Chapter 7

Proof of the Orbifold Theorem

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

7.1 Topological preliminaries

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

Theorem 7.1 (The Orbifold Theorem).
Suppose that $\mathcal{O}$ is a compact, orientable, orbifold-irreducible 3-orbifold with (possibly empty) orbifold-incompressible boundary consisting of Euclidean 2-orbifolds. Suppose that $\Sigma(\mathcal{O})$ is a non-empty graph. Then there is an incompressible Euclidean 2-suborbifold $\mathcal{T}$ (possibly empty) such that each component of $\mathcal{O} - \mathcal{T}$ is a geometric orbifold.

To begin the proof of the Orbifold Theorem, we need an orbifold version of the torus decomposition of a 3-manifold. This will provide the incompressible Euclidean 2-suborbifold $\mathcal{T}$ along which the orbifold is decomposed. The statement of the decomposition theorem given in 2.55 has been specialized to the orbifolds that arise in the version of the Orbifold Theorem given above.

Theorem 7.2 (Euclidean Decomposition Theorem).
Suppose that $\mathcal{O}$ is a compact, orientable, orbifold-irreducible 3-orbifold with (possibly empty) boundary consisting of orbifold-incompressible Euclidean 2-orbifolds. Then there is a (possibly empty) closed, orientable, incompressible Euclidean 2-suborbifold $\mathcal{T} \subset \mathcal{O}$ such that if P is the closure of a component of $\mathcal{O} - \mathcal{T}$, then P is either an orbifold Seifert fibre space or it is orbifold-atoroidal. If P has non-empty boundary, that boundary will be orbifold-incompressible.

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

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

Theorem 7.3 (Thurston's Theorem for Manifolds with Boundary).
Suppose that M is a compact, orientable, irreducible, atoroidal 3-manifold with $\partial M \neq \phi$. Either the interior of M admits a complete hyperbolic structure or M is Seifert fibred.

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

Proposition 7.4 (Complement Seifert Fibred).
Suppose that $\mathcal{O}$ is a compact, orientable, orbifold-irreducible 3-orbifold and

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

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

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

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

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

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

Theorem 7.5 (Doubling Trick).
Suppose that $\mathcal{O}$ is a compact, orientable, orbifold-irreducible, orbifold-atoroidal 3-orbifold with orbifold-incompressible Euclidean boundary. Furthermore, suppose that $\mathcal{O}$ is not an orbifold Seifert fibre space or an I-bundle over a Euclidean 2-orbifold. As above, let $D\mathcal{O}_{\{v_i\}}$ denote $\mathcal{O}$, with neighbourhoods of some of its vertices removed, doubled along its non-torus boundary components. Then $D\mathcal{O}_{\{v_i\}}$ has orbifold-incompressible boundary and every orbifold-incompressible pillowcase or turnover in $D\mathcal{O}_{\{v_i\}}$ is orbifold-isotopic to $\partial_{NT}\mathcal{O}_{\{v_i\}}$. If $\{v_i\}$ consists of all the vertices in $\mathcal{O}$, then $D\mathcal{O}_{\{v_i\}} - \Sigma(D\mathcal{O}_{\{v_i\}})$ admits a finite volume, complete hyperbolic structure.

7.2 Deforming hyperbolic structures

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

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

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

If Q_0 was obtained by removing vertices and doubling, one type of degeneration that can occur will lead us to replace one or more of the vertices that were removed. The result will be a hyperbolic cone-manifold structure on a new orbifold Q_1 whose singular locus will be a graph, not just a link. The cone angles on Q_1 will then be increased, with the deformation still parametrized so that at time t the cone angles will be $t\theta_i$, where the θ_i correspond to the orbifold angles. We then need to analyze the possible degenerations of this new family of structures. In order to include such families as well, the theorems below 7.6 (trouble at $t < 1$) and 7.7 (trouble at $t = 1$) are stated for orbifolds that may have vertices and for parametrization that may begin with some value of t bigger than 0.

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

Theorem 7.6 (trouble at $t < 1$). *Suppose that Q is a compact, orientable 3-orbifold with singular locus a graph Σ and (possibly empty) boundary consisting of tori. Suppose that $\epsilon < t_\infty < 1$ and that there is a continuous family of hyperbolic cone-manifold structures on Q for $t \in [\epsilon, t_\infty)$. Then one of the following happens:*
1 (hyperbolic) There is a hyperbolic cone-manifold structure on Q with cone angles corresponding to $t = t_\infty$.
2 (Euclidean) There is a Euclidean cone-manifold structure on Q with cone angles corresponding to $t = t_\infty$.
3 (vertex filling) Q contains an open subset which has the structure of a finite-volume complete hyperbolic 3-dimensional cone-manifold M' with cone angles corresponding to $t = t_\infty$. The 2-dimensional cross-sections of the ends of M', when given the orbifold angles of Q, are orbifold-incompressible in Q.

Case 3 (vertex filling) only occurs for $t_\infty < 1$ when the original orbifold had vertices, neighbourhoods of which were removed; Q was obtained by doubling. Any boundary components of M' that don't come from the boundary of Q itself, arise from the boundary of some of these vertex neighbourhoods. They appear as cusps in the hyperbolic structure on M'. In this case the cone angles can still be increased a small amount in M' and the boundary components from these vertices can be filled in to give a cone-manifold structure on the new orbifold Q' that includes these vertices. (Topologically Q' is obtained by cutting Q open along some of the turnovers created during the vertex removal and doubling process and then adding cones to the resulting boundary turnovers.) The family of hyperbolic structures on this new cone-manifold can be extended by increasing the cone angles, at least a small amount. Once all the vertices have been filled back in, or if there were none to begin with, this case can no longer occur. The underlying topological structure at this stage will be that of the original orbifold, doubled along any non-torus Euclidean boundary components it might have had.

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

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

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

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

Theorem 7.7 (trouble at $t = 1$). *Suppose that Q is a compact, orientable, orbifold-irreducible 3-orbifold with singular locus a graph and (possibly empty) boundary consisting of tori. Suppose that there is a continuous family of hyperbolic cone-manifold structures on Q for $t \in [\epsilon, 1)$, for some $\epsilon < 1$. Then one of the following happens:*
1 (hyperbolic) Q contains a finite-volume complete hyperbolic 3-suborbifold whose ends have orbifold-incompressible cross-sections.
2 (Euclidean) Q is a compact Euclidean 3-orbifold.
3 (graph) Q is a graph orbifold.
4 (bundle) Q is an orbifold bundle with generic fibre a pillowcase or a turnover and base a 1-orbifold.

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

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

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

Using theorem 7.6 (trouble at $t < 1$) and the discussion after it, we can analyze the types of degeneration that can occur as $t \to t_\infty < 1$. In case 1 (hyperbolic) of 7.6 there is no degeneration; the family can be extended. So we can assume this case doesn't occur. If case 2 (Euclidean) occurs, Theorem 7.12 (Euclidean/spherical transition) (which depends on the work of Hamilton and whose proof is outlined in Section 7.4) implies that Q_0 is spherical. If $Q_0 = \mathcal{O}$, then $\mathcal{O}$ is spherical, hence geometric as desired. If Q_0 has been obtained by any vertex removal and/or doubling, it can be shown to have infinite orbifold fundamental group which is not possible for a spherical orbifold. Again, $Q_0 = \mathcal{O}$ is the only possibility in this case.

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

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

The cone angles of the cone-manifold structure on Q_1 can now be increased, forming a new continuous family of hyperbolic cone-manifolds, beginning with parameter $t = \epsilon$, where $0 < \epsilon$. We can then apply the same arguments to this family. If Q_1 has also been obtained from $\mathcal{O}$ by removing some vertices and/or doubling, then, again, the only degeneration possible as $t \to t_\infty < 1$ is case 3 (vertex filling). The only other possibility would be for there to be degeneration as $t \to t_\infty$, where $t_\infty = 1$. But this is not possible until all the vertices have been filled in. To see this, note that, if Q_1 has also been obtained by removing some vertices and doubling, it will contain incompressible spherical turnovers. For each such spherical turnover, there will be a value of t strictly less than 1 for which the angles correspond to a Euclidean cone structure on a turnover. One can show that there would actually be a totally geodesic Euclidean turnover in the hyperbolic cone structure on Q_1; this is impossible.

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

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

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

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

Case 4 (bundle) can't occur as a double since all the incompressible, non-peripheral Euclidean 2-suborbifolds have the property that cutting along them leads to an I-bundle over a Euclidean 2-orbifold, a case that was ruled out since it could be handled directly. It can be shown that the Euclidean 2-suborbifolds created in Q by doubling $\mathcal{O}$ are represented by totally geodesic 2-dimensional sub-cone-manifolds in the approximating hyperbolic cone-manifold structures. If Q is Euclidean (case 2), the doubling suborbifolds will be totally geodesic so $\mathcal{O}$ will also be Euclidean. Similarly, in case 3 (graph), if Q is Seifert fibred, it can be shown that $\mathcal{O}$ is Seifert fibred. In case 3 (graph), if Q is not Seifert fibred, it is a union of Seifert fibred orbifolds glued along incompressible Euclidean 2-suborbifolds. By 7.5 (doubling trick) these must have come from the boundary components of $\mathcal{O}$ since $\mathcal{O}$ was orbifold-atoroidal. Thus the Seifert fibred pieces of the graph orbifold must equal $\mathcal{O}$, possibly doubled along a subset of its boundary components. By the previous argument, $\mathcal{O}$ itself is Seifert fibred.

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

7.3 Controlling degenerations

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

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

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

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

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

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

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

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

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

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

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

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

Theorem 7.8 (Bounded Decay of Injectivity Radius).
Given $\epsilon, \delta, r > 0$ there is $\eta > 0$ such that if M is any complete 3-dimensional cone-manifold of constant curvature $\kappa \in [-1, 0]$ and with all cone angles in $[\delta, \pi]$ and if $x, z \in M$ with $inj(z) > \epsilon$ and $d(z, x) < r$ then $inj(x) > \eta$.

In particular, this theorem implies that no pieces of the singular locus can come together and no sets can collapse within any bounded distance of the basepoint. It follows that the limit will again be a complete 3-dimensional cone-manifold.

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

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

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

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

Proposition 7.10 (ends at $t = 1$). *Suppose that Q is a complete, finite-volume, hyperbolic 3-orbifold. There is a compact (non-convex) core C of Q such that each component, E, of $Q - C$ is isometric to the quotient of a* *torus cusp by a finite group of isometries. Thus the closure of E is orbifold isomorphic to $F \times [0, \infty)$ where F is an orientable, closed, Euclidean 2-orbifold: a turnover, pillowcase or torus.*

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

Proposition 7.10 (ends at $t = 1$) follows easily from the fact that Q is finitely orbifold-covered by a hyperbolic manifold. Proposition 7.11 (ends at $t < 1$) requires knowledge of the possible non-compact 3-dimensional Euclidean cone-manifolds. A discussion of this topic appears in the last section of this chapter.

If such a cusp develops as t approaches t_∞, then the Gromov-Hausdorff topology implies that there are almost isometric maps of large compact pieces of the geometric limit into the cone-manifolds. If $t_\infty < 1$ and the geometric limit has an end with a non-torus cross-section, then 7.11 (ends at $t < 1$) implies that there are turnovers in the cone-manifolds whose angle sums approach 2π as $t \to t_\infty$. Since $t_\infty < 1$ the turnovers must be spherical in the orbifold; hence they must be the result of removing vertices in the original orbifold and doubling.

If $t_\infty = 1$ (or if there are only torus cross-sections when $t_\infty < 1$), all the cross-sections of ends in the geometric limit will be Euclidean orbifolds. We claim that, for t sufficiently close to t_∞, the images of these cross-sections will be **incompressible** in the orbifold Q. Assuming this claim, we can finish case 2 (inj bd at base pt).

Using the Gromov-Hausdorff topology, we obtain embeddings of the compact core C, as described in 7.10 (ends at $t = 1$) and 7.11 (ends at $t < 1$), of the limit cone-manifold or orbifold. If $t_\infty < 1$, the image of the boundary of C consists of tori that are incompressible in Q, and turnovers that, with their angles replaced by the orbifold angles, are spherical in Q. The tori, since they are incompressible, must be boundary parallel and, by 7.5 (doubling trick), the turnovers must be orbifold-isotopic to those created by removing vertices. It follows that Q contains a finite-volume complete hyperbolic 3-dimensional cone-manifold M', (homeomorphic to the interior of C), with angles corresponding to $t = t_\infty$. This is case 3 (vertex filling) of 7.6 (trouble at $t < 1$). Note that, as discussed in the previous section, if the original orbifold $\mathcal{O}$ had no vertices or if they have all been filled in,

Q will not contain any spherical turnovers. Thus, all the boundary of C is boundary parallel in Q and M' will be homeomorphic to Q. Contrary to the choice of t_∞, Q has a hyperbolic cone-manifold structure with cone angles corresponding to $t = t_\infty$. No degeneration has occurred and the family can be extended.

Similarly, if $t_\infty = 1$, the image of the compact core, C, will have boundary consisting of incompressible tori (which must be boundary parallel) and incompressible Euclidean turnovers and pillowcases (which are the result of doubling). It follows that Q contains an orbifold-incompressible finite-volume complete 3-dimensional hyperbolic suborbifold. This is case 1 (hyperbolic) in 7.7 (trouble at $t = 1$).

It remains to be seen why the tori, pillowcases, and Euclidean turnovers are *incompressible* in Q. The turnovers are trivially incompressible because every simple closed curve in them bounds a (singular) disk in the turnover itself. We will first consider the case when there are no vertices and when there are only tori in the boundary of C. We then sketch the changes necessary when there are vertices or pillowcases.

The complement of the singular locus in Q is irreducible so, if an embedded torus is compressible in the complement of the singular locus, it is either contained in a ball or bounds a solid torus. Since $Q - \Sigma(Q)$ is also atoroidal, an incompressible torus must be boundary parallel.

The holonomy of elements of the fundamental group of the boundary of the core C are all parabolic in the geometric limit, so, by convergence of holonomy, any given element must become arbitrarily close to parabolic in the approximates. This is not possible if the torus is contained in a ball, in which case the holonomies are all trivial. Thus, if the torus is compressible in Q, it must bound a (singular) solid torus. The meridian curve will be represented by either the trivial element, if the solid torus is non-singular, or an elliptic element with rotation bounded away from 0, if the solid torus is singular. Therefore, by convergence of holonomy, in a sequence of approximates any given curve on a torus boundary of C can be a meridian at most a finite number of times.

The subtlety here is that, a priori, in a sequence of approximations, C could be mapped into Q in topologically distinct ways so that an infinite sequence of distinct curves bound (singular) disks in Q. Let Q_n denote the image of C under the nth approximating map union the (singular) solid tori bounded by the compressible tori. Then Q_n is obtained from C by Dehn filling along the nth meridians, μ_n^i. Viewed as a suborbifold of Q, the boundary of Q_n is incompressible. It is not difficult to show that there are only a finite number of 3-dimensional sub-orbifolds of Q with incompressible boundary, up to isotopy. After taking a subsequence, we can assume that the Q_n are all diffeomorphic to the same orbifold, Q_∞.

If $t_\infty < 1$ the cone angles will be less than the orbifold angles in Q and in Q_∞. To avoid using Mostow-Kojima rigidity for hyperbolic cone-manifolds 5.11 (see [53]) which depends on arguments similar to those in the proof of the Orbifold Theorem (including those in the next few paragraphs) and on local rigidity of cone-manifolds 5.10 (see [45]), we remove a neighbourhood of the singular locus.

Let N_∞ denote Q_∞ with a regular neighbourhood of its singular locus removed and let $\hat{C}$ denote C with a regular neighbourhood of its singular locus removed. It is not hard to see that the complement of the singular locus of a finite volume hyperbolic cone-manifold can be given a complete metric of strictly negative curvature. (See, e.g., [53]; a further argument is required in the case with vertices.) Hence, $\hat{C}$ is irreducible and atoroidal and by 7.3 (Thurston's Haken theorem) it has a finite volume hyperbolic metric.

Suppose that N_∞ is homeomorphic to $\hat{C}$; i.e., that the only solid tori added are singular. An infinite number of distinct curves on the boundary of $\hat{C}$ are mapped to the curves on the boundary of N_∞ that bound singular disks in Q. This implies that $\hat{C}$ has an infinite number of self-homeomorphisms which are homotopically distinct. But, since $\hat{C}$ can be given a complete finite volume hyperbolic metric, Mostow rigidity implies that the group of self-homotopy equivalences which are homeomorphisms on the boundary is finite, a contradiction.

Thus N_∞ is obtained from $\hat{C}$ by an infinite sequence of Dehn fillings where, on each filled-in torus, the same curve is a meridian at most a finite number of times. But, by Thurston's hyperbolic Dehn surgery theorem, all but a finite number of these Dehn fillings result in finite volume hyperbolic manifolds. Furthermore, the hyperbolic structures on such a sequence of fillings has arbitrarily short closed geodesics. By Mostow rigidity there is a unique hyperbolic structure on N_∞; by discreteness and finite volume, there is a shortest geodesic. This gives a contradiction. Thus, all the images of the boundary tori of $\hat{C}$, hence of C, are incompressible in Q as claimed.

When C has pillowcase boundary components, the argument is very similar. If the image of such a pillowcase in Q is compressible, it either is contained in a singular ball with a single unknotted arc of singular locus or it bounds a folded ball. In the first case the holonomy of the image pillowcase would be a single elliptic element. As before, this is impossible by convergence of holonomy for approximating maps sufficiently far out in the sequence.

If the image of a pillowcase bounds a folded ball, a simple closed curve, called the *meridian*, in the pillowcase which does not bound a (singular) disk in the pillowcase does bound a non-singular disk in the folded ball. Convergence of holonomy again implies that a single curve can be a meridian at most a finite number of times since, in C, the holonomy of every element in the orbifold fundamental group of the pillowcase is parabolic and non-trivial. If there are pillowcases in C, then $t_\infty = 1$ and C is an orbifold, not just a cone-manifold. It has a complete, finite volume hyperbolic structure on its interior. We again conclude that the same orbifold, Q_∞, is obtained by orbifold Dehn fillings on infinitely many distinct meridians on each component. The theory of hyperbolic Dehn filling, as extended to orbifolds by Dunbar-Meyerhoff ([64]), leads, as before, to a contradiction, using Mostow rigidity applied to Q_∞.

When there are vertices, the argument is essentially the same. However, when the limit, C, is not an orbifold but only a cone manifold, a further argument beyond that contained in [53] is required in order to show that there is a hyperbolic structure on the complement of the singular locus where the holonomies around the edges connecting the vertices are all parabolic. This fact is used to show it is not possible to obtain the same manifold by Dehn filling on infinitely many distinct curves on each torus boundary component of $\hat{C}$. The argument then proceeds as before. This completes the outline of the proof of Case (2 inj bd at base pt).

Case 3 (inj $\rightarrow$ 0 everywhere). The injectivity radius goes to 0 for all $x \in M_t$ as $t \rightarrow t_\infty$.

In this case, the diameter of M_t may actually go to 0. If so, we rescale so that the diameter is 1. If not, we don't rescale. There are 2 subcases here, depending on whether or not the injectivity radius in the (possibly) rescaled metric goes to 0 at all points.

Case 3a (rescaled inj bd). The injectivity radius does not go to 0 all $x \in M_t$ when the diameter is scaled to equal $\max(1, \operatorname{diam} M_t)$.

We are assuming that we are not in Case (2 inj bd at base pt), so the injectivity radius goes to 0 everywhere in the unscaled metric. The metric must have been scaled to have diameter 1 in this case. By 7.8 (decay of inj) , since the diameter is bounded above, the injectivity radius must be uniformly bounded below at all points in the rescaled metric. By 7.9 (3d limit) , the limit as $t \to t_\infty$ will be a compact Euclidean cone-manifold.

If $t_\infty = 1$, we are at the orbifold angles and Q has a Euclidean structure. This is case 2 (Euclidean) of 7.7 (trouble at $t = 1$).

If $t_\infty < 1$, we will argue, using the work of Hamilton, that Q has a spherical structure. This is case 2 (Euclidean) of 7.6 (trouble at $t < 1$).

We record this step as the following theorem. The argument will be outlined in the next section.

Theorem 7.12 (Euclidean/spherical transition). *Let Q be a compact orbifold with a Euclidean cone structure with some cone angles strictly less than the orbifold angles. Then Q has either a spherical structure or a $S^2 \times \mathbb{R}$ structure. If the Euclidean cone structure arises as a rescaled limit (in the Gromov-Hausdorff topology) of hyperbolic cone structures on Q, then Q has a spherical structure.*

Case 3b (collapsing). The injectivity radius goes to 0 for all $x \in M_t$ when the diameter is scaled to equal $\max(1, \operatorname{diam} M_t)$.

This is the most complicated case in the analysis, which we refer to as the "collapsing case". In the manifold context there has been considerable analysis (see [17], [18], [66], [29]) of the topology of manifolds that admit a sequence of metrics with curvature bounds where the injectivity radius goes to 0 at every point. Such manifolds are shown to possess a generalized Seifert fibred structure called an "F-structure". A 3-dimensional manifold with an F-structure is a graph manifold.

The theorems below may be viewed as a generalization to cone-manifolds of these theorems. However, it is not apparent at this time that the techniques in the manifold context generalize directly.

We say that a 3-dimensional orbifold Q has an *ϵ-collapse* if there is a 3-dimensional hyperbolic cone-manifold M with $\operatorname{inj}(x) < \epsilon \cdot \min(1, \operatorname{diam}(M))$ for all $x \in M$, and a homeomorphism $f : Q - \partial Q \longrightarrow M$ such that $f(\Sigma(Q - \partial Q)) = \Sigma(M)$. In addition, for every edge e of $\Sigma(Q)$, the difference between the cone angle on e in Q and the cone angle on $f(e)$ in M is less than ϵ. It is often convenient to use orbifold terminology when referring to M so we will pass back and forth between M and Q.

Theorem 7.13 (Collapsing Theorem for Cone-manifolds). *Suppose that Q is a compact, orientable, orbifold-irreducible 3-orbifold with non-empty 1-dimensional singular locus and with orbifold-incompressible Euclidean boundary. Then there is $\epsilon > 0$ such that if Q has an ϵ-collapse*

then either
(1) Q is a graph orbifold or
(2) Q is an orbifold bundle with generic fibre a turnover or pillowcase and base a 1*-orbifold.*
Furthermore there is an edge of $\Sigma(Q)$ *labelled* 2.

We discuss the theorems used in the proof of this theorem in the final section of this chapter.

Cases (1) and (2) of the collapsing theorem correspond to cases 3(graph) and 4 (bundle) in 7.7 (trouble at $t = 1$) . Assuming the theorems stated in this section, this concludes the outline of the proofs of 7.6 (trouble at $t < 1$) and 7.7 (trouble at $t = 1$) and, hence, of the Orbifold Theorem.

We summarize the logic in this section in the following flow diagram:

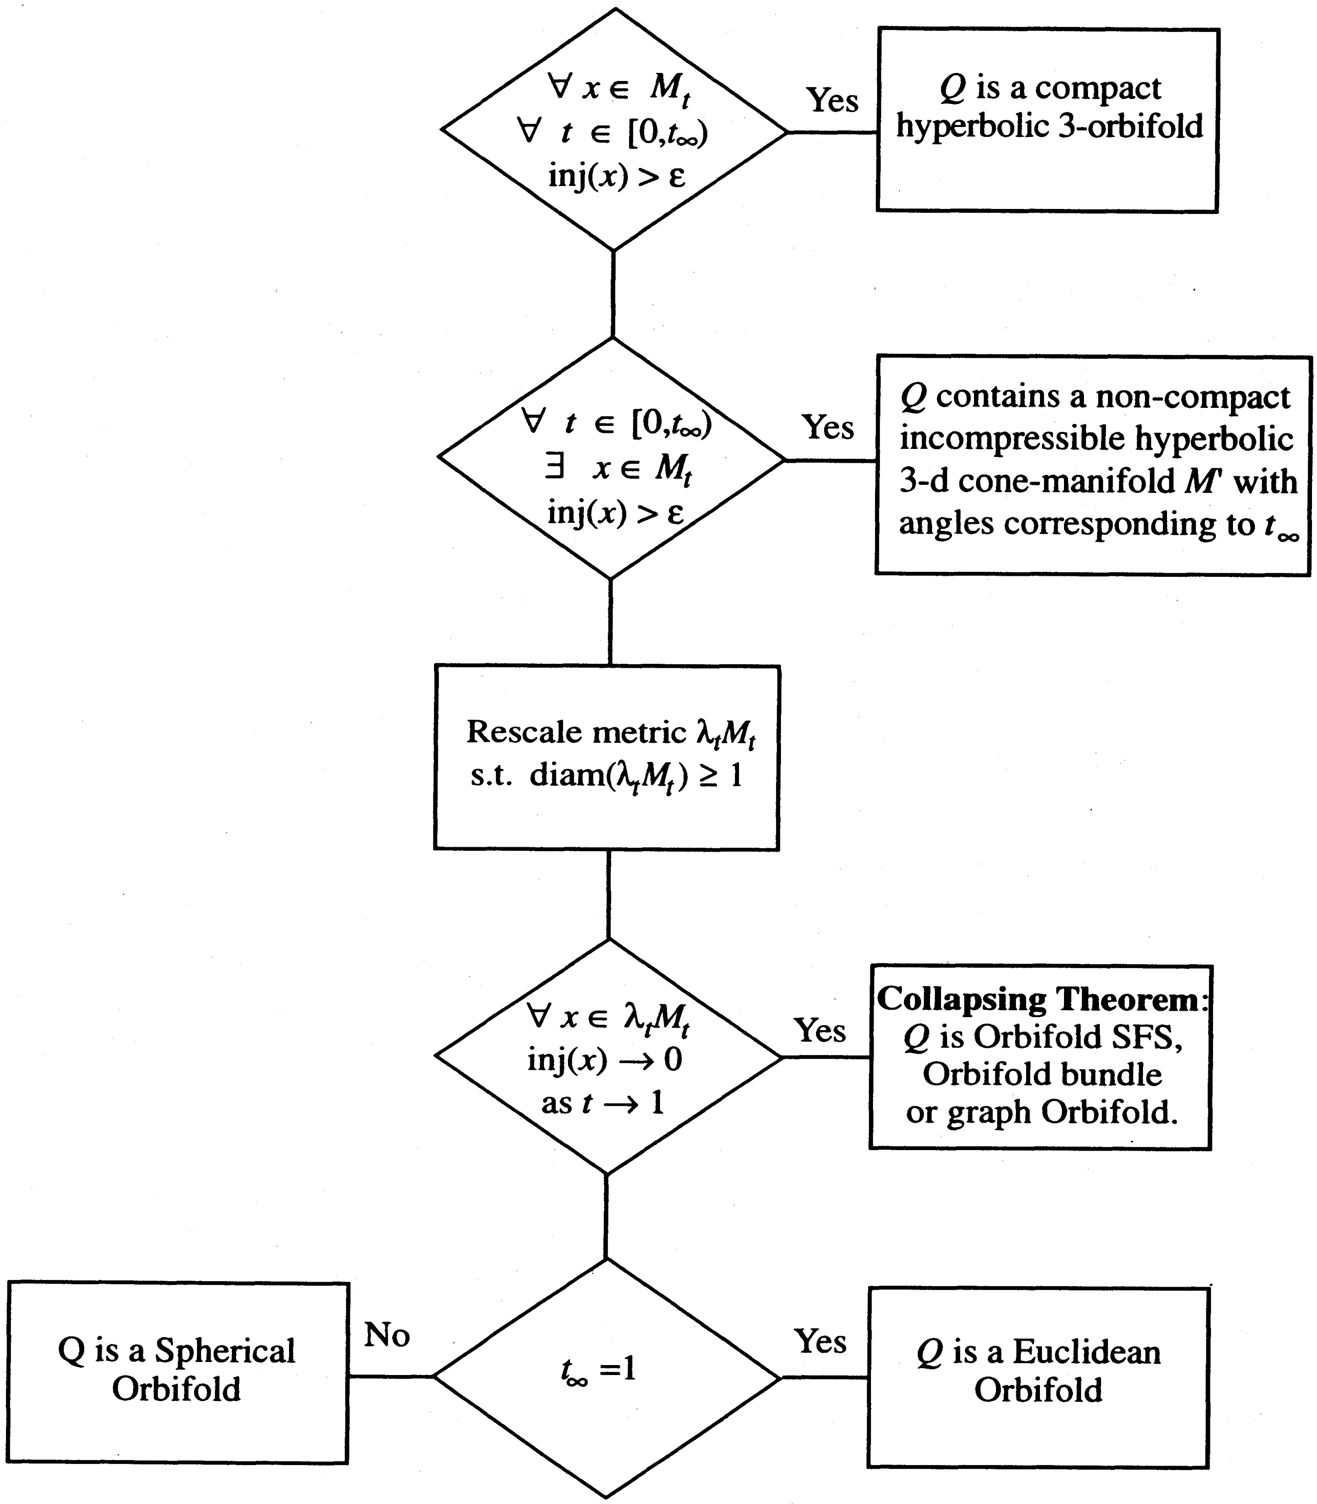

$\forall x \in M_t$
$\forall t \in [0,t_\infty)$
$\mathrm{inj}(x) > \varepsilon$
Yes
Q is a compact hyperbolic 3-orbifold
$\forall t \in [0,t_\infty)$
$\exists x \in M_t$
$\mathrm{inj}(x) > \varepsilon$
Yes
Q contains a non-compact incompressible hyperbolic 3-d cone-manifold M' with angles corresponding to t_∞
Rescale metric $\lambda_t M_t$
s.t. $\mathrm{diam}(\lambda_t M_t) \geq 1$
$\forall x \in \lambda_t M_t$
$\mathrm{inj}(x) \to 0$
as $t \to 1$
Yes
Collapsing Theorem:
Q is Orbifold SFS, Orbifold bundle or graph Orbifold.
Q is a Spherical Orbifold
No
$t_\infty = 1$
Yes
Q is a Euclidean Orbifold

7.4 Euclidean to spherical transition

In this section we outline the proof of 7.12 (Euclidean/spherical transition).

The basic idea behind this theorem is that, if there is a Euclidean cone-manifold structure on Q with angles strictly less than the orbifold angles, then one should be able to spread out some of the concentrated curvature away from the singular locus to obtain a metric with the orbifold angles and some positive curvature on the smooth part of the orbifold. The Ricci flow on Q with this metric should either lead to a spherical orbifold metric or imply that there is an $S^2 \times \mathbb{R}$ structure on Q.

When the singular locus is a link, this process of "smoothing" the metric to the orbifold angles can be done simply and explicitly. When there are vertices, it is less clear how to reach the orbifold angles while maintaining positive curvature on the smooth part of the orbifold so a more ad hoc argument is used. Furthermore, Hamilton's results are proved only for manifolds so, in all cases, a device for finding an orbifold cover which is a manifold is required. (Hamilton has an unpublished manuscript [36] which generalizes his results to orbifolds, but we won't use that here.)

The following theorem of Hamilton is in [37].

Theorem 7.14. *A compact 3-manifold, M, with non-negative Ricci curvature which is not everywhere flat is diffeomorphic to a quotient of either S^3 or $S^2 \times \mathbb{R}$ by a group of fixed point free isometries in the standard metrics. Furthermore, if the original metric with non-negative Ricci curvature has a non-trivial group of symmetries, the homogeneous metric on M will possess the same group of symmetries.*

We suppose that Q has a Euclidean cone structure with cone angles strictly less than the orbifold angles. Assume that the singular locus is a link; i.e., assume that there are no vertices. Consider disjoint singular solid tori, each containing a single component of the singular locus. Assume that each torus consists of points a constant distance from the component it contains. The cross-sections perpendicular to the singular loci are Euclidean disks with a single cone point. The metrics on these can be smoothed in a rotationally symmetric way to obtain a smooth metric with non-negative Ricci curvature on the underlying manifold X of Q.

This smoothed metric on X is not flat so 7.14 (Hamilton's theorem) implies that X is finitely covered by S^3 or $S^2 \times S^1$. Assume that we are in the S^3 case. Viewing Q as X with a link determining the singular locus, we can lift to the topological universal cover of X which is homeomorphic to S^3. This defines an orbifold cover $\tilde{Q}$ of Q whose underlying space is S^3

and whose singular locus is a link. The homology of the complement of the link is a direct product of infinite cyclic groups, each of which is generated by a meridian around a component of the link. We map onto a product of finite cyclic groups, sending the meridian to a generator and killing the pth power if the singular component has local group $\mathbb{Z}_p$. The kernel of this homomorphism defines an orbifold covering of $\tilde{Q}$ which is a manifold. We denote this manifold by M. It also is an orbifold cover of Q.

Return to the Euclidean cone-manifold structure on Q. This time we "smooth" the metric so that the cone angles along the singular locus equal the orbifold angles. This is done in the same way as before, in tubular neighbourhoods of each singular component, in a radially symmetric fashion on each transverse disk. The new metric on each disk still has a singular point at the centre of the disk but it has the orbifold angle; the smooth portion of the disk has some positive curvature near the singular point.

This metric lifts to a smooth metric on the manifold cover M with non-negative Ricci curvature. Hamilton's theorem implies that the Ricci flow converges to a spherical metric possessing any symmetries that the original metric on M had. Thus the spherical metric descends to Q and Q has a spherical structure as desired.

If the underlying manifold X is finitely covered by $S^2 \times S^1$, then Hamilton's proof shows that the algebraic splitting of the curvature operator that appears in an $S^2 \times \mathbb{R}$ will exist in all the metrics that occur in the Ricci flow for all *positive* times. Since the smoothed metric on X has concentrated positive curvature orthogonal to the original singular locus, this must be compatible with the splitting. From this, it follows that Q is, up to a 2-fold cover, homeomorphic to a bundle over S^1 with fibre S^2 with finitely many singular points. Hamilton's proof also provides a metric with positive curvature on the fibres for all positive times, so using Gauss-Bonnet and the fact that all the cone angles are less than π, there will be at most 3 singular points in the fibres.

From this description, it is apparent that Q is, in fact, Seifert fibred with the singular locus contained in the fibres. This is easily seen to have a geometric structure. However, it is also not hard to see that such an orbifold cannot have a hyperbolic metric in the complement of its singular locus, so this case does not actually arise in our context.

When there are vertices, it is less clear how to do the smoothing parts of the argument so we resort to a more ad hoc argument which we hope to simplify in the near future.

We will call a vertex whose link is $S^2(2,2,n)$ a *dihedral vertex*; an edge, labelled n, joining two distinct dihedral vertices is called a *dihedral edge* of order n. Such an edge has a neighbourhood whose 2-fold branched cover over the edges labelled 2 is a tubular neighbourhood of a closed curve labelled n. The process of smoothing such a neighbourhood (both to the orbifold angle and to angle 2π) described above was radially symmetric on each disk cross-section so it is symmetric with respect the order 2 symmetry on the cover.

If Q has a Euclidean cone structure with some of its angles less than the orbifold angles, it can be seen to have finite orbifold fundamental group. Otherwise, the orbifold universal cover is non-compact and would contain an bi-infinite ray. This can be ruled out, using triangle comparison theory. The same argument shows that any orbifold obtained from Q by decreasing the labels on some edges (i.e. increasing the desired orbifold cone angles) will also have finite orbifold fundamental group. Since the only orbifolds with a geometric structure that can have finite orbifold fundamental group are spherical, it suffices to show that such an orbifold is geometric in order to conclude that it is spherical.

To guarantee the existence of a dihedral edge in our orbifold, we change all of the labels to 2's. Denote this orbifold by Q_2. We claim that it is geometric, hence spherical.

Assuming that Q_2 is spherical, its topological type, including the singular graph, is an OSFS and belongs to one of a few known families (see [26], [28]). The original orbifold Q has the same topology with some of the labels increased. By looking at each family, it is then possible to show that changing the labels leads either to an orbifold with infinite orbifold fundamental group, which is impossible for Q, or to a spherical orbifold.

In order to obtain a spherical structure on Q_2, we note that there must be a dihedral edge (of order 2) if there are any vertices. We first attempt to find a hyperbolic structure on Q_2 where the holonomy around the dihedral edge remains parabolic. This parabolicity requirement has the effect of removing a neighbourhood of the edge, creating a sphere with 4 cone points on the boundary. The cone angles begin at 0; we attempt to increase them to π. Throughout the deformation the meridian curve will be parabolic. If we reach cone angles π, the boundary will become a pillowcase cusp.

The only way this sequence of hyperbolic structures can degenerate is for it to collapse at time $t = 1$, in which case Q_2 with the neighbourhood removed is a graph manifold. By the arguments in the last section (using the latitude hypothesis) on collapsing, the fibration can be extended over the folded ball that the pillowcase bounds in Q_2. Thus Q_2 is a graph manifold; since it has finite orbifold fundamental group it is a spherical OSFS.

If we reach the final angles, we can begin to increase the cone angle around the dihedral edge. (At the angles of the form $\frac{2\pi}{k}$, this amounts to doing hyperbolic Dehn filling along a meridian of the pillowcase.) Since Q_2 can't be hyperbolic, this must degenerate at some stage. Either we again conclude that Q_2 is a graph manifold, hence a spherical OSFS or we obtain a Euclidean structure on Q_2 with cone angles π along all edges except the dihedral edge.

One can then find an orbifold cover which unfolds the angle π edges, leaving one with a link singularity in the cover. This can be smoothed symmetrically with respect to the covering maps as described above. The argument now proceeds as before in the link singularity case. The spherical structure obtained from the Ricci flow will be symmetric and descend to one on Q_2.

7.5 Analysis of the thin part

We have seen that as long as the injectivity radius is bounded below, controlling the degeneration of the family of hyperbolic cone-manifolds is not difficult. However, when the injectivity radius goes to 0, controlling the topology of the geometric limits becomes more difficult. In the outline, there were two key theorems concerning the topology of the regions where the injectivity radius is small. The first 7.11 (ends at $t < 1$) described the ends of a finite volume hyperbolic cone-manifold. The cross-sections of these ends provided us with 2-dimensional Euclidean sub-cone-manifolds that put strong topological limitations on the orbifolds Q that could degenerate when the injectivity radius was bounded below at the basepoint but went to 0 elsewhere. The second theorem 7.13 (Collapsing theorem) provided a topological classification for those orbifolds Q that could admit a family of metrics where the injectivity radius went to 0 everywhere.

One way to understand the topological and metric structure near a point where the injectivity radius is going to 0 is to rescale the metric so the injectivity radius is 1, using that point as the basepoint. Since the injectivity radius at the basepoint goes to zero, the sequence of scale factors will go to infinity. Applying 7.8 (decay of inj) , one sees that the limit of the scaled structures will be a complete Euclidean cone-manifold. If the diameter of the original sequence doesn't go to zero (or goes to zero at a slower rate than the injectivity radius), the diameter of the geometric limit will be infinite and the cone-manifold will be non-compact.

One of the most important tools in the proof of the Orbifold Theorem is the classification of non-compact 3-dimensional cone-manifolds whose cone angles are at most π. The restriction on the cone angles is crucial to this theorem. If the cone angles are allowed to lie between 0 and 2π, the number of possibilities becomes unbounded.

Theorem 7.15 (Bieberbach-Soul Theorem). *A non-compact, orientable, 3-dimensional Euclidean cone-manifold with cone angles in* $(0,\pi]$ *is isometric to one of the following.*
(1) A cone.
(2) A (possibly singular) solid torus, possibly with a twisted product metric.
(3) The product of a compact orientable 2-dimensional Euclidean cone-manifold with a line:
(i) torus$\times\mathbb{R}$*,*
(ii) pillowcase$\times\mathbb{R}$*, or*
(iii) turnover$\times\mathbb{R}$*.*

(4) (i) A folded ball, or
(ii) a singular folded ball.

(5) (i) $D^2(\pi,\pi)\times S^1$*,*
(ii) a twisted product $D^2(\pi,\pi)\tilde{\times}S^1$*, or*
(iii) a twisted line bundle over a Klein bottle.

(6) (i),(ii),(iii) $\mathbb{R}^3$ *with the singular locus shown, or*
(iv) a twisted line bundle over $\mathbb{R}P^2(\pi,\pi)$*.*

(7) (i),(ii) $\mathbb{R}^3$ *with the singular locus shown.*

Remark: These are illustrated in the following figure. The reader may wish to refer to Chapter 2 for an explanation of some of the terms in the theorem. In particular, definition 2.48 extends to cone-manifolds in the obvious way.

This theorem is actually a special case of a general theorem about non-compact, orientable, n-dimensional Euclidean cone-manifolds with cone angles at most π. That theorem states that such a Euclidean cone-manifold is, up to a 2-fold branched cover, isometric to a normal bundle of a lower dimensional, compact Euclidean cone-manifold.

This general theorem is analogous to the Bieberbach Theorems for Euclidean manifolds. In particular, it reduces the classification of non-compact,

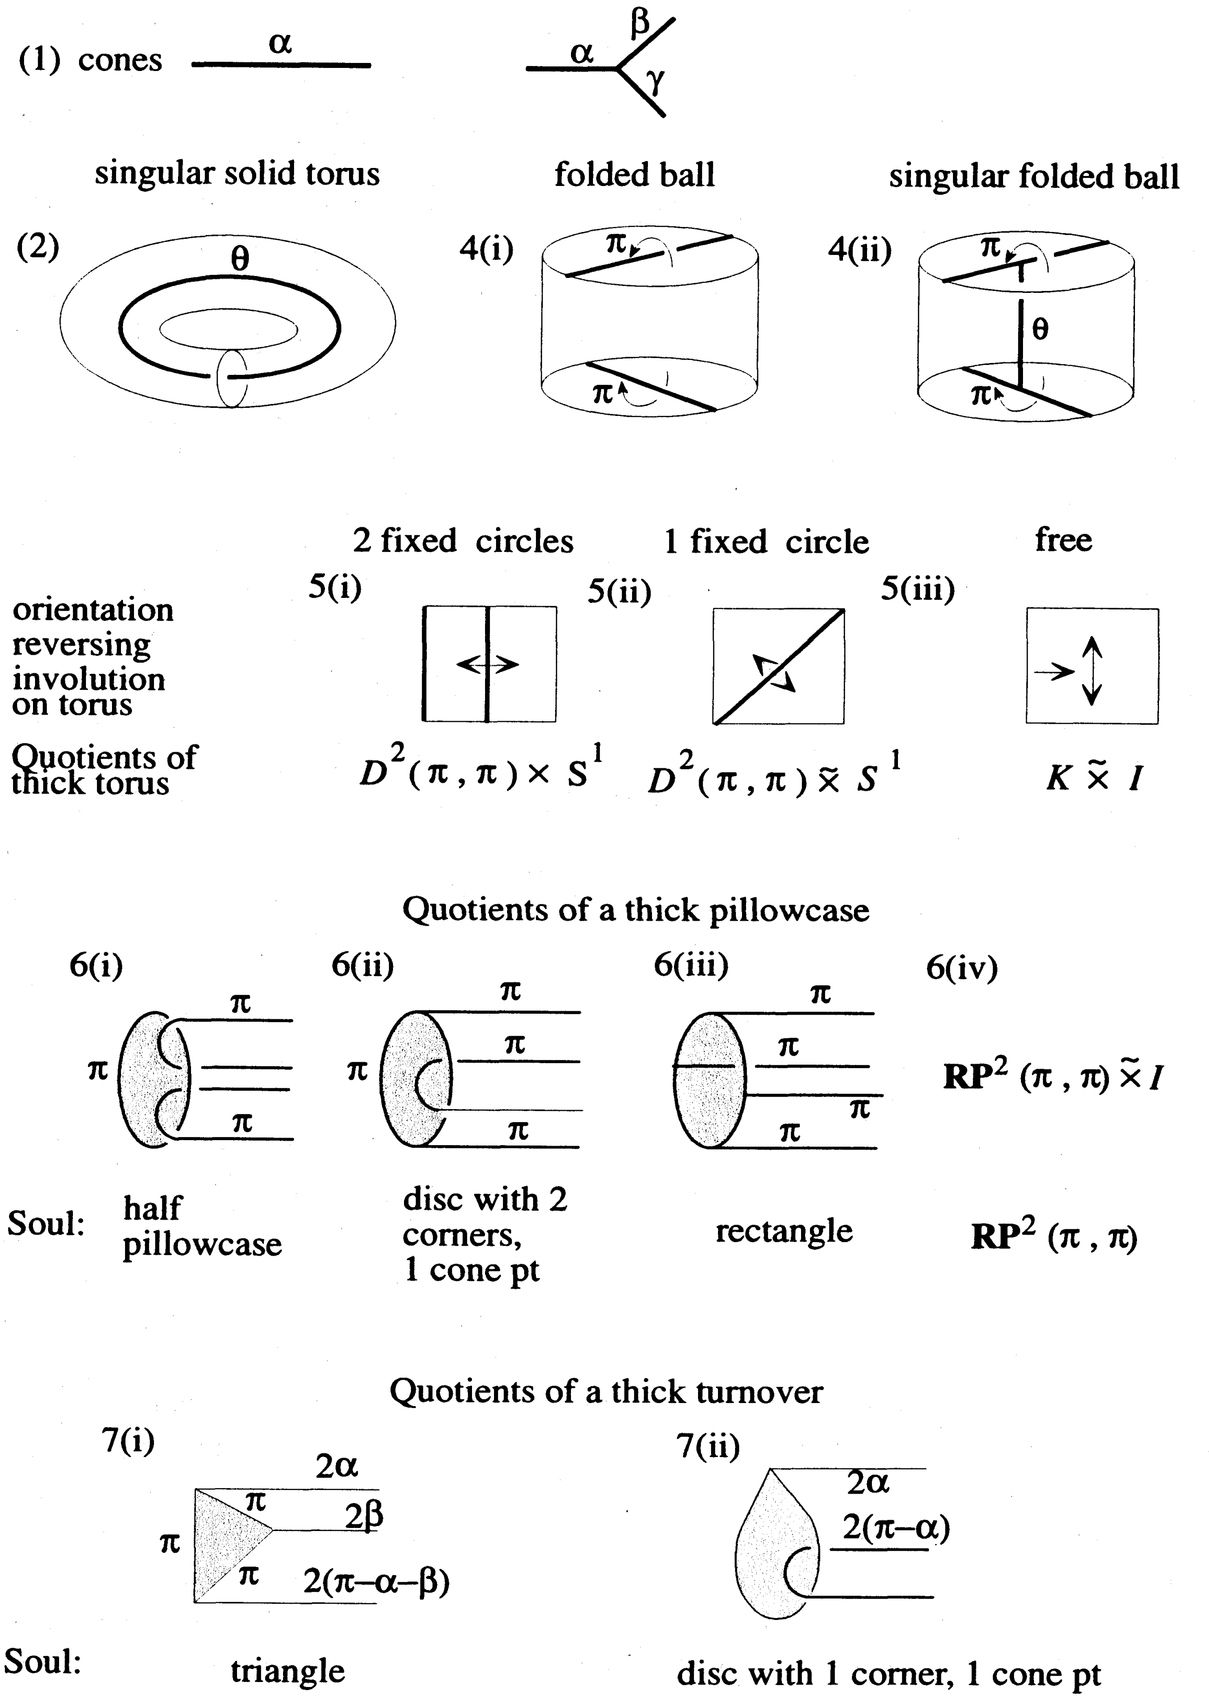

(1) cones
α
α β γ
singular solid torus
folded ball
singular folded ball
(2)
θ
4(i)
π
π
4(ii)
π
θ
π
2 fixed circles
1 fixed circle
free
5(i)
5(ii)
5(iii)
orientation reversing involution on torus
Quotients of thick torus
$D^2(\pi,\pi)\times S^1$
$D^2(\pi,\pi)\tilde{\times} S^1$
$K\tilde{\times} I$
Quotients of a thick pillowcase
6(i)
6(ii)
6(iii)
6(iv)
π
π
π
$\mathbf{RP}^2(\pi,\pi)\tilde{\times} I$
Soul:
half pillowcase
disc with 2 corners, 1 cone pt
rectangle
$\mathbf{RP}^2(\pi,\pi)$
Quotients of a thick turnover
7(i)
7(ii)
2α
2β
2(π–α–β)
2α
2(π–α)
Soul:
triangle
disc with 1 corner, 1 cone pt

Euclidean cone-manifolds to that of compact, lower dimensional Euclidean cone-manifolds and involutions on them. For example, the possible non-compact, orientable, 2-dimensional Euclidean cone-manifolds are a cone (normal bundle over a point with some angle; this includes the plane), an infinite cylinder (normal bundle over a circle), and an infinite pillow-case (normal bundle over an interval with angle π attached to its endpoints, which is the circle divided out by an involution). The 3-dimensional theorem above follows from the classification of 0, 1, and 2 dimensional Euclidean cone-manifolds with angles at most π and involutions on them.

The proof of the general theorem starts by following the outline of the proof of the Soul Theorem, due to Cheeger and Gromoll ([16]) which gives a structure theorem for non-compact manifolds with non-negative Ricci curvature. Indeed, in some cases the topology of the underlying space of the Euclidean cone-manifolds can be inferred directly from the Soul Theorem if one can smooth the metric to obtain a positively curved one. (The topology of the soul may change, however.) This argument works for all cone angles at most 2π. However, the more precise isometric description as a normal bundle is only true for cone angles at most π and requires further analysis.

The compact set, C, for which the Euclidean cone-manifold, B, is the normal bundle is called the **soul** of B. It can be described in terms of Busemann functions on B. A Busemann function, b_γ, is determined by an infinite ray γ; it is defined by $b_\gamma(x) = \lim_{t\to\infty} d_B(x, \gamma(t)) - t$, where $d_B(\cdot, \cdot)$ denotes distance in B. The soul C is derived from the level set for the maximum value of the function obtained by taking the infimum over all rays emanating from a chosen point, $p \in B$. This construction is at the centre of [16].

The 3-dimensional theorem above gives a list of the possible geometric limits under the scaling process, at least when the limit is non-compact. The Gromov-Hausdorff topology implies that neighbourhoods of points with small injectivity radius in hyperbolic cone-manifolds can be approximated by these. This leads to a structure theorem for the topology of the set with small injectivity radius that generalizes the Margulis Lemma for hyperbolic manifolds. The following theorem, which is a consequence of the generalized Margulis lemma, says that, if a point in a 3-dimensional hyperbolic cone-manifold, M, has a small injectivity radius when M is scaled to have diameter at least 1, it has a neighbourhood that is almost isometric to a neighbourhood of the soul in a non-compact Euclidean cone-manifold. The possible list for such models comes from 7.15 (Bieberbach-Soul theorem), where a few cases have been eliminated using the finite volume hypothesis.

Theorem 7.16 (Local Margulis for Cone-Manifolds).
Let M be a finite volume 3-dimensional hyperbolic cone-manifold with cone angles in the range $(0, \pi]$. Then there is an $\epsilon > 0$ so that if $x \in M$ with $inj(x) < \epsilon \min(1, \mathrm{diam(M)})$ then x has a compact neighbourhood, containing $N(x, 1000\mathrm{inj}(x))$, which is almost isometric to one of the following:
(1) A (singular) solid torus,
(2) A (singular) folded ball,
(3) A thick torus, pillowcase, or turnover,
(4) A folded thick torus, pillowcase or turnover.

The main idea in the proofs of both the Collapsing Theorem for cone-manifolds 7.13 (Collapsing theorem) and the structure of the ends of cone-manifolds 7.11 (ends at $t < 1$) is to use the neighbourhoods of points with small injectivity radius whose topology is described by this theorem and analyze how they can be glued together.

The analysis of the ends of hyperbolic cone-manifolds in 7.11 (ends at $t < 1$) is simplified by the fact that, since $t_\infty < 1$ and we have not reached the orbifold angles, all the cone angles are strictly less than π. This reduces the list of possible local models from 7.16 (local Margulis) to a (singular) solid torus, a thick torus, or a thick turnover. The (singular) solid tori can be incorporated into the compact part. The remainder of the proof involves showing that any pair of standard neighbourhoods homeomorphic to a thick turnover (torus) that intersect can be amalgamated into a larger neighbourhood homeomorphic to a thick turnover (torus). The product ends are created in this manner.

7.6 Outline of the Collapsing Theorem

In this section we will outline the proof of the Collapsing Theorem 7.13, which states that, if an orbifold Q has an ϵ-collapse, for sufficiently small ϵ, then it is either a graph orbifold or an orbifold bundle with generic fibre a turnover or a pillowcase.

The proof begins with the local Margulis theorem 7.16 (local Margulis) which provides the local models from which the orbifold is built. There is an analogy to a child's construction kit. The construction kit contains pieces which are Euclidean models that are almost isometric to standard neighbourhoods and we can build orbifolds by fitting together pieces from this kit. The pieces are metric spaces which must be glued by almost isometries along parts of their boundaries. With one exception, which is easily analyzed separately, the pieces have OSFS structures. The main issue is to show that these structures can be glued together to give the structure of a graph orbifold, except in special cases when Q is an orbifold bundle. The geometry of the collapse provides extra information, called the latitude hypothesis, on the relation between the fibres of the different pieces. This plays a key role in the argument.

With the exception of (folded) thick turnovers, every standard model admits at least one OSFS structure. However, it is easy to see that if there are any (folded) thick turnovers then Q is a bundle.

The idea is that (folded) thick turnovers are the only standard neighbourhoods with "triangular" shaped boundaries. The standard neighbourhoods are almost isometric to Euclidean models, and are glued together by isometries. Thus the corresponding Euclidean model neighbourhoods have almost isometric boundaries. Hence the only standard neighbourhood that can be connected to a (folded) thick turnover is another such. This implies that Q is a union of a finite sequence of (folded) thick turnovers, arranged in either a linear fashion (giving a bundle over an interval with generic fibre a turnover) or a circular fashion (giving a bundle over a circle). This is essentially the same argument that provides the structure of the ends of a cone-manifold whose angles are strictly less than π 7.11 (ends at $t < 1$).

Having dealt with (folded) thick turnovers, we may assume that every standard neighbourhood admits at least one OSFS structure. One of these neighbourhood types, folded thick tori, is easily analyzed. The boundary of a folded thick torus, V, is a torus T which is incompressible in V. If T is incompressible in Q, it must be boundary parallel and Q equals V, which is an OSFS. If T compresses in Q, then $Q - V$ is a (singular) solid torus. Every every folded thick torus admits two OSFS structures, given by the two eigenvectors of the involution that does the folding. By 2.46, a fibration on the boundary of a (singular) solid torus extends over the interior unless the fibre is isotopic to a meridian. This can occur for at most one of the two fibrations of V, so Q is an OSFS.

Thus we are reduced to the case that the only standard neighbourhoods in Q are (folded) thick pillowcases, (singular) folded balls, (singular) solid tori and thick tori. This corresponds to moving down the flowchart at the end of this section past the first two boxes.

The basic strategy is to attempt to fit the OSFS structures on the standard neighbourhoods together to give a single OSFS structure on Q when it is orbifold-atoroidal or, more generally, when Q arises as a double, to give an OSFS structure to each component after Q is cut along incompressible

Euclidean 2-dimensional sub-orbifolds.

The boundaries of these standard neighbourhoods consist of tori or pillowcases. By 2.46 they both admit countably many orbifold Seifert fibrations, parametrized by the slope of a regular fibre. The main tool used to extend an OSFS defined on the boundary over the interior of a (singular) solid torus or (singular) folded ball is the lemma 2.46 which states this can be done unless a regular fibre is a meridian.

In a general topological setting, it is quite possible to have a manifold or orbifold that is the union of two SFS glued along their boundaries which is not a SFS or a graph manifold. For example the exterior, X, of the trefoil knot is a SFS. Glue a solid torus, which also is a SFS, to X along their boundaries so that a meridian curve of the solid torus is glued to a regular fibre in the boundary of X. This is ± 6 Dehn filling depending on whether the trefoil is left- or right-handed. The resulting closed manifold is $L(2,1)\#L(3,1)$. It is not difficult to show that it is not a graph manifold.

In our situation there is an additional piece of information coming from the geometry of an almost collapsed orbifold called the **latitude hypothesis.** A *latitude* is the isotopy class of any shortest closed geodesic on a pillowcase or torus. There may be up to three such isotopy classes, though generically there is exactly one. When we attempt to construct an OSFS structure on Q at each stage we will have a finite number of suborbifolds each of which has been given an OSFS structure such that a regular fibre is isotopic to a latitude of the boundary pillowcase (or torus) of some (singular) folded ball (or solid torus) standard neighbourhood. We discuss the pillowcase case here; the torus case is similar. The latitude hypothesis is the statement that if α is a latitude of a standard neighbourhood which is a (singular) folded ball then α is not homotopic in $Q-\Sigma$ to either a point or to a meridian of Σ. Thus all regular fibres appearing in our construction satisfy this condition. This prevents the phenomenon of killing the homotopy class a regular fibre when gluing together two OSFS.

The latitude hypothesis is proved by estimating the holonomy of a latitude. A standard neighbourhood in Q which is a (singular) folded ball is almost isometric to a compact Euclidean cone-manifold. The first step is to show that the Euclidean holonomy of a latitude is the composition of two rotations through π around almost parallel axes. This uses the fact that the diameter of a (singular) folded ball standard neighbourhood is very large compared to the diameter of the soul, i.e. the distance between the two axes of rotation. Roughly speaking, by taking a large finite orbifold cover of the (singular) folded ball, one sees there are two almost parallel rotation axes in the cover which are not too far apart. Thus the hyperbolic holonomy of a latitude is almost the composition of two rotations through angles almost equal to π around almost parallel distinct axes. The hyperbolic holonomy is non-trivial; hence, the latitude is essential in $Q-\Sigma$. This estimate also shows that the latitude has very small complex translation length. Its holonomy can't be close to a rotation through an angle $2\pi/n$ around some edge of Σ because such an elliptic does not have "very small" rotation angle. Thus a latitude is not homotopic in $Q - \Sigma$ to a meridian of Σ.

The union of standard neighbourhoods meeting Σ is a suborbifold bounded by tori. We denote by $\mathcal{N}$ the union of those components of this orbifold that are not singular solid tori. A foliation argument is used to show that:

Theorem 7.17. *At least one component of the union of the standard neighbourhoods meeting Σ is not a singular solid torus. Hence $\mathcal{N}$ is non-empty.*

This is important because the boundaries of the components of $\mathcal{N}$ are incompressible in $\mathcal{N}$ so they must either be boundary parallel or compressible in Q. A compressible torus must bound a (singular) solid torus since Q is irreducible. It follows that if each component of $\mathcal{N}$ is an OSFS then (using the latitude hypothesis) this OSFS structure can be extended over all of Q. Similarly, if each component of $\mathcal{N}$ is a graph orbifold, then so is Q.

The standard neighbourhoods in $\mathcal{N}$ are of three types: thick pillowcases, (singular) folded balls, and folded thick pillowcases. All three have pillowcases as boundary components. By topological methods, it is possible to combine any neighbourhoods of the same type that intersect. Either it can be arranged that any two neighbourhoods of the same type are disjoint, or a component of $\mathcal{N}$ fibres over a 1-dimensional orbifold with Euclidean fibre, a case that can be easily handled separately. Furthermore, the same argument used previously to deal with folded thick tori shows that either all the folded thick pillowcases are incompressible in Q, or Q is a folded thick pillowcase union a (singular) folded ball. As with the folded thick torus, the latter two cases are readily seen to be OSFS. When a folded thick pillowcase, V, is incompressible, we cut along a pillowcase which is boundary parallel in V. This creates a component which is again homeomorphic to V and is used in the decomposition as a graph orbifold; the remaining component contains a new thick pillowcase with one component on the boundary. We continue to denote the latter piece by $\mathcal{N}$; it suffices to show that its components are graph orbifolds.

A component, X, of $\mathcal{N}$ is a union $X = A \cup B$ where A is a disjoint union of thick pillowcases and B is a disjoint union of (singular) folded balls. Furthermore, it is possible to arrange that each (singular) folded ball component of B intersects A in exactly two components each of which is a $D^2(2,2)$. The combinatorics of this arrangement may be complicated, so, to simplify the picture, we cut along a (possibly compressible) pillowcase inside each thick pillowcase which does not already have one boundary component on the boundary of X. In the resulting pieces exactly one boundary component of each thick pillowcase is on the boundary; the other boundary component is connected to (singular) folded balls along $D^2(2,2)$'s.

Specifically, each piece is of the form $A \cup B$ where $B = \bigcup_i B_{2i}$ is a disjoint union of (singular) folded balls and $A = \bigcup_i A_{2i-1}$ is a disjoint union of thick pillowcases. Furthermore B_{2i} intersects only $A_{2i\pm1}$ and each component of the intersection is isomorphic to $D^2(2,2)$. Hence each component of X is the union of a finite number of thick pillowcases and (singular) folded balls arranged alternately in a circular way like hollow beads on a string. Each hollow bead is a thick pillowcase. Each piece of string between two beads is a (singular) folded ball. We call this cutting up process "disassembling" $\mathcal{N}$.

"Disassembly" picture

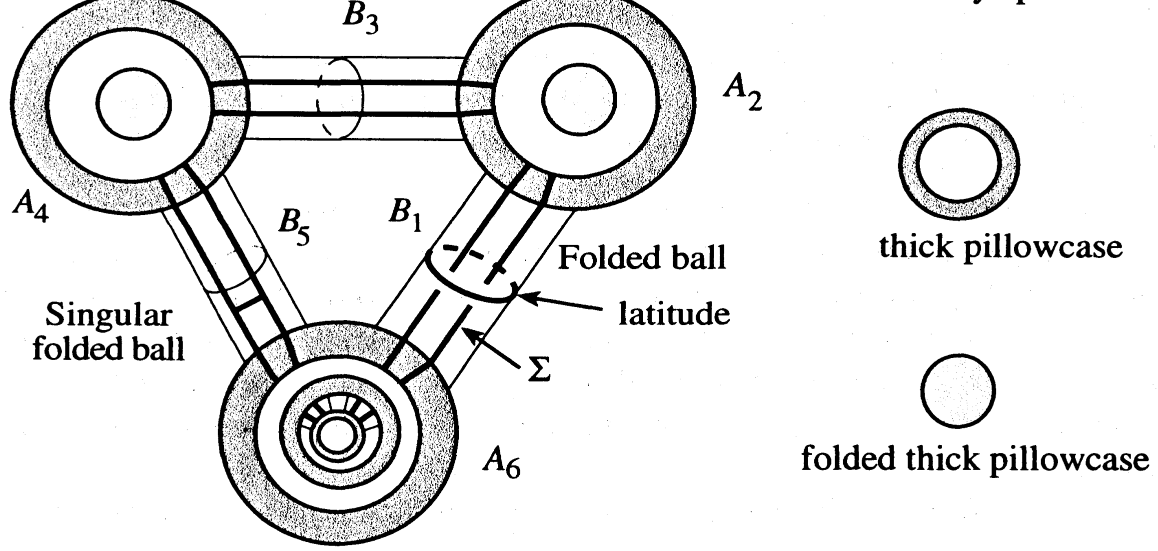

We will now show that each piece X of the disassembled orbifold admits a OSFS structure. Choose a (singular) folded ball B_1 in X. Then $\partial B_1 \cap \partial X$ is an annulus, as shown in the figure. A core curve of this annulus is a latitude of ∂B_1. Since this latitude does not compress in B_1 there is a OSFS structure on B_1 with this latitude a regular fibre. Now $C = A_2 \cap B_1 \cong D^2(2,2)$. The OSFS structure on B_1 may be isotoped to give an orbifold fibration of C with ∂C one of the fibres. This fibration extends productwise over A_2 to an OSFS structure on A_2. This may be isotoped and then extended over the next (singular) folded ball B_3 attached to A_2. In this way we can extend the OSFS structure around the string of beads which makes up X.

We now use these OSFS structures on the pieces to define a graph orbifold structure on Q. The boundary components of each piece has an induced fibration. One boundary component of each piece of the disassembled orbifold is a torus. It is either boundary parallel or bounds a (singular) solid torus in Q. The induced fibration extends over the (singular) solid tori by the latitude hypothesis; fill in these (singular) solid tori with the extended fibration. We now glue back along the pillowcases used to disassemble $\mathcal{N}$. If such a pillowcase, P, is incompressible, then we do not try to match the OSFS structures on the two sides of P, since P will be part of a set of incompressible pillowcases decomposing Q in OSFS pieces.

If a pillowcase, P, becomes compressible when two pieces are glued together, it must bound a (singular) folded ball in Q (see 2.47). (That it can't be contained in a (singular) ball follows from the Seifert fibres on the pieces.) By the latitude hypothesis, the fibration on P extends over the (singular) folded ball. This may change the fibration that was already on the (singular) folded ball. Continuing in this manner the fibrations are matched up, piece by piece along all of the compressible pillowcases. Thus, every component of $\mathcal{N}$, cut along incompressible pillowcases, is an OSFS. Hence $\mathcal{N}$ is a graph manifold as desired.

It is useful to note that when Q is the union of two (folded) thick pillowcases glued along their mutual boundary, the process of cutting along incompressible pillowcases decomposes Q into two (folded) thick pillowcases. This is a degenerate version of a graph orbifold where the pieces are each (possibly the quotient by an involution of) a product. In this case Q is a bundle with generic fibre a pillowcase and base a circle or interval. This happens if Q is a Solv orbifold.

The following flowchart summarizes the overall structure of this proof.

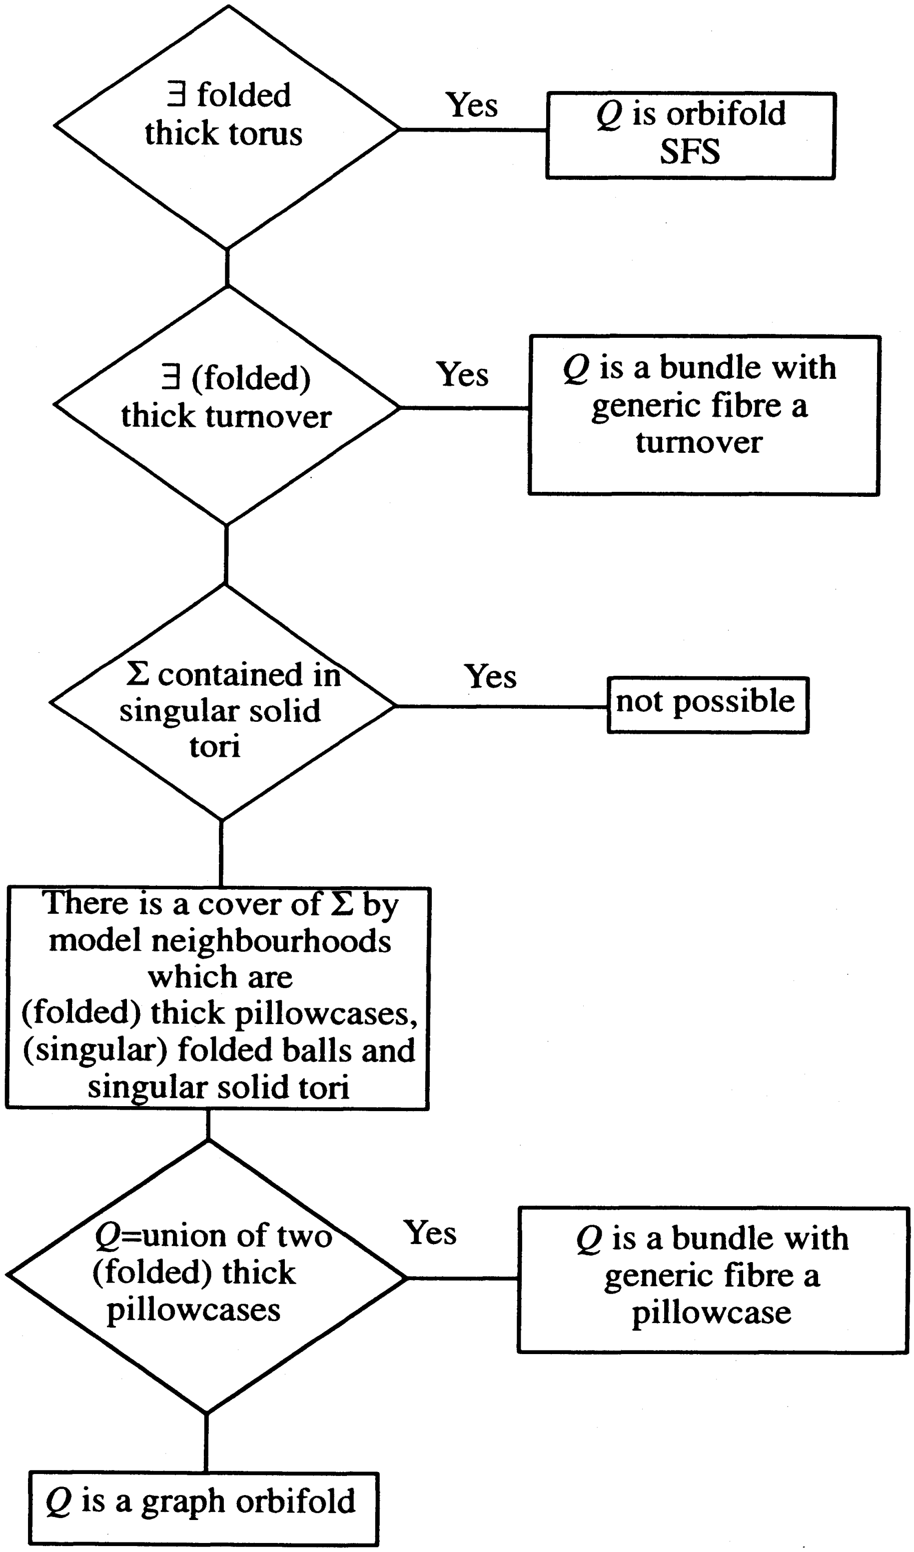
∃ folded thick torus
Yes
Q is orbifold SFS
∃ (folded) thick turnover
Yes
Q is a bundle with generic fibre a turnover
Σ contained in singular solid tori
Yes
not possible
There is a cover of Σ by model neighbourhoods which are (folded) thick pillowcases, (singular) folded balls and singular solid tori
Q=union of two (folded) thick pillowcases
Yes
Q is a bundle with generic fibre a pillowcase
Q is a graph orbifold