}) - N_{i_{l-1}/2}(O^{(l-1)})$, the ball $B_{i_l/2}(x)$ is contained in a standard ball with radius i_l. Thus the assertion holds with $\rho := \frac{i_n}{2}$ and $\rho' := i_0$. As for Proposition 4.1.4, the estimates can be localized. □

4.2 Compactness of thick orbifolds

4.2.1 Gromov-Hausdorff and smooth convergence of orbifolds

We recall, e.g. from [Gr81, ch. 3] and [BBI01, ch. 7, 8.1], that a sequence of pointed metric spaces (X_k, x_k), $k \in \mathbb{N}$, converges in the *Gromov-Hausdorff* sense to a pointed metric space (X_∞, x_∞),

$$(X_k, x_k) \xrightarrow{GH} (X_\infty, x_\infty), \tag{4.2.1}$$

if for all $r, \epsilon > 0$ there exist for sufficiently large k (not necessarily continuous) maps $f_k : B_r^{X_k}(x_k) \to X_\infty$ such that $f_k(x_k) = x_\infty$, $|d^{X_\infty}(f_k(y_k), f_k(z_k)) - d^{X_k}(y_k, z_k)| < \epsilon$ for all $y_k, z_k \in B_r^{X_k}(x_k)$, and $B_{r-\epsilon}^{X_\infty}(x_\infty) \subseteq N_\epsilon(f_k(B_r^{X_k}(x_k)))$.

An ϵ-*net* in a metric space Y is an ϵ-separated subset $\nu \subseteq Y$ (i.e. any two points in ν have distance $> \epsilon$ from each other) with $d_H(\nu, Y) \leq \epsilon$. For a function $N : (0,\infty)^2 \to \mathbb{N}$ we denote by $\mathcal{L}_N$ the space of isometry classes of pointed complete length spaces (X, x) such that for all $r, \epsilon > 0$ the ϵ-nets in $B_r^X(x)$ have cardinality $\leq N(r, \epsilon)$. Note that the metric spaces in $\mathcal{L}_N$ are *proper* and *geodesic*.

It is a basic fact (see e.g. [BBI01, Thm. 8.1.10]) that $\mathcal{L}_N$ is *compact* with respect to the Gromov-Hausdorff topology, i.e. any sequence of pointed metric spaces in $\mathcal{L}_N$ has a convergent subsequence. The pointed limit metric space is unique up to isometry if it is complete and boundedly compact, i.e. if all closed bounded sets in it are compact (see [BBI01, Thm. 8.1.7]). Given a convergent sequence (4.2.1) in $\mathcal{L}_N$, the balls of a fixed radius around the base points converge in the sense of pointed Gromov-Hausdorff *distance*,

$$d_{GH}((B_r^{X_k}(x_k), x_k), (B_r^{X_\infty}(x_\infty), x_\infty)) \to 0$$

for all $r > 0$. That is, for every $r > 0$ there exist isometric embeddings $\iota_k : B_r^{X_k}(x_k) \hookrightarrow Z$ and $\iota_\infty : B_r^{X_\infty}(x_\infty) \hookrightarrow Z$ into some auxiliary metric space Z with the property that $d_H^Z(B_r^{X_k}(x_k), B_r^{X_\infty}(x_\infty)) \to 0$ and $x_k \to x_\infty$ in Z, where d_H^Z denotes the Hausdorff distance of subsets of Z. (Here we use that in a length space concentric metric balls of radii $r_1, r_2 > 0$ have Hausdorff distance $\leq |r_1 - r_2|$.) One may think of the family of embeddings ι_k and ι_∞ as a local *realization* of the convergence (4.2.1). Once a realization is fixed, it makes sense to speak of the convergence of sequences (x_k) of points $x_k \in B_r^{X_k}(x_k)$. A diagonal argument shows that one can choose compatible realizations of the convergences of r-balls simultaneously for all $r > 0$.

When considering metric spaces with additional structure, one may be interested in promoting Gromov-Hausdorff convergence to stronger forms of convergence which establish a closer tie between the approximators and the limit space.

Definition 4.2.2. We say that a sequence of pointed connected Riemannian n-orbifolds (O_k, p_k) converges *smoothly* (or in the Cheeger-Gromov sense) to a pointed complete connected Riemannian n-orbifold (O_∞, p_∞),

$$(O_k, p_k) \xrightarrow{\mathcal{C}^\infty} (O_\infty, p_\infty), \tag{4.2.3}$$

if for every relatively compact open subset $W_\infty \subseteq O_\infty$ containing the base point p_∞ there are for sufficiently large k diffeomorphisms $\Phi_k : W_\infty \to W_k$ onto open subsets $W_k \subset O_k$ such that $\Phi_k^{-1}(p_k) \to p_\infty$ in O_∞ and $\Phi_k^* g_k \stackrel{\mathcal{C}^\infty}{\to} g_\infty$.

Note that we require the approximating diffeomorphisms to be only *almost* base point preserving instead of exactly base point preserving as in [Lu01, Def. 3.4]. In the manifold case this would make no difference, but in the orbifold case it does because orbifold diffeomorphisms preserve the singular stratifications. It seems natural to choose the definition of smooth convergence such that if $p_k \to p_\infty$ is a convergent sequence in a connected Riemannian orbifold O then $(O, p_k) \overset{\mathcal{C}^\infty}{\to} (O, p_\infty)$. We observe that smooth convergence implies convergence of stratifications because local orbifold diffeomorphisms preserve the stratification of orbifolds.

As in the manifold case one can show that under appropriate conditions, smooth limits are unique up to Riemannian isometry, for instance, if one requires them to be complete. Also, if O_∞ is complete (and hence also the O_k), then smooth convergence (4.2.3) implies Gromov-Hausdorff convergence with respect to the associated path metrics.

4.2.2 Orbifold structures on Gromov-Hausdorff limits

For an integer $n \geq 2$, a nondecreasing function $\kappa : (0,\infty) \to [0,\infty)$ and constants $r_0, v_0 > 0$, we denote by $\mathcal{O}_{n,\kappa,r_0,v_0}$ the space of isometry classes of pointed complete connected smooth[1] Riemannian n-orbifolds (O,p), such that $|\sec| \leq \kappa \circ d(p,\cdot)$, i.e. such that for every $r > 0$ the sectional curvature on $B_r(p)$ is bounded by $|\sec| \leq \kappa(r)$, and such that $\mathrm{vol}(B_{r_0}(p)) \geq v_0$.

The local lower curvature bounds imply via Bishop-Gromov volume comparison (Prop. 3.1.4) that the thickness (volume non-collapse) in the base point propagates across the orbifold and can be bounded below in terms of the distance from the base point, i.e. there are uniform estimates $\mathrm{vol}(B_r(x)) \geq v(n,\kappa,r_0,v_0,d(p,x),r)$, cf. Lemma 4.0.1. Moreover, the size of ϵ-nets in metric balls can be bounded above, and hence $\mathcal{O}_{n,\kappa,r_0,v_0} \subset \mathcal{L}_N$ for a suitable counting function $N(r,\epsilon) = N_{n,\kappa,r_0,v_0}(r,\epsilon)$. In particular, the space $\mathcal{O}_{d,\kappa,r_0,v_0}$ is Gromov-Hausdorff *precompact*.

We are interested in the Gromov-Hausdorff closure of $\mathcal{O}_{n,\kappa,r_0,v_0}$, more precisely, in the local structure of limit spaces in the Gromov-Hausdorff boundary $\partial^{GH}\mathcal{O}_{n,\kappa,r_0,v_0}$, and in compactness properties of $\mathcal{O}_{n,\kappa,r_0,v_0}$ and related spaces with respect to the finer topologies of smooth convergence. Since by construction lower Alexandrov curvature bounds pass to Gromov-Hausdorff limits, we already know that the limit spaces in $\partial^{GH}\mathcal{O}_{n,\kappa,r_0,v_0}$ have again local lower curvature bounds in the Alexandrov sense.

If we impose more regularity on a space of orbifolds, we can also expect the limit spaces to be more regular and to obtain stronger compactness properties. For an integer $n \geq 2$, a function $D : \mathbb{N}_0 \times (0,\infty) \to [0,\infty)$ and constants $r_0, v_0 > 0$, we denote by $\mathcal{O}_{n,D,r_0,v_0}$ the space of isometry classes of pointed complete connected smooth Riemannian n-orbifolds (O,p) for which $\mathrm{vol}(B_{r_0}(p)) \geq v_0$ and $\|\nabla^l R\| \leq D(l,r)$ on $B_r(p)$ for all $l \geq 0$ and $r > 0$. Clearly, $\mathcal{O}_{d,D,r_0,v_0} \subset \mathcal{O}_{d,\kappa,r_0,v_0}$ for suitable $\kappa(r) = \kappa_{D(0,\cdot)}(r)$ and $\mathcal{O}_{n,D,r_0,v_0}$ is also Gromov-

[1]Since we are only interested in curvature at the moment, it would actually suffice to require only regularity of class $\mathcal{C}^4$, say.

Hausdorff *precompact.*

We consider a sequence of pointed Riemannian orbifolds $(O_k, p_k) \in \mathcal{O}_{n,\kappa,r_0,v_0}$. Due to the Gromov-Hausdorff precompactness of $\mathcal{O}_{n,\kappa,r_0,v_0}$, we may assume after passing to a subsequence that the (O_k, p_k) converge to a limit space,

$$(O_k, p_k) \xrightarrow{GH} (X_\infty, p_\infty).$$

When investigating the local structure of X_∞ near some point x_∞, we fix a constant $R >> d(p_\infty, x_\infty)$ and a realization of the Gromov-Hausdorff convergence of balls $(B_R^{O_k}(p_k), p_k) \overset{GH}{\to} (B_R^{X_\infty}(p_\infty), p_\infty)$ in some auxiliary ambient metric space Z.

We now apply our uniform lower injectivity radius bounds obtained in section 4.1 to the orbifolds O_k.

Proposition 4.2.4. *Any point in X_∞ is contained in an open metric ball which is a Gromov-Hausdorff limit of* standard *balls in the O_k.*

Proof. Given a point $x_\infty \in X_\infty$, Corollary 4.1.6 implies that there exists a point $x'_\infty \in X_\infty$ (arbitrarily close to x_∞) and a sequence of points $x'_k \in O_k$ such that $x'_k \to x'_\infty$ (in Z) and $\liminf_{k\to\infty} \text{inj}(x'_k) > d(x_\infty, x'_\infty)$. Then for $d(x_\infty, x'_\infty) < r < \liminf_{k\to\infty} \text{inj}(x'_k)$ we have

$$(B_r^{O_k}(x'_k), x'_k) \xrightarrow{GH} (B_r^{X_\infty}(x'_\infty), x'_\infty). \tag{4.2.5}$$

and the $B_r^{O_k}(x'_k)$ are standard balls for large k. □

Since there are uniform curvature bounds on the balls $B_r^{O_k}(x'_k)$ and since their volumes are bounded below uniformly in k, we have that $\liminf_{k\to\infty} \text{vol}(\Sigma_{x'_k} O_k) > 0$. Therefore the links $\Sigma_{x'_k} O_k$ fall into finitely many isometry types, compare Remark 3.2.4, and after passing to a subsequence we may assume that they are isometric to each other. Then there exist a finite subgroup $\Gamma \subset O(n)$ and isometric identifications $i_k : \mathbb{R}^n/\Gamma \overset{\cong}{\to} T_{x'_k} O_k$ (preserving origins).

Let us denote by $e_k : B_r^{\mathbb{R}^n/\Gamma}(0) \overset{\cong}{\to} B_r^{O_k}(x'_k)$ the diffeomorphisms obtained from restricting $\exp_{x'_k} \circ i_k$, by $\pi_\Gamma : \mathbb{R}^n \to \mathbb{R}^n/\Gamma$ the quotient projection, and $\hat{e}_k := e_k \circ \pi_\Gamma : B_r^{\mathbb{R}^n}(0) \to B_r^{O_k}(x'_k)$. Pulling back the orbifold Riemannian metrics g_k from O_k via $\hat{e}_k$ yields smooth Riemannian metrics $\hat{h}_k := \hat{e}_k^* g_k$ on the ball $B_r(0) \subset \mathbb{R}^n$. If r is chosen sufficiently small, say $r < \pi/2\sqrt{\kappa(R)}$, the Rauch estimates imply that the $\hat{h}_k$ are uniformly bilipschitz equivalent to the euclidean metric.

Depending on the additional regularity assumptions which one imposes on the O_k, one can deduce regularity for the limit space X_∞. We assume now that $(O_k, p_k) \in \mathcal{O}_{n,D,r_0,v_0}$. The function $\kappa(r)$ can be chosen so that $\mathcal{O}_{n,D,r_0,v_0} \subset \mathcal{O}_{n,\kappa,r_0,v_0}$. Then there are uniform bounds for the curvature tensors of the metrics $\hat{h}_k$ and all their covariant derivatives (for sufficiently small r as above).

Lemma 4.2.6 (cf. [Ba07, Lem. 3.2.1]). *For $r < \infty$, $l, n \in \mathbb{N}$ and constants $0 < D_0, \dots, D_l < \infty$ there are constants $0 < D'_0, \dots, D'_l < \infty$ such that the following holds: Let (M, g) be an n-dimensional smooth Riemannian manifold with $x \in M$ and suppose that the exponential map is defined on $B_r(0) \subset T_xM$. If $\|\nabla^k R\| < D_k$ on M for all $0 \le k \le l$, then $\|d^l exp^* g\| < D'_l$ on $B_r(0)$.*

Based on the lemma, we may apply Arzelà-Ascoli and pass to a subsequence so that the $\hat{h}_k$ converge to a Γ-invariant smooth Riemannian limit metric $\hat{h}_\infty$ on $B_r(0)$,

$$\hat{h}_k \stackrel{\mathcal{C}^\infty}{\to} \hat{h}_\infty.$$

The smooth convergence implies Gromov-Hausdorff convergence with respect to the associated intrinsic path metrics, $(B_r(0), d_{\hat{h}_k}) \stackrel{GH}{\to} (B_r(0), d_{\hat{h}_\infty})$, and equally for the Γ-quotients, $(B_r(0), d_{\hat{h}_k})/\Gamma \stackrel{GH}{\to} (B_r(0), d_{\hat{h}_\infty})/\Gamma$.

Note that the balls $B_r^{O_k}(x'_k)$ may not be convex, but their intrinsic path metrics are at least *locally* isometric to their extrinsic metrics obtained by restricting the path metrics d_{g_k} on the O_k. More precisely, for a point $y'_k \in B_r^{O_k}(x'_k)$, both metrics coincide on the ball around y'_k with radius $\frac{1}{2}(r - d_{g_k}(x'_k, y'_k))$. Hence the ball quotients $(B_r(0), d_{\hat{h}_k})/\Gamma$ are in a uniform way locally isometric to the balls $B_r^{O_k}(x'_k)$ equipped with their extrinsic metrics. We conclude with (4.2.5) that $B_r^{X_\infty}(x'_\infty)$ is locally isometric to the Riemannian ball quotient $(B_r(0), d_{\hat{h}_\infty})/\Gamma$. It follows that the limit space X_∞ is a smooth Riemannian n-orbifold O_∞, and that O_∞ inherits from the O_k the same rate of local non-collapsedness and the same bounds for the curvature tensor along with all its derivatives. We therefore have completed the proof of

Theorem 4.2.7 (Gromov-Hausdorff compactness). *$\mathcal{O}_{n,D,r_0,v_0}$ is Gromov-Hausdorff compact.*

Using an isometric identification $i_\infty : \mathbb{R}^d/\Gamma \stackrel{\cong}{\to} T_{x'_\infty}O_\infty$, we convert the i_k into isometries $\iota_k = i_k \circ i_\infty^{-1} : T_{x'_\infty}O_\infty \stackrel{\cong}{\to} T_{x'_k}O_k$ (preserving origins). For the corresponding diffeomorphisms of standard balls

$$\phi_k := \exp_{x'_k}^{O_k} \circ \iota_k \circ (\exp_{x'_\infty}^{O_\infty})^{-1} : B_r^{O_\infty}(x'_\infty) \longrightarrow B_r^{O_k}(x'_k) \tag{4.2.8}$$

then holds

$$\phi_k^* g_k \stackrel{\mathcal{C}^\infty}{\longrightarrow} g_\infty|_{B_r^{O_\infty}(x'_\infty)} \tag{4.2.9}$$

where we denote by g_∞ the Riemannian metric on O_∞. The Gromov-Hausdorff convergence

$$(B_r^{O_k}(x'_k), x'_k) \stackrel{GH}{\longrightarrow} (B_r^{O_\infty}(x'_\infty), x'_\infty), \tag{4.2.10}$$

compare (4.2.5), and the smooth convergence (4.2.9) can be coordinated so that

$$\phi_k \to \mathrm{id}_{B_r^{O_\infty}(x'_\infty)} \tag{4.2.11}$$

uniformly. A priori one has only that for large k the ϕ_k are close to isometries ψ_k of the limit ball $B_r^{O_\infty}(x'_\infty)$, but (4.2.11) can be achieved by replacing the ϕ_k with $\phi_k \circ \psi_k^{-1}$.

4.2.3 Riemannian center of mass

We recall from [GK73, Ka77] the construction of a *center of mass* for sufficiently concentrated measures with finite total mass on a Riemannian manifold. We will apply it in the next section where we will use it to interpolate between (equivariant lifts of) local orbifold diffeomorphisms.

Let M be a Riemannian manifold. For a point $p \in M$, the function $d(p, \cdot)^2$ is smooth and strictly convex near p. Hence, for a measure μ on M with finite mass $0 < \|\mu\| < \infty$ and supported on a sufficiently small ball $B_\rho(p)$, the function $f_\mu := \int_M d^2(q, \cdot)\, d\mu(q)$ is smooth and strictly convex near p, and it has a unique global minimum near p which one calls the *center of mass* $C(\mu)$ of μ. The gradient vector field $\operatorname{grad} f_\mu$, viewed as a smooth section of the tangent bundle $TM \to M$ defined near p, intersects the zero section transversally in $C(\mu)$.

Let $\lambda_i : M \to [0, \infty)$ be smooth (weight) functions, $1 \leq i \leq m$, such that $\sum_{i=1}^m \lambda_i > 0$. Then the map

$$(q_1, \dots, q_m) \overset{c}{\mapsto} C\big(\lambda_1(q_1) \cdot \delta_{q_1} + \dots + \lambda_m(q_m) \cdot \delta_{q_m}\big) \tag{4.2.12}$$

is well-defined on a sufficiently small (open) neighbourhood U of the diagonal Δ in $M^{\times m} = M \times \dots \times M$, and $c|_\Delta : \Delta \to M$ is the canonical identification.

That c is smooth, can be seen by a transversality argument as follows. Let $\pi : U \times M \to M$ denote the projection onto the second factor. Then

$$(q_1, \dots, q_m, x) \overset{s}{\mapsto} \Big(\operatorname{grad} \sum_{i=1}^m \lambda_i(q_i) \cdot d(q_i, \cdot)^2\Big)(x)$$

is a smooth section of the vector bundle $\pi^*TM \to U \times M$. It intersects the zero section precisely along $\operatorname{graph}(c)$ and the intersection is transversal. Thus $\operatorname{graph}(c)$ is a submanifold of $U \times M$ with codimension $\dim(M)$. Moreover, it is transversal to the fibration $U \times M \to U$ due to the strict convexity of the functions $\sum_{i=1}^m \lambda_i(q_i) \cdot d(q_i, \cdot)^2$. This shows that c is indeed smooth.

Our discussion and in particular the smoothness of c imply:

Lemma 4.2.13. *Let $\tau : N \to M$ be a smooth map from a manifold N into a Riemannian manifold M. Furthermore, for $1 \leq i \leq m$, let $\tau_k^i : N \to M$ be sequences of smooth maps, $k \in \mathbb{N}$, such that $\tau_k^i \overset{C^\infty}{\to} \tau$. Then for any compact subset K of N, the map $c \circ (\tau_k^1, \dots, \tau_k^m)$ is defined and smooth on K for sufficiently large k, and $c \circ (\tau_k^1, \dots, \tau_k^m) \overset{C^\infty}{\to} \tau$.*

Proof. $c \circ (\tau_k^1, \dots, \tau_k^m) \xrightarrow{\mathcal{C}^\infty} c \circ (\tau, \dots, \tau) = \tau$. □

4.2.4 Promoting Gromov-Hausdorff convergence to smooth convergence

Following the argument in [Ba07, Sec. 3.2], we will now strengthen the Gromov-Hausdorff convergence of Theorem 4.2.7 to smooth convergence. More precisely, suppose that

$$(O_k, p_k) \xrightarrow{GH} (O_\infty, p_\infty) \tag{4.2.14}$$

is a Gromov-Hausdorff convergent sequence in $\mathcal{O}_{n,D,r_0,v_0}$. We want to show that Gromov-Hausdorff convergence can be promoted to smooth convergence by passing to a suitable subsequence. We proved this locally in section 4.2.2 and explain now how to interpolate local almost isometric smooth approximations to obtain global ones.

Let $W_\infty \subseteq O_\infty$ be a relatively compact open subset. As in section 4.2.2 we fix a radius R such that $\overline{W}_\infty \subset B_R(p_\infty)$ and a realization of the Gromov-Hausdorff convergence of balls $(B_R^{O_k}(p_k), p_k) \xrightarrow{GH} (B_R(p_\infty), p_\infty)$ in an auxiliary ambient metric space Z.

According to Proposition 4.2.4, $\overline{W}_\infty$ can be covered by finitely many open standard balls $B_{r_i}(y_i) \subset O_\infty$, $1 \le i \le m$, such that the larger balls $B_{16r_i}(y_i) \subset O_\infty$ are standard and Gromov-Hausdorff limits of open standard balls $B_{16r_i}^{O_k}(y_i^k) \subset O_k$ for all $1 \le i \le m$. Moreover, after passing to a suitable subsequence there exist sequences of diffeomorphisms

$$\phi_i^k : B_{16r_i}(y_i) \to B_{16r_i}^{O_k}(y_i^k) \tag{4.2.15}$$

such that $(\phi_i^k)^* g_k \xrightarrow{\mathcal{C}^\infty} g_\infty|_{B_{16r_i}(y_i)}$ and $\phi_i^k \to \mathrm{id}_{B_{16r_i}(y_i)}$ uniformly for every i as $k \to \infty$, cf. (4.2.8, 4.2.9, 4.2.10). Let $\Gamma_i := \Gamma_{y_i} = \Gamma_{y_i^k}$ be the corresponding local groups. For technical reasons, we choose the radii r_i pairwise distinct and arrange them in increasing order $r_1 < \dots < r_m$.

For $i \le j \in \{1, \dots, m\}$ we say that $i \prec j$ if $\overline{B_{4r_i}(y_i)} \subset B_{16r_j}(y_j)$. For sufficiently large k, i.e. $k \ge k_0$ for some value of k_0, we have the implication that if $\phi_i^k(B_{r_i}) \cap \phi_j^k(B_{r_j}) \neq \emptyset$ for some $i \le j$, then $i \prec j$.

For an index pair $i \prec j$ and sufficiently large k, the *transition map* $\psi_{ij}^k := (\phi_k^j)^{-1} \circ \phi_k^i : B_{4r_i}(y_i) \to B_{16r_i}(y_i)$ is a well-defined orbifold embeddings. (Because there are only finitely many index pairs $i \prec j$, this works for all of them for sufficiently large k.) By the metric properties of the diffeomorphisms ϕ_i^k and ϕ_j^k, we have $\psi_{ij}^k \to \mathrm{id}_{B_{4r_i}(y_i)}$ uniformly.

By Lemma 2.1.6, we can lift the transition maps to $|\Gamma_i|$ distinct embeddings $\widetilde{B}_{4r_i}(y_i) \to \widetilde{B}_{16r_i}(y_i)$ which are equivariant with respect to group isomorphisms $\Gamma_i \to \Gamma_i$.

Lemma 4.2.16. *For any sufficiently small $\epsilon > 0$ there is $k_0 \in \mathbb{N}$ such that for all $i < j \in \{1, \dots, m\}$ with $i \prec j$ and $k \geq k_0$ there is a unique lift $\widetilde{\psi}_{ij}^k : \widetilde{B}_{4r_i}(y_i) \to \widetilde{B}_{16r_i}(y_i)$ of ψ_{ij}^k such that $d(\widetilde{\psi}_{ij}^k(x), x) < \epsilon$ for all $x \in \widetilde{B}_{4r_i}(y_i)$. This lift is Γ_i-equivariant, i.e. $\widetilde{\psi}_{ij}^k \circ \gamma = \gamma \circ \widetilde{\psi}_{ij}^k$ for all $\gamma \in \Gamma_i$.*

Proof. Let $1 \leq i, j \leq m$ with $i \prec j$. In $\widetilde{B}_{4r_i}(y_i)$, let V denote the open subset of all points with trivial Γ_i-stabilizer. It is connected unless Γ_i contains reflections at hyperplanes. In this case, the different components of V are mapped to each other by reflections in Γ_i. Let $V_0 \subset V$ be a connected component of V.

For small $\epsilon' > 0$, we have the open set $U_{\epsilon'} \subset V \subset \widetilde{B}_{4r_i}(y_i)$ consisting of all points $x \in \widetilde{B}_{4r_i}(y_i)$ such that $d(\gamma \cdot x, x) > \epsilon'$ for all nontrivial $\gamma \in \Gamma_i$. By construction, $U_{\epsilon'}$ is Γ_i-invariant. Let $U_{\epsilon',0} = U_{\epsilon'} \cap V_0$. By choosing ϵ' sufficiently small, we can assume that $U_{\epsilon',0}$ is connected. (This follows from the fact that every point in V_0 lies in $U_{\epsilon',0}$ for sufficiently small ϵ'.)

For sufficiently large k the map $\psi_{ij}^k|_{B_{4r_i}(y_i)}$ is ϵ'-close to $id_{B_{4r_i}(y_i)}$. Thus for every point $x \in U_{\epsilon',0}$ one lift $\widetilde{\psi}_{ij}^k$ of ψ_{ij}^k is uniquely characterized by the property that $d(\widetilde{\psi}_{ij}^k(x), x) < \epsilon'$. (We lift a minimizing segment between the projection $p_i(x) \in B_{4r_i}(y_i)$ and its image under ψ_{ij}^k; this segment is entirely located inside $B_{16r_i}(y_i)$.) Since $U_{\epsilon',0}$ is connected, this lift does not depend on the choice of x. The embedding $\widetilde{\psi}_{ij}^k : \widetilde{B}_{4r_i}(y_i) \to \widetilde{B}_{16r_i}(y_i)$ is equivariant with respect to some group monomorphism $\alpha_{ij}^k : \Gamma_i \to \Gamma_i$.

We now claim that $\widetilde{\psi}_{ij}^k$ satisfies $d(\widetilde{\psi}_{ij}^k, \mathrm{id}) < \epsilon'$ on the whole of $U_{\epsilon'}$ even if Γ_i contains reflections at hypersurfaces. To verify this claim, let $\gamma \in \Gamma_i$ be a reflection at a hyperplane H which intersects ∂V_0. After reducing ϵ' further if necessary, we can find points $z \in H$ and $x \in U_{\epsilon',0}$ such that z is much closer to x than to all the other hypersurfaces corresponding to reflections in Γ_i.

Because ϕ_i^k and ϕ_j^k are almost local Riemannian isometries for sufficiently large k, we know that the map $\widetilde{\psi}_{ij}^k$ cannot change the length of a curve much, say by more than a factor $\in (\frac{1}{2}, 2)$. If we apply this result to a minimizing geodesic segment joining z to x and recall that by construction $d(\widetilde{\psi}_{ij}^k(x), x) < \epsilon'$, we obtain that $\widetilde{\psi}_{ij}^k(z)$ must be contained in the same hyperplane H as z.

Since $\alpha_{ij}^k(\gamma)$ must fix $\widetilde{\psi}_{ij}^k(z)$, this implies that $\alpha_{ij}^k(\gamma) = \gamma$. Because

$$d(\widetilde{\psi}_{ij}^k \circ \gamma(x), \gamma(x)) = d(\gamma \circ \widetilde{\psi}_{ij}^k(x), \gamma(x)) = d(\widetilde{\psi}_{ij}^k(x), x) < \epsilon'$$

for all $x \in U_{\epsilon',0}$, it follows that $d(\widetilde{\psi}_{ij}^k, \mathrm{id}) < \epsilon'$ is also valid on $\gamma \cdot U_{\epsilon',0}$. The claim now follows by repeating above argument a finite number of times.

Consider now a point $x \in U_{\epsilon'}$ and an element $\gamma \in \Gamma_i$. By a computation as above, we conclude that both $\gamma \circ \widetilde{\psi}_{ij}^k(x)$ and $\alpha_{ij}^k(\gamma) \circ \widetilde{\psi}_{ij}^k(x)$ are ϵ'-close to $\gamma(x) \in U_{\epsilon'}$ which is only possible if $\alpha_{ij}^k(\gamma) = \gamma$. Thus, α_{ij}^k is precisely the standard inclusion $\Gamma_i \to \Gamma_j$.

We can make $\widetilde{B}_{4r_i}(y_i) \setminus U_{\epsilon'}$ arbitrarily small as $\epsilon' \to 0$, and again use that $\widetilde{\psi}_{ij}^k$ almost preserves the lengths of geodesic segments to deduce that $d(\widetilde{\psi}_{ij}^k(x), x) < c(\epsilon')$ for all $x \in \widetilde{B}_{4r_i}(y_i)$ for some constant $c(\epsilon') \to 0$ as $\epsilon' \to 0$. This completes the proof of the lemma for the index pair i, j and there are only finitely many such pairs. □

By definition, the transition maps ψ_{ij}^k have the compatibility property that $\psi_{ij}^k = \psi_{i'j}^k \circ \psi_{ii'}^k$ wherever they are defined. We will now show that the unique lifts $\widetilde{\psi}_{ij}^k$ have a similar property:

Let $1 \leq i < i' < j \leq m$ be an index triple with $i \prec i'$, $i' \prec j$ and $i \prec j$. This implies that $B_{r_i}(y_i) \subset B_{4r'_i}(y_{i'}) \subset B_{16r_j}(y_j)$. Let us now fix a Γ_i-equivariant embedding ι of $\widetilde{B}_{4r_i}(y_i)$ into $\widetilde{B}_{16r_{i'}}(y_{i'})$. For sufficiently large k (and some sufficiently small $\epsilon > 0$) we now have the two uniquely determined Γ_i-equivariant embeddings $\widetilde{\psi}_{ij}^k$ and $\iota^{-1} \circ \widetilde{\psi}_{i'j}^k \circ \iota \circ \widetilde{\psi}_{ii'}^k : \widetilde{B}_{r_i}(y_i) \to \widetilde{B}_{4r_i}(y_i)$. Since they both lift ψ_{ij}^k and are close to the identity map $\mathrm{id}_{\widetilde{B}_{r_i}(y_i)}$, we can conclude as in the proof of Lemma 4.2.16 that these two maps are identical for sufficiently large $k \in \mathbb{N}$. (We chose $\epsilon' > 0$ sufficiently small to guarantee $U_{\epsilon'} \cap \widetilde{B}_{r_i}(y_i) \neq \emptyset$.) Again, we can simultanously perform this construction for all the finitely many triples i, i', j as above.

After passing to a subsequence if necessary, Lemma 4.2.16 implies that we can arrange that for all $1 \leq i < j \leq m$ with $i \prec j$ we have uniquely determined lifts $\widetilde{\psi}_{ij}^k$ which converge uniformly to the identity on the uniformizations $\widetilde{B}_{4r_i}(y_i)$. We can improve the quality of this convergence even further:

Lemma 4.2.17. *After passing to a subsequence, we have* $\widetilde{\psi}_{ij}^k \overset{\mathcal{C}^\infty}{\to} \mathrm{id}_{\widetilde{B}_{4r_i}(y_i)}$ *uniformly for all* i, j *with* $i \prec j$ *as* $k \to \infty$.

Proof. (cf. [Ba07, Sec. 3.2.3]) For a pair of indices $i \prec j$, consider a radial normal geodesic $\gamma : [0, 4r_i) \to B_{4r_i}(y_i)$ emanating from y_i. By the construction of the orbifold diffeomorphism ϕ_i^k, $\phi_i^k \circ \gamma$ is a normal geodesic in $B_{4r_i}^{O_k}(y_i^k) \subset O_k$ emanating from y_i^k. It follows that $\psi_{ij}^k \circ \gamma$ is a geodesic in $B_{16r_i}(y_i)$ with respect to the pull-back metric $(\phi_j^k)^* g_k$ defined on some open neighbourhood of $\overline{B_{4r_i}(y_i)}$ (which we can assume to be independent of k).

Let $\tilde{y}_i$ be the unique lift of y_i in $\widetilde{B}_{16r_i}(y_i)$ and fix an isometric identification $f : \mathbb{R}^n \to T_{\tilde{y}_i}\widetilde{B}_{16r_i}(y_i)$. Let $\exp : B_{4r_i}^{\mathbb{R}^n}(0) \to \widetilde{B}_{4r_i}(y_i)$ be the exponential map with $d_0 \exp = f$ and $(\widetilde{\phi}_i^k)^* g_k$ the Γ_i-invariant lifts of the pull-back metrics $(\phi_j^k)^* g_k$. They are defined on some open neighbourhood of the closure of $\widetilde{B}_{4r_i}(y_i)$ which is independent of k. From the corresponding property of the diffeomorphisms ϕ_i^k, it follows that

$$(\widetilde{\phi}_i^k)^* g_k \overset{\mathcal{C}^\infty}{\to} g_\infty|_{\widetilde{B}_{4r_i}(y_i)}. \tag{4.2.18}$$

Our discussion above now implies that the maps $\widetilde{\psi}_{ij}^k \circ \exp$ are exponential maps with respect to the lifted pull-back metrics $(\widetilde{\phi}_i^k)^* g_k$ centered at the points $\widetilde{\psi}_{ij}^k(\tilde{y}_i)$. By construction, these points converge to $\tilde{y}_i$ as $k \to \infty$.

The differentials $d_0(\widetilde{\psi}_{ij}^k \circ \exp)$ are uniformly bounded with respect to g_∞ because they are isometric with respect to metrics which become arbitrarily close to g_∞. Thus, after passing to a subsequence we can assume that $d_0(\widetilde{\psi}_{ij}^k \circ \exp) \to f_{ij}$ for some isometry $f_{ij} : \mathbb{R}^n \to T_{\tilde{y}_i} \widetilde{B}_{16r_i}(y_i)$.

We now use 4.2.18 to conclude that as $k \to \infty$, the maps $\widetilde{\psi}_{ij}^k \circ \exp$ converge smoothly uniformly to an exponential map $\exp_{ij} : B_{4r_i}^{\mathbb{R}^n}(0) \to \widetilde{B}_{4r_i}(y_i)$ with respect to the metric g_∞ and with differential $d_0 \exp_{ij} = f_{ij}$. But since the maps $\widetilde{\psi}_{ij}^k$ also converge uniformly to the identity on $\widetilde{B}_{4r_i}(y_i)$, we must have $f_{ij} = f$ and hence $\exp_{ij} = \exp$.

This concludes the proof that $\widetilde{\psi}_{ij}^k$ converges to the identity on $\widetilde{B}_{4r_i}(y_i)$ smoothly uniformly. The lemma follows once more from the finiteness of our index set. □

In particular, the lemma implies that $\psi_{ij}^k \overset{\mathcal{C}^\infty}{\to} \mathrm{id}_{B_{4r_i}(y_i)}$ as $k \to \infty$.

By interpolating the "almost compatible" local approximations ϕ_i^k, we wish to produce $\mathcal{C}^\infty$-almost isometric smooth embeddings

$$\Phi^k : W_\infty \to O_k$$

for large k such that $(\Phi^k)^* g_k \overset{\mathcal{C}^\infty}{\to} g_\infty|_{W_\infty}$ and $\Phi^k \to \mathrm{id}_{W_\infty}$ uniformly as $k \to \infty$.

Let $\lambda_i : O_\infty \to [0, \infty)$ be smooth functions supported on the $B_{r_i}(y_i)$ such that $\sum_i \lambda_i > 0$ on $\overline{W}_\infty$. We extend the function $\lambda_i \circ (\phi_i^k)^{-1}$ by zero outside $B_{r_i}^{O_k}(y_i^k)$ to a smooth function on all of O_k. Then for all $x_k \in O_k$ the finitely supported measures

$$\mu_k(x_k) := \sum_i (\lambda_i \circ (\phi_k^i)^{-1})(x_k) \cdot \delta_{(\phi_k^i)^{-1}(x_k)}$$

are well-defined, because the weight coefficient of one of the point masses is positive only when its location is well-defined. Note that $\mathrm{diam}(\mathrm{supp}(\mu_k(x_k))) \to 0$ uniformly as $k \to \infty$.

We will construct Φ^k by applying a center of mass construction to the measure valued map μ_k. Since the measures $\mu_k(x_k)$ have small supports this could be done pointwise using the center of mass construction for Riemannian *manifolds* introduced in [GK73, Ka77], see section 4.2.3, as long as the supports do not come too close to the singular locus. However, measures supported near the singular locus have in general no well-defined center of mass. We circumvent this obstacle by averaging measure valued *maps* instead of single measures. This works because the maps under consideration by Lemma 4.2.16 have distinguished lifts to local charts which then can be averaged using the center of mass for manifolds.

We will first construct averaged maps from open subsets of O_k to O_∞ and then show that these maps coincide on the overlap of the sets ranges. The inverse of these maps will then be the desired map $\Phi^k : W_\infty \to O_k$ for all sufficiently large k.

We first construct a family of open subsets $U_i \subset B_{y_i}(r_i) \subset O_\infty$ with the property that $d(\overline{U}_i, \mathrm{supp}(\lambda_j)) > 0$ for all $j < i$ and $W_\infty \subset \bigcup_i U_i$. In order to do so, we first find open sets $\mathrm{supp}(\lambda_i) \subset V_i \subset B_{r_i}(y_i)$ with $\overline{V}_i \subset B_{r_i}(y_i)$ and $d(\overline{V}_i, \mathrm{supp}(\lambda_i)) > 0$ and then set $U_i := B_{r_i}(y_i) \setminus (\overline{V}_1 \cup \cdots \cup \overline{V}_{i-1})$. Let $\widetilde{U}_i$ be the Γ_i-invariant lifts of U_i to $\widetilde{B}_{r_i}(y_i)$.

Let $\tilde{\lambda}_j$ denote the pull-backs of the functions λ_j from $B_{4r_i}(y_i)$ to $\widetilde{B}_{4r_i}(y_i)$ using the natural quotient projection. The center of mass map

$$(\tilde{q}_1, \ldots, \tilde{q}_m) \overset{c}{\mapsto} C\big(\tilde{\lambda}_1(\tilde{q}_1) \cdot \delta_{\tilde{q}_1} + \cdots + \tilde{\lambda}_m(\tilde{q}_m) \cdot \delta_{\tilde{q}_m}\big)$$

is defined and smooth close to the diagonal of $(\widetilde{B}^\infty_{4r_i}(y_i))^{\times m}$, cf. (4.2.12). If the function $\tilde{\lambda}_i$ vanishes identically, the function c is well-defined even if its ith entry is not.

We observe that

$$\mu_k \circ \phi_i^k = \sum_{i \prec j} (\lambda_j \circ \psi_{ij}^k) \cdot \delta_{\psi_{ij}^k}$$

on U_i for sufficiently large k: For $1 \le j \le m$ we know that $\phi_i^k(B_{r_i}) \cap \phi_j^k(B_{r_j}) \ne \emptyset$ implies $i \prec j$ or $j \prec i$, but for all $1 \le j < i$ we have by construction $\phi_i^k(U_i) \cap \phi_j^k(\mathrm{supp}(\lambda_j)) = \emptyset$.

To define its average on U_i, we note that

$$C\Big(\sum_{i \prec j} (\widetilde{\lambda}_j \circ \widetilde{\psi}_{ij}^k) \cdot \delta_{\widetilde{\psi}_{ij}^k}\Big) = c \circ (\widetilde{\psi}_{i1}^k, \ldots, \widetilde{\psi}_{im}^k) : \widetilde{U}_i \to \widetilde{B}_{4r_i}(y_i) \tag{4.2.19}$$

is Γ_i-equivariant and descends to a smooth orbifold map $U_i \to B_{4r_i}(y_i)$ which we define to be the average $C(\mu_k \circ \phi_i^k)$.

By Lemma 4.2.13, the maps (4.2.19) converge smoothly to the identity as $k \to \infty$, and hence also $C(\mu_k \circ \phi_i^k) \xrightarrow{\mathcal{C}^\infty} \mathrm{id}_{U_i}$. We now define F_i^k on $\phi_i^k(U_i)$ to be the smooth map

$$C(\mu_k \circ \phi_i^k) \circ (\phi_i^k)^{-1}.$$

To complete our construction, it suffices to verify that the maps F_i^k defined on the open sets $\phi_i^k(U_i)$ coincide on the overlaps $\phi_i^k(U_i) \cap \phi_{i'}^k(U_{i'})$ for sufficiently large k. To verify this, consider a nonempty intersection $\phi_i^k(U_i) \cap \phi_{i'}^k(U_{i'}) \ne \emptyset$ with $i \prec i'$. For sufficiently large k we observe that $\phi_i^k(U_{i'}) \cap \phi_l^k(\mathrm{supp}(\lambda_l)) = \emptyset$ for all $1 \le l < i'$. Hence, on $(\phi_i^k)^{-1}(\phi_i^k(U_i) \cap \phi_{i'}^k(U_{i'}))$ we have

$$\mu_k \circ \phi_i^k = \sum_{i' \prec j} (\lambda_j \circ \psi_{ij}^k) \cdot \delta_{\psi_{ij}^k}.$$

We now choose a Γ_i-invariant isometric embedding $\iota : \widetilde{B}_{4r_i}(y_i) \hookrightarrow \widetilde{B}_{16r_{i'}}(y_{i'})$ and use the compatibility of the lifts $\widetilde{\psi}_{ij}^k$ (cf. the discussion after Lemma 4.2.16) to compute

$$\iota \circ C\Big(\sum_{i' \prec j} (\widetilde{\lambda}_j \circ \widetilde{\psi}_{ij}^k) \cdot \delta_{\widetilde{\psi}_{ij}^k}\Big) = c \circ (\iota \circ \widetilde{\psi}_{i1}^k, \ldots, \iota \circ \widetilde{\psi}_{im}^k) = c \circ (\widetilde{\psi}_{i'1}^k \circ \iota \widetilde{\psi}_{ii'}^k, \ldots, \widetilde{\psi}_{i'm}^k \circ \iota \widetilde{\psi}_{ii'}^k) =$$

$$= c \circ (\widetilde{\psi}_{i'1}^k, \ldots, \widetilde{\psi}_{i'm}^k) \circ \iota \circ \widetilde{\psi}_{ii'}^k = C\Big(\sum_{i' \prec j} (\widetilde{\lambda}_j \circ \widetilde{\psi}_{i'j}^k) \cdot \delta_{\widetilde{\psi}_{i'j}^k}\Big) \circ \iota \circ \widetilde{\psi}_{ii'}^k$$

and hence (after projecting to O_∞)

$$C(\mu_k \circ \phi_i^k) = C(\mu_k \circ \phi_{i'}^k) \circ \psi_{ii'}^k$$

on $(\phi_i^k)^{-1}(\phi_i^k(U_i) \cap \phi_{i'}^k(U_{i'}))$, respectively

$$C(\mu_k \circ \phi_i^k) \circ (\phi_i^k)^{-1} = C(\mu_k \circ \phi_{i'}^k) \circ (\phi_{i'}^k)^{-1}.$$

This completes the proof that $F_i^k = F_{i'}^k$ on the overlap $\phi_i^k(U_i) \cap \phi_{i'}^k(U_{i'})$ for sufficiently large k. By the finiteness of our index set $\{1, \ldots, n\}$ we therefore have for sufficiently large k constructed well-defined orbifold maps F_k from an open subset of O_k into O_∞ satisfying

$$F_k \circ \phi_i^k|_{U_i} \xrightarrow{\mathcal{C}^\infty} \mathrm{id}_{U_i} \tag{4.2.20}$$

uniformly as $k \to \infty$.

For sufficiently large k, it follows from 4.2.20 that the maps F_k are embeddings with $W_\infty \subset \mathrm{Im}(F_k)$. We now define $\Phi^k := (F^k)^{-1}|_{W_\infty}$ to obtain orbifold embeddings $\Phi^k : W_\infty \to O_k$ such that $(\Phi^k)^* g_k \overset{\mathcal{C}^\infty}{\to} g_\infty|_{W_\infty}$ and $\Phi^k \to \mathrm{id}_{W_\infty}$ uniformly, as desired. It is automatic that also $\Phi_k^{-1}(p_k) \to p_\infty$. This completes the proof of

Theorem 4.2.21 (Smooth compactness). *$\mathcal{O}_{n,D,r_0,v_0}$ is compact with respect to the topology of smooth convergence.*

Remark 4.2.22. An alternative approach to the proof of Theorem 4.2.21 is to consider the frame bundle of orbifolds, thus reducing the problem to studying the $O(n)$-equivariant convergence of manifolds. This strategy is pursued in [KL08, Sec. 4].

Later on, we will also require a version of Theorem 4.2.21 for a situation where we only have uniform bounds on finitely many covariant derivatives of the curvature tensors. In other words, we are also interested in the space O_{n,l,C,r_0,v_0} where $r_0, v_0 > 0$, l is a (sufficiently large) natural number, $C : (0,\infty) \to (0,\infty)$ a function and a complete connected smooth pointed Riemannian orbifold (O,p) is contained in O_{n,l,C,r_0,v_0} if $\mathrm{vol}\, B(p,r_0) \geq v_0$ and on every ball $B(p,r)$ around p the curvature tensor and its covariant derivatives up to order l are bounded by the constant $C(r)$.

The proofs presented above for the compactness of $\mathcal{O}_{n,D,r_0,v_0}$ can also be applied to O_{n,l,C,r_0,v_0}, with the sole exception that the resulting orbifold structure on a Gromov-Hausdorff limit and the approximating diffeomorphisms are no longer smooth, but only have lower regularity $\mathcal{C}^{l'}$ (with some $l' < l$). More precisely, consider a sequence $(O_k, p_k) \in O_{n,l,C,r_0,v_0}$ with Gromov-Hausdorff limit (O_∞, p_∞). When we apply 4.2.6 to locally obtain a limit metric in geodesic coordinates on (O_∞, p_∞), we now have only control on the first l derivatives of this metric. Thus, we can now only conclude that O_∞ has the structure of an orbifold of regularity class $\mathcal{C}^{l-2}$. (Isometries between Riemannian manifolds of class $\mathcal{C}^l$ are only local diffeomorphisms of class $\mathcal{C}^{l-2}$.) Consequently, the diffeomorphisms which we interpolate in the proof of Theorem 4.2.21 are only of this lower regularity class, and we lose two further degrees of regularity in the proof of Lemma 4.2.17. As a special case of these considerations, we obtain

Theorem 4.2.23. *Let $r_0, v_0 > 0$ and let $C : (0, \infty) \to (0, \infty)$ be a function. Then $\mathcal{O}_{n,20,C,r_0,v_0}$ has the following precompactness property: every sequence of orbifolds $(O_k, p_k) \in O_{n,20,C,r_0,v_0}$ subconverges $\mathcal{C}^5$-smoothly to a $\mathcal{C}^{10}$-smooth Riemannian n-orbifold.*

5. Coarse stratification of roughly ≤ 2-dimensional Alexandrov spaces

5.1 Preliminaries

5.1.1 Alexandrov balls

Throughout this section, by a *segment* we mean more precisely a distance minimizing geodesic segment. Given two points x and y, then xy denotes one of the possibly several segments connecting these points.

All arguments from Alexandrov geometry used in this text will be local; we accordingly work in an appropriate class of local Alexandrov spaces.

Definition 5.1.1 (Alexandrov ball). An *Alexandrov ball* of curvature $\geq k$ is a local Alexandrov space (cf. 3.4 with curvature $\geq k$ of the form $X = B(x, \rho)$, $\rho > 0$, with the additional properties that the closed balls $\overline{B}(x, r)$ for $r \in (0, \rho)$ are metrically complete, and that for any two points $y, z \in X$ with $d(y, z) + d(x, y) + d(x, z) < 2\rho$ there is a segment yz joining y and z.

The first property can be viewed as metrical completeness "up to radius $< \rho$" and the second is a global form of the length space condition. We call x the *center* of the Alexandrov ball $B(x, \rho)$ and the minimal number $r \in [0, \rho]$ such that $\overline{B}(x, r) = B(x, \rho)$ its *radius* (with respect to x). An Alexandrov space may be regarded as an Alexandrov ball with infinite radius. For any pair of points in $B(x, \frac{\rho}{2})$ or, more generally, in a ball $B(y, r)$ with $d(x, y) + 2r \leq \rho$ there is at least one segment connecting them. For any triple of vertices in $B(x, \frac{\rho}{3})$ or, more generally, in a ball $B(y, r)$ with $d(x, y) + 3r \leq \rho$ geodesic triangles exist and satisfy triangle comparison, again due to Perelman's version of Toponogov's Theorem for Alexandrov spaces [BGP92, Thm. 3.2].

5.1.2 Strainers and cross sections

Let $X = B(x, R)$ be an Alexandrov ball with curvature ≥ -1. All points occuring in our discussion below are supposed to lie in $B(x, \frac{R}{3})$.

For a (small) constant $\theta > 0$, a θ-straight *n-strainer* of length l $(> l)$ in a point $x \in X$ consists of n pairs of points a_i, b_i at distance l $(> l)$ from x such that $\tilde{\angle}_p(a_i, b_i) \geq \pi - \theta$, $\tilde{\angle}_p(a_i, a_j) \geq \frac{\pi}{2} - \theta$ and $\tilde{\angle}_p(b_i, b_j) \geq \frac{\pi}{2} - \theta$ for all $i \neq j$, and $\tilde{\angle}_p(a_i, b_j) \geq \frac{\pi}{2} - \theta$ for all i, j. (All comparison angles are taken with respect to the comparison curvature value -1, i.e. in the hyperbolic plane.) We call the strainer $< \theta$-straight if it is θ'-straight for some $\theta' < \theta$. (Compare the definition of *burst points* in [BGP92, §5.2] and the definition of strainers [BBI01, §10.8.2].) We say that a strainer is *equilateral* if all points a_i, b_i have the same distance from x. We analogously define strainers with respect to other comparison curvature values.

Similarly, we define a θ-straight $n\frac{1}{2}$-*strainer* of length l $(> l)$ in x as such an n-strainer together with an additional point a_{n+1} at distance l $(> l)$ from x such that $\tilde{\angle}_p(a_{n+1}, a_i) \geq \frac{\pi}{2} - \theta$ and $\tilde{\angle}_p(a_{n+1}, b_i) \geq \frac{\pi}{2} - \theta$ for all $i \leq n$. We define an *infinitesimal* strainer as a configuration of directions in $\Sigma_x X$ satisfying analogous inequalities.

Due to the monotonicity of comparison angles, the existence of a strainer of length l at x implies the existence of strainers at x of the same type and straightness with any length $l' < l$.

For an n-strainer $(a_1, b_1, \ldots, a_n, b_n)$ in x we put

$$f_i := f_{a_i,b_i} := \frac{1}{2}(d(a_i, \cdot) - d(b_i, \cdot)) - \frac{1}{2}(d(a_i, x) - d(b_i, x))$$

and $f := f_{a_1,b_1,\ldots,a_n,b_n} := (f_1, \ldots, f_n)$, normalized to vanish in x. The functions f_i are 1-Lipschitz. We call the level sets $f^{-1}(t)$ the *cross sections* of the strainer and denote by $\Sigma_{y;a_i,b_i} = f_i^{-1}(f_i(y))$ and $\Sigma_{y;a_1,b_1,\ldots,a_n,b_n} = f^{-1}(f(y))$ the cross sections through the point y.

An important property of Alexandrov spaces and balls (see [BGP92, Cor. 10.8.21]) is that their (Hausdorff) *dimension* (of Alexandrov spaces or balls) can be characterized in terms of strainers as follows. There exists a constant $\bar{\theta}_{d\frac{1}{2}} > 0$ such that X has dimension $> d$ if and only if some point in X admits a $\bar{\theta}_{d\frac{1}{2}}$-straight $d\frac{1}{2}$-strainer of some positive length. (This in turn is implied by the existence of a $< \bar{\theta}_{d\frac{1}{2}}$-straight infinitesimal $d\frac{1}{2}$-strainer in some point.) These constants need not be extremely small, e.g. $\bar{\theta}_{1\frac{1}{2}}$ may be chosen arbitrarily in $(0, \frac{\pi}{2})$ and $\bar{\theta}_{2\frac{1}{2}}$ in $(0, \frac{\pi}{10})$.

Thus, if a sequence of d-dimensional pointed Alexandrov balls (X_i, x_i) with curvature ≥ -1 *collapses* to a pointed Alexandrov ball (X_∞, x_∞) of dimension $\leq k < d$, $(X_i, x_i) \to (X_\infty, x_\infty)$, then for any radius $r > 0$ the supremum of the lengths of $\bar{\theta}_{k\frac{1}{2}}$-straight $k\frac{1}{2}$-strainers in points of $\overline{B}(x_i, r)$ tends to zero as $i \to \infty$.

5.1.3 Comparing comparison angles

We will need to compare the comparison angles of (equilateral) 1-strainers with respect to different model spaces of constant curvature with curvature values in the interval $[-1, 0]$. We observe that for a triangle with fixed side lengths comparison angles increase monotonically with the comparison curvature value. We are interested in 1-strainers of bounded length. Consider a triangle Δ in euclidean space with two sides of length 2 and angle $\pi - \theta$ between them. We define the angle $\alpha(\theta)$ by letting $\pi - \alpha(\theta) < \pi - \theta$ be the comparison angle of Δ with respect to hyperbolic space of curvature -1, i.e. the corresponding angle of a triangle in hyperbolic space with the same side lengths as Δ.

Lemma 5.1.2. *The function $\alpha(\theta)$ is differentiable in $\theta = 0$ with $\alpha'(0^+) = (2\coth 2)^{\frac{1}{2}} < \frac{3}{2}$.*

Proof. Let $l = 4 - h$ denote the length of the third side of Δ, i.e. $l = 4\cos\frac{\theta}{2}$ and $h = \frac{1}{2}\theta^2 + O(\theta^3)$.

By the hyperbolic law of cosines, we have $\cosh l = (\cosh 2)^2 + (\sinh 2)^2 \cos\alpha = \cosh 4 - (\sinh 2)^2(1 - \cos\alpha) = \cosh 4 + \frac{1}{2}(\sinh 2)^2\alpha^2 + O(\alpha^3)$. Moreover, $\cosh l = \cosh(4 - h) = \cosh 4 - \sinh 4\ h + O(h^2) = \cosh 4 - \frac{1}{2}\sinh 4\ \theta^2 + O(\theta^3)$. Combining these equations we obtain that $\lim_{\theta\to 0}\alpha^2/\theta^2 = 2\coth 2$ and the lemma follows. □

We will frequently use the following application of the lemma. Let (X, x) be an Alexandrov space or ball of curvature ≥ -1. Suppose that (a, b) is a 1-strainer at x of length ≤ 2 with comparison angle $\tilde{\angle}_x(a, b) \geq \pi - \theta$ with respect to some comparison curvature value $k \in [-1, 0]$; and hence in particular with *euclidean* comparison angle (with respect to $k = 0$) $\geq \pi - \theta$. Then for sufficiently small θ, i.e. $\theta \in (0, \theta_0)$ for some universal $\theta_0 > 0$, Lemma 5.1.2 implies that the strainer (a, b) has comparison angle $\geq \pi - \frac{3}{2}\theta$ with respect to the comparison curvature value -1. For future reference, we also compute that $\alpha(\frac{\pi}{2}) < \frac{3\pi}{4}$ by solving the above equations with $l = 8^{\frac{1}{2}}$ for α. In other words, if we have an equilateral 1-strainer of length ≤ 2 which is $\frac{\pi}{2}$-straight with respect to some comparison curvature value in $[-1, 0]$, it is still $\frac{3\pi}{4}$-straight with respect to all other comparison curvature values in $[-1, 0]$.

5.2 Uniform local approximation by cones

Alexandrov spaces can in every point be arbitrarily well locally approximated by their tangent cone if one zooms in sufficiently far (see [BBI01, Thm. 10.9.3]). We need a quantitative version of this infinitesmial conelikeness, that is, we need uniform scales on which one can well approximate by cones. (Compare the scaling argument in [MT08, 3.5].)

Definition 5.2.1 (Local approximation by cones). We say that the Alexandrov ball $B(y,1)$ is in the point $z \in B(y,\frac{1}{2})$ on the scale $s \leq \frac{1}{2}$ μ-well *approximated by a cone* if the rescaled pointed ball $s^{-1} \cdot (B(z,s),z)$ has Gromov-Hausdorff distance $< \mu$ from the euclidean cone of radius 1 over some Alexandrov space with curvature ≥ 1 and with base point in the tip of the cone.

The base of the approximating cone may be empty, in which case the cone is just a point. The following result says that Alexandrov balls of dimension $\leq d$ or, more generally, which are *roughly* $\leq$ *d-dimensional* in the sense that $\bar{\theta}_{d\frac{1}{2}}$-straight $d\frac{1}{2}$-strainers must be very short, can be arbitrarily well locally approximated by cones of dimension $\leq d$.

Proposition 5.2.2 (Local approximation by cones on uniform scales). *For $d \in \mathbb{N}$ and $\sigma, \mu > 0$ exist scales $0 < s_{d\frac{1}{2}} = s_{d\frac{1}{2}}(\sigma,\mu) << s_1 = s_1(d,\sigma,\mu) << \sigma$ such that any Alexandrov ball $B(y,1)$ with curvature ≥ -1 and without $\bar{\theta}_{d\frac{1}{2}}$-straight $d\frac{1}{2}$-strainers of length $\geq s_{d\frac{1}{2}}$ can in every point $z \in B(y,\frac{1}{2})$ on some scale $s(z) \in [s_1,\sigma]$ be μ-well approximated by a cone of dimension $\leq d$.*

Proof. Consider a sequence of Alexandrov balls $B(y_k,1)$ with curvature ≥ -1 and without $\bar{\theta}_{d\frac{1}{2}}$-straight $d\frac{1}{2}$-strainers of length $\geq \frac{\sigma}{k}$, and suppose that there exist points $z_k \in B(y_k,\frac{1}{2})$ such that $B(y_k,1)$ can in z_k not be μ-well approximated by a cone of dimension $\leq d$ on any scale $s \in [\frac{\sigma}{k},\sigma]$. The $(B(y_k,1),z_k)$ Gromov-Hausdorff converge to a pointed Alexandrov ball $(B(y_\infty,1),z_\infty)$ with curvature ≥ -1 and dimension $\leq d$, and $z_\infty \in B(y_\infty,\frac{1}{2})$. Now $B(y_\infty,1)$ can on a sufficiently small scale $s' < \sigma$ be μ-well approximated in z_∞ by (the truncation at radius 1 of) its tangent cone. Since $\frac{\sigma}{k} < s'$ for large k, we obtain a contradiction. □

Throughout the remainder of this text, we fix some small value for σ, say $\sigma = \frac{1}{1000}$.

5.3 Islands without strainers

Building on 5.2.2 we will now divide our roughly $\leq d$-dimensional Alexandrov balls $B(y,1)$ with curvature ≥ -1 into two regions according to the (non-)existence of good 1-strainers on a uniform scale. We will show that the points without such strainers accumulate in "islands" which are uniformly separated from each other.

Definition 5.3.1 (Hump). For small $\theta > 0$ we call a point $z \in B(y,\frac{1}{2})$ a (θ,μ)-*hump* if the base of the approximating cone provided by 5.2.2 has diameter $< \pi - \frac{\theta}{2}$.

Let $H = H_{\theta,\mu} \subset B(y,\frac{1}{2})$ denote the subset of (θ,μ)-humps, and let $S = S_{\theta,\mu} \subset B(y,1)$ denote the subset of points admitting $< \theta$-straight 1-strainers of length $> \frac{1}{11}s_1(d,\sigma,\mu)$, cf. 5.2.2. We are interested in the distribution of the set $H - S$ for small θ and μ. (Our notation suppresses the dependence on d. Later we will only need the case $d = 2$.)

If the approximation accuracy μ is sufficiently small, then humps and non-humps have the following properties.

Lemma 5.3.2. *For sufficiently small $\theta > 0$ (i.e. $0 < \theta \leq \theta_0$ for some positive θ_0) there exists $\mu_0(\theta) > 0$ such that for $\mu \in (0, \mu_0(\theta)]$ the following holds:*

(i) A (θ, μ)-hump z admits no $\frac{\theta}{4}$-straight 1-strainers of length $\frac{1}{111}s(z)$, but

(ii) all points in the closed annulus $\overline{A}(z; \frac{1}{10}s(z), \frac{9}{10}s(z))$ do admit $\frac{\theta}{11}$-straight 1-strainers of length $> \frac{1}{11}s(z)$, i.e. $\overline{A}(z; \frac{1}{10}s(z), \frac{9}{10}s(z)) \subset S_{\theta,\mu}$. (iii) Moreover, if $y' \in B(y, \frac{1}{2})$ is no (θ, μ)-hump, then it admits $< \theta$-straight 1-strainers of length $> \frac{99}{100}s(z) > \frac{1}{11}s_1(d, \sigma, \mu)$, i.e. $B(y, \frac{1}{2}) \subset H_{\theta,\mu} \cup S_{\theta,\mu}$.

Proof. Property (i) follows from the fact that euclidean comparison angles (with respect to comparison curvature value 0) are larger than their hyperbolic equivalents (with respect to comparison curvature value -1).

For the proof of (ii), every point in $\overline{A}(z; \frac{1}{10}s(z), \frac{9}{10}s(z))$ can be made arbitrarily close to the midpoint of a segment of length $> \frac{2}{11}s(z)$ if μ is chosen sufficiently small. This implies the existence of a $< \theta$-straight 1-strainer of length $> \frac{1}{11}s(z)$ at every point in $\overline{A}(z; \frac{1}{10}s(z), \frac{9}{10}s(z))$ for sufficiently small μ.

To prove (iii), we observe that for every $\theta > 0$ and sufficiently small $\mu > 0$ every point y' which is not a (θ, μ)-hump z admits a 1-strainer (a, b) of length $> \frac{1}{11}s(y')$ which is $\frac{2}{3}\theta$-straight as a *euclidean* 1-strainer, i.e. with respect to the comparison curvature value 0. By Lemma 5.1.2 and the discussion afterwards this strainer then is also $< \theta$-straight with respect to comparison curvature value -1 for sufficiently small $\theta > 0$. □

Now we show that humps can be grouped into finitely many islands.

Proposition 5.3.3. *Let $d \in \mathbb{N}$ and $\sigma, \mu, \theta > 0$ such that $\mu \leq \mu_0(\theta)$. Suppose that $B(y, 1)$ is an Alexandrov ball with curvature ≥ -1 and without $\bar{\theta}_{d\frac{1}{2}}$-straight $d\frac{1}{2}$-strainers of length $\geq s_{d\frac{1}{2}}(\sigma, \mu)$. Then there exist finitely many (θ, μ)-humps $z_j \in H_{\theta,\mu}$ such that*

$$B(y, \frac{1}{2}) \subset \left(\bigcup_j B(z_j, \frac{1}{10}s(z_j)) \right) \cup S_{\theta,\mu}.$$

Moreover, $d(z_j, z_k) > \frac{9}{10}s(z_j)$ for $j \neq k$.

Proof. Let $z, z' \in H - S$. Then $z' \notin \overline{A}(z; \frac{1}{10}s(z), \frac{9}{10}s(z))$ and $z \notin \overline{A}(z'; \frac{1}{10}s(z'), \frac{9}{10}s(z'))$, i.e. $\frac{1}{s(z)}d(z, z'), \frac{1}{s(z')}d(z, z') \notin [\frac{1}{10}, \frac{9}{10}]$. If $z \in B(z', \frac{1}{10}s(z'))$ but $z' \notin B(z, \frac{1}{10}s(z))$, then $\frac{9}{10}s(z) < d(z, z') < \frac{1}{10}s(z')$, and so $B(z, \frac{1}{10}s(z)) \subset B(z', \frac{1}{9}s(z')) \subset B(z', \frac{1}{10}s(z')) \cup S$. There can be no infinite sequence of points $z, z', z'', \dots \in H - S$ such that $s(z) < \frac{1}{9}s(z') <$

$\frac{1}{9^2}s(z'') < \ldots$ because the scales take values in the bounded interval $[s_1, \sigma]$. Let $H_1 \subset H - S$ denote the subset of all points $z \in H - S$ for which no point $z' \in H - S$ exists with $\frac{9}{10}s(z) < d(z,z') < \frac{1}{10}s(z')$. We note that $\cup_{z\in H-S}B(z, \frac{1}{10}s(z)) \subset \cup_{z\in H_1}B(z, \frac{1}{10}s(z)) \cup S$.

By construction, the relation on H_1 defined by $z \sim z' :\Leftrightarrow z' \in B(z, \frac{1}{10}s(z))$ is reflexive and symmetric, and we have that $z \not\sim z' \Rightarrow d(z,z') > \frac{9}{10}s(z)$ and $> \frac{9}{10}s(z')$. To verify that the relation is also transitive, suppose that $z' \sim z \sim z''$ and $z' \not\sim z''$. Then $\frac{9}{20}(s(z') + s(z'')) < d(z', z'') \leq d(z,z') + d(z,z'') < \frac{1}{10}(s(z') + s(z''))$, a contradiction. Thus, "$\sim$" is an equivalence relation on H_1. We call an equivalence class a (θ, μ)-*island*. Note that the island containing z is contained in the intersection $\cap_{z'\sim z}B(z', \frac{1}{10}s(z'))$. Since inequivalent points $z', z'' \in H_1$ satisfy $d(z', z'') > \frac{9}{20}(s(z')+s(z''))$, islands are separated, $B(z', \frac{9}{20}s(z'))\cap B(z'', \frac{9}{20}s(z'')) = \emptyset$. In particular, there are only finitely many islands.

If $z \sim z'$ and $s(z') < 4s(z)$, then $\subset B(z, \frac{1}{10}s(z)) \cup S$. Let $R \subset H_1$ be a subset which contains exactly one representative z from each island with almost maximal scale value $s(z)$ (among its fellow islanders). Then $\cup_{z\in H_1}B(z, \frac{1}{10}s(z)) \subset \cup_{z\in R}B(z, \frac{1}{10}s(z)) \cup S$. R is finite and for any two distinct points $z', z'' \in R$ holds $d(z', z'') > \frac{9}{10}s(z')$. Altogether we obtain

$$H \subset \cup_{z\in H-S}B(z, \frac{1}{10}s(z)) \cup S \subset \cup_{z\in H_1}B(z, \frac{1}{10}s(z)) \cup S \subset \cup_{z\in R}B(z, \frac{1}{10}s(z)) \cup S$$

□

5.4 The 1-strained region

We will now study the geometry of the region $S = S_{\theta,\mu}$ of points admitting good ($< \theta$-straight) 1-strainers which was introduced in section 5.3.

In the estimates provided in this section we will abstain from giving explicit constants although this could be done in each case. Instead we use the symbols $c, c', \ldots$ to denote generic positive constants, i.e. constants which constantly change from estimate to estimate and which in each estimate (or assertion) take some fixed value independent of the other parameters $\theta, l, l', r, \lambda, \cdots > 0$ involved. The estimates hold (or become nontrivial) for sufficiently small values of the parameter θ, i.e. there exists some $\theta_0 > 0$ such that they hold for all $\theta \in (0, \theta_0]$. By decreasing the upper bound θ_0 for θ as we go along, we can also guarantee that the frequently occuring terms of the form $c\theta$ are as small as we wish. We will always assume that the upper bound θ_0 is sufficiently small such that the conclusions from Lemma 5.3.2 hold.

Throughout this section, let $X = B(x, 10)$ be an Alexandrov ball with curvature ≥ -1.

5.4.1 Local almost product structure

Let (a,b) be a $< 2\theta$-straight 1-strainer of length $> (1-\theta)$ at x. We want to apply the following considerations to $< \theta$-straight 1-strainers on scale $\frac{1}{11}s_1(d,\sigma,\mu)$ which we rescale to length ≈ 1. They are then not necessarily anymore $< \theta$-straight with respect to comparison curvature value -1, but by Lemma 5.1.2 they still are $< 2\theta$-straight. In this way, we avoid that our constants depend on the scale $s_1(d,\sigma,\mu)$.

The estimates given below express that near x there is on a certain small scale an almost product structure with a one-dimensional factor in the direction of the strainer. To begin with, the points near x admit 1-strainers close to (a,b) of comparable quality and almost the same length. More precisely, we have

$$\tilde{\angle}_{\cdot}(a,b) > \pi - c\theta \qquad \text{on } B(x,\theta) \tag{5.4.1}$$

with a certain constant $c > 1$. For future reference, let us denote by $C_0 > 1$ a constant such that (5.4.1) holds with $c = C_0$. To verify (5.4.1), note that the function $d(a,\cdot)+d(b,\cdot)$ along xx' has first derivative $< \theta$ in x and second derivative $< c'$ (in a barrier sense) along the whole segment. Thus $d(a,x') + d(x',b) < d(a,x) + d(x,b) + c''\theta^2 < d(a,b) + c'''\theta^2$ which implies $\tilde{\angle}_{x'}(a,b) > \pi - c\theta$. It follows that the function $f_{a,b} + d(b,\cdot) + \text{const} = -(f_{a,b} - d(a,\cdot)) + \text{const} = \frac{1}{2}(d(a,\cdot) + d(b,\cdot)) + \text{const}$ is $c\theta$-Lipschitz on $B(x,\theta)$.

One can define near x a coarse *flow* in the strainer direction which replaces the gradient flows of $d(a,\cdot)$ and $d(b,\cdot)$. For $t \in (-\theta,\theta)$ and $x' \in B(x,3\theta)$ let $\Phi_t^{a,b}(x')$ be the intersection point $ax'b \cap f_{a,b}^{-1}(f_{a,b}(x') + t)$. The resulting maps $\Phi_t^{a,b} : B(x,3\theta) \to X$ are well-defined only up to small ambiguity because the broken segments $ax'b$ need not be unique. Correspondingly, these maps are in general not continuous.

However, they are almost distance non-decreasing, $d(\Phi_t^{a,b}x_1, \Phi_t^{a,b}x_2) > (1-c\theta)d(x_1,x_2) - c'\theta|t|$. This follows from triangle comparison applied to $\Delta(x_1,x_2,a)$ or $\Delta(x_1,x_2,b)$ and the fact that $|t| \leq d(x_i, \Phi_t^{a,b}x_i) < (1+c'\theta)|t|$ because $f_{a,b}$ has slope ≈ 1 along ax_ib. They are also almost inverse to each other, i.e. $\Phi_{-t}^{a,b}\Phi_t^{a,b}$ is $c\theta|t|$-close to the identity (where it is defined). To see this, consider the triangle $\Delta(x', \Phi_t^{a,b}x', \Phi_{-t}^{a,b}\Phi_t^{a,b}x')$ and note that $\tilde{\angle}_{\Phi_t^{a,b}x'}(x', \Phi_{-t}^{a,b}\Phi_t^{a,b}x') < c''\theta$ by (5.4.1) and again $|t| \leq d(\Phi_t^{a,b}x', x'), d(\Phi_t^{a,b}x', \Phi_{-t}^{a,b}\Phi_t^{a,b}x') < (1+c'\theta)|t|$. Hence we also have an upper bound $d(\Phi_t^{a,b}x_1, \Phi_t^{a,b}x_2) \leq (1+c'''\theta)d(x_1,x_2) + c''''\theta|t|$, i.e. the $\Phi_t^{a,b}|_{B(x,\theta)}$ are $(1+c'''\theta t, c''''\theta t)$-quasi-isometric.

Let $x_1, x_2 \in B(x,\theta)$. Since $\angle \geq \tilde{\angle}$ and geodesic triangles in $\Sigma_{x_i}X$ have circumference $\leq 2\pi$, (5.4.1) yields $\angle_{x_1}(a,x_2) + \angle_{x_1}(b,x_2), \angle_{x_2}(a,x_1) + \angle_{x_2}(b,x_1) < \pi + c\theta$. Since also $\tilde{\angle}_{x_1}(a,x_2) + \tilde{\angle}_{x_2}(a,x_1), \tilde{\angle}_{x_1}(b,x_2) + \tilde{\angle}_{x_2}(b,x_1) > \pi - c'\theta$, we obtain that angles and comparison angles with a strainer direction almost coincide,

$$\angle_{x_1}(a,x_2) - \tilde{\angle}_{x_1}(a,x_2) < c\theta, \tag{5.4.2}$$

and

$$|\tilde{\angle}_{x_1}(a, x_2) + \tilde{\angle}_{x_1}(b, x_2) - \pi| < c\theta, \qquad |\tilde{\angle}_{x_1}(a, x_2) - \tilde{\angle}_{x_2}(b, x_1)| < c\theta. \tag{5.4.3}$$

As a consequence, $d(a, \cdot)$ is near x *almost affine* along segments in the sense that its slope is almost constant. More precisely,

$$d(a, x_1) - d(a, x_2) = d(x_1, x_2) \cos \alpha \tag{5.4.4}$$

with some angle α satisfying $|\alpha - \tilde{\angle}_{x_1}(a, x_2)| < c\theta$, as follows from (5.4.2, 5.4.3) and the monotonicity of the cosine by integrating the derivative of $d(a, \cdot)$ along x_1x_2. Estimates of the same form hold for $d(b, \cdot)$ and $f_{a,b}$. (Note that (5.4.4) and the corresponding estimate $d(b, x_1)-d(b, x_2) = d(x_1, x_2) \cos \beta$ with $\beta = \pi-\alpha'$ satisfying $|\alpha'-\tilde{\angle}_{x_1}(a, x_2)| < c'\theta$ yield that $f_{a,b}(x_1) - f_{a,b}(x_2) = d(x_1, x_2) \cos \alpha''$ with $\cos \alpha'' = \frac{1}{2}(\cos \alpha + \cos \alpha')$, i.e. $|\alpha'' - \tilde{\angle}_{x_1}(a, x_2)| < c''\theta$.) It follows that there exists a constant $L > 0$ such that the function

$$f_{a,b} - \frac{f_{a,b}(x_2) - f_{a,b}(x_1)}{d(x_2, x_1)} d(x_1, \cdot) \tag{5.4.5}$$

and the analogous functions derived from $d(a, \cdot)$ and $d(b, \cdot)$ are $L\theta$-Lipschitz continuous along the segment x_1x_2.

In particular, the *cross sections* of a strainer are topological hypersurfaces *almost perpendicular* to it: If $x_1, x_2 \in f_{a,b}^{-1}(t) \cap B(x, \theta)$, then (5.4.2, 5.4.3) imply

$$\frac{\pi}{2} - c\theta < \tilde{\angle}_{x_1}(a, x_2) \leq \angle_{x_1}(a, x_2) < \frac{\pi}{2} + c\theta \tag{5.4.6}$$

because the comparison angles of the triangles $\Delta(x_1, x_2, a)$ amd $\Delta(x_1, x_2, b)$ almost agree. Moreover, the functions $d(a, \cdot)$, $d(b, \cdot)$ and $f_{a,b}$ are $L\theta$-Lipschitz continuous on any segment with endpoints in the same (piece of) cross section $f_{a,b}^{-1}(t) \cap B(x, \theta)$.

The maps $\Phi_t^{a,b}$ can be used to compare cross sections, since by their definition they satisfy $\Phi_t^{a,b}(f_{a,b}^{-1}(t') \cap B(x, 2\theta)) \subset f_{a,b}^{-1}(t' + t)$. Regarding the *size* of cross sections of n-strainers $(a_1, b_1, \dots, a_n, b_n)$ we obtain from (5.4.1, 5.4.6): If points $x_1, x_2 \in f_{a_1,b_1,\dots,a_n,b_n}^{-1}(t) \cap B(x, \theta)$ have distance l, then they admit $c\theta$-straight $n\frac{1}{2}$-strainers of length l.

The cross sections of the 1-strainer (a, b) are *connected*: Let $x_1, x_2 \in f_{a,b}^{-1}(t) \cap B(x, \theta)$. For the midpoint m of (a segment) x_1x_2 holds $|f_{a,b}(m)-t| < c\theta d(x_1, x_2)$. As above, by *amb* with $f_{a,b}^{-1}(t)$ we find an almost midpoint $y \in f_{a,b}^{-1}(t)$ at distance $< (\frac{1}{2} + c\theta)d(x_1, x_2)$ from x_1 and x_2. Iterating this procedure yields a continuous curve in $f_{a,b}^{-1}(t) \cap B(x, 3\theta)$ connecting x_1 and x_2. (Here one uses the earlier estimates with the parameter 3θ instead of θ.)

To simplify notation, let us put $\Sigma_{y;a,b}^o := \Sigma_{y;a,b} \cap B(x, \theta)$.

Lemma 5.4.7 (Projecting to cross sections). *Let $y, y_1 \in B(x,\theta)$ and let z_1 be the intersection point $ay_1b \cap \Sigma_{y;a,b}$. Then*

$$\left| d(y_1,z_1) - d(y,y_1) \cdot |\cos \tilde{\angle}_y(a,y_1)| \right| < c\theta d(y,y_1) \tag{5.4.8}$$

and

$$|d(y,z_1) - d(y,y_1) \sin \tilde{\angle}_y(a,y_1)| < c'\theta d(y,y_1). \tag{5.4.9}$$

In particular,

$$|d(y,z_1)^2 + d(z_1,y_1)^2 - d(y,y_1)^2| < c''\theta d(y,y_1)^2. \tag{5.4.10}$$

Proof. We put $l = d(y,y_1)$ and $\alpha_1 = \tilde{\angle}_y(a,y_1)$.

We have $|f_{a,b}(y_1) - f_{a,b}(y)| \leq d(y_1,z_1) \leq (1+c\theta)|f_{a,b}(y_1) - f_{a,b}(y))|$ and, due to (5.4.4) and the remark thereafter, $f_{a,b}(y_1) - f_{a,b}(y) = -l\cos\alpha_1'$ with $|\alpha_1' - \alpha_1| < c\theta$. This yields (5.4.8).

To estimate $d(y,z_1)$, we consider a comparison triangle for $\Delta(y,y_1,z_1)$. In view of (5.4.3), we may exchange a and b, and therefore assume without loss of generality that $z_1 \in y_1a$. Then $\alpha_1 > \frac{\pi}{2} - c\theta$, cf. (5.4.6). Regarding $\tilde{\angle}_{y_1}(z_1,y)$, we have $\tilde{\angle}_{y_1}(a,y) \leq \tilde{\angle}_{y_1}(z_1,y) \leq \angle_{y_1}(z_1,y) = \angle_{y_1}(a,y)$ and hence

$$|\tilde{\angle}_{y_1}(z_1,y) - (\pi - \alpha_1)| < c\theta$$

because of (5.4.2) and $\pi - c\theta < \tilde{\angle}_{y_1}(a,y) + \alpha_1 \leq \pi$. This information implies (5.4.9). (Whether we use a hyperbolic or euclidean comparison triangle to compute the length of the side $y_1'z'$ corresponding to y_1z causes only a difference by a factor $< \frac{\sinh l}{l} < 1 + \theta^2 < 1 + c\theta$ (due to the distortion of the exponential map for hyperbolic plane up to radius l) and we may therefore work with a euclidean one. Then $|\frac{1}{l}d(y',z_1') - \sin\tilde{\angle}_{y_1}(z_1,y)| < |\cos(\pi-\alpha_1) - \cos\tilde{\angle}_{y_1}(z_1,y)| + c\theta < c'\theta$.)

Finally, (5.4.10) is a direct consequence. □

5.5 The roughly ≤ 2-dimensional case

We assume now in addition that $B(x,10)$ is *roughly ≤ 2-dimensional* in the sense that there are no $\bar{\theta}_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainers of length λ for some (very small) $\lambda > 0$.

5.5.1 Cross sections of 1-strainers

We continue our discussion in section 5.4. The next two results express that the cross sections of good 1-strainers are now roughly ≤ 1-dimensional.

Lemma 5.5.1. *Let (a,b) be a $C_0\theta$-straight $>(1-\theta^{\frac{1}{2}})$-long 1-strainer at a point $y \in B(x,5)$. Let y_1y_2 be a segment of length l with endpoints in $B(y,\frac{\theta}{3})$ and with midpoint m. Let y' be a point with $d(m,y') > \theta l$. Suppose that $f_{a,b}$ is $3L\theta$-Lipschitz on the segments y_1y_2 and $my' \cap B(m,\theta l)$. Then $\angle_m(y_i,y') < c\theta$ for $i=1$ or 2, if $\lambda = \lambda(\theta,l)$ is sufficiently small.*

If also $d(y_1,y') \leq \frac{l}{3}$, then $\angle_m(y_1,y') < c\theta$.

Proof. Let $z' \in my'$ be the point at distance θl from m. Since f_{y_1,y_2} has slope $\equiv 1$ on y_1y_2, there exists a point $z \in y_1y_2 \cap \overline{B}(m,\theta l)$ with $f_{y_1,y_2}(z) = f_{y_1,y_2}(z')$. It satisfies

$$|f_{a,b}(z) - f_{a,b}(z')| \leq 6L\theta^2 l \tag{5.5.2}$$

because $f_{a,b}$ is $3L\theta$-Lipschitz on zm and mz'.

The Lipschitz assumption also implies that $|\angle.(y_i,a) - \frac{\pi}{2}|, |\angle.(y_i,b) - \frac{\pi}{2}| < c\theta$ on (the interior of) y_1y_2 and hence $|\tilde\angle.(y_i,a) - \frac{\pi}{2}|, |\tilde\angle.(y_i,b) - \frac{\pi}{2}| < c'\theta$ by (5.4.2). Thus the quadrupel (a,b,y_1,y_2) is a $c'\theta$-straight 2-strainer at z. The $1\frac{1}{2}$-strainer (y_1,y_2,z') at z is $c\theta$-straight by (5.4.6). We consider now the $1\frac{1}{2}$-strainer (a,b,z') at z and estimate $\tilde\angle_z(a,z')$. Since $d(z,z') \leq 2\theta l$ and $f_{a,b} - d(a,\cdot)$ is $c\theta$-Lipschitz on $B(y,\theta)$, (5.5.2) translates to $|d(a,z) - d(a,z')| < c'\theta^2 l$. With (5.4.4) it follows that $d(z,z')|\cos\tilde\angle_z(a,z')| < c'\theta^2 l + c''\theta d(z,z') < c'''\theta^2 l$.

If $d(z,z') > 2c'''\bar\theta_{2\frac{1}{2}}^{-1}\theta^2 l$ (with the constant c''' from the last estimate), then $|\cos\tilde\angle_z(a,z')| < \frac{1}{2}\bar\theta_{2\frac{1}{2}}$ and $|\tilde\angle_z(a,z') - \frac{\pi}{2}| < \bar\theta_{2\frac{1}{2}}$. Similarly, $|\tilde\angle_z(b,z') - \frac{\pi}{2}| < \bar\theta_{2\frac{1}{2}}$, and it follows that (a,b,y_1,y_2,z') is a $<\bar\theta_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainer at z with length $> 2c'''\bar\theta_{2\frac{1}{2}}^{-1}\theta^2 l$. This is a contradiction if $\lambda = \lambda(\theta,l)$ is sufficiently small, and we conclude that $d(z,z') < c''''\theta^2 l$ and consequently $\tilde\angle_m(z,z') < c\theta$.

Suppose without loss of generality that $z \in my_1$. Then $\angle_m(z,z') - \tilde\angle_m(z,z') \leq \angle_m(y_1,z') - \tilde\angle_m(y_1,z') < c'\theta$, where the last inequality follows from (5.4.2) after rescaling by the factor l^{-1}. Thus $\angle_m(y_1,y') = \angle_m(z,z') < c''\theta$, which shows the first assertion.

For the second assertion, suppose that $d(y_1,y') \leq \frac{l}{3}$ but $\angle_m(y_2,y') < c\theta$. Then y' is close to my_2 or y_2 is close to my'. Since $d(y_2,y') \geq \frac{2l}{3}$, only the second alternative can occur and $d(m,y') \gtrapprox \frac{l}{2} + \frac{2l}{3} > l$. On the other hand $d(m,y') \leq d(m,y_1) + d(y_1,y') < l$, a contradiction. Thus $\angle_m(y_1,y') < c\theta$. □

Lemma 5.5.3 (Roughly one-dimensional cross section). *Let $l' \leq l \leq \frac{\theta}{50}$, $y \in B(x,5)$ and let (a,b) be a $C_0\theta$-straight $>(1-\theta^{\frac{1}{2}})$-long 1-strainer at y. Suppose that* $\operatorname{diam}(\Sigma^o_{y;a,b}) >$ $40l$.

(i) Then there exists a point $y' \in \Sigma_{y;a,b} \cap B(y,l)$ and a segment $y'm$ of length $\geq 10l$, such that $f_{a,b}$ is $3L\theta$-Lipschitz on $y'm$ and $\Sigma_{y;a,b} \cap B(y,l) \subset N_{c\theta l}(y'm)$, if $\lambda = \lambda(\theta,l)$ is sufficiently small.

(ii) Moreover, if $\mathrm{diam}(\Sigma_{y;a,b}\cap B(y,l)) < \frac{199}{100}l$*, then* y' *can be chosen in* $B(y,\frac{199}{200}l)$ *so that* $\Sigma_{y;a,b}\cap B(y',r)\subset N_{c\theta r}(y'm)$ *for all* $r\in[l',l]$*, if* $\lambda=\lambda(\theta,l')$ *is sufficiently small.*

Proof. (i) By assumption, there is $q\in\Sigma_{y;a,b}$ with $d(y,q)=20l$. (Recall from section 5.4.1 that $\Sigma_{y;a,b}$ is path connected near y.) Let m be the midpoint of yq. Since $f_{a,b}$ is $L\theta$-Lipschitz on yq, we have that $|f_{a,b}(m)-f_{a,b}(y)|<10L\theta l$. Thus for every point $z\in\Sigma_{y;a,b}\cap B(y,l)$ the function $f_{a,b}$ has along the segment mz Lipschitz constant $<\frac{10}{9}L\theta+L\theta=\frac{19}{9}L\theta<3L\theta$, cf. (5.4.5), and 5.5.1 yields that $\angle_m(y,z)<c\theta$. It follows that the segments mz have pairwise angles $<2c\theta$ and are all contained in the $c'\theta l$-neighbourhood of an almost longest one among them. We choose y' as its endpoint.

(ii) By part (i), the piece of cross section $\Sigma_{y;a,b}\cap B(y,l)$ is $c\theta l$-close to a subsegment $y'z\subset y'm$ where $z\in y'm$ is a point with $d(y,z)=l$ and $|d(z,m)+l-d(y,m)|<c'\theta l$. Hence $d(y,y')<(\frac{99}{100}+c''\theta)l$. The points in $\Sigma_{y;a,b}\cap B(y,l)$, which are further away from m than y', must be $2c\theta l$-close to y', i.e. $d(m,\cdot)|_{\Sigma_{y;a,b}\cap B(y,l)}$ assumes a maximum in $\Sigma_{y;a,b}\cap\overline{B}(y',2c\theta l)$. We replace y' by this maximum and then have $d(y,y')<\frac{199}{200}l$. Also by (i), $\Sigma_{y;a,b}$ contains no $\frac{l}{1000}$-long $\frac{\pi}{2}$-straight 1-strainer at y'.

Suppose that $\Sigma_{y;a,b}$ contains a $\theta l'$-long $\frac{\pi}{2}$-straight 1-strainer at y'. Using (i) at the point y' and on the scale $\theta l'$, it follows that for sufficiently small $\lambda=\lambda(\theta,l')$ there exists a segment τ of length almost $2\theta l'$, say, of length $\frac{199}{100}\theta l'$ along which $f_{a,b}$ is $c\theta$-Lipschitz and such that y' lies at distance $<c'\theta^2l'$ from the midpoint of τ. (When applying part (i) to points nearby y, the 1-strainer (a,b) may only be $c\theta$-straight for some constant $c>C_0$ and $(1-2\theta^{\frac{1}{2}})$-long at these points, and we use a version of part (i) with appropriate different constants.)

We consider $d(m,\cdot)$ along the middle third τ' of τ. The function $f_{a,b}$ is $c\theta$-Lipschitz along τ and along all segments zm initiating in interior points z of τ'. Lemma 5.5.1 implies for sufficiently small $\lambda=\lambda(\theta,l')$ that these segments zm have angles $<c'\theta$ with τ'. This means that $d(m,\cdot)$ has slope $\approx\pm1$ along τ', i.e. the directional derivatives of $d(m,\cdot)$ in directions tangent to τ' take values in $[-1,-1+c''\theta^2)\cup(1-c''\theta^2,1]$. (See e.g. [BGP92, Sec. 11] for a discussion of directional derivatives of distance functions.) If at some interior point z_0 of τ' the directional derivatives in the two antipodal directions tangent to τ' are both negative, then $d(m,\cdot)$ decays with slope ≈ -1 along both subsegments of τ' with initial point z_0. In particular, z_0 is a maximum of $d(m,\cdot)|_{\tau'}$. If such a point z_0 does not exist, then $d(m,\cdot)|_{\tau'}$ is almost affine, i.e. with respect to an appropriate orientation of τ' it increases with slope ≈ 1 along the whole segment.

Since y' is a maximum of $d(m,\cdot)|_{\Sigma_{y;a,b}\cap B(y',\frac{l}{200})}$, it follows that $d(m,\cdot)|_{\tau'}$ attains a maximum at a point $y''\in\tau'$ close to the midpoint of τ', more precisely, at distance $<c''\theta^2l'$ from y'. Furthermore, there are at least two segments σ_1 and σ_2 connecting y'' to m whose initial directions $\dot\sigma_i(0)$ (with respect to unit speed parametrizations starting at y'') are close

to the two antipodal directions of τ' at y''. Each endpoint of τ lies at distance $< c'''\theta^2 l'$ from one of the two segments σ_1 and σ_2, and hence $d(\sigma_1(\theta l'), \sigma_2(\theta l')) > \frac{19}{10}\theta l'$.

Applying part (i) at y'' on the scales between $\theta l'$ and l yields that for $\lambda(\theta, l')$ sufficiently small the continuous function $t \mapsto \frac{1}{t} d(\sigma_1(t), \sigma_2(t))$ on $[\theta l', l]$ takes values close to 0 and 2, i.e. in $[0, c\theta) \cup (2 - c\theta, 2]$. However, by the above, it has value ≈ 0 for $t = \frac{l}{1000}$ and value ≈ 2 for $t = \theta l'$, a contradiction. (Note that the smaller the scale, the smaller λ has to be, and there exists a λ which serves simultaneously for all scales $s \in [\theta l', l]$.)

Thus $\Sigma_{y;a,b}$ contains no $\theta l'$-long $\frac{\pi}{2}$-straight 1-strainer at y'. It follows, again by part (i) on the scale $\theta l'$, that $\operatorname{diam}(\Sigma_{y;a,b} \cap \partial B(y', \theta l')) < c\theta^2 l'$, and hence every point $z \in \Sigma_{y;a,b}$ at distance $d \in [\theta l', l]$ from y' has distance $< c\theta d$ from $y'm$. This proves the assertion. □

Lemma 5.5.4 (Spreading 1-strainers). *Let l, l', y and (a, b) be as in 5.5.3. Let $z, z_1, z_2 \in \Sigma_{y;a,b} \cap B(y, l)$ such that (z_1, z_2) is a $\frac{\pi}{2}$-straight 1-strainer of length $r \in [l', \frac{1}{2}l]$ at z. Then there exists a constant $C_1 \geq C_0$ such that for sufficiently small $\lambda = \lambda(\theta, l')$ holds:*

(i) The 1-strainer (z_1, z_2) at z is $< C_1\theta$-straight. If moreover the 1-strainer (a, b) is $< 2\theta$-straight at y and has length $> (1-\theta)$, the 2-strainer (a, b, z_1, z_2) at z is $< C_1\theta$-straight.

(ii) All points $u \in azb \cap B(z, \min(100\theta^{-\frac{3}{4}} r, \theta))$ admit $< C_1\theta$-straight $\frac{r}{2}$-long 1-strainers contained in $\Sigma_{u;a,b}$.

Proof. (i) This is a consequence of 5.5.3(i) applied at the point z on the scale r, and of 5.4.6.

(ii) Now we use the $(1+c\theta, c'\theta t)$-quasi-isometry property of the maps $\Phi_t^{a,b}$ up to distance $\approx \theta$ from y. We may assume that $u = \Phi_t^{a,b} z$ with $|t| < \min(100\theta^{-\frac{3}{4}} r, \theta)$ and put $u_i = \Phi_t^{a,b} z_i$. Then $d(u, u_i) < (1 + c\theta)r + c'\theta t$ and $d(u_1, u_2) > (1 - c\theta)\frac{99}{50} r - c'\theta t$, using that $d(z_1, z_2) \approx 2r$ according to part (i). We obtain a $\frac{\pi}{3}$-straight 1-strainer (u_1, u_2) at u contained in $\Sigma_{u;a,b}$ and with length $\approx r$. Close to the midpoints of the segments uu_i we find a $\frac{\pi}{2}$-straight 1-strainer (u_1', u_2') at u contained in $\Sigma_{u;a,b}$ and with length $\frac{r}{2}$. (Compare the argument for the connectivity of cross sections.) Applying part (i) again on the scale $\frac{r}{2}$ yields the $c\theta$-straight $\frac{r}{2}$-long 1-strainer, once we can make sure that $\operatorname{diam}(\Sigma_{u;a,b}^o) > 20r$. But this follows from our assumption that $\operatorname{diam}(\Sigma_{y;a,b}^o) > 40l \geq 40r$ by using the maps $\Phi_t^{a,b}$ as before. □

Remark 5.5.5. As in part (i) of the lemma, we obtain that even if the 1-strainer (z_1, z_2) is only $\frac{3\pi}{4}$-straight, it is also $< C_1\theta$-straight and hence in particular $\frac{\pi}{2}$-straight.

5.5.2 Edges

In view of the local product structure, the points y' obtained in 5.5.3(ii) can be considered as points “near the edge” of our space. They are characterized by the property that they admit no long 1-strainers contained in the cross section. We will now investigate the geometry near the edge, keeping the assumption of rough 2-dimensionality from section 5.5. Good 1-strainers at points near the edge must be almost perpendicular to the cross section if they are not too short:

Lemma 5.5.6 (Almost unique 1-strainers near the edge). *Let l, y and (a, b) be as in 5.5.3. Suppose that (a', b') is a $c\theta^{\frac{1}{2}}$-straight 1-strainer of length $\geq l$ at y.*

(i) If $\Sigma_{y;a,b}$ contains no $\frac{\pi}{2}$-straight $c'l$-long 1-strainer at y and $\lambda = \lambda(\theta, l)$ is sufficiently small, then $\angle_y(a, a'), \angle_y(a, b'), \angle_y(b, a'), \angle_y(b, b') \notin [\frac{\pi}{100}, \frac{99}{100}\pi]$.

(ii) If $\Sigma_{y;a,b}$ contains no $\frac{\pi}{2}$-straight θl-long 1-strainer at y and $\lambda = \lambda(\theta, l)$ is sufficiently small, then $\angle_y(a, a'), \angle_y(a, b'), \angle_y(b, a'), \angle_y(b, b') \notin [c'\theta^{\frac{1}{2}}, \pi - c'\theta^{\frac{1}{2}}]$.

In both cases, y admits no $c\theta^{\frac{1}{2}}$-straight l-long 2-strainer.

Proof. There exists a $c\theta^{\frac{1}{2}}$-straight l-long 1-strainer (y_1, y_2) at y. (Choose $y_1 \in ya'$ and $y_2 \in yb'$.) We put $\alpha_i = \tilde{\angle}_y(a, y_i)$.

By 5.4.4 and the remark afterwards, $|f_{a,b}(y_i) - f_{a,b}(y) + l\cos\alpha_i| < c\theta l$. Thus for $u_1 = ay_1b \cap \Sigma_{y_2;a,b}$ holds $(1 - c\theta)d(y_1, u_1) < |f_{a,b}(y_1) - f_{a,b}(y_2)| \leq |f_{a,b}(y_1) - f_{a,b}(y)| + |f_{a,b}(y) - f_{a,b}(y_2)|$ and

$$\frac{1}{l}d(y_1, u_1) < |\cos\alpha_1| + |\cos\alpha_2| + c\theta,$$

compare the proof of 5.4.7.

Let $z_i = ay_ib \cap \Sigma_{y;a,b}$. According to 5.4.7, we have $|d(y, z_i) - l\sin\alpha_i| < c\theta l$. Now the rough one-dimensionality of the cross section $\Sigma_{y;a,b}$, i.e. 5.5.3(i) applied on the scale l, and our assumption in part (ii) yield that $d(y, y') < c'\theta l$ and $|d(z_1, z_2) - |d(y, z_1) - d(y, z_2)|| < c''\theta l$. Hence $|\frac{1}{l}d(z_1, z_2) - |\sin\alpha_1 - \sin\alpha_2|| < c'''\theta$. With the metric properties of the maps $\Phi_t^{a,b}$ it follows that

$$\left|\frac{1}{l}d(u_1, y_2) - |\sin\alpha_1 - \sin\alpha_2|\right| < c\theta.$$

Since (y_1, y_2) is $c\theta^{\frac{1}{2}}$-straight, we have $d(y_1, y_2) > (2 - c\theta)l$, and (5.4.10) implies

$$(|\cos\alpha_1| + |\cos\alpha_2|)^2 + (\sin\alpha_1 - \sin\alpha_2)^2 > 4 - c'\theta.$$

Writing $\alpha_i = \frac{\pi}{2} + \beta_i$, the last inequality becomes $4\sin^2\frac{|\beta_1|+|\beta_2|}{2} = (\sin|\beta_1| + \sin|\beta_2|)^2 + (\cos\beta_1 - \cos\beta_2)^2 > 4 - c'\theta$ and

$$\sin\frac{|\beta_1| + |\beta_2|}{2} > 1 - c\theta.$$

Thus $|\beta_i| > \frac{\pi}{2} - c'\theta^{\frac{1}{2}}$ and $\tilde{\angle}_y(a, y_1), \tilde{\angle}_y(a, y_2), \tilde{\angle}_y(b, y_1), \tilde{\angle}_y(b, y_2) \notin [c'\theta^{\frac{1}{2}}, \pi - c'\theta^{\frac{1}{2}}]$.

To pass from comparison angles to angles, we note that if e.g. $\tilde{\angle}_y(a, y_1) < c\theta^{\frac{1}{2}}$, then $\angle_y(a, b') = \angle_y(a, y_2) \geq \tilde{\angle}_y(a, y_2) > \pi - c'\theta^{\frac{1}{2}}$. This shows (ii).

Under the assumption of part (i) we obtain weaker but still useful estimates. In this case, $|d(z_1, z_2) - |d(y, z_1) - d(y, z_2)|| < (2c' + c''\theta)l < 3c'l$ and we obtain the estimates $\left|\frac{1}{l}d(u_1, y_2) - |\sin\alpha_1 - \sin\alpha_2|\right| < c''$ and $|\beta_i| > \frac{\pi}{2} - c'''$. The constant c''' depends on the constant c' in the hypothesis of (i) and can be made arbitrarily small by choosing c' small enough. Assertion (i) follows. □

Remark 5.5.7. Similarly, one shows e.g. that if $\Sigma_{y;a,b}$ contains no $\frac{\pi}{2}$-straight $\theta^{\frac{1}{2}}l$-long 1-strainer at y, then $\angle_y(a, a'), \angle_y(a, b'), \angle_y(b, a'), \angle_y(b, b') \notin [c'\theta^{\frac{1}{4}}, \pi - c'\theta^{\frac{1}{4}}]$, if $\lambda = \lambda(\theta, l)$ is sufficiently small.

We make the following choice of scales to quantify rough edges.

Definition 5.5.8 (Edgy points). A point $y \in B(x, 5)$ is called θ-*edgy* relative to a $< 2\theta$-straight 1-strainer (a, b) of length $> (1 - \theta)$ at y if $\operatorname{diam}(\Sigma^o_{y;a,b}) > \theta^{\frac{5}{2}}$ and $\Sigma_{y;a,b}$ contains no $\frac{\pi}{2}$-straight θ^4-long 1-strainer at y.

We say that $y \in B(x, 5)$ is θ-*weakly edgy* relative to a $< 2\theta$-straight 1-strainer (a, b) of length $> (1 - \theta)$ at y if $\operatorname{diam}(\Sigma^o_{y;a,b}) > \frac{1}{2}\theta^{\frac{5}{2}}$ and $\Sigma_{y;a,b}$ contains no $\frac{\pi}{2}$-straight $2\theta^4$-long 1-strainer at y.

We call $y \in B(x, 5)$ θ-*strongly edgy* relative to a $< 2\theta$-straight 1-strainer (a, b) of length $> (1 - \theta)$ at y if $\operatorname{diam}(\Sigma^o_{y;a,b}) > \frac{4}{3}\theta^{\frac{5}{2}}$ and $\Sigma_{y;a,b}$ contains no $\frac{\pi}{2}$-straight $\frac{3}{4}\theta^4$-long 1-strainer at y.

Lemma 5.5.9. *Suppose that every point $y \in B(x, \theta)$ admits a $< 2\theta$-straight 1-strainer (a_y, b_y) of length $> (1 - \theta)$ and that (a, b) is a $< 2\theta$-straight 1-strainer of length $> (1 - \theta)$ at x such that $\Sigma^o_{x;a,b}$ has diameter $\geq 2\theta^{\frac{5}{2}}$ and contains no $< C_1\theta$-straight 1-strainer of length θ^4 centered at x. If $\lambda = \lambda(\theta)$ is sufficiently small, then x is θ-edgy relative to the 1-strainer (a, b) and moreover, there is a point $z \in B(x, 2\theta^4)$ which is θ-strongly edgy relative to (a_z, b_z).*

Proof. That x is θ-edgy follows immediately from 5.5.4 with $r = l' = \theta^4$ and $l = 2\theta^4$.

Applying 5.5.3 on the same scales, we obtain that there is a point $z \in \Sigma_{x;a,b} \cap B(x, 2\theta^4)$ such that $\Sigma^o_{z;a,b} \subset \Sigma^o_{x;a,b}$ has diameter $\geq 2\theta^{\frac{5}{2}}$ and admits no $\frac{\pi}{2}$-straight 1-strainer of length $\frac{1}{2}\theta^4$. The 1-strainer (a, b) is still $< C_0\theta$-straight at z.

By assumption, the point z admits a $< 2\theta$-straight 1-strainer (a', b') of length $> (1 - \theta)$. By Lemma 5.5.6, the angles between the two 1-strainers (a, b) and (a', b') at z are (after changing the order of the strainer points if necessary) less than $c'\theta^{\frac{1}{2}}$.

For any point $z' \in \Sigma_{z;a',b'}$ we have $|\angle_z(a', z') - \frac{\pi}{2}| \leq c\theta$ (by 5.4.6), $|\angle_z(a', z') - \angle_z(a, z')| \leq$

$c'\theta^{\frac{1}{2}}$ and $|\angle_z(a,z') - \tilde{\angle}_z(a,z')| \le c''\theta$ (by 5.4.2). This implies that $|\tilde{\angle}_z(a,z') - \frac{\pi}{2}| \le c'''\theta^{\frac{1}{2}}$. We now apply 5.4.7 and project the cross section $\Sigma^o_{z;a',b'}$ to $\Sigma_{z;a,b}$. This implies that $\Sigma_{z;a',b'}$ contains no $\frac{\pi}{2}$-straight 1-strainer of length $\frac{3}{4}\theta^4$ centered at z, since such a strainer would project to a 1-strainer of length $\ge \frac{1}{2}\theta^4$ in $\Sigma_{z;a,b}$ with distance $\ge \frac{1}{2}\theta^4$ between its end points. Such a strainer must be $\frac{\pi}{2}$-straight by Lemma 5.5.4.

Similarly, we apply 5.4.7 to project $\Sigma^o_{z;a,b}$ to $\Sigma_{z;a',b'}$. This implies that $\operatorname{diam}\Sigma^o_{z;a',b'} \ge \frac{4}{3}\theta^{\frac{5}{2}}$. Thus, z is indeed θ-strongly edgy relative to (a',b'). □

By 5.5.6(ii), at a θ-edgy point y there are no θ^3-long $c\theta^{\frac{1}{2}}$-straight 2-strainers and any two θ^3-long $c\theta^{\frac{1}{2}}$-straight 1-strainers have angle $< c'\theta^{\frac{1}{2}}$ (in the sense that their pairs of directions are $c'\theta^{\frac{1}{2}}$-Hausdorff close subsets in $\Sigma_y X$).

The almost uniqueness of 1-strainers at edgy points extends to uniform neighbourhoods:

Lemma 5.5.10 (Almost uniqueness of 1-strainers extends). *Let y be θ-edgy relative to (a,b). Suppose that (a_i,b_i) are $C_0\theta$-straight 1-strainers of lengths $\in (1-\theta^{\frac{1}{2}}, 1+\theta^{\frac{1}{2}})$ at points $z_i \in B(y,\theta)$ for $i=1,2$. Then $\angle.(a_1,a_2), \angle.(a_1,b_2), \angle.(b_1,a_2)$, $\angle.(b_1,b_2) \notin [c'\theta^{\frac{1}{2}}, \pi - c'\theta^{\frac{1}{2}}]$ on $B(y,\theta)$, if $\lambda=\lambda(\theta)$ is sufficiently small.*

Proof. The strainers (a_i,b_i) are $c'\theta$-straight at y, cf. (5.4.1), and hence $c\theta^{\frac{1}{2}}$-straight with the constant c as in the hypothesis of 5.5.6, because $c'\theta < c\theta^{\frac{1}{2}}$. Applying 5.5.6(ii) with $l=\theta^3$ in the edgy point y yields up to switching a_1 and b_1 that $\angle_y(a_1,a_2), \angle_y(b_1,b_2) < c''\theta^{\frac{1}{2}}$. Due to our condition on the lengths of the 1-strainers (a_i,b_i) it follows that they are $c'''\theta^{\frac{1}{2}}$-Hausdorff close (as two point subsets), which in turn implies that $\angle.(a_1,a_2), \angle.(b_1,b_2) < c''''\theta^{\frac{1}{2}}$ on $B(y,\theta)$. □

Lemma 5.5.11 (Relative position of nearby edgy points). *Let x be θ-edgy relative to (a,b). If $\lambda=\lambda(\theta)$ is sufficiently small, the following hold:*

(i) All points in $B(x,\theta^3)$ which are θ-weakly edgy are contained in the $\theta^{\frac{15}{4}}$-neighbourhood of axb.

(ii) Suppose that every point $y \in B(x,\theta)$ admits a $< 2\theta$-straight 1-strainer (a_y,b_y) of length $> (1-\theta)$. Then for every point $y' \in axb \cap B(x,\theta^3)$ we have $\operatorname{diam}\Sigma_{y';a_{y'},b_{y'}} < \frac{3}{4}\theta^2$ or y' lies at distance $< \theta^{\frac{15}{4}}$ from a point z which is θ-strongly edgy relative to (a_z,b_z).

Proof. (i) Let $y \in \Sigma_{x;a,b} \cap A(x, \frac{1}{2}\theta^{\frac{15}{4}}, 2\theta^3)$. Then $\Sigma_{x;a,b}$ contains a $c\theta$-straight $\frac{1}{3}\theta^{\frac{15}{4}}$-long 1-strainer at y_1, cf. 5.5.3(i). By 5.5.4(ii), every point $u \in ayb \cap B(y, 2\theta^3)$ admits a $c\theta$-straight $\frac{1}{6}\theta^{\frac{15}{4}}$-long 1-strainer contained in $\Sigma_{u;a,b}$. By the metric properties of the maps $\Phi^{a,b}_t$, every point $z \in B(y,\theta^3)$ outside the $\theta^{\frac{15}{4}}$-neighbourhood of axb lies at distance $< c'\theta^4$ from such a point u (for some such y) and therefore admits a $c''\theta$-straight $> \frac{1}{7}\theta^{\frac{15}{4}}$-long 1-strainer contained in $\Sigma^o_{z;a,b}$.

Suppose that z is θ-weakly edgy with respect to some 1-strainer (a', b'). Then by Lemma 5.5.10 the two 1-strainers (a, b) and (a', b') have angles $\leq c'\theta^{\frac{1}{2}}$ at z, and we can project $\Sigma^o_{z;a,b}$ to $\Sigma_{z;a',b'}$ as in the proof of Lemma 5.5.9 to obtain a contradiction.

(ii) Let $y' \in axb \cap B(x, \theta^3)$. Suppose that $\Sigma^o_{y';a,b}$ contains a $\frac{\pi}{2}$-straight $\frac{1}{2}\theta^{\frac{15}{4}}$-long 1-strainer at y. Then $y' = ayb \cap \Sigma_{x;a,b}$ has distance $< c\theta^4$ from x. By 5.5.4, $\Sigma^o_{x;a,b}$ then contains a $c'\theta$-straight $\frac{1}{4}\theta^{\frac{15}{4}}$-long 1-strainer at y', which is also a $\frac{\pi}{2}$-straight $> \frac{1}{5}\theta^{\frac{15}{4}}$-long 1-strainer at x. This contradicts the θ-edgyness of x. Thus $\Sigma^o_{y';a,b}$ contains no $\frac{1}{2}\theta^{\frac{15}{4}}$-long $\frac{\pi}{2}$-straight 1-strainer at y'.

If $\operatorname{diam} \Sigma^o_{y';a,b} < \theta^{\frac{9}{8}}$, by assumption there is a $< 2\theta$-straight 1-strainer (a', b') of length $> (1 - \theta)$ at y'. Projecting $\Sigma_{y';a,b}$ to $\Sigma_{y';a',b'}$ yields that $\operatorname{diam} \Sigma^o_{y';a',b'} < \frac{3}{4}\theta^2$. Otherwise, 5.5.3(ii) applied on the scale $l = \theta^{\frac{15}{4}}$ yields that $\Sigma_{y';a',b'} \cap B(y, 2\theta^{\frac{15}{4}})$ contains a point z such that $\Sigma^o_{z;a',b'}$ has diameter $\geq \theta^{\frac{9}{4}}$ and admits no $\frac{\pi}{2}$-straight 1-strainer of length $\frac{1}{2}\theta^4$ centered at z. This implies that z is θ-strongly edgy relative to some 1-strainer (a_z, b_z) as in the proof of Lemma 5.5.9. □

Lemma 5.5.12 (Almost parallel cross sections of edges). *Let x be a θ-edgy point relative to a θ-straight 1-strainer (a, b) with length $> (1 - \theta)$, and let y be θ-weakly edgy relative to another 1-strainer (a', b'). We consider the truncated cross sections $\check{\Sigma}_x := \Sigma_{x;a,b} \cap B(x, \frac{1}{2}\theta^3)$ and $\check{\Sigma}_y := \Sigma_{y;a',b'} \cap B(y, \frac{1}{2}\theta^3)$. If $\lambda = \lambda(\theta)$ is sufficiently small, the following hold:*

(i) Suppose that $\check{\Sigma}_y$ and $\check{\Sigma}_y$ intersect. Then $d(x, y) < c\theta^{\frac{7}{2}}$.

(ii) Suppose that $d(x, y) \leq c\theta^{\frac{10}{3}}$. Then the Hausdorff distance $d_H(\check{\Sigma}_x, \check{\Sigma}_y)$ is less than $\theta^{\frac{99}{30}}$.

Proof. (i) Let z be one of the intersection points of the cross sections $\check{\Sigma}_x$ and $\check{\Sigma}_y$. By 5.5.10, $f_{a,b} - f_{a',b'}$ is $c\theta^{\frac{1}{2}}$-Lipschitz on $B(z, \frac{1}{2}\theta^3)$ (in fact, on $B(z, \theta)$), and hence $f_{a,b}$ is $c\theta^{\frac{1}{2}}$-Lipschitz on $\Sigma_{y;a',b'} \cap B(z, \frac{1}{2}\theta^3)$. It follows that $|f_{a,b}(x) - f_{a,b}(y)| = |f_{a,b}(z) - f_{a,b}(y)| < c\theta^{\frac{7}{2}}$. Since y is contained in the $\theta^{\frac{15}{4}}$-neighbourhood of axb due to 5.5.11, we obtain that $d(y_1, y_2) < c'\theta^{\frac{7}{2}}$.

(ii) Let $z = ayb \cap \Sigma_{x;a,b}$ and consider a point $u \in \check{\Sigma}_y$. By 5.5.10, we have $\angle_y(a, a') \leq c'\theta^{\frac{1}{2}}$. Thus, we can apply 5.4.7 to project $\check{\Sigma}_y$ to $\Sigma_{y;a,b}$. In particular, the point u projects to a point v with $d(u, v) \leq c''\theta^{\frac{7}{2}}$ and $|d(y, v) - d(y, u)| \leq c'''\theta^3$.

Next, we apply the coarse flow $\Phi^{a,b}_{-f_{a,b}(y)}$ to transport $\Sigma_{y;a,b}$ to $\Sigma_{x;a,b}$. We have $z = \Phi^{a,b}_{-f_{a,b}(y)}$ and set $w := \Phi^{a,b}_{-f_{a,b}(y)}(v)$. Our condition that $d(x, y) \leq c\theta^{\frac{10}{3}}$ and the metric properties of the flow yield that $d(v, w) \leq c''''\theta^{\frac{10}{3}}$ and $|d(w, z) - d(y, u)| \leq \theta^{\frac{7}{2}}$. Finally, Lemma 5.5.11 implies that $d(z, x) \leq \theta^{\frac{7}{2}}$. Thus, the distance between w and $\check{\Sigma}_x$ is at most $2\theta^{\frac{7}{2}}$. (Here, we use again that close to x the cross section $\Sigma_{x;a,b}$ is almost 1-dimensional, i.e. close to an interval.)

All in all, we conclude that $d(u, \check{\Sigma}_x) \leq \theta^{\frac{99}{30}}$. By switching the roles of $\check{\Sigma}_x$ and $\check{\Sigma}_y$, we similarly obtain that every point in $\check{\Sigma}_x$ has distance $\leq \theta^{\frac{99}{30}}$ to $\check{\Sigma}_y$. This completes the proof. □

For future reference, we observe that the lemma also holds true if we replace $\check{\Sigma}_y$ by $B(y, \frac{1}{2}\tau\theta^3)$ for some $\tau \in (1-\theta, 1+\theta)$. The proof for (i) goes through unchanged and for (ii) it suffices to observe that the almost 1-dimensionality of $\Sigma_{y;a',b'}$ near y (Lemma 5.5.4) implies that $d_H(\check{\Sigma}_y, B(y, \frac{1}{2}\tau\theta^3)) < \theta^{\frac{99}{30}}$.

5.6 Necks

A neck occurs where the connected component of a cross section has small diameter.

Definition 5.6.1 (Necklike points). A point $y \in B(x, 5)$ is called θ-*necklike* relative to a $< 2\theta$-straight 1-strainer (a, b) of length $> (1-\theta)$ (at y) if $\operatorname{diam}(\Sigma_{y;a,b}) < \theta^2$. We say that $y \in B(x, 5)$ is θ-*weakly necklike* relative to a $< 2\theta$-straight 1-strainer (a, b) of length $> (1-\theta)$ at y if $\operatorname{diam}(\Sigma_{y;a,b}) < 2\theta^2$ and that it is θ-*strongly necklike* relative to such a strainer if $\operatorname{diam}(\Sigma_{y;a,b}) < \frac{3}{4}\theta^2$.

Our new definition allows us to reformulate part (ii) of Lemma 5.5.11: Suppose that for an edgy point x, every point in $B(x, \theta)$ admits a $< 2\theta$-straight 1-strainer of length $> (1-\theta)$ Then every point $y \in axb \cap B(x, \theta^3)$ is θ-strongly necklike or has distance $< \theta^{\frac{15}{4}}$ from a point z which is θ-strongly edgy.

Suppose that x is θ-necklike relative to (a, b). Nearby cross sections have comparable diameters: Using the metric properties of the maps $\Phi_t^{a,b}$ one sees that

$$\operatorname{diam}(\Sigma_{z;a,b}) < (1 + c\theta)\theta^2 + c\theta|f_{a,b}(z) - f_{a,b}(y)| < c'\theta^2$$

for $z \in B(y, \theta)$.

Every segment of length $> \theta$ initiating in $B(y, \theta)$ must pass through one of the two cross sections $f_{a,b}^{-1}(f_{a,b}(y) \pm \frac{9}{10}\theta)$. Hence, triangle comparison and (5.4.2) imply for any point a' with $d(y, a') > \theta$ that

$$\angle_z(a, a'), \angle_z(b, a') \notin [c\theta, \pi - c\theta] \tag{5.6.2}$$

for $z \in B(y, \frac{\theta}{2})$. In particular, any $\frac{\pi}{2}$-straight 1-strainer (a', b') of length $> \theta$ at a point in $B(y, \frac{\theta}{2})$ is $c'\theta$-straight, and $f_{a',b'} - f_{a,b}$ is $c''\theta$-Lipschitz on $B(y, \frac{\theta}{2})$.

If x is θ-necklike relative to a 1-strainer (a, b) with $\operatorname{diam}\Sigma_{x;a,b} < \frac{1}{2}\theta^2$, and if all points in $B(x, \theta)$ admit $< 2\theta$-straight 1-strainers of length $> (1-\theta)$ we can conclude from the above observations as in the proof of Lemma 5.5.9 that all points in $B(x, \theta)$ are also θ-strongly necklike. We now deduce that cross sections of nearby necklike points are almost parallel.

Lemma 5.6.3 (Almost parallel cross sections of necks). *Let x be θ-necklike relative to (a,b). Furthermore, let $y \in B(x,\theta)$ be θ-weakly necklike with respect to a 1-strainer (a',b').*

(i) If $\Sigma_{x;a,b}$ and $\Sigma_{y;a',b'}$ have nonempty intersection, then $d(x,y) < c'\theta^3$.

(ii) If $d(x,y) < \theta^{\frac{11}{6}}$, then the Hausdorff distance of $\Sigma_{x;a,b}$ and $\Sigma_{y;a',b'}$ is less that $\theta^{\frac{5}{3}}$.

Proof. The proof is closely related to the one for edges, i.e. 5.5.12.

(i) Let z be one of the intersection points of the cross sections $\Sigma_{x;a,b}$ and $\Sigma_{y;a',b'}$. By our discussion above, $f_{a,b} - f_{a',b'}$ is $c\theta$-Lipschitz on $B(z,\theta^2)$, and hence $f_{a,b}$ is $c\theta^{\frac{1}{2}}$-Lipschitz on $\Sigma_{y;a',b'}$. It follows that $|f_{a,b}(x) - f_{a,b}(y)| = |f_{a,b}(z) - f_{a,b}(y)| < c'\theta^3$.

(ii) Consider a point $u \in \check{\Sigma}_y$. By 5.6.2, we have $\angle_y(a,a') \le c'\theta$. When projecting $\Sigma_{y;a',b'}$ to $\Sigma_{y;a,b}$, we map u to a point v with $d(u,v) \le c''\theta^3$ by 5.4.7. The coarse flow $\Phi^{a,b}_{-f_{a,b}(y)}$ transports $v \in \Sigma_{y;a,b}$ to some point $w \in \Sigma_{x;a,b}$ with $d(v,w) \le c'''\theta^{\frac{11}{6}}$.

This shows that $d(u,\Sigma_{x;a,b}) \le \theta^{\frac{5}{3}}$. Again, we switch the roles of $\Sigma_{x;a,b}$ and $\Sigma_{y;a',b'}$ to complete the proof. □

6. Locally volume collapsed 3-orbifolds are graph

6.1 Setup and formulation of main result

Let (O, g) be a closed connected smooth Riemannian 3-orbifold which does *not* have non-negative sectional curvature, $sec \ngeq 0$.

Definition 6.1.1 (Curvature scale). For $-b^2 \in [-1, 0)$ we define the $-b^2$-*(sectional) curvature scale* in a point $x \in O$ as the maximal radius $\rho_{-b^2}(x) \in (0, \infty)$ such that the rescaled ball $B_{\rho_{-b^2}(x)^{-2}g}(x, 1) = \rho_{-b^2}(x)^{-1} \cdot B_g(x, \rho_{-b^2}(x))$ has sectional curvature $sec \geq -b^2$.

Note that

$$b\rho_{-1} \leq \rho_{-b^2} \leq \rho_{-1}. \tag{6.1.2}$$

The function ρ_{-b^2} is continuous on O. More precisely, ρ_{-b^2} *does not oscillate too fast* in the sense that for $0 < \lambda < 1$ one has

$$(1 - \lambda)\rho_{-b^2}(x) \leq \rho_{-b^2} \leq (1 + \lambda)\rho_{-b^2}(x) \tag{6.1.3}$$

on $B(x, \lambda\rho_{-b^2}(x))$.

The rescaled balls $B_{\rho_{-b^2}(x)^{-2}g}(x, 1)$ are Alexandrov balls with curvature $\geq -b^2$ and radius ≤ 1 in the sense of definition 5.1.1.

We study the geometry and topology of 3-orbifolds which are locally collapsed relative to the curvature scale.

Definition 6.1.4 (Local volume collapse). Let $v > 0$ and let $\sigma : O \to (0, \infty)$ be some (not necessarily continuous) function. We say that (O, g) is v-*collapsed* at the scale σ, if for all points x we have $\mathrm{vol}(B_{\sigma(x)^{-2}g}(x, 1)) < v$, equivalently, $\mathrm{vol}(B_g(x, \sigma(x))) < v\sigma(x)^3$.

If $sec \ngeq 0$, we say that (O, g) is $(v, -b^2)$-*collapsed* if it is v-collapsed at the scale ρ_{-b^2}.

Note that if (O, g) is locally v-collapsed at some scale $\sigma \leq \rho_{-b^2}$, Bishop-Gromov volume comparison (Proposition 3.1.4) yields that it is locally $(v', -b^2)$-collapsed with $v' = \frac{vol(B_{-b^2}(1))}{vol(B_0(1))} v \leq \frac{vol(B_{-1}(1))}{vol(B_0(1))} v$ (independent of $-b^2$!). Here $B_{-b^2}(1)$ denotes the unit 3-ball with $sec \equiv -b^2$.

Strongly volume collapsed Riemannian 3-orbifolds are, on the scale of their collapse, close to Alexandrov spaces of dimension ≤ 2 and with curvature $\geq -b^2$, and the volume collapse translates into the shortness of $2\frac{1}{2}$-strainers, cf. section 5.1.2).

Lemma 6.1.5. *For $\lambda > 0$ exists $v = v(\lambda) > 0$ such that the following holds: If (O, g) is $(v, -b^2)$-collapsed and $x \in O$, then $\bar{\theta}_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainers in the Alexandrov ball $\rho_{-b^2}(x)^{-1} B(x, \rho_{-b^2}(x))$ of curvature $\geq -b^2 \geq -1$ have length $< \lambda$.*

Proof. Suppose that the Riemannian 3-orbifolds (O_i, g_i) are $(\frac{1}{i}, -b_i^2)$-collapsed but contain points x_i which admit such that there are $\bar{\theta}_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainers of length $\geq \lambda$ in the balls $\rho_{-b_i^2}(x)^{-1} B(x, \frac{1}{2}\rho_{-b_i^2}(x))$. Then the rescaled balls $\rho_{-b^2}(x_i)^{-1} B_(x_i, 1)$ Gromov-Hausdorff subconverge to an Alexandrov ball with dimension ≤ 2 and curvature ≥ -1 which admits $\bar{\theta}_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainers of length λ, a contradiction. □

We will require some additional regularity for our Riemannian orbifolds. The conclusions on the global topology of collapsed 3-orbifolds are valid without this regularity condition, but it is technically convenient because it avoids the use of (an orbifold version of) Perelman's Stability Theorem for Alexandrov spaces and is expected to be satisfied by the output of the Ricci flow on 3-orbifolds. In the next definition, ω_3 denotes the volume of the euclidean unit 3-ball (cf. [Pe03, 7.4] and [KL10, Thm. 1.3]).

Definition 6.1.6 (Local curvature control). Fix numbers $s_0 \in \mathbb{N}$, $v_0 \in (0, \omega_3)$, a function $K : (0, \omega_3) \to (0, \infty)$ and a scale function $\sigma : O \to (0, \infty)$. We say that (O, g) has (v_0, s_0, K)-*curvature control* below scale σ, if the following holds: If $\mathrm{vol}\, B(x, r) \geq vr^3$ for $v \in [v_0, \omega_3)$ and $r \in (0, \sigma(x)]$, then $\|\nabla^s R\| \leq K(v) r^{-2-s}$ on $B(x, r)$, equivalently, $\|\nabla^s R\| \leq K(v)$ on the rescaled ball $r^{-1} \cdot B(x, r)$ for $s = 0, \ldots, s_0$.

We will apply this notion in the following situation.

Lemma 6.1.7. *Let (O_i, g_i) be a sequence of Riemannian 3-orbifolds as above with (v_i, s_0, K)-curvature control below scale ρ_{-b^2}, where $v_i \to 0$. Furthermore, let $x_i \in O_i$ be points and $\lambda_i \to 0$ positive numbers. Then, if s_0 is sufficiently large, the sequence of rescaled pointed orbifolds $(\lambda_i \rho_{-b^2}(x_i))^{-1} \cdot (O_i, x_i)$ subconverges either in the Gromov-Hausdorff sense to an Alexandrov space with curvature ≥ 0 and dimension ≤ 2, or in the $\mathcal{C}^5$-topology to a $\mathcal{C}^{10}$-smooth complete Riemannian 3-orbifold with* $\sec \geq 0$.

Proof. Clearly, we have Gromov-Hausdorff subconverge to a pointed Alexandrov space (X, x_0) with curvature ≥ 0 and dimension ≤ 3. If $\dim(X) = 3$, the convergence can be improved using our assumption of local curvature control. The approximating pointed 3-orbifolds $(\lambda_i\rho_{-b^2}(x_i))^{-1}\cdot(O_i,x_i)$ are uniformly noncollapsed. Indeed, for any $r>0$, we have $\operatorname{vol} B_{\lambda_i^{-2}\rho_{-b^2}(x_i)^{-2}g_i}(x_i,r) > \frac{1}{2}\operatorname{vol}_3 B(x_0,r)>0$ for large i, where volume in X is measured with respect to the 3-dimensional Hausdorff measure. Thus the (v_i,s_0,K)-curvature control on $(\lambda_i\rho_{-b^2}(x_i))^{-1}\cdot O_i$ below the scales $\lambda_i^{-1}\to\infty$ applies for large i and yields on the balls $B_{\lambda_i^{-2}\rho_{-b^2}(x_i)^{-2}g_i}(x_i,r)$ uniform bounds on the curvature tensor and its covariant derivatives up to order s_0. The smoothness of the limit and the convergence now follow from Theorem 4.2.23. □

The main result of this paper is the following (cf. [Pe03, Thm. 7.4], [MT08, Thm. 0.2] and [KL10, Thm. 1.3]):

Theorem 6.1.8. *Let $s_0\in\mathbb{N}$ and let $K:(0,\omega_3)\to(0,\infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0=v_0(s_0,K)\in(0,\omega_3)$ such that the following holds: If (O,g) is $(v_0,-1)$-collapsed, has (v_0,s_0,K)-curvature control below the scale ρ_{-1} and contains no bad 2-suborbifolds, then O admits a $\mathcal{C}^5$ Riemannian metric with $\sec\geq 0$, or can be decomposed by finitely many surgeries into components which are spherical or graph.*

Remark 6.1.9. Unlike in the manifold case we cannot conclude that O is always graph, because there are non-graph 3-orbifolds admitting nonnegatively curved (e.g. spherical or euclidean) metrics.

One can reduce to the case of a lower diameter bound relative to a curvature scale by suitably rescaling. Theorem 6.1.8 follows from:

Theorem 6.1.10. *Let $s_0\in\mathbb{N}$ and let $K:(0,\omega_3)\to(0,\infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0=v_0(s_0,K)\in(0,\omega_3)$ such that the following holds: If for some $-b^2\in[-1,0)$ the orbifold (O,g) is $(v_0,-b^2)$-collapsed, satisfies $\operatorname{rad}(O,\cdot)\geq\frac{1}{2}\rho_{-b^2}$, has (v_0,s_0,K)-curvature control below the scale ρ_{-b^2}, and contains no bad 2-suborbifolds, then O can be decomposed by finitely many surgeries into components which are spherical or graph.*

Proof that Theorem 6.1.10 implies Theorem 6.1.8. Suppose that (O_i) is a sequence of $(v_i,-1)$-collapsed orbifolds with (v_i,s_0,K)-curvature control below the scale ρ_{-1} where $v_i\to 0$. Then we must show that the O_i satisfy the conclusion of Theorem 6.1.8 for infinitely many i.

Note that for all $b\in(0,1]$ the O_i are $(b^{-3}v_i,-b^2)$-collapsed and have (v_i,s_0,K)-curvature control below the scale ρ_{-b^2}, cf. (6.1.2). Hence we are done if for some $-b^2\in[-1,0)$ we have $\operatorname{rad}(O_i)\geq\frac{1}{2}\rho_{-b^2}$ for all sufficiently large i.

Otherwise, after passing to a subsequence, there exist sequences of numbers $-b_i^2\to 0$

and points $x_i \in O_i$ such that $\mathrm{rad}(O_i, x_i) < \frac{1}{2}\rho_{-b_i^2}(x_i)$. It follows that $\rho_{-b_i^2} \equiv const_i \geq \mathrm{diam}(O_i)$ and we have collapse to the point in the sense of $\mathrm{diam}(O_i) \cdot (-\min \sec_{O_i})^{\frac{1}{2}} \to 0$. We rescale and increase the $-b_i^2$ so that $\mathrm{diam}(O_i) = 1$ and $\min \sec_{O_i} = -b_i^2 \to 0$. Then $\rho_{-b_i^2} \equiv 1$.

If $v_i' = \mathrm{vol}(O_i) \to 0$, then the O_i are $(v_i', -b_i^2)$-collapsed with (v_i, s_0, K)-curvature control below the scales $\rho_{-b_i^2}$ and with $\mathrm{rad}(O_i) \geq \frac{1}{2} \equiv \frac{1}{2}\rho_{-b_i^2}$, and we are done by Theorem 6.1.10.

Otherwise, after passing to a subsequence, we have a lower volume bound $\mathrm{vol}(O_i) \geq v' > 0$ and, due to the curvature control, uniform (global) bounds on the curvature tensor and its covariant derivatives up to order s_0. If s_0 is large enough, it follows that the O_i subconverge, say, in the $\mathcal{C}^5$-topology to a $\mathcal{C}^5$-Riemannian 3-orbifold O_∞ with $\sec \geq 0$. In particular, infinitely many O_i are diffeomorphic to O_∞ and therefore admit $\mathcal{C}^5$ metrics with $\sec \geq 0$. □

We fix some arbitrary (more than) sufficiently large value for s_0, say $s_0 := 1000$. For the rest of this section we make the following assumption on the orbifolds we work with:

Assumption 6.1.11. *(O, g) is a smooth closed connected Riemannian 3-orbifold such that sec $\not\geq 0$ and $\mathrm{rad}(O, \cdot) \geq \frac{1}{2}\rho_{-b^2}$ where $-b^2 \in [-1, 0)$.*

Together with Corollary 2.3.3, Theorem 6.1.8 implies the following

Theorem 6.1.12. *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds: If (O, g) is closed and $(v_0, -1)$-collapsed, has (v_0, s_0, K)-curvature control below the scale ρ_{-1} and contains no bad 2-suborbifolds, then O either admits a $\mathcal{C}^5$ metric with $\sec \geq 0$, or satisfies Thurston's Geometrization Conjecture.*

We can deal with the first alternative in the theorem by the application of the following orbifold version of Hamilton's corresponding application of the Ricci flow to 3-manifolds with sectional curvature ≥ 0 (see [KL11, Prop. 5.6] and also [Ha03]; for a more detailed discussion of the lower regularity case cf. [KL10, Lem. 3.11]):

Theorem 6.1.13. *Let O be a smooth closed connected 3-orbifold and g a $\mathcal{C}^5$ Riemannian metric on O with sectional curvature ≥ 0. Then O is geometric with model geometry S^3, $\mathbb{R}^3$ or $\S^2 \times \mathbb{R}$.*

Theorems 6.1.12 and 6.1.13 together yield

Corollary 6.1.14. *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds: If (O, g) is closed and $(v_0, -1)$-collapsed, has (v_0, s_0, K)-curvature control below the* *scale ρ_{-1} and contains no bad 2-suborbifolds, then O satisfies Thurston's Geometrization Conjecture.*

6.2 Conical approximation and humps

Sufficient local volume collapse leads to good local approximation by cones of dimension ≤ 2 on scales comparable to the curvature scale. We now apply our discussion of the (coarse) geometry of Alexandrov balls in sections 5.2, 5.3 and 5.4 to the orbifold setting and globalize it relative to the curvature scale. Proposition 5.2.2 for $d = 2$ and Lemma 6.1.5 imply:

Proposition 6.2.1 (Uniform conical approximation relative to the curvature scale). *For $0 < \sigma, \mu < 1$ there are a scale $0 < s_1 = s_1(\sigma, \mu) << \sigma$ and a rate of collapsedness $v = v(\sigma, \mu) > 0$ such that the following holds: If (O, g) is $(v, -b^2)$-collapsed, then it can in every point x be μ-well approximated (cf. Definition 5.2.1) on some scale $s(x) \in [s_1\rho_{-b^2}(x), \sigma\rho_{-b^2}(x)]$ by a cone of dimension 1 or 2, i.e. by the open interval $(-1, 1)$, by the half-open interval $[0, 1)$, or by the cone of radius 1 over a circle or an interval of diameter $\leq \pi$.*

Proof. The assertion follows with $s_1 = s_1(\sigma, \mu) := s_1(2, \sigma, \mu)$ from Proposition 5.2.2 and $v = v(\sigma, \mu) := v(s_{2\frac{1}{2}}(\sigma, \mu))$ from Lemma 6.1.5. □

We fix some small value $\sigma \in (0, \frac{1}{10})$ for the upper bound on the local scales of approximation.

We next wish to divide our orbifold into regions with and without good 1-strainers. The region with good 1-strainers has locally almost product geometry. It consists of the points where the bases of approximating cones provided by 6.2.1 have diameter $\approx \pi$.

For small $\theta > 0$, let $\mu_0(\theta) > 0$ be the constant given by Lemma 5.3.2. For $\mu \in (0, \mu_0(\theta)]$, we call a point $x \in O$ a $(\theta, \mu, -b^2)$-*hump*, if O can in x be μ-well approximated on the scale $s(x)$ by a cone with base of diameter $< \pi - \frac{\theta}{2}$, i.e. by the half-open interval $[0, 1)$ or the cone of radius 1 over a circle or an interval with diameter $< \pi - \frac{\theta}{2}$. In other words, x is a (θ, μ)-hump of the rescaled ball $B_{\rho_{-b^2}(x)^{-2}g}(x, 1)$ in the sense of Definition 5.3.1 for $d = 2$. If x is no $(\theta, \mu, -b^2)$-hump, then O can in x be μ-well approximated on the scale $s(x)$ by a cone with base of diameter $\geq \pi - \frac{\theta}{2}$, i.e. by the open interval $(-1, 1)$ or the cone of radius 1 over a circle or an interval with diameter $\geq \pi - \frac{\theta}{2}$.

We denote by $H = H_{\theta,\mu,-b^2} \subset O$ the subset of $(\theta, \mu, \rho_{-b^2}(x))$-humps and by $S = S_{\theta,\mu,-b^2} \subset O$ the open subset of points x which admit $< \theta$-straight equilateral 1-strainers with length in $(\frac{1}{11}s_1(\sigma, \mu), \frac{3}{22}s_1(\sigma, \mu))$ in the Alexandrov ball $(\rho_{-b^2}(x))^{-1}B(x, \rho_{-b^2}(x))$ of curvature ≥ -1.

Throughout the following chapters, we will abbreviate this property by saying that the points $x \in S_{\theta,\mu,-b^2}$ admit $< \theta$-straight (equilateral) 1-strainers with length in the interval $(\frac{1}{11}s_1(\sigma,\mu)\rho_{-b^2}(x), \frac{3}{22}s_1(\sigma,\mu)\rho_{-b^2}(x))$.

Due to our bound on the approximation accuracy μ, the implications of Lemma 5.3.2 hold: A $(\theta,\mu,-b^2)$-hump x admits no $\frac{\theta}{4}$-straight 1-strainers of length $\frac{1}{111}s(x)$, but all points in the closed annulus $\overline{A}(x; \frac{1}{10}s(x), \frac{9}{10}s(x))$ do admit $\frac{\theta}{11}$-straight 1-strainers of length $> \frac{1}{11}s(x) \geq \frac{1}{11}s_1(\sigma,\mu)\rho_{-b^2}(x)$, i.e. $\overline{A}(x; \frac{1}{10}s(x), \frac{9}{10}s(x)) \subset S_{\theta,\mu,-b^2}$. On the other hand, if x is no $(\theta,\mu,-b^2)$-hump, then it admits $< \theta$-straight 1-strainers of length $> \frac{99}{100}s(x) > \frac{1}{11}s_1(\sigma,\mu)\rho_{-b^2}(x)$. Hence $O = H_{\theta,\mu,-b^2} \cup S_{\theta,\mu,-b^2}$.

Proposition 6.2.2 (cf. Proposition 5.3.3). *If $\mu \leq \mu_0(\theta)$ and (O,g) is $(v(\sigma,\mu),-b^2)$-collapsed (cf. Proposition 6.2.1), then there exist finitely many $(\theta,\mu,-b^2)$-humps $x_j \in H_{\theta,\mu,-b^2}$ such that*

$$O = \left(\bigcup_j B(x_j, \frac{1}{10}s(x_j))\right) \cup S_{\theta,\mu,-b^2}.$$

Moreover, $d(x_j,x_k) > \frac{9}{10}s(x_j)$ for $j \neq k$.

Proof. We proceed as in the proof of Proposition 5.3.3.

Again, there can be no infinite sequence of points $x_1, x_2, \dots \in H_{\theta,\mu,-b^2} - S_{\theta,\mu,-b^2}$ such that for all k we have $x_k \in B(x_{k+1}, \frac{1}{10}s(x_{k+1}))$ and $x_{k+1} \notin B(x_k, \frac{1}{10}s(x_k))$, because then $s(x_1) < \frac{1}{9^{k-1}}s(x_k)$. But due to the continuity of the curvature scale ρ_{-b^2}, cf. (6.1.3), and the compactness of O the scales take also in this situation values in a bounded interval, namely in $[s_1 \min \rho_{-b^2}, \sigma \min \rho_{-b^2}]$.

The rest of the proof goes through without change. □

6.3 The Shioya-Yamaguchi blow-up

We recall the Shioya-Yamaguchi blow-up argument, see [SY00, Sec. 3, Key Lemma 3.6]. To simplify matters, we restrict ourselves to certain special situations. Some of our arguments are different; we also treat some additional cases not mentioned there explicitely.

We will use blow-up arguments in section 6.6 to determine the *topological type* of subsets of smooth orbifolds. By this we mean the equivalence class of an orbifold with (possibly non-smooth) boundary up to *homeomorphism* in the orbifold sense, i.e. up homeomorphism respecting the singular stratifications of the two orbifolds. Thus, such a homeomorphism restricts to a homeomorphism on the various strata of the two orbifolds.

The advantage of studying only the topological type of orbifolds is that we will not have to concern ourselves with questions of regularity on topological suborbifolds with boundary (i.e. subsets which locally lift to topological submanifolds with boundary) of smooth orbifolds. At least in dimension ≤ 3, if two such suborbifolds are homeomorphic in the above sense and have nonempty interior, their interiors are diffeomorphic: By our discussion in section 2.1.3, sufficiently small neighbourhoods of their singular loci are diffeomorphic to each other, and their regular parts are (homeomorphic and hence) diffeomorphic manifolds of dimension ≤ 3.

6.3.1 General discussion

Consider the following situation. Let $B(p_i, 1)$ be a sequence of n-dimensional Riemannian orbifold balls (i.e. open metric balls of radius 1 in complete Riemannian n-orbifolds without boundary) with curvature $sec \geq -1$ which collapse to an Alexandrov ball of strictly smaller dimension $1 \leq m < n$,

$$B(p_i, 1) \longrightarrow X = B(x, 1). \tag{6.3.1}$$

We suppose furthermore that the collapse limit X is (0-)*conelike* in the sense that every segment initiating in x extends to length 1, and that the closed balls $\overline{B}(p_i, \frac{1}{2})$ are *not discal*. Note that the conelikeness of X implies that for any fixed $\epsilon > 0$ the annulus $A(p_i, \epsilon, 1-\epsilon)$ contains for large i no critical points of the distance function $d(p_i, \cdot)$ such that in particular $\overline{B}(p_i, \frac{1}{2})$ is a topological suborbifold with boundary.

We recall that a point y in an Alexandrov space Y is critical for a distance function $d(p, \cdot)$ if for every direction in $\sigma\Sigma_y Y$ there is a minimizing segment from y to p making angle $\leq \frac{\pi}{2}$ with σ. On a smooth orbifold the absence of critical points (on an open subset) implies that one can construct gradient-like vector fields using a partition of unity (cf. the discussion in section 3.5).

Let $\hat{p}_i \in B(p_i, 1)$ be any sequence of points with $d(\hat{p}_i, p_i) \to 0$. The distance function $d(\hat{p}_i, \cdot)$ must have critical values in $(0, \frac{1}{2})$, because $\overline{B}(p_i, \frac{1}{2})$ is not discal.

Let δ_i be the maximal critical value in $(0, \frac{1}{2})$, and let $q_i \in B(p_i, \frac{1}{2})$ be a critical point at distance $d(\hat{p}_i, q_i) = \delta_i$ from $\hat{p}_i$. Then $\delta_i \to 0$ because X is conelike. For any constant $c > 1$ one has

$$\overline{B}(p_i, \frac{1}{2}) \cong \overline{B}(\hat{p}_i, c\delta_i) \tag{6.3.2}$$

for sufficiently large i, i.e. the topology of the balls $B(p_i, 1)$ is concentrated near their centers. (Note that, also due to the conelikeness of X, there exists a common gradient-like vector field for $d(p_i, \cdot)$ and $d(\hat{p}_i, \cdot)$, and so $\overline{B}(p_i, \frac{1}{2}) \cong \overline{B}(\hat{p}_i, \frac{1}{2})$.)

To help revealing the local topology at the p_i, we form (modulo passing to a subsequence) the *blow-up limit*

$$(\delta_i^{-1}B(p_i,1),\hat{p}_i) \longrightarrow (Y,y_0). \tag{6.3.3}$$

The limit space Y is a noncompact Alexandrov space with dimension $\geq m$ and curvature ≥ 0. Moreover $q_i \to z$ with z a critical point of $d(y_0,\cdot)$ at distance $d(y_0,z)=1$.

If no collapse happens any more in the blow-up limit (6.3.3), i.e. if $\dim Y = n$, then we need topological stability in order to relate the topologies of the balls $B(p_i,1)$ and Y. For instance, if Y and the convergence in (6.3.3) are sufficiently smooth as will be the case in the situations considered later in the paper, say Y is $\mathcal{C}^{10}$-regular and the convergence is $\mathcal{C}^5$-regular, then one can argue as follows:

There are $r,\epsilon>0$ and a smooth vector field V on $Y-\overline{B}(y_0,\frac{r}{2})$ such that for all $y\notin \overline{B}(y_0,\frac{r}{2})$ the vector $V(y)$ has angles $\geq \frac{\pi}{2}+\epsilon$ with all segments yy_0, compare the proof of the Soul Theorem in section 3.5. We regard $\delta_i^{-1}B(\hat{p}_i,2r\delta_i)$ as embedded in Y for large i. In particular, $\hat{p}_i\to y_0$ and V is on a neighbourhood of $\partial B(y_0,r)$ gradient-like not only for $d(y_0,\cdot)$ but also for $\delta_i^{-1}d_{B(p_i,1)}(\hat{p}_i,\cdot)$, viewed as a function on part of Y. Hence $\delta_i^{-1}\overline{B}(\hat{p}_i,r\delta_i)$ is isotopic (preserving singular strata) to $\overline{B}(y_0,r)$ in Y, and with (6.3.2) we see that $\overline{B}(p_i,\frac{1}{2})$ is for large i homeomorphic to the closed disc bundle in the normal bundle of the soul of Y, in other words, to a (small) closed tubular neighbourhood of the soul of Y.

The blow-up limit (6.3.3) is in general still a collapse to lower dimension. The aim of the following discussion is to find situations when the drop of dimension is strictly smaller than for the original collapse (6.3.1).

The following construction proceeds as in the proof of Soul Theorem: Let $\xi\in\Sigma_x X$ be a direction at x. Due to conelikeness it is represented by a segment σ_ξ emanating from x, which unique in view of the lower curvature bound. Fix some $t_0\in(0,1)$, say $t_0=\frac{1}{10}$, and let $\sigma_\xi^i\in B(p_i,1)$ be a sequence of points converging to the point $\sigma_\xi(t_0)$ on σ_ξ at distance t_0 from x. (Our choice of the $\sigma_\xi(t_0)$ is independent of the choice of the $\hat{p}_i$.) The segments $\hat{p}_i\sigma_\xi^i$ subconverge to a (not necessarily unique) ray ρ_ξ in Y emanating from y_0. Moreover, the normalized distance functions $\delta_i^{-1}(d(\sigma_\xi^i,\cdot)-d(\sigma_\xi^i,\hat{p}_i))$ on the rescaled balls $\delta_i^{-1}B(p_i,1)$ subconverge (due to Arzelà-Ascoli) to a concave 1-Lipschitz function β_ξ on Y with $\beta_\xi(y_0)=0$ which decays along ρ_ξ with extremal slope -1, $\beta_\xi(\rho_\xi(t))=-t$, where we use a unit speed parametrization $\rho_\xi(t)$ starting at $\rho_\xi(0)=y_0$. In fact, every point $y\in Y$ is the initial point of a ray ρ_ξ^y along which β_ξ decays with slope -1.

In particular, the level sets of β_ξ have no interior points. The comparison of β_ξ with the Busemann function $b_\xi=\lim_{t\to\infty}(d(\rho_\xi(t),\cdot)-d(\rho_\xi(t),y_0))$ associated to the ray ρ_ξ is given by the inequality

$$\beta_\xi \leq b_\xi.$$

To verify this, let $\sigma_\xi^i(a)$ denote the point on $\hat{p}_i\sigma_\xi^i$ at (unrescaled) distance $a\delta_i$ from $\hat{p}_i$. Then $d(\sigma_\xi^i,\cdot)-d(\sigma_\xi^i,\hat{p}_i)\leq d(\sigma_\xi^i(a),\cdot)-d(\sigma_\xi^i(a),\hat{p}_i)$ for large i and hence $\beta_\xi\leq d(\rho_\xi(a),\cdot)-a$ for all $a > 0$. Letting $a \to \infty$ yields the inequality. As a consequence, the convex suplevel sets of β_ξ are smaller (not larger) than the corresponding suplevels of b_ξ.

Since the q_i are critical for $d(\hat{p}_i, \cdot)$, we have $\tilde{\angle}_{q_i}(\hat{p}_i, \sigma^i_\xi) \le \frac{\pi}{2}$ and $\liminf_{i\to\infty} \delta_i^{-1}(d(\sigma^i_\xi, q_i) - d(\sigma^i_\xi, \hat{p}_i)) \ge 0$. So

$$\beta_\xi(z) \ge 0$$

and z is contained in the totally convex subset $\cap_\xi \{\beta_\xi \ge 0\} = \{\min_\xi \beta_\xi \ge 0\}$.

The blow-up *expands* $\Sigma_x X$ in the following sense: When performing the above construction for two directions ξ and ξ' at the same time, one obtains for every point $y \in Y$ a pair of rays ρ^y_ξ and $\rho^y_{\xi'}$ satisfying

$$\angle_{Tits}(\rho^y_\xi, \rho^y_{\xi'}) := \lim_{t\to\infty} \tilde{\angle}_y(\rho^y_\xi(t), \rho^y_{\xi'}(t)) \ge \angle_x(\xi, \xi'). \tag{6.3.4}$$

This construction can be performed for any finite subset $A \subset \Sigma_x X$ and hence yields weakly expanding maps $\epsilon_{y,A} : A \to \operatorname{tits} Y$.

We will use the following observations:

Lemma 6.3.5. *(i) The blow-up limit Y is not isometric to a euclidean space.*

(ii) If $\Sigma_x X$ contains an embedded unit l-sphere (i.e. with $\sec \equiv 1$*), then so does* $\operatorname{tits} Y$ *and Y splits off an $\mathbb{R}^{l+1}$-factor. If $\Sigma_x X$ contains an embedded unit l-hemisphere, then so does* $\operatorname{tits} Y$ *and Y contains an isometrically embedded copy of the $(l+1)$-dimensional euclidean halfspace (and in particular splits off an $\mathbb{R}^l$-factor).*

Proof. (i) $d(y_0, \cdot)$ has a critical point (at distance 1).

(ii) Choose A as the union of $l+1$ pairs of antipodes which span the embedded unit sphere (corresponding to coordinate axes). Then the expanding map $A \to \operatorname{tits} Y$ must be an isometric embedding and the assertion follows from the Splitting Theorem for Alexandrov spaces. The second assertion follows similarly by applying the first assertion to the boundary $(l-1)$-sphere of the embedded l-hemisphere. □

6.3.2 The case of flat conical limits with dimension ≤ 2

We apply the general discussion above in certain special situations. Note that always $\dim Y \ge \dim X$. We aim now to achieve that $\dim Y > \dim X$ by making a good choice of the $\hat{p}_i$.

Collapse to a flat k-disc. Suppose that X is isometric to the euclidean unit m-disc, $m \ge 1$. By Lemma 6.3.5, Y splits off an $\mathbb{R}^m$-factor and $Y \not\cong \mathbb{R}^m$. Hence we always obtain $\dim(Y) > m$, independently of the choice of the $\hat{p}_i$.

Collapse to the half-open interval (Noses). Suppose that $X = [0, 1)$ with $x = 0$. There is a unique direction ξ at $x = 0$. We choose $\hat{p}_i$ as a "tip" of the nose, i.e. as a maximum of $d(\sigma^i_\xi, \cdot)$. Then $\hat{p}_i \to x$ and the choice of the $\hat{p}_i$ is admissible in the sense that $d(p_i, \hat{p}_i) \to 0$. The base point y_0 is a maximum of β_ξ and hence $\beta_\xi(z) = \beta_\xi(y_0) = 0$. If $\dim(Y) = 1$, then Y is a halfline since $Y \ncong \mathbb{R}$ by Lemma 6.3.5(i), and $\beta_\xi = b_\xi$ has a unique maximum. This contradicts $z \neq y_0$. Thus $\dim(Y) \geq 2$.

Note that Y contains a flat half-strip, but that nevertheless its geometry is in general not rigid.

Collapse to a flat 2-disc with cone point or a sector (Humps). Suppose that X is the cone of radius 1 over a circle or interval with diameter $< \pi$. We generalize the cases of noses to humps by adapting the argument in [SY00, Sec. 3] to this case.

Let $A \subset \Sigma_x X$ be a finite subset such that $\sum_{\xi \in A} d(\sigma_\xi(t_0), \cdot)$ has a unique maximum in x. We choose the $\hat{p}_i$ as maxima of the corresponding functions $\sum_{\xi \in A} d(\sigma^i_\xi, \cdot)$. Then $\hat{p}_i \to x$. The point y_0 is a maximum of $\sum_{\xi \in A} \beta_\xi$. It follows that $\beta_\xi(z) = 0$ for all $\xi \in A$, and the totally convex subset $\cap_{\xi \in A}\{\beta_\xi \geq 0\} = \cap_{\xi \in A}\{\beta_\xi = 0\}$ containing y_0 and z has positive dimension. In particular, $\dim Y \geq 2$.

Suppose that $\dim Y = 2$. Then $\cap_{\xi \in A}\{\beta_\xi = 0\}$ is one-dimensional. Let y be an interior point of a segment $y_0 z$. Since the rays ρ^y_ξ for $\xi \in A$ are perpendicular to $y_0 z$, there can be at most two of them, $|A| \leq 2$. Since we are free to choose $|A|$ with any cardinality, we obtain a contradiction. Thus $\dim Y \geq 3$.

An argument analogous to the last one shows furthermore that the soul of Y (if it exists) must have codimension ≥ 2. In particular, if $\dim Y = 3$, then $\dim \mathrm{soul}(Y) \leq 1$.

Collapse to the flat 2-halfdisc. Suppose that X is the flat unit halfdisc in $\{u \in \mathbb{R}^2 : u_2 \leq 0\}$ centered at $x = 0$. (This case has not been treated explicitely in [SY00, Sec. 3]. There, blow-up limits have been obtained under the assumption that $\mathrm{diam}(\Sigma_x X) < \pi$.)

Lemma 6.3.5 implies that Y contains a flat halfplane, but $Y \ncong \mathbb{R}^2$. In particular, Y splits metrically as $Y \cong \mathbb{R} \times W$. If $\dim Y = 2$, then W is a halfline and Y a flat halfplane. If $\dim Y = 3$, then W is a noncompact Alexandrov surface with curvature ≥ 0. We may assume that $y_0 \in 0 \times W$. The critical points of $d(y_0, \cdot)$ lie on $0 \times W$.

If we denote by $\eta^\pm \in \Sigma_x X$ the directions pointing to $(\pm 1, 0)$, then for any $y \in Y$ the rays $\rho^y_{\eta^\pm}$ have angle π at y, cf. (6.3.4), and their union is the line $\mathbb{R} \times w$ through y. Moreover, $\{\beta_{\eta^+} = \beta_{\eta^-}\} = 0 \times W$, and it is the Gromov-Hausdorff limit of the bisectors $\{d(\sigma^i_{\eta^+}, \cdot) = d(\sigma^i_{\eta^-}, \cdot)\}$.

For any direction $\xi \in \Sigma_x X$ we have $\angle_y(\rho^y_\xi, \rho^y_{\eta^\pm}) = \angle_x(\xi, \eta^\pm)$ because $\pi = \angle_{\mathrm{Tits}}(\rho^y_{\eta^+}, \rho^y_\xi) + \angle_{\mathrm{Tits}}(\rho^y_\xi, \rho^y_{\eta^-}) \geq \angle_y(\rho^y_{\eta^+}, \rho^y_\xi) + \angle_y(\rho^y_\xi, \rho^y_{\eta^-}) \geq \angle_x(\eta^+, \xi) + \angle_x(\xi, \eta^-) = \pi$, also by (6.3.4). Let $\eta \in \Sigma_x X$ denote the (bisector) direction pointing to $(0, -1)$. Then the rays ρ^y_η are orthogonal to $\rho^y_{\eta^\pm}$ and contained in layers $t \times W$.

We now choose $\hat{p}_i$ as a maximum of $d(\sigma^i_\eta, \cdot)$ on the bisector $\{d(\sigma^i_{\eta^+}, \cdot) = d(\sigma^i_{\eta^-}, \cdot)\}$. Again $\hat{p}_i \to x$, and y_0 is a maximum of β_η on $\{\beta_{\eta^+} = \beta_{\eta^-}\}$. If Y is a flat halfplane, then $\{\beta_{\eta^+} = \beta_{\eta^-}\}$ is a halfline and y_0 its endpoint. This is a contradiction because $d(y_0, \cdot)$ has critical points and y_0 cannot lie on the boundary of the halfplane Y. Thus $\dim Y \geq 3$.

The next observation narrows down the possibilities for a 3-dimensional blow-up limit Y. Again, we consider the case where e.g. Y is a $\mathcal{C}^{10}$-regular orbifold of non-negative sectional curvature, and hence has a soul by Proposition 3.5.1.

Lemma 6.3.6. *If* $\dim Y = \dim W + 1 = 3$ *and* $\dim \mathrm{soul}(W) = 1$ *(and hence* W *is a quotient of the flat cylinder), then* W *must be one-ended.*

Proof. To see this, assume the contrary. Then W splits off a line, i.e. $W \cong \mathbb{R} \times F^1$ with a connected closed 1-orbifold F^1, and we may assume that $y_0 \in 0 \times F^1$. Since y_0 is a maximum of β_η, we have $\beta_\eta(s, f) = -|s|$. There are two unit speed rays $\rho_i : [0, \infty) \to 0 \times W$ starting from y_0 in antipodal directions, $\angle_y(\dot\rho_1(0), \dot\rho_2(0)) = \pi$, such that $\beta_\eta(\rho_i(t)) = -t$. From every point in $Y - 0 \times F^1$ a unique ray starts along which β_η decays with slope 1, and thus for $s > 0$ we have $\rho^{\rho_i(s)}_\eta(t) = \rho_i(s + t)$. It follows that there are points $x_{ij} \in \{d(\sigma^i_{\eta^+}, \cdot) = d(\sigma^i_{\eta^-}, \cdot)\}$ such that, with respect to the rescaled metrics, the segments $x_{ij}\sigma^i_\eta$ converge to the ray ρ_j. In particular, $\delta_i^{-1} d(\hat{p}_i, x_{ij}) \to 0$. On the other hand, without rescaling, the two sequences of segments converge to the same segment $x\sigma_\eta(t_0)$.

It follows by continuity that there exist points $z_{ij} \in x_{ij}\sigma^i_\eta$ such that $d(\sigma^i_\eta, z_{i1}) = d(\sigma^i_\eta, z_{i2})$ and $\tilde\angle_{\hat{p}_i}(z_{i1}, z_{i2}) = \frac{\pi}{3}$. We put $l_i = d(z_{i1}, z_{i2})$. Then $d(\hat{p}_i, z_{ij}) \to 0$ and $\delta_i^{-1} d(\hat{p}_i, z_{ij}) \to \infty$. Moreover, $\delta_i^{-1} |d(\hat{p}_i, z_{ij}) - l_i| \to 0$. Let m_i be the midpoints of segments $z_{i1}z_{i2}$.

Triangle comparison applied to the triangles $\Delta(z_{i1}, z_{i2}, \sigma^i_\eta)$ yields $\angle_{z_{ij}}(\sigma^i_\eta, m_i) \gtrapprox \frac{\pi}{2}$, and $d(m_i, z_{ij}\sigma^i_\eta) \gtrapprox \frac{l_i}{2}$, whereas comparison at $\Delta(z_{i1}, z_{i2}, \hat{p}_i)$ yields $\angle_{z_{ij}}(\hat{p}_i, m_i) \gtrapprox \frac{\pi}{3}$ and $d(m_i, z_{ij}\hat{p}_i) \gtrapprox \frac{\sqrt{3}}{4} l_i$. It follows that $l_i^{-1} d(m_i, x_{ij}\sigma^i_\eta)$ is bounded away from 0 and $\tilde\angle_{\hat{p}_i}(z_{ij}, m_i) \geq \phi_0 > 0$ for large i. The segments $\hat{p}_i m_i$ subconverge to a ray ρ in Y with initial point y_0 and $\angle_{y_0}(\rho_j, \rho) \geq \phi_0$.

Comparison at the triangles $\Delta(x_{ij}, \sigma^i_\eta, \sigma^i_{\eta^\pm})$ yields that $\liminf \tilde\angle_{x_{ij}}(z_{ij}, \sigma^i_{\eta^\pm}) \geq \frac{\pi}{2}$. Rescaling with the factors $l_i^{-1} \to \infty$ and taking into account that a Gromov-Hausdorff limit splits off a line shows that in fact $\tilde\angle_{x_{ij}}(z_{ij}, \sigma^i_{\eta^\pm}) \to \frac{\pi}{2}$ and furthermore $\tilde\angle_{x_{ij}}(m_i, \sigma^i_{\eta^\pm}) \to \frac{\pi}{2}$. Thus $\angle_{y_0}(\rho_{\eta^\pm}, \rho) = \frac{\pi}{2}$. In view of $\angle_{y_0}(\rho_j, \rho) \geq \phi_0$, this is a contradiction. □

Combining the discussion of collapse in the various special cases, we obtain:

Proposition 6.3.7 (Blow-up limits of strictly larger dimension). *If* $X = B(x, 1)$ *is a flat cone of dimension* ≤ 2*, then the base points* $\hat{p}_i$ *can be chosen so that* $\dim Y > \dim X$*.*

In fact, the arguments in the special cases above only used the conelikeness of X and the geometry of $\Sigma_x X$, and hence the conclusion of Proposition 6.3.7 holds more generally whenever X is conelike of dimension ≤ 2.

6.4 Strainers

6.4.1 Position relative to the singular locus

A Riemannian orbifold can be considered a local Alexandrov space with a very special singularity structure. The existence of a strainer in a point has implications for the singular type of this point. More precisely, for sufficiently small $\theta > 0$, a point admits a θ-straight m-strainer if and only if its link splits off a join factor isometric to the $(m-1)$-dimensional unit sphere, i.e. if and only if the singular stratum containing it has dimension $\geq m$. Similarly, a boundary point admits an m-strainer if and only if the singular stratum containing it has dimension $\geq m+1$. Such a strainer must be almost tangent to the singular stratum.

Thus an interior point in a Riemannian 3-orbifold O admits 3-strainers if and only if it is regular, admits 2-strainers if and only if it is regular or a reflector boundary point, and admits 1-strainers if and only if it is no singular vertex. For instance, $S_{\theta,\mu,-b^2} \cap O^{(0)} = \emptyset$.

6.4.2 Gradient-like vector fields

We recall that for a point $p \in O$ the distance function $d(p,\cdot)$ has directional derivatives. The derivative $\partial_v d(p,\cdot)$ in the direction of a unit tangent vector $v \in \Sigma_x O$ equals $-\cos(v, D_{x,p})$ where $D_{x,p} \subset \Sigma_x O$ is the compact subset of the directions of all segments xp. The function $v \mapsto \partial_v d(p,\cdot)$ on the unit tangent bundle of O is lower semicontinuous outside $\Sigma_p O$, and there its suplevel sets are open.(This argument also works for the distance function from a compact subset of O, e.g. from a component of ∂O if O has boundary.)

For $x \neq p$ and $c \in [0,1]$ the subset $\{v \in \Sigma_x O : \partial_v d(p,\cdot) \geq c\}$ is totally convex and has diameter $\leq 2\arccos c$ (e.g. because it has angular distance $\geq \pi - \arccos c$ from $D_{x,p}$). Such a subset for $c > 0$ can be nonempty only if $x \notin O^{(0)}$, because this requires that $\operatorname{diam}(\Sigma_x O) \geq \pi - \arccos c > \frac{\pi}{2}$. If it is nonempty for a singular point $x \in O^{(1)} \cup O^{(2)}$, then it contains a singular direction at x.

By the usual construction using a partition of unity, it follows that for $c \in [0,1)$ there exists a gradient-like vector field X for $d(p,\cdot)$ on the open subset $\{x \neq p : \exists v \in \Sigma_x O \text{ with } \partial_v d(p,\cdot) > c\}$ which is tangent to the singular locus and of course satisfies $\partial_X d(p,\cdot) > c$.

For distinct points a and b there exist gradient-like vector fields for $d(a,\cdot)$ and $d(b,\cdot)$ on the open set $S_{a,b} = \{\angle.(a,b) > \frac{\pi}{2}\} \subset O - \{a,b\}$. (Note that the function $\angle.(a,b)$ is lower semicontinuous on $O - \{a,b\}$.) More precisely, for $\phi \in (0, \frac{\pi}{2}]$ there exists a gradient-like vector field X for $d(a,\cdot)$ on $\{\angle.(a,b) > \pi - \phi\}$ with $\angle.(a,X) > \pi - \phi$. Such a vector field satisfies $\angle.(b,-X) \geq \angle.(a,b) - \angle.(a,-X) > \pi - 2\phi$ and $\angle.(b,X) < 2\phi$. (Note that $-X$ is defined since X is tangent to the singular locus.) Thus, if $\phi \leq \frac{\pi}{4}$ then $\partial_X d(b,\cdot) < 0$ and $\partial_X f_{a,b} > 0$, i.e. X is gradient-like for $f_{a,b}$.

If (a,b) is a θ-straight 1-strainer at p for sufficiently small θ, then its cross section $\Sigma_{x;a,b}$ is near p a topological 2-suborbifold, because a gradient-like vector field for $f_{a,b}$ has local cross sections through p which are smooth 2-suborbifolds, and any two local cross sections can be isotoped to each other using the flow.

6.4.3 Local bilipschitz charts and fibrations by cross sections

Suppose that $(a_1,b_1,a_2,b_2,a_3,b_3)$ is a $<\theta$-straight 3-strainer at a regular point $x \in O$ for some small $\theta > 0$. Then there exist unit vectors $v_1,v_2,v_3 \in \Sigma_x O$ with $\angle_x(a_i,v_i) > \pi - \theta$. It follows that $\angle_x(b_i,-v_i) > \pi - 2\theta$, $|\angle_x(a_i,v_j) - \frac{\pi}{2}|, |\angle_x(b_i,v_j) - \frac{\pi}{2}| < 2\theta$ and $|\angle_x(v_i,v_j) - \frac{\pi}{2}| < 3\theta$ for $i \neq j$. Let X_i be arbitrary commuting smooth vector fields near x with $X_i(x) = v_i$. By continuity, on a sufficiently small neighbourhood of x they have length ≈ 1 and satisfy the same angle inequalities $\angle.(a_i,X_i) > \pi-\theta$ and their implications. Thus, they are almost orthogonal, $|\angle(X_i,X_j) - \frac{\pi}{2}| < 3\theta$ for $i \neq j$, and gradient-like for the functions f_{a_i,b_i} associated to the 1-substrainers, $\partial_{X_i} f_{a_i,b_i} > \frac{1}{2}(\cos\phi + \cos 2\phi) > 1 - c\theta^2$, and $|\partial_{X_j} f_{a_i,b_i}| < \sin 2\theta < 2\theta$ for $i \neq j$. The X_i are the coordinate vector fields for some local coordinates, and it follows that $(f_{a_1,b_1}, f_{a_2,b_2}, f_{a_3,b_3})$ restricts to a bilipschitz homeomorphism from a neighbourhood of x onto an open subset of $\mathbb{R}^3$. (Compare the discussion in [BGP92, Sec. 11.8] and [MT08, 2.4.1].)

An analogous argument can be carried out for a $2\frac{1}{2}$-strainer (a_1,b_1,a_2,b_2,a_3) at a reflector boundary point $x \in O^{(2)}$. Then the 2-substrainer (a_1,b_1,a_2,b_2) is almost tangent to the reflector boundary. One uses an orbifold chart at x, lifts the functions f_{a_1,b_1}, f_{a_2,b_2} and constructs f_{a_3,b_3} by lifting a_3 to a 1-strainer in the chart. Furthermore, one adapts the smooth vector fields X_i to the reflection ι on the chart, i.e. constructs them so that $\iota^*X_1 = X_1$, $\iota^*X_2 = X_2$ and $\iota^*X_3 = -X_3$. One obtains a local bilipschitz homeomorphism to the 3-dimensional halfspace with reflector boundary.

Consider now a $<\theta$-straight 2-strainer (a_1,b_1,a_2,b_2) at a point x. That its cross section $\Sigma_{x;a_1,b_1,a_2,b_2}$ is near x a bilipschitz 1-suborbifold can be seen as follows: If x is a regular point, then one can choose an $\approx\theta$-straight 1-strainer (a',b') contained in $\Sigma_{x;a_1,b_1,a_2,b_2}$ and, with respect to the local bilipschitz coordinates near x provided by the 3-strainer (a_1,b_1,a_2,b_2,a',b'), the cross section $\Sigma_{x;a_1,b_1,a_2,b_2}$ is a coordinate line. If x is singular, then it must be a reflector boundary point. One chooses a point $a' \in \Sigma_{x;a_1,b_1,a_2,b_2}$ near x and works with the bilipschitz coordinates provided by the $2\frac{1}{2}$-strainer (a_1, b_1, a_2, b_2, a').

Suppose that C is a compact connected component of $\Sigma_{x;a_1,b_1,a_2,b_2}$ such that (a_1, b_1, a_2, b_2) is a $c\theta$-straight 2-strainer at all points of C. Then C is a closed 1-suborbifold and hence homeomorphic to S^1 or the mirrored interval I^1. The map $(f_{a_1,b_1}, f_{a_2,b_2})$ yields a product fibration of a neighbourhood of C by cross sections of the 2-strainer. This can be seen as follows using the bilipschitz coordinates near the points $y \in C$. The distance functions f_{a_1,b_1}, f_{a_2,b_2} given by the 2-strainer and the auxiliary distance function f_{a_3,b_3} near y are normalized so that $f_{a_i,b_i}(y) = 0$. For sufficiently small $\epsilon_i > 0$ the map $(f_{a_1,b_1}, f_{a_2,b_2})$ yields near y a product fibration of the box $\{|f_{a_i,b_i}| \leq \epsilon_i \,\forall i\}$ over the rectangle $[-\epsilon_1, \epsilon_1] \times [-\epsilon_2, \epsilon_2]$. By covering C with finitely many such boxes one obtains the fibration of a neighbourhood.

In the above discussion, one or both of the functions f_{a_1,b_1} and f_{a_2,b_2} can be replaced by $d(a_i, \cdot)$, because their directional derivatives differ only slightly. (Recall that $f_{a_i,b_i} - d(a_i, \cdot)$ is $c'\theta$-Lipschitz in the region where (a_i, b_i) is a $c\theta$-straight 1-strainer, cf. section 5.4.1).

6.5 A decomposition according to the coarse stratification

We start by formulating the collapse assumption on the orbifolds (O, g) needed in this section and the quantities involved in it.

The parameter θ (straightness of 1-strainers) is required to be small, $\theta \in (0, \theta_0]$, where θ_0 is sufficiently small for the arguments in sections 5.4 and 5.5 to apply (cf. the discussion at the beginning of section 5.4). The upper bound for θ will be decreased several times during our later arguments. The parameter μ (accuracy of conical approximation) needs to be sufficiently small so that the conclusions of 5.3.2 regarding the existence of θ-straight 1-strainers apply, $\mu \leq \mu_0(\theta)$ with the constant $\mu_0(\theta)$ from there. The parameter μ determines (together with the fixed parameter σ) via Proposition 6.2.1 the bound $s_1(\sigma, \mu)$ and the scale $\hat{s}_{\mu,-b^2} := \frac{1}{11} s_1(\sigma, \mu)\rho_{-b^2}$. Conical approximation in all points $x \in O$ on scales $s(x) \in [s_1(\sigma, \mu)\rho_{-b^2}(x), \sigma\rho_{-b^2}(x)]$ holds if (O, g) is $(v(\sigma, \mu), -b^2)$-collapsed with the constant $v(\sigma, \mu) > 0$ from Proposition 6.2.1.

In order to make the results on edgy points in section 5.5 available, we need to rule out the existence of $\bar{\theta}_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainers with length $\lambda(\theta)\hat{s}_{\mu,-b^2}(x)$ at all points x for a constant $\lambda(\theta)$ which is sufficiently small so that Lemmas 5.5.10, 5.5.11 and 5.5.12 hold. This is achieved by requiring (O, g) to be $(v(\theta, \mu), -b^2)$-collapsed for a suitable small constant $v(\theta, \mu) > 0$, cf. Lemma 6.1.5. We make it so small that it is smaller than the constant $v(\sigma, \mu)$ mentioned above.

Furthermore, in order to obtain sufficient collapse on the scale $\theta^4\hat{s}_{\mu,-b^2}(x)$ for small θ, we ask that $v(\theta,\mu) \le (\theta^5 s_1(\sigma,\mu))^3$. We will also assume that (O,g) has $(v(\theta,\mu), s_0, K)$-curvature control below scale ρ_{-b^2}.

6.5.1 The 2-strained region

We define the *2-strained region* $R_{\theta,\mu,-b^2} \subset O$ as the open set consisting of all points x which admit $< C_1\theta$-straight 2-strainers of length $> \theta^4\hat{s}_{\mu,-b^2}(x)$, where C_1 is the constant from Lemma 5.5.4. Note that $R_{\theta,\mu,-b^2} \cap O^{(sing)} \subset O^{(2)}$.

We will show that for sufficiently small θ the 2-strained region admits metrically an almost product fibration with short fibers almost orthogonal to the 2-strainers. (The smallness of the parameter μ is not important at this point, because we are not yet using the conical approximation from Proposition 6.2.1.)

We begin with a local approximation result:

Proposition 6.5.1. *For $\epsilon > 0$ there is $\theta_1 = \theta_1(\epsilon) > 0$ such that the following holds: Let $\theta \le \theta_1$ and $\mu \le \mu_0(\theta)$. If (O,g) is $(v(\theta,\mu), -b^2)$-collapsed with $(v(\theta,\mu), s_0, K)$-curvature control below scale ρ_{-b^2}, and if (a_1,b_1,a_2,b_2) is a $C_1\theta$-straight $\theta^4\hat{s}_{\mu,-b^2}(x)$-long 2-strainer at x, then $\mathrm{diam}(\Sigma^o_{x;a_1,b_1,a_2,b_2})^{-1}\cdot(O,x)$ is ϵ-close in the pointed $\mathcal{C}^5$-topology to the product $\mathbb{R}^2 \times F^1$ of the euclidean plane with a connected closed 1-orbifold of diameter 1.*

Proof. Let $-b_i^2 \in [-1,0)$ and $\theta_i, \mu_i > 0$ such that $\theta_i \to 0$ and $\mu_i \le \mu_0(\theta_i)$. Suppose that the orbifolds (O_i,g_i) are $(v(\theta_i,\mu_i), -b_i^2)$-collapsed with $(v(\theta_i,\mu_i), s_0, K)$-curvature control below the scales $\rho_{-b_i^2}$, and that the $(a_1^i,b_1^i,a_2^i,b_2^i)$ are $C_1\theta_i$-straight $\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i)$-long 2-strainers at the points x_i. We have to show that the conlusion of the proposition holds for large i.

Consider the cross sections $\Sigma_i^o = \Sigma^o_{(x_i;a_1^i,b_1^i,a_2^i,b_2^i)}$. For any point $x_i \neq y_i \in \Sigma_i^o$, the $2\frac{1}{2}$-strainer $(a_1^i,b_1^i,a_2^i,b_2^i,y_i)$ at x_i is $\bar\theta_{2\frac{1}{2}}$-straight for large i. Since $v(\theta,\mu) \le (\theta^5 s_1(\sigma,\mu))^3$, it follows that $\lambda_i := (\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}\,\mathrm{diam}(\Sigma_i^o) \to 0$, cf. Lemma 6.1.5. The rescaled pointed orbifolds $(\lambda_i\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}\cdot(O_i,x_i)$ Gromov-Hausdorff subconverge to a pointed Alexandrov space (X,x_0) with dimension ≤ 3 and curvature ≥ 0. Moreover, the broken segments $a_1^ix_ib_1^i$ and $a_2^ix_ib_2^i$ (sub)converge to two perpendicular lines through x_0, and hence X splits metrically as a product $\mathbb{R}^2\times\Sigma$.

To see that the rescaled cross sections $(\lambda_i\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}\cdot\Sigma_i^o$ subconverge to Σ, we note that $(\lambda_i\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}(d(a_j^i,\cdot)-d(a_j^i,x_i))$ and $(\lambda_i\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}(d(b_j^i,\cdot)-d(b_j^i,x_i))$ subconverge (due to Arzelà-Ascoli) to concave 1-Lipschitz functions α_j and β_j with $\alpha_j(x_0) = \beta_j(x_0) = 0$. The concavity of the sums $\alpha_j+\beta_j$ implies together with the triangle inequality that the functions α_j and β_j are constant on fibers $pt\times\Sigma$. Since for any sequence points $y_i\in\Sigma_i^o$ we have $(\lambda_i\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}(d(a_j^i,y_i)-d(a_j^i,x_i))$, $(\lambda_i\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}(d(b_j^i,y_i)-$

$d(b_j^i, x_i)) \to 0$, compare (5.4.6), it follows that $\Sigma_i^o \to \Sigma$. Therefore Σ is a compact 1-dimensional Alexandrov space of diameter 1.

Since in the blow-up limit $(\lambda_i \theta_i^4 \hat{s}_{\mu_i,-b_i^2}(x_i))^{-1} \cdot (O_i, x_i) \to (X, x_0)$ we have no dimension drop, after passing to another subsequence, the convergence can be improved to $\mathcal{C}^5$-smooth convergence and $F^1 = \Sigma$ is a connected closed 1-orbifold, cf. Lemma 6.1.7. □

We observe that in the proof, the local fibrations induced by 2-strainers as discussed in section 6.4 locally Gromov-Hausdorff converge to the natural product fibration of X. They are therefore isotopic by *small* isotopies to the product approximations in Corollary 6.5.1.

Note that in the situation of the proposition, $\text{diam}(\Sigma_{x;a_1,b_1,a_2,b_2}) << \theta^4 \hat{s}_{\mu,-b^2}(x)$. Moreover, for $x \in R_{\theta,\mu,-b^2}$ the cross sections $\Sigma^o_{x;a_1,b_1,a_2,b_2}$ of different $C_1\theta$-straight $\theta^4 \hat{s}_{\mu,-b^2}(x)$-long 2-strainers (a_1, b_1, a_2, b_2) at x have Hausdorff distance $<< \text{diam}(\Sigma^o_{x;a_1,b_1,a_2,b_2})$ and almost equal diameters, and we define the *width* $w(x)$ at x as the infimum of these diameters.

The fiber direction of a local approximation as in Corollary 6.5.1 yields a smooth line field which is *almost vertical* in the sense that it is perpendicular to the stratum $O^{(2)}$ and almost perpendicular to sufficiently long segments. Any two such local line fields almost agree on the overlaps of their domains of definition, and using a partition of unity we can combine such local line fields to a global almost vertical line field $L = L_{\theta,\mu,-b^2}$ on $R_{\theta,\mu,-b^2}$. More precisely, for (small) $\nu > 0$ there is $\theta_1 = \theta_1(\nu) > 0$ such that the following holds: If $\theta \leq \theta_1$, $\mu \leq \mu_0(\theta)$ and if (O, g) is $(v(\theta, \mu), -b^2)$-collapsed with $(v(\theta, \mu), s_0, K)$-curvature control below scale ρ_{-b^2}, then for $x \in R_{\theta,\mu,-b^2}$ the line $L(x)$ has angle $> \frac{\pi}{2} - \nu$ with any segment of length $10\nu^{-1}w(x)$ initiating in x, and in particular with any segment of length $\theta^4 \hat{s}_{\mu,-b^2}(x) >> w(x)$. In fact, a line field $L_{\theta,\mu,-b^2}$ with these properties can be constructed on the slightly larger open set $\hat{R}_{\theta,\mu,-b^2} := \cup_{x \in R_{\theta,\mu,-b^2}} B(x, \frac{1}{\nu} w(x))$. The reflector boundary $O^{(2)}$ is, where it meets $R_{\theta,\mu,-b^2}$, *almost horizontal* in the sense that it is almost tangent to sufficiently long segments and, in particular, to sufficiently long 2-strainers.

The trajectories of L starting in a point $x \in R_{\theta,\mu,-b^2}$ move almost orthogonally to sufficiently long segments and therefore remain close to the cross section $\Sigma^o_{x;a_1,b_1,a_2,b_2}$ of a suitable 2-strainer as above for length at least $>> w(x)$.

If $x \in O^{(2)} \cap R_{\theta,\mu,-b^2}$, then the trajectory is orthogonal to $O^{(2)}$ in x and reaches $O^{(2)}$ again after length $\approx w(x)$. More generally, all trajectories intersecting $B(x, 10w(x))$ have length $\approx w(x)$ and connect reflector boundary points. The construction of an almost product fibration by mirrored intervals close to $O^{(2)} \cap R_{\theta,\mu,-b^2}$ is therefore immediate.

Away from the reflector boundary, the L-trajectories starting in points $x \in R_{\theta,\mu,-b^2}$ almost close up after length $\approx 2w(x)$. The question whether L can be globally perturbed to an integrable line field with closed trajectories of lengths $\approx 2w(x)$ has been treated in

[MT08, Sec. 4.2] in a very similar setting. The discussion there uses only the control on finitely many derivatives of the curvature tensor and goes through without change in the situation considered here. One obtains the following result which can be considered as a version of a special case of Yamaguchi's Fibration Theorem [Ya91] "without a priori given base".

Proposition 6.5.2 (Almost vertical fibration of the 2-strained region, cf. [MT08, Prop. 4.4]). *For $\nu > 0$ there is $\theta_1 = \theta_1(\nu) > 0$ such that the following holds: If $\theta \leq \theta_1$, $\mu \leq \mu_0(\theta)$ and if (O, g) is $(v(\theta,\mu), -b^2)$-collapsed with $(v(\theta,\mu), s_0, K)$-curvature control below scale ρ_{-b^2}, then there exists an open subset U, $R_{\theta,\mu,-b^2} \subseteq U \subseteq \hat{R}_{\theta,\mu,-b^2}$, such that every connected component of U is the total space of a smooth orbifold fibration with fiber S^1 or the mirrored interval I, and all fibers have angle $< \nu$ with the almost vertical line field $L_{\theta,\mu,-b^2}$.*

The local fibrations provided by the product approximations in Corollary 6.5.1 have only a finite degree of regularity, but the global fibration of U obtained by interpolating these local fibrations can be smoothed. The smoothness will however not be important to us.

From now on, we fix some small positive value of ν (e.g. $\nu = \frac{1}{1000}$) and set $\theta_1 = \theta_1(\nu)$. Moreover, whenever an orbifold is sufficiently collapsed, we implicitely fix a fibration as in Proposition 6.5.2.

6.5.2 Edges

Definition 6.5.3. We define a point $x \in S_{\theta,\mu,-b^2}$ to be $(\theta, \mu, -b^2)$-*edgy* relative to an equilateral $< \theta$-straight 1-strainer (a, b) at x with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$ if $\operatorname{diam}(\Sigma^o_{x;a,b}) > \theta^{\frac{5}{2}}\hat{s}_{\mu,-b^2}(x)$ and if $\Sigma_{x;a,b}$ contains no $\frac{\pi}{2}$-straight $\theta^4\hat{s}_{\mu,-b^2}(x)$-long 1-strainers at x, compare Definition 5.5.8.

Let us briefly discuss the correspondence between the two definitions 5.5.8 and 6.5.3. First, we recall from our discussion in section 5.1.3 that a $(\theta, \mu, -b^2)$-edgy point x relative to a $< \theta$-straight 1-strainer (a, b) with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$ is also θ-edgy in the space $(\hat{s}_{\mu,-b^2}(x))^{-1} \cdot B(x, \rho_{-b^2}(x))$ of curvature ≥ -1 in the sense of Definition 5.5.8: In the rescaled space, the strainer (a, b) is $< 2\theta$-straight of length $> (1 - \theta)$ and the cross section $\Sigma_{x;a,b}$ cannot contain any $\frac{\pi}{2}$-straight 1-strainer of length θ^4.

Conversely, suppose that an equilateral 1-strainer (a, b) is $< \theta$-straight at x with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$ and that in the rescaled space $\hat{s}_{\mu,-b^2}(x)^{-1} \cdot B(x, \rho_{-b^2}(x))$ the point x is θ-edgy relative to (a, b) (again in the sense of Definition 5.5.8). Then Remark 5.5.5 together with the discussion at the end of section 5.1.3 implies that x is also $(\theta, \mu, -b^2)$-

edgy. For suppose that $\Sigma_{x;a,b}$ contains a $\frac{\pi}{2}$-straight $\theta^4\hat{s}_{\mu,-b^2}(x)$-long 1-strainer at x. Then after rescaling this strainer is still $< \frac{3\pi}{4}$-straight and hence in fact $< C_1\theta$-straight which is a contradiction.

By construction, for any $x \in O$ there are no $\bar{\theta}_{2\frac{1}{2}}$-straight $2\frac{1}{2}$-strainers of length $\lambda(\theta)$ in the rescaled space $\hat{s}_{\mu,-b^2}(x)^{-1}B(x, \frac{1}{2}\rho_{-b^2}(x))$ Moreover, estimate 6.1.3 implies that on the ball $B(x, \theta\hat{s}_{\mu,-b^2}(x))$ we have the estimate $\hat{s}_{\mu,-b^2}/\hat{s}_{\mu,-b^2}(x) \in (1-\theta, 1+\theta)$.

This allows us to generalize the above arguments to the following results: If $y \in B(\theta\hat{s}_{\mu,-b^2}(x))$ is $(\theta,\mu,-b^2)$-edgy relative to a 1-strainer (a,b) then it is θ-weakly in the rescaled space $\hat{s}_{\mu,-b^2}(x)^{-1} \cdot B(x,\rho_{-b^2}(x))$ of curvature ≥ -1. Conversely, suppose that $y \in B(\theta\hat{s}_{\mu,-b^2}(x))$ admits an equilateral $< \theta$-straight 1-strainer (a,b) with length in $(\hat{s}_{\mu,-b^2}(y), \frac{3}{2}\hat{s}_{\mu,-b^2}(y))$ such that y is θ-strongly edgy relative to (a,b) in the rescaled space $\hat{s}_{\mu,-b^2}(x)^{-1} \cdot B(x,\rho_{-b^2}(x))$. Then y is also $(\theta,\mu,-b^2)$-edgy.

Hence our results from section 5.5.2 can be applied to our present situation and used to control e.g. the relative position of edgy points.

We will globally construct tubes along the "coarse edges". Our previous discussion applies provided that θ and μ are sufficiently small (i.e. $\theta \leq \theta_1$ and $\mu \leq \mu_0(\theta)$), and that (O,g) is $(v(\theta,\mu),-b^2)$-collapsed with $(v(\theta,\mu),s_0,K)$-curvature control below scale ρ_{-b^2}.

For every $(\theta,\mu,-b^2)$-edgy point $x \in O$, let (a_x,b_x) be some equilateral $< \theta$-straight 1-strainer with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$ relative to which x is edgy. (Since 1-strainers at x are almost unique by Lemma 5.5.6, it does not matter which one we use.) Associated to it is the truncated cross section $\check{\Sigma}_x = \Sigma_{x;a_x,b_x} \cap B(x, \frac{1}{2}\theta^3\hat{s}_{\mu,-b^2}(x))$.

For $j = 1,2,3$ we consider the fibers $\gamma_x^j = \Sigma_{x;a_x,b_x} \cap B(x, \frac{j}{8}\theta^3\hat{s}_{\mu,-b^2}(x))$ of the partial (topological) product fibration of $\Sigma_{x;a_x,b_x}$ induced by $d(x,\cdot)$.(Compare the discussion in section 6.4.3.) The γ_x^j and also their $\frac{1}{16}\theta^3\hat{s}_{\mu,-b^2}(x)$-neighbourhoods are contained in the 2-strained region $R_{\theta,\mu,-b^2}$, and they are almost vertical in the sense that they are isotopic to a fiber of the fibration given by Proposition 6.5.2 by a small isotopy, say, supported on the $\theta^4\hat{s}_{\mu,-b^2}(x)$-neighbourhood of γ_x^j.

An almost unique 1-strainer at a $(\theta,\mu,-b^2)$-edgy point x locally defines a vector field on the ball $B(x,\theta\hat{s}_{\mu,-b^2})$ as in section 6.4.2. By interpolating these fields by a partition of unity, we obtain a smooth vector field $L = L_{\text{edge}}$ tangent to the singular locus whose open domain of definition contains the balls $B(x,\theta\hat{s}_{\mu,-b^2}(x))$ around the $(\theta,\mu,-b^2)$-edgy points x, and which is on these balls almost parallel to the 1-strainers (a_x,b_x), meaning that $\angle.(L_{\text{edge}},a_x), \angle.(L_{\text{edge}},b_x) \notin [c\theta^{\frac{1}{2}}, \pi - c\theta^{\frac{1}{2}}]$ on $B(x,\theta\hat{s}_{\mu,-b^2}(x))$, see Lemma 5.5.10. In particular, we have $|\angle.(L_{\text{edge}},x) - \frac{\pi}{2}| < c'\theta^{\frac{1}{2}}$ on the $\frac{1}{7}\theta^3\hat{s}_{\mu,-b^2}(x)$-neighbourhood of γ_x^2.

We choose a maximal subfamily of pairs $(x,\check{\Sigma}_x)$ such that the corresponding subset ϵ of $(\theta,\mu,-b^2)$-edgy points x is separated in the sense that for any two distinct points $x_1,x_2 \in \epsilon$

holds $d(x_1, x_2) > \theta^{\frac{10}{3}} \hat{s}_{\mu,-b^2}(x_1)$. By Lemma 5.5.12, the $\check{\Sigma}_x$ for $x \in \epsilon$ are pairwise disjoint.

We call $x_1, x_2 \in \epsilon$ *adjacent*, if $d(x_1, x_2) < 4\theta^{\frac{10}{3}} \hat{s}_{\mu,-b^2}(x_1)$ and if the arc in the L_{edge}-trajectory (leaf) connecting x_1 to its intersection point with $\check{\Sigma}_{x_2}$ does not meet the other cross sections $(x, \check{\Sigma}_x)$ for $x \in \epsilon - \{x_1, x_2\}$. By Lemma 5.5.11, we have $d(x_1, x_2) \approx \theta^{\frac{10}{3}} \hat{s}_{\mu,-b^2}(x_1)$ unless points on the segment x_1x_2 have no θ-straight $\hat{s}_{\mu,-b^2}$-long 1-strainers or have such strainers with cross sections of diameter $\approx \theta^{\frac{5}{2}}$. The relation of adjacency generates an equivalence relation on ϵ, and we call an equivalence class a *chain* of $(\theta, \mu, -b^2)$-edgy points. Each chain can be given a linear or cyclic order.

Consider two adjacent edgy points $x_1, x_2 \in \epsilon$. The cross sections $\check{\Sigma}_{x_i}$ have Hausdorff distance $< \theta^{\frac{99}{30}} \hat{s}_{\mu,-b^2}(x_1)$ by Lemma 5.5.12. Using the integral curves of the line field L_{edge}, we flow the 1-orbifolds $\gamma^j_{x_1}$ from $\check{\Sigma}_{x_1}$ into $\check{\Sigma}_{x_2}$. Their images $\gamma'^j_{x_1,x_2}$ in $\check{\Sigma}_{x_2}$ are $< \theta^{\frac{9}{20}+\frac{99}{30}} \hat{s}_{\mu,-b^2}(x_1) < \theta^{\frac{7}{2}} \hat{s}_{\mu,-b^2}(x_1)$-close to the $\gamma^j_{x_2}$, compare the discussion of the maps $\Phi^{a,b}_t$ in section 5.4.1. Moreover, since the isotopy of $\gamma^j_{x_1}$ to $\gamma'^j_{x_1,x_2}$ takes place inside a $\theta^{\frac{98}{30}} \hat{s}_{\mu,-b^2}(x_1)$-ball contained in $R_{\theta,\mu,-b^2}$, $\gamma'^j_{x_1,x_2}$ is isotopic by a small isotopy to the almost vertical fibers of the fibration of $R_{\theta,\mu,-b^2}$, and therefore it must inside $\check{\Sigma}_{x_2}$ be homotopic to $\gamma^j_{x_2}$, and hence isotopic by a small isotopy. Thus the trace of the isotopy of $\gamma^j_{x_1}$ to $\gamma'^j_{x_1,x_2}$ can be adjusted (by a small isotopy supported near $\gamma^j_{x_2}$) to a 2-suborbifold $S^j_{x_1,x_2}$ homeomorphic to $\cong \gamma^j_{x_1} \times [0,1]$, contained in $A(x; (\frac{j}{8} - \frac{1}{100})\theta^3 \hat{s}_{\mu,-b^2}(x), (\frac{j}{8} + \frac{1}{100})\theta^3 \hat{s}_{\mu,-b^2}(x))$, lying between $\check{\Sigma}_{x_1}$ and $\check{\Sigma}_{x_2}$, and with boundary $\gamma^j_{x_1} \cup \gamma^j_{x_2}$. By concatenating the $S^j_{x_1,x_2}$, we obtain three disjoint embedded 2-suborbifolds S^j following along the chains of edgy points. We call the region contained between S^1 and the two final cross section of the chain a *tube* along the coarse edge.

We observe that by Lemma 5.5.11 all edgy points close to an edgy point $x \in \epsilon$ are contained in the corresponding tube along the coarse edge. Again by 5.5.11, a chain can only end inside a $(\theta, \mu, -b^2)$-hump or if the cross sections to $C_0\theta$-straight $\hat{s}_{\mu,-b^2}$-long 1-strainers have diameter $\approx \theta^{\frac{5}{2}}$. It is also possible that an edgy point has no adjacent edgy points. We simply discard such isolated cross sections.

The simplest *interface* between chains of edgy points (the "coarse edge") and the rest of O arises for a *cyclic chain* $\kappa \subseteq \epsilon$ of $(\theta, \mu, -b^2)$-edgy points. The parts $S^j_\kappa \subset S^j$ corresponding to κ are then closed topological 2-suborbifolds. Let T^j_κ denote the tube containing κ and bounded by S^j_κ, and let $A^{i,j}_\kappa := T^j_\kappa - int(T^i_\kappa)$ for $1 \leq i < j \leq 3$. The T^j_κ and $A^{i,j}_\kappa$ are compact topological 3-suborbifolds compatibly fibering over the circle, and $A^{1,3}_\kappa \subset R_{\theta,\mu,-b^2}$. The fiber of T^j_κ is (homeomorphic to) a compact 2-orbifold with one boundary component. Note that S^2_κ separates S^1_κ and S^3_κ.

We wish to replace S^2_κ by a smooth 2-suborbifold which is *vertically saturated*, i.e. saturated with respect to the fibration of $R_{\theta,\mu,-b^2}$, cf. Proposition 6.5.2. To do so, we take a vertically saturated compact connected 3-suborbifold W such that $S^2_\kappa \subset \text{int}(W)$ and

$W \subset \mathrm{int}(A_\kappa^{1,3})$. Then W separates S_κ^1 and S_κ^3, and according to Lemma 6.5.4 below, one of the boundary components of W separates S_κ^1 and S_κ^3. We denote this boundary component by $S_\kappa^{2,v}$.

Then as a consequence of Lemma 6.5.5, $S_\kappa^{2,v}$ is isotopic to the S_κ^j. In other words, we can isotope S_κ^2 by an isotopy supported in $\mathrm{int}(A_\kappa^{1,3})$ so that it becomes vertically saturated and the fibrations on it induced by $R_{\theta,\mu,-b^2}$ and T_κ^2 match. (We observe that the boundary of every truncated cross ection $\check{\Sigma}_x$, $x \in \epsilon$ can be isotoped by a small isotopy to a fiber of $R_{\theta,\mu,-b^2}$ by the remark after the proof of Proposition 6.5.1.)

Lemma 6.5.4. *Let Σ be a connected closed 2-orbifold without singular points and let $W \subset \Sigma \times [0,1]$ be a compact connected 3-suborbifold disjoint from $\Sigma \times 0$ and $\Sigma \times 1$ and separating them. Then some component of ∂W separates them also.*

Proof. For the purpose of this lemma, we consider Σ and Σ' as compact manifolds, possibly with boundary.

Based on the existence and uniqueness of smooth structures on 3-manifolds and the uniqueness up to isotopy of smooth structures on 2-manifolds, we know that there exists a smooth structure on $\Sigma \times [0,1]$ with respect to which the embedded topological 2-submanifold ∂W is a smooth submanifold. (Cut along ∂W, put a smooth structure and glue again after adjusting the induced smooth structures on the boundaries by an isotopy in a collar. See e.g. [Mu60], [Wh61], [Ep66].)

Moreover, given a smooth structure on Σ (and hence on $\Sigma \times [0,1]$), there exists a homeomorphism of $\Sigma \times [0,1]$ carrying the embedded topological 2-submanifolds ∂W to a smooth submanifold. Hence we may work without loss of generality in the smooth category.

Let V denote the component of $\Sigma \times (0,1) - int(W)$ containing $\Sigma \times 0$. Then $\Sigma' = V \cap W$ separates $\Sigma \times 0$ and $\Sigma \times 1$, and we have to show that it is connected. Suppose the contrary, i.e. that it decomposes as the disjoint union $\Sigma' = \Sigma'_1 \cup \Sigma'_2$ of closed 2-submanifolds. Then there exists an embedded circle γ in $\Sigma \times (0,1)$ which intersects Σ'_1 once transversally. This is absurd because γ can be homotoped into $\Sigma \times 0$. □

The following fact is for tori a simple special case of a result of Waldhausen [Wa67, 2.8]. The arguments in the non-orientable and orbifold cases are similar.

Lemma 6.5.5. *(i) Let Σ and Σ' be closed surfaces, each of which is homeomorphic to the 2-torus T^2 or to the Klein bottle K^2. Suppose that $\Sigma' \subset \Sigma \times [0,1]$ is embedded so that it is disjoint from $\Sigma \times 0$ and $\Sigma \times 1$ and separates them. Then Σ' is isotopic to $\Sigma \times 0$ and $\Sigma \times 1$.*

(ii) The same conclusion holds if Σ and Σ' are closed 2-orbifolds, each of which is homeomorphic to the annulus Ann^2 or the Möbius strip $Möb^2$ with reflector boundary.

Proof. Again we can assume without loss of generality that Σ' is a smooth surface, respectively, suborbifold.

(i) Cut open $\Sigma \times [0,1]$ along an annulus A so as to obtain a solid torus $[0,1] \times [0,1] \times S^1$. After adjusting Σ' we can assume that it intersects A transversally in circles. Because Σ' separates $\Sigma \times [0,1]$, every circle which is null-homotopic in A bounds a 2-ball in Σ'. (Otherwise, by the orientability and irreducibility of the solid torus, such a circle would decompose Σ' into an annulus, and we could compress Σ' to a circle or a point.) Using irreducibility again, we can isotope Σ' such that it intersects A only in circles which are not null-homotopic and decompose Σ' into annuli. Moreover, we can assume that these circles are vertical in the fibration of the full torus by circles. Hence every annulus component of $\Sigma' - A$ can be isotoped to be vertical, too. (Cf. [Wa67, 2.4].) Thus, Σ' can be isotoped to be vertical in a fibration of $\Sigma \times [0,1]$ by circles. This clearly implies (i).

(ii) Without loss of generality, we can assume that every boundary component of $\partial\Sigma'$ is horizontal in the product $\Sigma \times [0,1]$. In the case where Σ is homeomorphic to Ann^2, so is Σ' and the two boundary components of Σ' lie above different boundary components of Σ.

As above, we cut open $\Sigma \times [0,1]$ along a 2-ball B to obtain a 3-ball $[0,1] \times [0,1] \times [0,1]$. Again, we arrange that Σ' is transversal to B and hence intersects it in null-homotopic circles or in intervals connecting opposite sides of $B \cong [0,1] \times [0,1]$ or one side to itself.

Circles in $B \cap \Sigma'$ must again bound 2-balls in Σ' and hence can be removed by suitable isotopies. (If Σ' is a Moebius band, this follows from the orientability of the 3-ball. If Σ' is an annulus which is decomposed into two annuli by a component of $\Sigma' \cap B$, we would obtain a compression disc for Σ which is equally impossible.) Similarly, it is impossible that an interval component of $\Sigma' \cap B$ connects one side of B to itself: This can only occur if Σ was a Moebius band and Σ' an annulus. However, it follows in this case that Σ' can be compressed to the $\partial\Sigma \times [0,1]$.

This implies that without loss of generality $\Sigma' \cap B$ consists of one or two intervals connecting opposite sides of the square B. Since they decompose Σ' into 2-balls, claim (ii) now follows. $\square$

6.5.3 Necks

Throughout this section, we assume that $\theta < \theta_1$ and $\mu < \mu_0(\theta)$ are chosen sufficiently small, and that (O, g) is $(v(\theta,\mu), -b^2)$-collapsed with $(v(\theta,\mu), s_0, K)$-curvature control below scale ρ_{-b^2}.

We define a point $x \in O$ to be $(\theta, \mu, -b^2)$-*necklike* relative to an equilateral $< \theta$-straight 1-strainer (a, b) at x with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$ if $\operatorname{diam}(\Sigma^o_{x;a,b}) < \theta^2 \hat{s}_{\mu,-b^2}(x)$, compare 5.6.1. We call the open subset $N_{\theta,\mu,-b^2} \subseteq S_{\theta,\mu,-b^2}$ of $(\theta, \mu, -b^2)$-necklike points the *necklike region* of O. It does not contain any singular vertices, $O^{(0)} \cap N_{\theta,\mu,-b^2} = \emptyset$.

As in the beginning of the previous section, we verify that a point $x \in O$ is $(\theta, \mu, -b^2)$-necklike relative to a 1-strainer (a, b) if and only if it is θ-necklike in the rescaled space $(\hat{s}_{\mu,-b^2}(x))^{-1} \cdot B(x, \rho_{-b^2}(x))$ in the sense of Definition 5.6.1.

Similarly, if $y \in B(x, \rho_{-b^2}(x))$ is $(\theta, \mu, -b^2)$-necklike it is θ-weakly necklike relative to (a, b) in $(\hat{s}_{\mu,-b^2}(x))^{-1} \cdot B(x, \rho_{-b^2}(x))$. If y admits an equilateral $< \theta$-straight 1-strainer (a, b) with length in $(\hat{s}_{\mu,-b^2}(y), \frac{3}{2}\hat{s}_{\mu,-b^2}(y))$ and is θ-strongly necklike relative to (a, b) in $(\hat{s}_{\mu,-b^2}(x))^{-1} \cdot B(x, \rho_{-b^2}(x))$, it is also $(\theta, \mu, -b^2)$-necklike.

Thus we can again use the results from section 5.5.2 to control the existence and relative position of necklike points.

As in section 6.5.2, we construct a smooth line field L_{neck} tangent to the singular locus whose open domain of definition contains the balls $B(x, \theta^{\frac{3}{2}}\hat{s}_{\mu,-b^2}(x))$ around all $(\theta, \mu, -b^2)$-necklike points x, and which is on every such ball almost parallel to the equilateral 1-strainers (a, b) at x with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$, i.e. $\angle.(L_{\text{neck}}, a), \angle.(L_{\text{neck}}, b) < \theta^{\frac{1}{2}}$ on that ball. The line fields L_{neck} and L_{edge} can be matched in the overlap of their domains of definition.

For a $< \theta$-straight equilateral 1-strainer (a, b) with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x), \Sigma_{x;a,b}$ is a closed topological 2-suborbifold almost perpendicular to L_{neck}, i.e. it has angle $> \frac{\pi}{2} - c''\theta$ with it. We call $\Sigma_{x;a,b}$ a *neck cross section* through the point x.

If two neck cross sections $\Sigma_{x_1;a_1,b_1}$ and $\Sigma_{x_2;a_2,b_2}$ intersect, then they have Hausdorff distance $< c\theta^3\hat{s}_{\mu,-b^2}(x)$ by Lemma 5.6.3. If they are disjoint but also not too far apart from each other, say if they have Hausdorff distance $< \theta^{\frac{5}{3}}\hat{s}_{\mu,-b^2}(x)$, then one can move one of the cross sections to the other along the trajectories of L_{neck}, and therefore the $\Sigma^o_{x_i;a_i,b_i}$ are in this case topologically *parallel*, i.e. they bound a product suborbifold $\cong \Sigma^o_{x_i;a_i,b_i} \times [0, 1]$.

Among all neck cross sections, we choose a maximal subfamily ν such that any two distinct cross sections $\Sigma_{x_1;a_1,b_1}$ and $\Sigma_{x_2;a_2,b_2}$ in ν have Hausdorff distance $> \theta^{\frac{5}{2}}\hat{s}_{\mu,-b^2}(x_1)$. In particular, they are disjoint. Due to the compactness of O, ν is finite. Inside ν, we form equivalence classes of topologically parallel cross sections. Each equivalence class has a linear or cyclic order. Any two successive cross sections in it are topologically parallel and bound a cylinder ("segment") homeomorphic to the product of one of them with the compact interval. By concatenating these pieces, the equivalence class yields an embedded *neck* in O which fibers over the interval or over the circle, unless the equivalence class consists only of a single neck cross section. Such an isolated neck cross section has diameter $\approx \theta^2\hat{s}_{\mu,-b^2}(x)$ (with respect to some point x in it); it is contained in the union of the 2-strained and edgy regions, and we simply disregard it. Any two necks are disjoint. The union N of all necks is a compact (topological) 3-suborbifold.

A *cyclic* neck is a closed 3-orbifold and hence fills out O entirely. The topology of cyclic necks will be determined later (in section 6.6.1).

A *linear* neck has two boundary components which are neck cross sections. We call an end of the neck *thick* if its boundary $\Sigma_{x;a,b}$ has diameter $> \theta^{\frac{49}{20}}\hat{s}_{\mu,-b^2}(x)$, and *thin* otherwise. If the end is thin, then nearby $\Sigma_{x;a,b}$, according to our construction e.g. at distance $< \theta^{\frac{49}{20}}\hat{s}_{\mu,-b^2}(x)$, there must be points outside the 1-strained region $S_{\theta,\mu,-b^2}$, i.e. a $(\theta,\mu,-b^2)$-hump, cf. Proposition 6.2.2. The interface between a thin end of a neck and a hump will be discussed in the next section.

Let $\Sigma_{x;a,b}$ be a neck cross section with $\mathrm{diam}(\Sigma_{x;a,b}) > \theta^{\frac{49}{20}}\hat{s}_{\mu,-b^2}(x)$. Then every point in $\Sigma_{x;a,b}$ is $(\theta,\mu,-b^2)$-edgy or belongs to the 2-strained region $R_{\theta,\mu,-b^2}$.

Let $y \in \Sigma_{x;a,b}$ be a point which is not $(\theta,\mu,-b^2)$-edgy. Then $\Sigma_{x;a,b}$ contains a $C_1\theta$-straight 1-strainer (z_1,z_2) at y with length $\theta^4\hat{s}_{\mu,-b^2}(y)$, cf. Lemmas 5.5.3 and 5.5.4(i). As discussed in section 6.4.3, the portion $A_y = \Sigma_{x;a,b}\cap\{|f_{z_1,z_2}| \leq \frac{1}{10}\theta^4\hat{s}_{\mu,-b^2}(y)\}\cap B(y,\theta^4\hat{s}_{\mu,-b^2}(y))$ of the cross section fibers over a compact interval with fibers the f_{z_1,z_2}-level sets. These are embedded 1-suborbifolds $\cong S^1$ or I. Let $\gamma_y = \Sigma_{x;a,b}\cap f^{-1}_{z_1,z_2}(0) = \Sigma^o_{y;a,b,z_1,z_2}$ denote the central fiber.

To combine these local fibrations to a global one, we choose a maximal family F of γ_y's so that any two distinct $\gamma_{y_1},\gamma_{y_2} \in F$ have Hausdorff distance $> \frac{1}{100}\theta^4\hat{s}_{\mu,-b^2}(y_1)$. The family F is finite, since $\Sigma_{x;a,b}$ is compact. If γ_{y_1} and γ_{y_2} have Hausdorff distance $< \frac{9}{100}\theta^4\hat{s}_{\mu,-b^2}(y_1)$, then γ_{y_2} separates A_{y_1} and is isotopic inside A_{y_1} (by a small isotopy) to a fiber of the above fibration of A_{y_1}. We call γ_{y_1} and γ_{y_2} *adjacent*, if they are not separated inside A_{y_1} by another $\gamma_y \in F$. In this case, they have Hausdorff distance $\approx \frac{1}{100}\theta^4\hat{s}_{\mu,-b^2}(y_1)$. It follows that the A_y for all $\gamma_y \in F$ can be simultaneously isotoped (by small isotopies) so that their fibrations match afterwards. This yields a fibration of part of $\Sigma^o_{x;a,b}$ and, if $\Sigma_{x;a,b}$ contains no edgy points, a global fibration.

If $e \in \Sigma_{x;a,b}$ is a $(\theta,\mu,-b^2)$-edgy point, then the discussion in section 6.5.2 implies that $\partial B(e,\frac{1}{2}\theta^3\hat{s}_{\mu,-b^2}(e))\cap\Sigma_{x;a,b}$ contains no edgy points and can be slightly isotoped inside $\Sigma_{x;a,b}$ to match the fibration obtained so far or, vice versa, the fibration can be adapted so that $\partial B(e,\frac{1}{2}\theta^3\hat{s}_{\mu,-b^2}(e))\cap\Sigma_{x;a,b}$ becomes a fiber. Let us call $\overline{B}(e,\frac{1}{2}\theta^3\hat{s}_{\mu,-b^2}(e))\cap\Sigma_{x;a,b}$ a *cap* of $\Sigma_{x;a,b}$. Since $\Sigma_{x;a,b}$ is compact and connected, it must have two disjoint caps which contain all $(\theta,\mu,-b^2)$-edgy points.

In the case when the neck cross section contains no $(\theta,\mu,-b^2)$-edgy points, we have $\Sigma_{x;a,b} \subset R_{\theta,\mu,-b^2}$ and can sandwich it between two nearby neck cross sections and proceed as in section 6.5.2 (for S^2_κ) to isotope it by an isotopy supported nearby (i.e. in the sandwich) so that it becomes vertically saturated.

If the neck cross section $\Sigma_{x;a,b}$ contains edgy points and hence two caps, then we can coordinate the fibration of $\Sigma_{x;a,b}$ with the fibration of the tubes T^j along the coarse edges, cf. section 6.5.2, so that the intersection $T^j\cap\Sigma_{x;a,b}$ consists of two fibers of T^j, namely one for each cap of $\Sigma_{x;a,b}$. (Here, we refer to the fibration of T^j by compact 2-orbifolds with one boundary component.) This can e.g. achieved by perturbing $\Sigma_{x;a,b}$ by a small isotopy (using the flow of the vector field L_{neck}) based near $T^j \cap \Sigma_{x;a,b}$ until it coincides with the closest tube cross sections of the T^j. Moreover, by a small isotopy of the fibration on $\Sigma_{x;a,b}$ minus the two caps, we can arrange that the intersections $S^j \cap \Sigma_{x;a,b}$ are fibers of the fibration of $\Sigma_{x;a,b}$. (In the latter step, we just use that any two noncontractible simple closed curves in an annular 2-orbifold are isotopic.) Alternatively, we could perturb the tubes T^j by a small isotopy near $\Sigma_{x;a,b}$.

6.5.4 Humps

We keep our assumption that (O, g) is $(v(\theta, \mu), -b^2)$-collapsed with $(v(\theta, \mu), s_0, K)$-curvature control below scale ρ_{-b^2} for sufficiently small $\theta, \mu > 0$. Then the discussion of sections 6.5.1, 6.5.2 and 6.5.3 applies.

Let $x \in O$ be a $(\theta, \mu, -b^2)$-hump as defined after Proposition 6.2.1. This means that on the scale $s(x) \in [s_1(\sigma, \mu)\rho_{-b^2}(x), \sigma\rho_{-b^2}(x)]$ of uniform conical approximation provided by Proposition 6.2.1, O is μ-well approximated in x by a flat disc of radius 1 with cone point of angle $\leq 2\pi - \theta$ or by a flat sector of radius 1 with angle $\leq \pi - \frac{\theta}{2}$. (This includes the half-open interval $[0, 1)$ as the degenerate case of the disc with cone angle 0.)

The closed ball $\overline{B}(x_i, \frac{1}{2}s(x_i))$ is a compact 3-suborbifold, since its boundary is almost orthogonal to radial (with respect to x) $\frac{\theta}{11}$-straight 1-strainers of length $\frac{1}{11}s(x)$, cf. Lemma 5.3.2, and hence a closed topological 2-suborbifold.

If $\text{diam}(\partial B(x, \frac{1}{2}s(x)))$ is not too small, e.g. if $\text{diam}(\partial B(x, \frac{1}{2}s(x))) > \theta^{\frac{49}{20}}\hat{s}_{\mu,-b^2}(x)$, then all points in $\partial B(x, \frac{1}{2}s(x))$ are edgy or 2-strained and, as above in section 6.5.3 for neck cross sections, we can construct a 1-dimensional fibration on most or all of $\partial B(x, \frac{1}{2}s(x))$. More precisely, at every point $z \in \partial B(x, \frac{1}{2}s(x))$ which is not edgy we can find an almost radial θ-straight 1-strainer (x, z') of length $\frac{1}{2}s(x)$ and another 1-strainer (a, b) of length $\theta^4\hat{s}_{\mu,-b^2}(z)$ such that the 2-strainer (x, z', a, b) is $c\theta$-straight (for some positive constant c, cf. section 5.4.1). This gives us a fibration of $\partial B(x, \frac{1}{2}s(x)$ on a neighbourhood of z as section 6.4.3. If $\partial B(x, \frac{1}{2}s(x))$ contains edgy points, we integrate the corresponding tube cross sections into its fibration as described in section 6.5.3.

If the conical approximation in x is by a disc with a cone point, then $\partial B(x, \frac{1}{2}s(x)) \subset R^O_{\theta,\mu,-b^2}$ and the fibration is global. If the approximation is by a sector, then $\partial B(x, \frac{1}{2}s(x))$ contains edgy points close to the edges of the sector; these and the complement of the fibered region are covered by two caps whose boundaries are fibers. Since the edge cross sections ($\Sigma^o_{x;a,b}$, cf. section 6.5.2) associated to edgy points in $\partial B(x, \frac{1}{2}s(x))$ are also almost orthogonal to the radial direction (with respect to x), they can be embedded into

$\partial B(x, \frac{1}{2}s(x))$ using an almost radial gradient like flow. As in section 6.5.3, we can coordinate the fibration of $\partial B(x, \frac{1}{2}s(x))$ with the fibration of the tubes T^j along the coarse edge so that the intersection $T^j \cap \partial B(x, \frac{1}{2}s(x))$ consists of two fibers of T^j, one for each cap of $\partial B(x, \frac{1}{2}s(x))$, and the intersections $\partial B(x, \frac{1}{2}s(x))$ are fibers of the fibration of $\partial B(x, \frac{1}{2}s(x))$. We say that the hump x has a *thick end*.

On the other hand, consider the case where $B(x, \frac{3}{4}s(x))$ contains a $(\theta, \mu, -b^2)$-necklike point. This occurs in particular if $\operatorname{diam}(\partial B(x, \frac{1}{2}s(x)))$ is not too large, e.g. if we have $\operatorname{diam}(\partial B(x, \frac{1}{2}s(x))) < \theta_i^{\frac{401}{200}} \hat{s}_{\mu,-b^2}(x)$, because then we have a $\frac{\theta}{11}$-straight 1-strainer of length $> \frac{1}{11}s(x) \geq \frac{1}{11}\hat{s}_{\mu,-b^2}(x)$ at a point in $\partial B(x, \frac{1}{2}s(x))$. In this case $\partial B(x_i, \frac{1}{2}s(x))$ is via an almost radial flow isotopic to a cross section of the associated neck, and we have a *neck-hump* interface. Such an interface corresponds to a *thin* end of a neck (as defined in section 6.5.3).

We now have constructed a covering of the orbifold O by finitely many humps $B(x_i, \frac{1}{2}s(x_i)$, necks, tubes and the fibration of $R_{\theta,\mu,-b^2}$. This follows from Lemma 5.5.9: Consider a point x which is not contained in the balls $B(x_i, \frac{1}{4}s(x_i)$ for the humps x_i; it admits an equilateral $< \theta$-straight 1-strainer with length $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$. If the diameter of its cross section is $\geq \theta^{\frac{9}{4}}$ and $x \notin R_{\theta,\mu,-b^2}$, the lemma implies that x is contained in a tube with no end near x (cf. our discussion in section 5.5.2).

We now adjust the boundaries of humps, necks and tubes to our fibration of the 2-strained part $R_{\theta,\mu,-b^2}$. We have already done this in sections 6.5.2 and 6.5.3 for cyclic chain of edgy points and for thick ends of necks which contain no edgy points. We now proceed analogously for $\partial B(x_i, \frac{1}{2}s(x_i))$ for all $(\theta, \mu, -b^2)$-humps with $\operatorname{diam}(\partial B(x, \frac{1}{2}s(x_i))) > \theta^{\frac{49}{20}}\hat{s}_{\mu,-b^2}(x_i)$ and conical approximation by a disc with a cone point.

At this point, we discard all necks which are entirely contained in the union of humps, tubes and $R_{\theta,\mu,-b^2}$. We deal with all remaining neck-hump interfaces as described above. By our previous discussion, every thick end of a hump or a neck meets $R_{\theta,\mu,-b^2}$. If they contain no edgy points, we have already isotoped them (by an isotopy supported nearby) to a vertically saturated 2-suborbifold.

Every thick end (of a hump or a neck) containing edgy points meets precisely two linear tubes along the coarse edge. By our discussion in section 6.5.2, these tubes can only end deep inside a neck or a hump, and hence a finite time after leaving the end intersect another hump or neck in a thick end. By the compactness of (O, g), such a sequence of tubes and thick ends must eventually close up. The union C of all humps, necks and tubes contained in such a closed chain is a topological 3-suborbifold with one boundary component $\partial_0 C$ which is a topological 2-suborbifold homeomorphic to T^2, K^2, Ann^2 or $\mathrm{Möb}^2$, as follows from our discussion in sections 6.5.2 and 6.5.3. Of course, C may have other boundary components if it contains at least one neck.

For every such chain C, we now extend the suborbifolds S_j^i (for the different tubes $T_j \subset C$) to closed 2-suborbifolds S_C^i isotopic (in C) to $\partial_0 C$, e.g. by forming suitable unions with the first three neck cross sections of a neck in C, and similarly for humps. The 2-suborbifold S_C^2 is then contained in $R_{\theta,\mu,-b^2}$, and we can apply our Waldhausen-like arguments from section 6.5.2 to isotope it (in the region between S_C^1 and S_C^3) so that it becomes vertically saturated with respect to the fibration of $R_{\theta,\mu,-b^2}$.

After performing this isotopy, we cut off the tubes T_j by a suitable cross section such that the fibrations on S_j^2 induced by T_j and $R_{\theta,\mu,-b^2}$ match. The complement of all humps, necks and tubes is now a saturated subset of the fibration of $R_{\theta,\mu,-b^2}$ (see remark after Definition 5.5.8). By performing a suitable isotopy inside the chain C, we can arrange that the cross section cutting off the tubes becomes a smooth 2-suborbifold with boundary intersecting the singular locus and the new boundary of the chain transversally. Similarly, we can isotope neck-hump interfaces to smooth 2-suborbifolds without boundary. Then the components of our decomposition of a 3-orbifold (O, g) have piecewise smooth boundary and their interiors are disjoint open smooth 3-suborbifolds.

Let us sum up our progress so far: For every $0 < \theta < \theta_1$ and $0 < \mu, < \mu_0(\theta)$ the following holds: If a 3-orbifold (O, g) is $(v(\theta, \mu), -b^2)$-collapsed with $(v(\theta, \mu), s_0, K)$-curvature control, it admits a decomposition (according to its coarse stratification) into topological 3-suborbifolds with disjoint interiors and piecewise smooth boundary, namely into $(\theta, \mu, -b^2)$-humps, necks, tubes and total spaces of orbifold fibrations with 1-dimensional fibers.

We also control how the different components of our decomposition intersect each other: If an end (of a hump or a neck) is thin, a hump ends in a neck and vice versa. If the end is thick and meets no tubes, it intersects one of the components with 1-dimensional fibration. In this case the end is toric and vertically saturated with respect to this fibration. Finally, there is the possibility that a thick end meets precisely two tubes and one of the components with 1-dimensional fibration. The boundary component of such an end can then be further decomposed into cross sections of the two tubes and an annular part between them which again is vertically saturated with respect to the 1-dimensional fibration. A tube which does not meet any hump or neck is cyclic; its boundary then is toric and vertically saturated.

6.6 Local topology

In this section, we determine the topological structures of the components of the decomposition we constructed in the previous section. More precisely, we will determine the topology of cross sections to tubes and necks and the topological type of humps.

6.6.1 Tube and neck cross sections

In this section we prove that after decreasing θ further if necessary, we can control the topological type of the cross sections to tubes and necks in sufficiently collapsed 3-orbifolds.

The following proposition is related to an argument in the appendix of [FY92]; see also [MT08, 4.24] for a simplification of the special case needed here.

Proposition 6.6.1 (Topology of edge cross sections). *There exists $\theta_2 > 0$ such that for $\theta \in (0, \theta_2]$ and $\mu \in (0, \mu_0(\theta)]$ the following holds: If (O, g) is $(v(\theta, \mu), -b^2)$-collapsed with $(v(\theta, \mu), s_0, K)$-curvature control below scale ρ_{-b^2}, and if $x \in O$ is $(\theta, \mu, -b^2)$-edgy relative to a $< \theta$-straight 1-strainer (a, b) with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$, then the truncated cross section $\Sigma_{x;a,b} \cap \overline{B}(x, \frac{1}{2}\theta^3\hat{s}_{\mu,-b^2}(x))$ is a connected compact 2-suborbifold with one boundary component and Euler characteristic $\chi \geq 0$.*

Proof. Let $-b_i^2 \in [-1, 0)$ and let θ_i, μ_i be sequences of small positive numbers $\theta_i \to 0$ and $\mu_i \leq \mu_0(\theta)$. We suppose that the orbifolds (O_i, g_i) are $(v(\theta_i, \mu_i), -b_i^2)$-collapsed with $(v(\theta_i, \mu_i), s_0, K)$-curvature control below scale $\rho_{-b_i^2}$, and that the points $x_i \in O_i$ are $(\theta_i, \mu_i, -b_i^2)$-edgy relative to $< \theta_i$-straight 1-strainers (a_i, b_i) with lengths in the interval $(\hat{s}_{\mu_i,-b_i^2}(x_i), \frac{3}{2}\hat{s}_{\mu_i,-b_i^2}(x_i))$.

We consider the neighbourhoods of the points x_i on the scales $\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i)$. The rescaled pointed orbifolds $(\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1} \cdot (O_i, x_i)$ Gromov-Hausdorff subconverge (collapse) to a 2-dimensional Alexandrov space with curvature ≥ 0 which splits off a line. In view of Lemma 5.5.3, this limit is the pointed flat halfplane with base point on the boundary. The cross sections $\Sigma_{x_i;a_i,b_i}$ converge to the cross sectional ray of the halfplane through the base point.

Let $z_i \in \Sigma_{x_i;a_i,b_i}$ with $d(x_i, z_i) = \theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i)$. Then (a_i, b_i, x_i, z_i) is a $c\theta_i$-straight 2-strainer near the intersection $\gamma_i = \Sigma_{x_i;a_i,b_i} \cap \partial B(x_i, \frac{1}{2}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$, for instance in the $\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i)$-neighbourhood of it. Note that $(\theta_i^4\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1}\operatorname{diam}(\gamma_i) \to 0$ because we have $v(\theta_i, \mu_i) \leq (\theta_i^5 s_1(\sigma, \mu_i))^3$.

The following consideration applies for sufficiently large i. From the discussion in section 6.4.3 we know that a neighbourhood (of at least comparable size) of γ_i is fibered by the level sets of the $\mathbb{R}^2$-valued map $(f_{a_i,b_i}, d(x_i, \cdot))$. In particular, $\gamma_i \subset \Sigma_{x_i;a_i,b_i}$ is a connected closed 1-suborbifold. This fibration exists in fact on a larger region, for instance on a neighbourhood of $\Sigma_{x_i;a_i,b_i} \cap \overline{A}(x_i, \frac{1}{100}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i), \frac{99}{100}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$. In particular, $d(x_i, \cdot)$ yields a product fibration (topologically) of $\Sigma_{x_i;a_i,b_i} \cap \overline{A}(x_i, \frac{1}{100}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i), \frac{99}{100}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$ over a compact interval. We note that as a consequence e.g. $\Sigma_{x_i;a_i,b_i} \cap B(x_i, \frac{3}{4}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$ (deformation) retracts onto $\Sigma_{x_i;a_i,b_i} \cap \overline{B}(x_i, \frac{1}{2}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$. Using the flow of a gradient like vector field X_i for the 1-strainer (a_i, b_i) as constructed in section 6.4.2, we obtain a homotopy of $B_i = \overline{B}(x_i, \frac{1}{2}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$ into $\Sigma_{x_i;a_i,b_i} \cap B(x_i, \frac{3}{4}\theta_i^3\hat{s}_{\mu_i,-b_i^2}(x_i))$ relative $\Sigma_i = \Sigma_{x_i;a_i,b_i} \cap B_i$, and together with the retraction $\Sigma_{x_i;a_i,b_i} \cap B(x_i, \frac{3}{4}\theta_i^3 \hat{s}_{\mu_i,-b_i^2}(x_i)) \to \Sigma_i$ a retraction $r_i : B_i \to \Sigma_i$. It is a retraction in the orbifold sense because X_i is tangential to the singular locus. We have $\partial\Sigma_i = \gamma_i$.

If B_i is homeomorphic to the closed 3-ball, then $\pi_1(\Sigma_i) \cong 1$ due to the retraction r_i, and hence Σ_i is a closed 2-disc. More generally, if B_i is discal then r_i lifts to an equivariant retraction $\tilde{r}_i : \tilde{B}_i \to \tilde{\Sigma}_i$ of manifold covers. As before, it follows that $\tilde{\Sigma}_i$ is a 2-disc and thus Σ_i is discal.

Suppose now that (after passing to a subsequence) none of the B_i is discal. We then determine the possible topological types of the B_i using the Shioya-Yamaguchi blow-up argument. Since the O_i also have $(v(\theta_i, \mu_i), s_0, K)$-curvature control below scale $\rho_{-b_i^2}$, according to our discussion in section 6.3.2, B_i is for large i homeomorphic to the product $[0,1] \times \Sigma_i'$ of the compact interval with a connected compact 2-orbifold with Euler characteristic $\chi \geq 0$ and one boundary component. Namely, after a suitable choice of base points, all blow-up limits are 3-dimensional of the form $(Y, y_0) = (\mathbb{R} \times W, (0, w_0))$ with a noncompact $\mathcal{C}^{10}$-smooth 2-orbifold W with $\sec \geq 0$, cf. Proposition 6.3.7 and Lemma 6.1.7. The topology of W is restricted by the orbifold Soul Theorem. If $\mathrm{soul}(W)$ is a point, then W is discal. If $\dim \mathrm{soul}(W) = 1$, then W is a one-ended quotient of the flat cylinder $S^1 \times \mathbb{R}$, cf. Lemma 6.3.6. The relation between the topologies of Σ_i' (for a subsequence yielding the blow-up limit) and W is that W is homeomorphic to the interior of Σ_i'.

Knowing that $B_i \cong [0,1] \times \Sigma_i'$ for large i, we derive the topology of the truncated cross sections Σ_i using the embeddings $\Sigma_i \subset B_i \cong [0,1] \times \Sigma_i'$ and the retractions r_i as before. If Σ_i' is discal, then we saw above that also Σ_i is discal (and $\Sigma_i \cong \Sigma_i'$). Otherwise, Σ_i' is finitely covered by an annulus $\tilde{\Sigma}_i'$ and r_i lifts to an equivariant retraction $\tilde{r}_i : [0,1] \times \tilde{\Sigma}_i' \to \tilde{\Sigma}_i$ of smooth finite covers. Since the composition $\pi_1(\tilde{\Sigma}_i) \to \pi_1(\tilde{\Sigma}_i') \cong \mathbb{Z} \stackrel{(\tilde{r}_i)_*}{\to} \pi_1(\tilde{\Sigma}_i)$ of induced maps of fundamental groups is the identity, it follows that $\pi_1(\tilde{\Sigma}_i) \cong \mathbb{Z}$ or 0 and Σ_i is finitely covered by a 2-disc or an annulus. (We do not worry about excluding the case of the disc here.) □

It follows from the proposition and our discussion in section 6.5.3 that for sufficiently small $\theta, \mu > 0$, *thick* ends of necks in $(v(\theta, \mu), -b^2)$-collapsed orbifolds with $(v(\theta, \mu), s_0, K)$-curvature control have Euler characteristic $\chi \geq 0$. Thus, we alredy control the topological structure of necks with at least one thick end. The following result generalizes this to arbitrary necks, i.e. cyclic necks or linear necks with two thin ends.

Proposition 6.6.2 (Topology of neck cross sections). *There exists $\theta_3 > 0$ such that for $\theta \in (0, \theta_3]$ and $\mu \in (0, \mu_0(\theta)]$ the follwoing holds: If (O, g) is $(v(\theta, \mu), -b^2)$-collapsed with $(v(\theta, \mu), s_0, K)$-curvature control below scale ρ_{-b^2}, and if $x \in O$ is $(\theta, \mu, -b^2)$-necklike relative to a $<\theta$-straight 1-strainer (a, b) with length in $(\hat{s}_{\mu,-b^2}(x), \frac{3}{2}\hat{s}_{\mu,-b^2}(x))$, then the corresponding cross section $\Sigma_{x;a,b}$ is a closed 2-suborbifold with one boundary component* *and Euler characteristic* $\chi \geq 0$.

Proof. Let $-b_i^2 \in [-1, 0)$ and let θ_i, μ_i be sequences of small positive numbers $\theta_i \to 0$ and $\mu_i \leq \mu_0(\theta)$. We suppose that the orbifolds (O_i, g_i) are $(v(\theta_i, \mu_i), -b_i^2)$-collapsed with $(v(\theta_i, \mu_i), s_0, K)$-curvature control below scale $\rho_{-b_i^2}$, and that the points $x_i \in O_i$ are $(\theta_i, \mu_i, -b_i^2)$-necklike relative to $< \theta_i$-straight 1-strainers (a_i, b_i) with lengths in the interval $(\hat{s}_{\mu_i,-b_i^2}(x_i), \frac{3}{2}\hat{s}_{\mu_i,-b_i^2}(x_i))$.

We let $d_i < \theta_i^2$ denote the diameter of the cross sections $(\hat{s}_{\mu_i,-b_i^2}(x_i))^{-1} \cdot \Sigma_{x_i;a_i,b_i}$ and rescale by $2d_i^{-1}$. Then after passing to a subsequence, the orbifolds $(\frac{1}{2}d_i\hat{s}_{\mu_i,-b_i^2})^{-1} \cdot (O_i, x_i)$ converge to an Alexandrov space (Y, y) of curvature ≥ 0 which splits off a line. The factor of Y orthogonal to the line is the limit of the rescaled cross sections to the 1-strainers (a_i, b_i), and hence has diameter 1 and is a compact Alexandrov space of curvature ≥ 0 and dimension 1 or 2.

If $\dim(Y) = 3$, by Lemma 6.1.7 Y is a $\mathcal{C}^{10}$-smooth orbifold and the convergence can be improved to $\mathcal{C}^5$-smooth. It follows that Y is isometric to $\Sigma \times \mathbb{R}$ for some closed 2-orbifold Σ of Euler characteristic $\chi(\Sigma) \geq 0$ and diameter 2. For sufficiently large i, we have an embedding of Σ into O_i which is transversal to the gradient-like vector field X_i for the strainer (a_i, b_i). This implies that Σ is isotopic to $\Sigma_{x_i;a_i,b_i}$.

If $\dim(Y) = 2$, Y must be isotopic to $[-1, 1] \times \mathbb{R}$ or $S^1 \times \mathbb{R}$. Hence there are constants $\phi_i \to 0$ such that every point in $(\frac{1}{2}d_i\theta_i\hat{s}_{\mu_i,-b_i^2})^{-1} \cdot (O_i, x_i) \cap B(x_i, 100)$ either admits a $C_1\phi_i$-straight 2-strainer of length ϕ_i^4 or lies within ϕ_i^3 of a point z admitting a $C_0\phi_i$-straight 1-strainer (a', b') of length 1 such that $\Sigma_{z;a',b'}$ has diamter $\geq \phi_i^{\frac{5}{2}}$ and admits no $\frac{\pi}{2}$-straight 1-strainer of length ϕ_i^4.

If Y is isometric to $[-1, 1] \times \mathbb{R}$, we construct two tubes of diameter $\approx \phi_i^3$ along $(\frac{1}{2}d_i\theta_i\hat{s}_{\mu_i,-b_i^2})^{-1} \cdot (O_i, x_i) \cap B(x_i, 100)$ as in section 6.5.2 corresponding to the two edges of Y. As described in section 6.5.3, we can now decompose $\Sigma_{x_i;a_i,b_i}$ into an annular part admitting a fibration by 1-dimensional orbifolds and two caps isotopic to cross sections of the two tubes. Note that for any $v > 0$ and sufficiently large i, the balls $(\frac{1}{2}d_i\theta_i\hat{s}_{\mu_i,-b_i^2})^{-1} \cdot (O_i, x_i) \cap B(z, 1)$ centered at points $z \in (\frac{1}{2}d_i\theta_i\hat{s}_{\mu_i,-b_i^2})^{-1} \cdot (O_i, x_i) \cap B(x_i, 10)$ uniformly have volume $< v$ with (v, s_0, K)-curvature control on scale 1 because of $1 << \rho_{-b_i^2}$. Thus we can proceed as in the proof of Proposition 6.6.1 to deduce that for sufficiently large i the cross sections to both tubes are compact 2-orbifolds with Euler characteristic $\chi \geq 0$ and one boundary component. This implies $\chi(\Sigma_{x_i;a_i,b_i}) \geq 0$.

If on the other hand Y is isometric to $S^1 \times \mathbb{R}$, it follows as in section 6.5.3 that for sufficiently large i the cross sections $\Sigma_{x_i;a_i,b_i}$ admit a global fibration by embedded 1-dimensional orbifolds and hence are toric (have Euler characteristic $\chi = 0$). □

6.6.2 Humps

For the remaining part of the proof, we fix $\bar{\theta} > 0$ such that the results of section 6.5 and Propositions 6.6.1 and 6.6.2 apply. Thus whenever $\mu < \mu_0(\bar{\theta})$ and a 3-orbifold (O, g) is $(v(\bar{\theta}, \mu), -b^2)$-collapsed with $(v(\bar{\theta}, \mu), s_0, K)$-curvature control, it admits a decomposition according to its coarse stratification and we have control over the cross sections of all tubes and necks.

In order to determine the local topology of humps we will improve the quality of our conical approximations, i.e. make μ sufficiently small. Again, we will adjust the upper bound for μ in several steps.

We say that a $(\bar{\theta}, \mu, -b^2)$-hump $x \in (O, g)$ is a *thick hump* if O can in x be μ-well approximated on scale $s(x)$ by a flat cone with a base of diameter $\in (\frac{\pi}{4}, \pi - \frac{\bar{\theta}}{2})$. In particular, this excludes the case of conical approximation by the 1-dimensional cones $(-1, 1)$ or $[0, -1)$. A thick hump must have a thick end in the sense of section 6.5.4.

Proposition 6.6.3 (Topological type of thick humps). *There exists $0 < \mu_1 < \mu_0(\bar{\theta})$ such that the follwing holds: Let $0 < \mu < \mu_1$. If (O, g) is $(v(\bar{\theta}, \mu), -b^2)$-collapsed with $(v(\bar{\theta}, \mu), s_0, K)$-curvature control and if $x \in O$ is a thick $(\bar{\theta}, \mu, -b^2)$-hump, then $B(x, \frac{1}{2}s(x))$ is discal or solid toric.*

Proof. This is an application of the Shioya-Yamaguchi blow-up discussed in section 6.3.2. Let $b_i^2 \in [-1, 0)$ and $\mu_i \to 0$. Suppose that the orbifolds (O_i, g_i) are $(v(\bar{\theta}, \mu_i), -b_i^2)$-collapsed with $(v(\bar{\theta}, \mu_i), s_0, K)$-curvature control and that the points $x_i \in O_i$ are thick $(\bar{\theta}, \mu_i, -b_i^2)$-humps, i.e. μ_i-well approximated on scale $s(x_i)$ by flat cones C_i with bases of diameter $\in (\frac{\pi}{4}, \pi - \frac{\bar{\theta}}{2})$.

A subsequence of the cones C_i converges to a flat cone C_∞ with a base of diameter $\in [\frac{\pi}{4}, \pi - \frac{\bar{\theta}}{2}]$. It follows that a subsequence of the rescaled balls $s(x_i)^{-1} \cdot B(x_i, s(x_i))$ also converges to C_∞ in the Gromov-Hausdorff sense.

Unless infinitely many of the balls $B(x_i, s(x_i))$ are discal, by our results in section 6.3.2 there is a sequence of rescaling factors $\delta_i \to 0$ such that the sequence $(\delta_i s(x_i))^{-1} \cdot B(x_i, s(x_i)$ Gromov-Hausdorff subconverges to a 3-dimensional limit space (Y, y) of curvature ≥ 0. As discussed in section 6.3.1, (Y, y) is actually a $\mathcal{C}^{10}$-smooth 3-orbifold and the convergence can be improved to $\mathcal{C}^5$-smooth.

The soul of the blow-up limit Y must be a point or 1-dimensional since it cannot be 2-dimensional by 6.3.2. Hence Y must be either discal or solid toric. Again by our discussion in section 6.3.1, this implies that for sufficiently large i, the balls $B(x_i, \frac{1}{2}s(x_i))$ are also either discal or solid toric. □

We can "read off" the topological type of a thick hump from the components of the decomposition it intersects. Let $\mu > 0$ be sufficiently small and suppose that the orbifold (O, g) is sufficiently volume collapsed with curvature control such that it admits a decomposition according to its coarse stratification and that Proposition 6.6.3 holds. Let $x \in O$ be a thick hump of this decomposition.

If O is in x μ-well approximated on scale $s(x)$ by a flat cone over a circle, equivalently if $\partial B(x, \frac{1}{2}s(x))$ is contained in $R_{\bar{\theta},\mu,-b^2}$, it follows that $\partial B(x, \frac{1}{2}s(x))$ admits a fibration by 1-dimensional fibers and hence cannot be spherical. This implies that the hump x is a solid toric 3-suborbifold bounded by a vertically saturated component of a 1-fibered component of the decoposition of O.

If the conical approximation of O in x is by a flat sector, $\partial B(x, \frac{1}{2}s(x))$ intersects precisely two tubes with cross sections Σ_1, Σ_2. Both cross sections have Euler characteristic $\chi \geq 0$. Since $\partial B(x, \frac{1}{2}s(x))$ can be decomposed into a union of Σ_1, Σ_2 and an annular (1-fibered) component, we have $\chi(\partial B(x, \frac{1}{2}s(x)) = \chi(\Sigma_1) + \chi(\Sigma_2)$. Thus if $\chi(\Sigma_1) = \chi(\Sigma_2) = 0$, it follows that $\chi(\partial B(x, \frac{1}{2}s(x)) = 0$ and hence that the hump x is again solid toric. Conversely, if at least one of the Σ_i is 2-discal, the hump x must be 3-discal.

We cannot expect that the arguments from the proof of the last proposition also work for general (not necessarily thick) humps, i.e. humps with conical approximation by cones with a base of arbitrarily small diameter.

In this case, it is possible that the $s(x_i)^{-1} \cdot B(x_i, s(x_i))$ collapse to a 1-dimensional cone and that the rescaled blow-ups $(\delta_i s(x_i))^{-1} \cdot B(x_i, s(x_i))$ only Gromov-Hausdorff converge to a 2-dimensional Alexandrov space (Y, y) of non-negative curvature (see section 6.3.2). We will however see that in this case we can again apply our arguments from section 6.5 to obtain a decomposition of the humps $B(x_i, s(x_i))$ with respect to the scale $\delta_i s(x_i)$ such that no thin humps or necks occur in the decomposition. When investigating collapse to the 2-dimensional space Y, we operate on the scale $\delta_i s(x_i)$ rather than on the natural curvature scale $\rho_{-b_i^2}$. Equivalently, the rescaled orbifolds $(\delta_i s(x_i))^{-1} \cdot (O_i, x_i)$ collapse to Y on scale 1. Note that we have already encountered a similar situation (in a very restricted setting) in the proof of Proposition 6.6.2.

Throughout the following considerations, we will always assume that $\rho_{-b^2} >> 1$. This implies $\sec \geq -b^2 \geq -1$ on balls of radius 1. Moreover, it means that (v, s_0, K)-curvature control on scale ρ_{-b^2} implies (v, s_0, K)-curvature control on scale 1.

We ommit $-b^2$ in our notation to indicate that we work on scale 1 rather than ρ_{-b^2}. For instance, we say that a 3-orbifold is v-collapsed at a point p if $\operatorname{vol} B(p, 1) < v$. Similarly, for $\theta, \mu > 0$ we define the set $R_{\theta,\mu}$ as the set of all points admitting $< C_1\theta$-straight 2-strainers of length $> \theta^4 s_1(\sigma, \mu)$. Similarly, we define (θ, μ)-edgy points, (θ, μ)-necklike points and (θ, μ)-humps.

We can adapt our previous results to our new setting of collapse at scale 1. More precisely, we have

Lemma 6.6.4. *There are $\hat{\theta} > 0$ and $0 < \hat{\mu} < \mu_0(\hat{\theta})$ such that the following holds: Let (O, g) be a 3-orbifold with* $\sec \not\geq 0$ *and $x \in O$. Suppose that for some $R > 0$ the orbifold O is on the ball $B(x, 6R)$ $(v(\hat{\theta}, \hat{\mu}))$-collapsed with $(v(\hat{\theta}, \hat{\mu}), s_0, K)$-curvature control on scale $1 << \rho_{-b^2}$, and that* $\operatorname{diam} O \geq 6R$. *Then*

(i) the orbifold O can in every $y \in B(x, 5R)$ be $\hat{\mu}$-well approximated on some scale $s(y) \in [s_1(\sigma, \hat{\mu}), \sigma]$ by a cone of dimension 1 or 2;

(ii) there exists an open subset U, $R_{\hat{\theta},\hat{\mu}} \cap B(x, 3R) \subseteq U \subset B(x, 4R)$ such that every connected component of U is the total space of a smooth orbifold fibration with fiber S^1 or the mirrored interval I, and all fibers have angle $< \nu$ with the almost vertical line field $L_{\hat{\theta},\hat{\mu}}$;

(iii) if $y \in B(x, 4R)$ is $(\hat{\theta}, \hat{\mu})$-edgy relative to an equilateral $\hat{\theta}$-straight 1-strainer (a, b) with length in $(s_1(\sigma, \hat{\mu}), \frac{3}{2}s_1(\sigma, \hat{\mu}))$, then the truncated cross section $\Sigma_{y;a,b} \cap \overline{B}(y, \frac{1}{2}\hat{\theta}^3 s_1(\sigma, \hat{\mu}))$ is a connected compact 2-suborbifold with one boundary component and Euler characteristic $\chi \geq 0$; and

(iv) if $y \in B(x, 4R)$ is a thick $(\hat{\theta}, \hat{\mu})$-hump, then $B(y, \frac{1}{2}s(y))$ is discal or solid toric.

Proof. The proof works exactly as for Propositions 6.2.1, 6.5.2, 6.6.1 and 6.6.3 since in all of these proofs we rescale by the collapse scale anyway. □

We now return to our original discussion of the topological structure of general humps in a $(v(\bar{\theta}, \mu), -b^2)$-collapsed 3-orbifold with $(v(\bar{\theta}, \mu), s_0, K)$-curvature control.

Proposition 6.6.5. *There exists $0 < \mu_2 < \mu_0(\bar{\theta})$ such that the following holds: Let $0 < \mu < \mu_2$. If (O, g) is $(v(\bar{\theta}, \mu), -b^2)$-collapsed with $(v(\bar{\theta}, \mu), s_0, K)$-curvature control and if $x \in O$ is any $(\bar{\theta}, \mu, -b^2)$-hump, then (at least) one of the three following cases occurs:*

(i) $B(x, \frac{1}{2}s(x))$ is discal or solid toric.

(ii) $B(x, \frac{1}{2}s(x))$ is has the topological type of $(\Sigma \times [-1, 1])/\mathbb{Z}_2$ with Σ a closed 2-orbifold with $\chi(\Sigma) \geq 0$ and $\mathbb{Z}_2$ operating as a reflection on $[-1, 1]$.

(iii) $B(x, \frac{1}{2}s(x))$ admits a decomposition as in section 6.5 into a 1-fibered part, tubes, humps and precisely one neck containing $A(x, \frac{1}{4}s(x), \frac{1}{2}s(x))$. The cross section of this neck is a closed 2-orbifold with $\chi \geq 0$. Finally, all the humps occuring in this decomposition are discal or solid toric.

Proof. Again, let $b_i^2 \in [-1, 0)$ and $\mu_i \to 0$. Suppose that the orbifolds (O_i, g_i) are $(v(\bar{\theta}, \mu_i), -b_i^2)$-collapsed with $(v(\bar{\theta}, \mu_i), s_0, K)$-curvature control and that the points $x_i \in O_i$ are (general) $(\bar{\theta}, \mu_i, -b_i^2)$-humps, i.e. μ_i-well approximated on scale $s(x_i)$ by flat cones C_i with bases of diameter $< \pi - \frac{\bar{\theta}}{2}$.

A subsequence of the cones C_i converges to some flat cone C_∞ with a base of diameter $< \pi - \frac{\bar{\theta}}{2}$, and thus a subsequence of the rescaled balls $s(x_i)^{-1} \cdot B(x_i, s(x_i))$ also converges to C_∞ in the Gromov-Hausdorff sense.

If the cone C_∞ is 2-dimensional, the proof proceeds as for Proposition 6.6.3 to show that for sufficiently large i we are in case (i) of the proposition.

We now suppose that C_∞ is 1-dimensional. It is then isometric to the half-open interval $[0, 1)$ with cone point $\{0\}$. Moreover we suppose that for infinitely many i the ball $B(x_i, s(x_i))$ is not discal.

By our discussion in 6.3.2 we therefore have rescaling factors $\delta_i \to 0$ such that the sequence $(\delta_i s(x_i))^{-1} \cdot B(x_i, s(x_i))$ Gromov-Hausdorff subconverges to a non-compact limit Alexandrov space (Y, y) of curvature ≥ 0 and dimension 2 or 3.

In case $\dim(Y) = 3$ Y is again a $\mathcal{C}^{10}$-smooth orbifold and we can improve the convergence to $\mathcal{C}^5$-smooth. Depending on the dimension of its soul, Y is discal, solid toric or diffeomorphic to $(\Sigma \times [-1, 1])/\mathbb{Z}_2$ with $\chi(\Sigma) \geq 0$ and $\mathbb{Z}_2$ operating on $[-1, 1]$ by a reflection. (We can exclude a product structure since the $B(x_i, \frac{1}{2}s(x_i))$ and hence Y are one-ended.) To finish this case, we note again that $B(x_i, \frac{1}{2}s(x_i))$ is homeomorphic to Y for sufficiently large i by our discussion in section 6.3.1.

We are now left with the case where C_∞ is isometric to $[0, 1)$ and the pointed blow-up limit (Y, y) is 2-dimensional. Remember from section 6.3.1 that we have a concave 1-Lipschitz function β_ξ on Y coming from the unique direction at $\{0\} \in [0, 1)$. By construction, y is a maximum of β with $\beta(y) = 0$.

We observe two important properties of the space Y with respect to its curvature bound ≥ 0:

(i) For every point $z \in Y$, there is a $\frac{\pi}{2}$-straight 1-strainer of length $\frac{1}{2}$ centered at z.

More precisely, let ρ_ξ^z be a ray of maximal β_ξ-decay emanating from z and let $y' \in Y$ be a maximum of β_ξ, i.e. $\beta_\xi(y') = 0$. Recall that by construction, there is a critical point x at distance 1 from y with $\beta_\xi(x) = 0$. Concavity of β_ξ implies $\beta_\xi = 0$ on the whole segment yx of length 1. Hence we can choose y' such that $d(z, y') \geq \frac{1}{2}$. Since β_ξ is 1-Lipschitz, we have $d(\rho_\xi^z(t), y') \geq |\beta_\xi(\rho_\xi^z(t))| \geq t$ which implies that for arbitrarily small $\epsilon > 0$ and sufficiently large $t > t_0(\epsilon)$ we have $\tilde{\angle}_z(y', \rho_\xi^z(t)) \geq \frac{\pi}{2} - \epsilon$. This implies property (i).

(ii) There is a radius R (depending on $\bar{\theta}$ and Y) such that for every point $x \in Y$ with $d(z, y) = r \geq R$ there is a $\bar{\theta}$-straight 1-strainer yzz' of length r.

Otherwise, we could find a sequence of points $z_i \to \infty$ with $d(z_{i+1}, y) \geq 2d(z_i, y)$ such that $\tilde{\angle}_{z_i}(y, z_j) \leq \pi - \frac{\bar{\theta}}{4}$ for $i < j$. After passing to a subsequence, we can assume that moreover $\tilde{\angle}_{z_j}(y, z_i) \leq \frac{\bar{\theta}}{8}$: We only have to make sure that $d(y, z_i)$ is growing sufficiently fast, i.e. $d(y, z_{i+1}) \geq \lambda d(y, z_i)$ for some $\lambda(\bar{\theta})$. But this implies that $\angle_y(x_i, x_j) \geq \tilde{\angle}_y(x_i, x_j) \geq \frac{\bar{\theta}}{4}$

for all $i \neq j$ which is absurd.

Since we have by construction $(\delta_i s(x_i))^{-1} \cdot B(x_i, s(x_i)) \to (Y, y)$, for sufficiently large i the $6R$-balls in $(\delta_i s(x_i))^{-1} \cdot B(x_i, s(x_i))$ are $v(\hat{\theta}, \hat{\mu})$-collapsed with $(v(\hat{\theta}, \hat{\mu}), s_0, K)$-curvature control on scale $1 << \rho_{-b^2}$. Thus Lemma 6.6.4 applies to these balls.

In particular, for these i every point $z \in B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 6R)$ can be $\hat{\mu}$-well be approximated on scale $s(z) \in [s_1(\sigma, \hat{\mu}), \sigma]$ by a flat cone.

If we make i sufficiently large (so that $d_{GH}(B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 6R), B(y, 6R))$ becomes sufficiently small), we also have conical approximation on the ball $B(y, 6R) \in Y$ of slightly lower quality, say $2\hat{\mu}$-good approximation. On the other hand, it follows from properties (i) and (ii) and the fact that Y contains a flat strip of width 1 (see 6.3.2) that the diameters of approximating cones must be $> \frac{\pi}{4}$ on $B(y, 6R)$ and $> \pi - \frac{\hat{\theta}}{2}$ on $A(y, R, 6R)$. Hence for i sufficiently large we can deduce that the same bounds on the diameters of approximating cones hold on the balls $B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 6R)$.

We can now apply Proposition 5.3.3 to obtain a decomposition of $B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 6R)$ into finitely many $(\hat{\theta}, \hat{\mu})$-humps and the set $S^i_{\hat{\theta}, \hat{\mu}}$ of points admitting $\hat{\theta}$-straight 1-strainers with length in $(\frac{1}{11} s_1(\sigma, \hat{\mu}), \frac{3}{22} \frac{1}{11} s_1(\sigma, \hat{\mu}))$. Note that all humps must lie in $B_{(\delta_i s(x_i))^{-1} g_i}(x_i, R)$ and are *thick*.

We now proceed as in the proof of Lemma 6.6.4 and construct a covering of the ball $B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 3R)$ by the total spaces of fibrations with 1-dimensional fibers, tubes and thick humps as in section 6.5. Note that all humps occuring in this covering are discal or solid toric.

By construction, the region $N = A_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R, 6R)$ is homeomorphic to a product of the interval and $\partial B(x_i, s(x_i))$. We add it as a "neck" to our covering of $B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 3R)$ and note that all points on its inner boundary $\partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R)$ are edgy or 2-strained. Moreover, at every point $z \in \partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R)$ we can by property (ii) find 1-strainers of length $2R$ which are almost orthogonal to $\partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R)$. As in section 6.5.4, this allows us to construct a 1-dimensional fibration on most or all of $\partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R)$, which is either global or capped off by the cross sections of two tubes. (For every point $z \in \partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R)$ which is not edgy, in the unrescaled hump $B_{g_i}(x_i, s(x_i))$ we can find a $c\theta$-straight 2-strainer $(x_i, z_i'; a_i, b_i)$ such that the 1-strainers (x_i, z_i') and (a_i, b_i) have lengths $2R\delta_i s(x_i)$ and $\theta^4 s_1(\sigma, \hat{\mu}) \delta_i s(x_i)$; these yield local fibrations of the boundary $\partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R) \cong \partial B_{g_i}(x_i, 2R\delta_i s(x_i))$.) In other words, we treat $\partial B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 2R)$ like a thick end of the "neck" N.

We now adjust the interfaces in $B_{(\delta_i s(x_i))^{-1} g_i}(x_i, 3R)$ as in section 6.5.4 and extend N radially away from x_i to $\partial B(x_i, \frac{1}{2} s(x_i))$ using an gradient-like vector field for $d(x_i, .)$. (Recall that there are no critical points for x_i at distance $> \delta_i s(x_i)$.) Thus for sufficiently large i we have constructed a decomposition as in case (iii) of the proposition. □

6.6.3 Proof of the main result

In this section we complete the proof of Theorem 6.1.10. As explained in section 6.1, Theorem 6.1.10 together with Corollary 2.3.3 implies our main result Theorem 6.1.12, which can be further reduced to Corollary 6.1.14 by the Ricci flow argument 6.1.13.

In addition to $\bar{\theta} > 0$ from section 6.6.2, we fix some $0 < \bar{\mu} \leq \mu_2$ with μ_2 as in Proposition 6.6.5. We set $v = v(\bar{\theta}, \bar{\mu})$. Then our discussion so far implies the following:

Let (O, g) be a closed connected 3-orbifold with sec $\not\geq 0$ and $\mathrm{rad}(O) \geq \frac{1}{2}\rho_{-b^2}$ for some $-b^2 \in [-1, 0)$ and which contains no bad 2-suborbifold. Suppose that (O, g) is $(v, -b^2)$-collapsed with (v, s_0, K)-curvature control on scale ρ_{-b^2}. Then O admits a decomposition according to its coarse stratification into finitely many components of the following kind: total spaces of orbifold fibrations with 1-dimensional fibers, tubes and necks with cross sections of Euler characteristic $\chi \geq 0$, and humps which are solid toric, 3-discal or homeomorphic to $(\Sigma \times [-1, 1])/\mathbb{Z}_2$ as in Proposition 6.6.5.

There are three possibilities for every end of a neck or hump: If the end is thin, a hump ends in a neck and vice versa. If the end is thick and meets no tubes, it intersects one of the components with 1-dimensional fibration. In this case the end is toric and vertically saturated with respect to this fibration. Finally, there is the possibility that a thick end meets precisely two tubes and one of the components with 1-dimensional fibration. The boundary component of such an end can then be further decomposed into cross sections of the two tubes and an annular part between them which again is vertically saturated with respect to the 1-dimensional fibration. The end can be spherical or toric depending on the Euler characteristic of the tube cross sections.

In order to complete the proof of Theorem 6.1.10 it therefore suffices to show that after performing a finite number of surgeries such a decomposition can be simplified to a graph decomposition.

Because O admits no bad 2-suborbifolds, the cross sections of all necks must be spherical or toric. We perform surgery across all necks with spherical cross section. If such a neck bounds a discal component or one of type $(\Sigma \times [-1, 1])/\mathbb{Z}_2$ for some closed spherical 2-orbifold Σ (corresponding to a hump at a thin end of the neck) the resulting summmand is a finite quotient of S^3. Similarly, if O is a cyclic neck with spherical cross section, it is decomposed by one surgery into a finite quotient of the 3-sphere. We also perform surgery along the boundaries of all components of type $(\Sigma \times [-1, 1])/\mathbb{Z}_2$ with spherical Σ which are not adjacent to a neck, i.e. coming from humps with thick ends.

The orbifold O may now be disconnected. We discard all spherical summands. Every remaining summand admits a decomposition as above without necks with spherical cross sections or humps of type $(\Sigma \times [-1,1])/\mathbb{Z}_2$ for spherical Σ.

Let V be a 3-discal component of this decomposition. It meets two coarse edges of O; let T and T' be the corresponding tubes. (We do not exclude the case $T = T'$.) Due to Euler characteristic reasons, at least one of the tube cross sections Σ_T, $\Sigma_{T'}$ must be discal. If both are discal, then the two cross sections must be homeomorphic because otherwise ∂V would be a bad 2-suborbifold of O. In this case, V is homeomorphic to $\Sigma_T \times [0,1]$ and we can replace the union $T \cup V \cup T'$ by a single tube with cross section Σ_T, thereby simplifying the decomposition.

Because of the finiteness of the decomposition of O, after repeating this step a finite number of times we can assume that no 3-discal component of the decomposition of O meets two tubes with discal cross section.

Consider now a tube T with discal cross section. If T is cyclic, it is homeomorphic to a fibration over S^1 with discal fiber and hence to a solid toric 3-orbifold with boundary. If T is linear, it ends in two 3-discal components V_1 and V_2 such that the other tubes ending in the V_i have *annular* cross section. In this case, the union $V_1 \cup T \cup V_2$ is again solid toric, cf. the discussion after Proposition 3.5.2. By considering these solid toric suborbifolds as components of our decomposition, we therefore can assume that all tubes occuring in the decomposition of O have annular cross section.

We recall from sections 6.5.2 and 6.5.4 that for each remaining tube T (which now must have annular cross section) intersecting the total space U of a fibration with 1-dimensional fiber, the two (smooth) fibrations of U and T can be matched on the 2-suborbifold $T \cap U$. Since every annular 2-orbifold inherits an orbifold fibration with 1-dimensional fiber from the fibration of the annulus by circles, we can extend the fibration of U to a Seifert fibration of $T \cup U$.

We have now obtained a decomposition along disjoint embedded toric 2-suborbifolds into components which are total spaces of orbifold Seifert fibrations, solid toric suborbifolds, necks with toric cross section and components of type $(\Sigma \times [-1,1])/\mathbb{Z}_2$ with toric Σ. This decomposition now is graph (cf. section 2.2.2). The proof of Theorem 6.1.10 is now complete. □

7. An extension to the case with boundary

In this section we extend the results of the previous one to a somewhat larger class of volume collapsed 3-orbifolds.

We define a *hyperbolic orbifold cusp* to be a complete 3-orbifold with boundary which is isometric to the quotient of a horoball in hyperbolic 3-space by a cocompact isometric group action. Thus, a hyperbolic orbifold cusp is diffeomorphic to $\Sigma^2 \times [0, \infty)$ for some toric orbifold Σ^2 (by Bieberbach's theroem). With the construction of the Ricci flow with surgery in mind (cf. [Pe03] and [KL10] for orientable manifolds), we will consider hyperbolic orbifold cusps with sectional curvature equal to $-\frac{1}{4}$.

Definition 7.0.1 (Almost cuspidal ends). A Riemannian 3-orbifold (O, g) with boundary has (v, s_0)-almost cuspidal ends if for every component $C \subset \partial O$ there is a hyperbolic orbifold cusp X_C such that the pairs $(N_{100}(C), C)$ and $(N_{100}(\partial X_C), \partial X_C)$ habe distance $\leq v$ in the $\mathcal{C}^{s_0}$-topology.

The following theorem generalizes Theorem 6.1.12 to locally volume collapsed 3-orbifolds with almost cuspidal ends (compare again [Pe03, Theorem 7.4], [MT08, Theorem 0.2] and [KL10, Theorem 1.3]).

Theorem 7.0.2. *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds: If (O, g) is closed or compact with (v_0, s_0)-almost cuspidal ends, is $(v_0, -1)$-collapsed, has (v_0, s_0, K)-curvature control below the scale ρ_{-1} and contains no bad 2-suborbifolds, then O is either closed and admits a $\mathcal{C}^5$ Riemannian metric with* $\sec \geq 0$*, or satisfies Thurston's Geometrization Conjecture.*

Again, we can use Theorem 6.1.13 to simplify the result of the theorem:

Corollary 7.0.3. *Let $s_0 \in \mathbb{N}$ and let $K : (0, \omega_3) \to (0, \infty)$ be a function. If s_0 is sufficiently large, then there exists a constant $v_0 = v_0(s_0, K) \in (0, \omega_3)$ such that the following holds:*

If (O, g) is closed or compact with (v_0, s_0)-almost cuspidal ends, is $(v_0, -1)$-collapsed, has (v_0, s_0, K)-curvature control below the scale ρ_{-1} and contains no bad 2-suborbifolds, then O is satisfies Thurston's Geometrization Conjecture.

Proof. Throughout the following proof, we choose s_0, $\bar{\theta}$, $\bar{\mu}$ and $v = v(\bar{\theta}, \bar{\mu})$ as in the proof of Theorem 6.1.10.

If (O, g) is closed, $(v, -1)$-collapsed, has (v, s_0, K)-curvature control below scale ρ_{-1} and contains no bad 2-suborbifolds, we have already shown that the theorem holds.

We therefore now suppose that (O, g) has at least one ((v, s_0)-cuspidal) end. In this case, we first observe that $\rho_{-1}(x) \approx 4$ near a boundary component C, say on $A(C, 10, 90)$. This means that there are points $x \in O$ with $\operatorname{diam} O >> 2\rho_{-1}(x)$. Hence collapse to a point cannot occur and we can work with $-b^2 = -1$. (In other words, we do not need to make use of the more general setting of Theorem 6.1.10.)

After decreasing v if necessary, we obtain that every point x close to a cuspidal end (again, say on $A(C, 10, 90)$ for a boundary component $C \subset \partial O$) admits a $< \bar{\theta}$-straight 1-strainer of length $\hat{s}_{\bar{\mu},-1}(x)$ almost orthogonal to level sets of $d(C, .)$. In particular, we conclude $A(C, 10, 90) \subset S_{\bar{\theta},\bar{\mu},-1}$.

We fix this new value of v and suppose from now on that (O, g) is compact with (v, s_0)-almost cuspidal ends and $(v, -1)$-collapsed, has (v, s_0, K)-curvature control below scale ρ_{-1} and contains no bad 2-suborbifolds.

We now define the *cusped necks* of O to be the closed sets $\overline{N}_{25}(C)$ for all boundary components $C \subset \partial O$. On the neighbourhoods $N_{90}(C)$ of the cusped necks, we have smooth gradient-like vector fields V_C for the distance function $d(C, .)$. Cusped necks are homeomorphic to $\Sigma^2 \times [0, 1]$ for some toric 2-orbifold Σ^2. Throughout the following discussion, we are only interested in the ends of cusped necks which are not boundary components of O.

By construction, cusped necks are disjoint from each other; they are also disjoint from the humps in $O \setminus \bigcup_C N_{10}(C)$ by 5.3.2 (i).

We will now show how cusped necks can be integrated in our decomposition of O according to its coarse stratification much like humps. As with humps, we call the end of a cusped neck *thin* if there are $(\bar{\theta}, \bar{\mu}, -1)$-necklike points in $\overline{N}_{50}(C)$. This occurs in particular if the diameter of $\{d(C, .) = 25\}$ is not too large, say $\leq \bar{\theta}^{\frac{401}{200}} s_1(\bar{\mu}, -1)$. (Remember that we have seen $\rho_{-1} \approx 4$ on a large neighbourhood of the end, so the above condition means that the diameter is of order $\theta^{\frac{401}{200}} \hat{s}_{\bar{\mu},-1}$ for all points in this neighbourhood.) A thin end of a cusped neck corresponds to the *thin end* of a neck (as defined in 6.5.3) and the interface can be matched up using he flow of V_C.

Similarly, we say that the end of a cusped neck is *thick* if the diameter of $\{d(C, .) = 25\}$

is sufficiently large, say $\geq \bar{\theta}^{\frac{49}{20}}$. In this case, we can proceed as we did for humps in section 6.5.4 and construct a 1-dimensional fibration on all or almost all of $\{d(C,.) = 25\}$, with the possible exception of two tube cross sections. We also can perturb $\{d(C,.) = 25\}$ such that it intersects the tubes (if there are any) in tube cross sections (again, using the flow of V_C). If the end of a cusped neck intersects two tubes, it follows immediately that both tubes have *annular* cross sections.

We now proceed to construct a decomposition of O as we did in the closed case. After adjusting the interfaces of the different components of the decomposition (using our Waldhausen-type arguments) and performing a finite number of surgeries we again obtain components which are spherical or admit a further decomposition along (piecewise smooth) toric suborbifolds into pieces which are orbifold Seifert fibrations, solid toric suborbifolds, necks with toric cross section or components of type $(\Sigma \times [-1,1])/\mathbb{Z}_2$ with toric Σ. (The new components coming from cusped necks of O are topologically of the same kind as necks with toric cross sections.) These decompositions are again graph, which by virtue of Corollary 2.3.3 completes the proof of the theorem. $\square$

Bibliography

[Ba07] R. Bamler, *Ricci flow with surgery*, diploma thesis, München 2007.

[BBI01] D. Burago, Y. Burago, S. Ivanov, *A course in metric geometry*, AMS 2001.

[BBBMP10] L. Bessières, G. Besson, M. Boileau, S. Maillot, J. Porti, *Collapsing irreducible 3-manifolds with nontrivial fundamental group*, Invent. Math. 179, No. 2, 435–460 (2010).

[BGP92] Y. Burago, M. Gromov, G. Perelman, *A. D. Alexandrov spaces with curvature bounded below*, Russ. Math. Surv. 47, No. 2, 1–58 (1992); translation from Usp. Mat. Nauk. 47, No. 2(284), 3–51 (1992).

[BLP05] M. Boileau, B. Leeb, J. Porti, *Geometrization of 3–dimensional orbifolds*, Ann. Math. (2) 162(1), 195–290 (2005).

[BMP03] M. Boileau, S. Maillot, J. Porti, *Three-dimensional orbifolds and their geometric structures*, Panoramas et Synthèses, vol. 15, SMF 2003.

[Bo92] J. Borzellino, *Riemannian Geometry of Orbifolds*, Ph.D. thesis, Los Angeles 1992.

[BS85] F. Bonahon, L.C. Siebenmann, *The classification of Seifert fibered 3-orbifolds*, in *Low dimensional topology* (ed. Roger Fenn), Cambridge 1985.

[BS87] F. Bonahon, L.C. Siebenmann, *The characteristic toric splitting of irreducible compact 3–orbifolds*, Math. Ann. 278, No. 1–4, 441–479 (1987).

[CC07] B. Chow, S.-C. Chu et alii, The Ricci flow: Techniques and applications. Part I: Geometric aspects, AMS 2007.

[CE75] J. Cheeger, D. Ebin, *Comparison theorems in Riemannian geometry*, North-Holland Mathematical Library 1975.

[CG72] J. Cheeger and D. Gromoll, *On the structure of complete manifolds of nonnegative curvature*, Ann. Math. 96, 413–443 (1972).

[CGT82] J. Cheeger, M. Gromov, M. Taylor, *Finite propagation speed, kernel estimates for functions of the Laplace operator, and the geometry of complete Riemannian manifolds*, J. Differ. Geom. 17, 15–53 (1982).

[Di08] J. Dinkelbach, *Equivariant Ricci flow with surgery*, Ph.D. thesis, München 2008.

[DL09] J. Dinkelbach, B. Leeb, *Equivariant Ricci flow with surgery and applications to finite group actions on geometric 3-manifolds*, Geometry & Topology 13, No. 2, 1129–1173 (2009).

[Du88] W.D. Dunbar, *Geometric orbifolds*, Rev. Math. Univ. Complutense Madr. 1, No. 1–3, 67–99 (1988).

[Ep66] D.B.A. Epstein, *Curves on 2-manifolds and isotopies*, Acta Math. 115, 83–107 (1966).

[Fa11] D. Faessler, *On the topology of locally volume collapsed Riemannian 3-orbifolds*, Mathematics arXiv, arXiv:math.DG/1101.3644, Jan. 19, 2011.

[FL] D. Faessler, B. Leeb, *A compactness property for Riemannian orbifolds*, unpublished notes.

[FY92] K. Fukaya, T. Yamaguchi, *The fundamental groups of almost nonnegatively curved manifolds*, Ann. Math. (2) 136, No. 2, 253–333 (1992).

[Fu86] K. Fukaya, *Theory of convergence for Riemannian orbifolds*, Jap. J. Math., New Series 12, 121–160 (1986).

[Gr81] M. Gromov, *Structures métriques pour les variétés riemanniennes*, Paris 1981.

[GK73] K. Grove, H. Karcher, *How to conjugate $\mathcal{C}^1$-close group actions*, Math. Z. 132, 11–20 (1973).

[Ha82] R. Hamilton, *Three-manifolds with positive Ricci curvature*, J. Diff. Geom. 17, 255-306 (1982).

[Ha95] R. Hamilton, *The formation of singularities in the Ricci flow*, Surv. in Diff. Geom. 2, 7–136 (1995).

[Ha03] R. Hamilton, *Three-orbifolds with positive Ricci curvature*, in *Collected papers on Ricci flow* (ed. H.D. Cao, B. Chow et alii), Somerville 2003.

[Ka77] H. Karcher, *Riemannian center of mass and mollifier smoothing*, Commun. Pure Appl. Math. 30, 509–541 (1977).

[Ka89] H. Karcher, *Riemannian comparison constructions*, Global differential geometry, MAA Stud. Math. 27, 170–222 (1989).

[KL08] B. Kleiner, J. Lott, *Notes on Perelman's papers*, Geometry & Topology 12, 2587–2855 (2008).

[KL10] B. Kleiner, J. Lott, *Locally collapsed 3-manifolds*, Mathematics arXiv, arXiv:1005.5106, 2010.

[KL11] B. Kleiner, J. Lott, *Geometrization of three-dimensional orbifolds via Ricci flow*, Mathematics arXiv, arXiv:math.DG/1101.3733, Jan. 19, 2011.

[Kn29] H. Kneser, *Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten*, Jahresbericht D. M. V. 38, 248–260 (1929).

[Lu01] P. Lu, *A compactness property for solutions of the Ricci flow on orbifolds*, Amer. J. Math. 123, 1103–1134 (2001),

[MY80] W.H. Meeks, S.T. Yau, *Topology of three dimensional manifolds and the embedding problems in minimal surface theory*, Ann. Math. (2) 112, 441–484 (1980).

[MT07] J. Morgan, G. Tian, *Ricci Flow and the Poincaré Conjecture*, Mathematics arXiv, arXiv:math.DG/0607607v2, 2007.

[MT08] J. Morgan, G. Tian, *Completion of the proof of the Geometrization Conjecture*, Mathematics arXiv, arXiv:0809.4040, 2008.

[Mu60] J. Munkres, *Obstructions to the smoothing of piecewise-differentiable manifolds*, Ann. of Math. 72, No. 2, 521–554 (1960).

[Ot88] J.-P. Otal, *Thurston's hyperbolization of Haken manifolds*, Surveys in Diff. Geom. III, 77–194 (1998).

[Pe03] G. Perelman, *Ricci flow with surgery on three-manifolds*, Mathematics arXiv, arXiv:math.DG/0303109, 2003.

[Sa56] I. Satake, *On a generalization of the notion of manifold*, Proc. Nat. Acad. Sci. USA 42, 359–363 (1956).

[Sc83] P. Scott, *The geometries of 3-manifolds*, Bull. London Math. Soc. 15, 401–487 (1983).

[SY00] T. Shioya, T. Yamaguchi, *Collapsing three-manifolds under a lower curvature bound*, J. Diff. Geom. 56, 1–66 (2000).

[SY05] T. Shioya, T. Yamaguchi, *Volume collapsed three-manifolds with a lower curvature bound*, Math. Ann. 333, No. 1, 131–155 (2005).

[Th86] W. Thurston, *Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds*, Ann. Math. (2) 124, 203–246 (1986).

[Th78] W. Thurston, *The geometry and topology of 3-manifolds*, Lecture Notes, Princeton 1978.

[Th97] W. Thurston, *Three-Dimensional Geometry and Topology*, Princeton 1997.

[Wa67] F. Waldhausen, *Eine Klasse von 3–dimensionalen Mannigfaltigkeiten: I, II*, Invent. Math. 3, 308–333; 4, 87–117 (1967).

[Wh61] J.H.C. Whitehead, *Manifolds with transverse fields in euclidean space*, Ann. of Math. 73, No. 2, 154–212 (1961).

[Ya91] T. Yamaguchi, *Collapsing and pinching under a lower curvature bound*, Ann. Math. 133, No. 2, 317–357 (1991).

Lebenslauf

Schulbildung

1993 – 1994	Christoph-Probst-Gymnasium in Gilching
1994 – 1997	Deutsche Schule Paris
1997 – 2002	Max-Born-Gymnasium in Germering

Zivildienst

2002 – 2003	Zivildienst in der neurologischen Klinik am Universitätsklinikum Großhadern

Studium

2003 – 2005	Physik an der LMU München
März 2005	Vordiplom Physik
2005 – 2006	Doppelstudium Mathematik und Physik
Sept. 2005	Vordiplom Mathematik
2006 – 2007	Masterstudiengang "M.Sc. in Mathematics and the Foundations of Computer Science" an der University of Oxford
Sept. 2007	Masterarbeit "Descents and Visualization Techniques for the Jacobians of Curves of Genus 2"

Promotion

2008 – 2011	Promotion am Lehrstuhl für Geometrie und Topologie (Prof. Leeb) an der LMU München
2008	Trimester am Institut Henri Poincaré in Paris
2008	Trimester am Hausdorff-Institut für Mathematik in Bonn
April 2011	Dissertation "The Topology of locally volume collapsed 3-Orbifolds"