

Heegaard gradient and virtual fibers

Joseph Maher*

February 12, 2019

Abstract

We show that if a closed hyperbolic 3-manifold has a sequence of finite covers of bounded Heegaard genus, then it is virtually fibered. This generalizes a theorem of Lackenby, removing restrictions needed about the regularity of the covers. Furthermore, we can replace the assumption that the covers have bounded Heegaard genus with the weaker hypotheses that the Heegaard splittings for the covers have Heegaard gradient zero, and also bounded width, in the sense of Scharlemann-Thompson thin position for Heegaard splittings.

1 Introduction

Inspired by Lubotzky's work on Property (τ) [12], Lackenby [10, 11] investigated the behaviour of Heegaard splittings in finite covers of hyperbolic 3-manifolds, and how this relates to other well-known conjectures about 3-manifolds, such as Thurston's conjecture that a hyperbolic 3-manifold has a finite cover which has positive Betti number, or Thurston's stronger conjecture that every hyperbolic 3-manifold has a finite cover which is fibered. Lackenby was able to show that if a hyperbolic manifold M has a sequence of covers M_i of bounded Heegaard genus, and bounded irregularity, then there is a finite cover which is fibered. Here, bounded irregularity means that the index of the normalizer of $\pi_1 M_i$ in $\pi_1 M$ is bounded. In this paper we show how to remove the assumptions about the regularity of the covers, and we can also replace the assumption about bounded genus with a weaker hypothesis. To be precise, we prove the following theorem:

Theorem 1.1. *Let M be a closed hyperbolic 3-manifold, with a sequence of finite covers $\{M_i \rightarrow M\}$, with Heegaard splittings H_i , of Euler characteristic χ_i . If the Heegaard gradient χ_i/d_i tends to zero, and the width $c_+(H_i)$ of the Heegaard splitting is bounded, then M is virtually fibered.*

Here $c_+(H_i)$ is the width of the Scharlemann-Thompson thin position for the Heegaard splitting, which in particular, is bounded by the genus of the splitting. In fact, the proof only requires the width $c_+(H_i)$ to grow sufficiently slowly, but the growth bound we produce is logarithmic, and not particularly convenient to write down.

Note that if a sequence of finite covers M_i of M are not Haken, then the width of the splitting is the same as the Heegaard genus, so if the widths are bounded, then the Heegaard gradient tends to zero. This gives the following corollary:

Corollary 1.2. *Let M be a closed hyperbolic manifold with a sequence of finite covers M_i for which $c_+(M_i)$ is bounded. Then M is virtually Haken.*

*email: maher@its.caltech.edu

1.1 Background

Every closed 3-manifold has a Heegaard splitting, which is a division of the manifold into two handlebodies. A handlebody is a compact 3-manifold homeomorphic to a regular neighbourhood of a graph in \mathbb{R}^3 . The graph is called a spine for the handlebody, and we say the genus of the handlebody is the genus of the surface forming its boundary.

A Heegaard splitting is a topological structure on a 3-manifold, but it is possible to relate a Heegaard surface to a minimal surface, which is a geometric object in the manifold. A Heegaard splitting S of a 3-manifold M gives rise to a sweepout, namely a degree one map $S \times I \rightarrow M$, such that $S \times 0$ and $S \times 1$ are taken to the two spines of the handlebodies on either side of S . We will normally write S_t for $S \times t$. The area of the surfaces S_t has some maximum value as t varies. The minimax value is the smallest maximal area over all possible sweepouts of the manifold. If we take a sequence of maximal area sweepout surfaces whose area tends to the minimax value, then it is a fundamental result of Pitts and Rubinstein [14] (see also Colding and De Lellis [5]) that there is a subsequence which converges to a minimal surface in the manifold, possibly of lower genus.

Recall that a minimal surface is a surface in a manifold with mean curvature zero. If the manifold is hyperbolic, then the intrinsic curvature of the surface is at most -1 , so the Gauss-Bonnet formula gives an upper bound for the area of the minimal surface, and hence for any sweepout surface in a minimax sweepout. As some sweepout surface must divide the manifold into two parts of equal volume, this gives an immediate link between the Heegaard genus of the manifold and the Cheeger constant, which in turn is closely related to the first eigenvalue of the Laplacian. Lubotzky [12] has shown that a closed hyperbolic 3-manifold has Property (τ) with respect to its congruence covers, which means that Cheeger constant of the covers is bounded away from zero. Lackenby [10] uses this in his proof that the Heegaard genus of the congruence covers of a hyperbolic 3-manifold grow linearly with their degree.

There is a bound on the diameter of a minimal surface in a hyperbolic 3-manifold M in terms of the injectivity radius of M , and the genus of the surface. As the injectivity radius of M is a lower bound for the injectivity radius of any cover, this gives a diameter bound for minimal surfaces in covers of M depending only on their genus. If the degree of the cover is very large with respect to the genus of the minimal surface, and the cover is fairly regular, then we can find many disjoint translates of the minimal surface in the cover. Lackenby is able to show the manifold is virtually fibered by showing that some of these translates must be parallel.

If we make no assumptions about the regularity of the covers, then we may not have any covering translations. However, if the volume of the manifold is large, and we have a sweepout of surfaces of bounded diameter, then we can find many disjoint surfaces in the manifold. The original sweepout surfaces need not have bounded diameter, but in Section 3 we define a generalized sweepout, and show that we can find one with sweepout surfaces of bounded diameter.

The diameter bound depends on the genus of the surface, so if a sequence of covers have Heegaard splittings of bounded genus, then they have sweepouts of bounded genus. However, there is a useful technique called Scharlemann-Thompson thin position for Heegaard splittings [15] which in some cases allows us to produce more general sweepouts with lower genus than the genus of the original Heegaard splitting. We now give a brief outline of this technique. We can think of a Heegaard splitting H as a handle structure for the manifold, with one of the handlebodies constructed from a zero handle and one-handles, and the other handlebody constructed from a three handle and two handles. It may be possible to rearrange the order of the handles, i.e. add some of the two handles before all of the one-handles have been added. We say that the maximum difference between the

number of one and two handles added so far is the width of a particular ordering of the handles. Let $c_+(H)$ be the minimum width over all possible orderings. The thin position width may just be the same as the genus of the original splitting. For example, a Heegaard splitting is strongly irreducible if any pair of essential discs on opposite sides of the Heegaard surface must intersect, so a strongly irreducible splitting is already in thin position as the order of the handles can not be changed. In [15], Scharlemann and Thompson show that if a Heegaard splitting is put in thin position, then the manifold is split up into pieces separated by incompressible surfaces, and each piece contains a strongly irreducible Heegaard splitting. So if a sequence of 3-manifolds M_i have splittings H_i of bounded width $c_+(H_i)$, then each manifold M_i can be split into pieces where each piece has a sweepout of bounded genus. If the Heegaard gradient χ_i/d_i tends to zero, then the average volume per handle becomes large, so there must be pieces of large diameter, and we can then argue that some of the generalized sweepout surfaces must be parallel.

1.2 Outline of the proof

Let M be a closed hyperbolic manifold. Let M_i be a sequence of finite covers of M , with degree d_i , and with Heegaard splittings H_i , of Euler characteristic χ_i . The volume of the cover is equal to the volume of M times the degree d_i of the cover, so if the Heegaard gradient χ_i/d_i tends to zero, then the average volume per handle becomes arbitrarily large. However, we have also assumed that the width of the Scharlemann-Thompson thin position for the Heegaard splittings H_i is bounded, so this means that there is a sequence of compression bodies C_i of bounded genus, whose volume tends to infinity.

A sweepout of a 3-manifold is a one-parameter family of surfaces which fill up the manifold in a degree one manner. A sweepout for a compression body can be constructed by homotoping ∂_+ down to a spine for the compression body. We can use the hyperbolic metric on the compression body to “straighten” the sweepout so that each sweepout surface is negatively curved, except for a single cone point of angle less than 2π . In Section 2 we review a precise construction of such sweepouts, called **simplicial sweepouts**, due to Bachman, Cooper and White [1]. A technical point is that it is convenient to straighten the sweepout in a complete manifold. We can construct a complete manifold from the compression body by gluing on negatively curved “flared ends” to the boundary components, so we end up working in a manifold homeomorphic to the interior of a compression body, with sectional curvature at most -1 , rather than the original hyperbolic manifold.

One of the main technical tools we will use is a **generalized sweepout**, which is a sweepout in which the genus of the surfaces in the one-parameter family is allowed to change. An ordinary sweepout can be defined as a degree one map from $S \times I$ to a 3-manifold, where S is a surface. The domain of the map, $S \times I$, is a 3-manifold with height function given by projection on to the I factor. In a generalized sweepout $S \times I$ is replaced by a 3-manifold with a Morse function, and the level sets of the Morse function are the sweepout surfaces, which are mapped across into M by a degree one map. We do not require the images of the level sets to be embedded.

The Gauss-Bonnet formula gives an area bound for negatively curved surfaces in terms of their genus, but such surfaces may have arbitrarily large diameter. However, a negatively curved surface has a thick-thin decomposition, so each surface has a “thick part”, where the injectivity radius is large, and a “thin part” where the injectivity radius is small. The thick part has bounded diameter, and the thin part consists of annuli, by the Margulis Lemma. In Section 3 we give an explicit construction for choosing a continuously varying family of thin parts for the sweepout surfaces, and show how to chop them off to produce a generalised sweepout whose leaves have bounded diameter.

The sweepout surfaces in the generalized sweepout have bounded diameter, so as the volume of the compression bodies becomes large, there are more and more disjoint sweepout surfaces inside the compression bodies. Although the sweepout surfaces are immersed, it still makes sense to say a sweepout surface S separates a disjoint sweepout surface S' from ∂_+ , if every path from a point on S' to ∂_+ has $+1$ algebraic intersection number with S . We show that as the volume of the compression body becomes large, we can find an arbitrarily large collection of sweepout surfaces which are **nested**, in the sense that for each pair of surfaces, one surface separates the other from ∂_+ .

We may also assume that the surfaces all have the same genus. As the surfaces are nested, and they are all compressions of ∂_+ , we can show that they are all homotopic. Furthermore, we can find a nested collection of surfaces S_1, \dots, S_n so that the homotopy from S_n to S_i is disjoint from S_j with $j < i$. Gabai [7] showed that the singular norm on homology is equal to the Thurston norm, so we can replace the immersed surfaces with embedded surfaces, which we can show are also homotopic. Homotopic surfaces in a compression body need not be isotopic; however if we change the isotopy class of a surface, we must do so by a homotopy that hits a spine of the compression body. As the homotopies between S_n and S_i are disjoint from S_1 for $i > 1$, they can't change the isotopy class of the surface, so the surfaces are in fact isotopic, and bound products in the compression body.

So we can find arbitrarily many disjoint parallel embedded surfaces in a compression body of large enough volume. Each compression body lives in a cover of M , and each cover is tiled by a choice of fundamental domain for M . As the surfaces have bounded diameter, each surface may only hit finitely many fundamental domains of M in the cover M_i . So there must be a pair of parallel surfaces which hit the same pattern of fundamental domains, so we can cut the compression body along the matching pair of fundamental domains and then glue them together to form a fibered cover of the original manifold M .

The remainder of this paper fills in the details of this argument. In Section 2 we describe the simplicial sweepout construction due to Bachman, Cooper and White [1], and give a minor generalization to sweepouts of compression bodies. In Section 3 we describe how to turn a simplicial sweepout of a compression body into a generalized sweepout with sweepout surfaces of bounded diameter. Finally in Section 4 we complete the proof by showing that as the volume of a compression body tends to infinity we can find arbitrarily many disjoint parallel surfaces of bounded diameter.

1.3 Acknowledgements

I would like to thank Marc Lackenby, Hyam Rubinstein, Nathan Dunfield, Jason Manning and Martin Scharlemann for helpful conversations. This paper was typeset in \LaTeX , and all the figures were drawn in `xfig`.

2 Simplicial sweepouts.

In this section we give a precise definition of a sweepout, and explain how to straighten it to a sweepout with leaves of bounded area.

A **handlebody** is a 3-manifold with boundary, homeomorphic to a regular neighbourhood of an embedded graph in \mathbb{R}^3 . A **Heegaard splitting** is a surface which divides a closed 3-manifold into two handlebodies. A **compression body** is the analogue of a handlebody for 3-manifolds with boundary, defined as follows. Take $S \times I$, where S is a closed connected orientable surface, and

attach 2-handles to $S \times 0$. If any of the resulting boundary components are 2-spheres, cap them off with 3-balls. The boundary component corresponding to $S \times 1$ is called the upper boundary, and will be denoted ∂_+ . The other boundary components are the lower boundary, denoted ∂_- , which need not be connected. A handlebody is a special case of a compression body in which ∂_- is empty.

A **spine** for a handlebody H is an embedded graph Γ in H so that the complement $H - \Gamma$ is a product $\partial H \times [0, 1)$. Similarly, a spine for a compression body is a 2-complex Γ which is the union of ∂_- and a properly embedded graph, so that $C - \Gamma$ is a product $\partial_+ \times [0, 1)$. We can construct a sweepout of a handlebody by shrinking the boundary down to the spine, using the product structure. More generally, we shall use the following definition of a sweepout of a compression body.

Definition 2.1. Sweepout. Let C be a compression body, and let Γ be a spine for C . Let S be a surface homeomorphic to ∂_+ , and let $\phi : S \times I \rightarrow C$ be a continuous map. We will write S_t for $S \times t$. We say the map ϕ is a **sweepout** if ϕ takes S_0 to $\partial_- \cup \Gamma$ and S_1 to ∂C , and ϕ is degree one, i.e. $\phi_* : H_3(S \times I, S \times \partial I) \rightarrow H_3(C, \partial C)$ is an isomorphism. \diamond

We will often think of the t variable as time. This definition does not require that the surfaces $\phi(S_t)$ be embedded in C , and surfaces at different times may intersect each other.

For our purposes, we would like to obtain some control over the geometry of the surfaces S_t in the sweepout. Roughly speaking, we shall do this by triangulating the surfaces, and then “straightening” the triangulation in the negatively curved metric on the manifold. The construction we use is that of a **simplicial sweepout**, and is described by Bachman, Cooper and White, in [1]. However, we provide a complete definition, as we need a slightly more general construction which allows for manifolds with boundary and change of basepoints.

When we give a surface a triangulation, it need not be a simplicial triangulation, as we wish to allow one-vertex triangulations. Our triangulations will be Δ -structures as defined by Hatcher [9].

In hyperbolic space the convex hull of three points is a geodesic triangle. In spaces of varying negative curvature this need not be the case, so instead we use coned simplices. A **coned 1-simplex** $\Delta^1 = (v_0, v_1)$ is a constant speed geodesic from v_0 to v_1 . We allow degenerate 1-simplices in which the speed is zero and the image of the 1-simplex is a point. A **coned n -simplex** is a map $\phi : \Delta^n \rightarrow M$ so that $\phi|_{\Delta^{n-1}}$ is a coned $n-1$ -simplex, and $\phi|_{\{tx + (1-t)v_n | t \in [0, 1]\}}$ is a constant speed geodesic for all $x \in \Delta^{n-1}$. Note that the map constructed in this way depends on the order of the vertices in the simplex, and the triangle need not be embedded in M .

Definition 2.2. Simplicial surface. Let S be a triangulated surface, and let M be a closed Riemannian manifold of sectional curvature at most -1 . Let $\phi : S \rightarrow M$ be a continuous map so that the map on each triangle $\phi : \Delta \rightarrow M$ is a coned 2-simplex. Then we say that $\phi : S \rightarrow M$ is a **simplicial surface**. \diamond

A simplicial sweepout is a sweepout in which every sweepout surface is a simplicial surface with a bounded number of triangles, and at most one cone point of positive curvature.

Definition 2.3. Simplicial sweepout. Let $\phi : \Sigma \rightarrow C$ be a sweepout, such that each surface S_t is mapped to a simplicial surface with at most $4g$ triangles, and with at most one vertex of angle sum less than 2π . Then we say that ϕ is a **simplicial sweepout** of C . \diamond

The following theorem follows from the proof of Theorem 2.3 in [1].

Theorem 2.4 (BCW). *Let M be a closed orientable Riemannian manifold of sectional curvature at most -1 . Let S_0 and S_1 be simplicial surfaces in M , with one vertex triangulations in which*

the vertices coincide in M . Let $\phi : S \times I \rightarrow M$ be a homotopy between S_0 and S_1 which does not move the vertex. Then there is a simplicial sweepout $\Phi' : S \times I \rightarrow M$, such that for all t , $\Phi(S \times t)$ consists of at most $4g$ triangles, and has at most one vertex at which the angle sum is less than 2π . Furthermore Φ is homotopic to Φ' , relative to the boundary of $S \times I$.

In fact we need a minor extension which enables us to change the vertex of the one-vertex triangulation.

Lemma 2.5. *Let M be a closed orientable Riemannian manifold of sectional curvature at most -1 . Let S_0 and S_1 simplicial surfaces with one vertex triangulations, which are homotopic by a homotopy $\Phi : S \times I \rightarrow M$. Then there is a simplicial sweepout $\Phi' : S \times I \rightarrow M$ which is homotopic to Φ , relative to $S \times \partial I$.*

Proof. We reduce to the case of Theorem 2.4 by showing how to homotop S_0 so that it shares a common basepoint with S_1 .

Let $\phi : S \times I \rightarrow M$ be the homotopy between the simplicial surfaces S_0 and S_1 . Let x_0 and x_1 be the basepoints for S_0 and S_1 respectively. We may assume that the pre-images of the basepoints lie in the same $x \times I$ fiber in $S \times I$. First homotop the homotopy ϕ to a new map, which by abuse of notation we will also call ϕ , in which ϕ restricted to $x \times I$ is a geodesic from x_0 to x_1 . Now homotop ϕ to a new map ϕ' so that $\phi'(x, t) = \phi(x, 2t)$, if $t \leq \frac{1}{2}$, and $\phi'(x, t) = \phi(x, 1)$ if $t \geq \frac{1}{2}$. This means that $\phi'(S \times \frac{1}{2})$ has the same basepoint as S_1 , and the homotopy from $S \times \frac{1}{2}$ to $S \times 1$ is constant on $x \times t$. Give $S \times \frac{1}{2}$ the same one vertex triangulation as $S \times 0$, and homotop ϕ' so that the edges get mapped to geodesics based at x_1 . The lower half of $S \times I$ now has a division into cells which are triangles $\Delta_i \times [0, \frac{1}{2}]$ in $S \times [0, \frac{1}{2}]$, and ϕ' maps the 1-skeleton of these triangular prisms to geodesic arcs in M . So we can homotop ϕ' so that ϕ' is simplicial on each $\Delta_i \times [0, \frac{1}{2}]$.

Now $\phi'(S \times \frac{1}{2})$ and $\phi'(S \times 1)$ share a common basepoint, and ϕ' restricted to $S \times [\frac{1}{2}, 1]$ fixes the basepoint, so we can apply Theorem 2.4. \square

3 Generalised sweepouts.

Simplicial sweepouts are sweepouts in which each surface has bounded area, but we would like to construct sweepouts in which each surface has bounded diameter. In order to do this, we use a more general notion of sweepout, called a generalized sweepout. Recall that a negatively curved surface has a thick-thin decomposition, in which the thick part has bounded diameter, so roughly speaking we wish to cut off the thin part of the sweepout surfaces to obtain bounded diameter surfaces. As the resulting surfaces may become disconnected, it will be convenient to work with the following definition of diameter.

Definition 3.1. ϵ -diameter. Let S be a set in a metric space M . The ϵ -**diameter** of S is the minimum number of balls of radius ϵ needed to cover S . \diamond

The advantage of this definition is that a surface which has two connected components, each of small diameter, will also be considered to have small diameter, even if the two components are far apart.

We now define a generalised sweepout, and give a method for construction one with bounded ϵ -diameter. Some of the following definitions and constructions are similar to those of Maher and

Rubinstein [13], but we reproduce them for the convenience of the reader. An important difference here is that the sweepout surfaces are not required to be embedded.

Definition 3.2. Generalised sweepouts.

A **generalised sweepout** of a 3-manifold M is a triple (Σ, f, h) , where

- Σ is an orientable 3-manifold.
- The map $h : \Sigma \rightarrow \mathbb{R}$ is a Morse function, which is constant on each boundary component of Σ , such that $h^{-1}(t)$ is a collection of surfaces, for all but finitely many t .
- The smooth map $f : (\Sigma, \partial\Sigma) \rightarrow (M, \partial M)$ is degree one, i.e. $f_* : H_3(\Sigma, \partial\Sigma) \rightarrow H_3(M, \partial M)$ is an isomorphism.

We will often write a sweepout as (Σ, ϕ) , where ϕ denotes the map $(f \times h) : \Sigma \rightarrow M \times \mathbb{R}$. We will think of $t \in \mathbb{R}$ as the time coordinate. ◇

We may think of the sweepout surfaces as a one-parameter family of immersed surfaces in M . There may be critical times at which the genus or number of components of the surface changes. We may change the sweepout in $B^3 \times I \subset M \times \mathbb{R}$ as follows. Replace $S_t \cap B^3$ with a different continuously varying family of surfaces, so that the new family agrees with the original sweepout on $\partial(B^3 \times I)$. This produces a new generalised sweepout as changing the map in $B^3 \times I$ doesn't change the degree of the map. We now give a precise definition.

Definition 3.3. Modifying sweepouts.

Let (Σ, ϕ) be a generalised sweepout. Let N be a 3-dimensional submanifold of Σ whose boundary is disjoint from the critical points of h . Let N' be a 3-manifold whose boundary is homeomorphic to ∂N , and let $\phi' : N' \rightarrow M \times \mathbb{R}$ be a map which is the same as the map ϕ on ∂N . Furthermore, assume $\pi_{\mathbb{R}} \circ \phi'$ is a Morse function on N' . Construct a new manifold from Σ by cutting out N and replacing it with N' , i.e. $\Sigma = (\Sigma - N) \cup N'$. If $\pi_M \circ \phi : N \rightarrow M$ is still degree one, then Σ' is a new generalised sweepout which we say is a **modification** of the original one. ◇

In fact, we will only use one type of modification, which will be to replace a one parameter family of annuli with a family of surfaces which start off as annuli, then get pinched into a pair of discs, and finally get pasted back to the original annulus, as illustrated below.




Figure 1: Pinching off a family of annuli.

The one parameter family of annuli is a solid torus in Σ , and this modification corresponds to doing $(0, 1)$ surgery on this solid torus. This does not change the degree of the map, as the solid torus represents zero in $H_3(\Sigma, \partial\Sigma)$, so the modified map is still a sweepout.

We wish to take the original simplicial sweepout and modify it to produce a generalised sweepout in which each sweepout surface S_t has bounded ϵ -diameter.

The simplicial sweepout surfaces S_t are composed of triangles of curvature at most -1 , with at most one vertex of angle sum less than 2π , which we shall call x_t . Let \bar{S}_t be the completion of the universal cover of $S_t - x_t$. This is a simply-connected 2-complex which is a union of triangles of curvature at most -1 , with all cone angles at vertices greater than 2π , so it is a complete CAT(-1) geodesic metric space, with the following well known properties, see Bridson and Haefliger [3].

- Geodesics are unique, and hence convex.
- If C is convex, then $N_r(C)$ is convex.
- If C is convex, then the closest point projection map onto C is distance decreasing.
- Isometries are either elliptic, parabolic or hyperbolic.

We remind the reader of the classification of isometries of CAT(-1) spaces. The **translation distance** $D(\phi)$ of an isometry ϕ of a metric space (X, d) is $\inf\{d(x, \phi(x)) | x \in X\}$. The **min set** of ϕ is $\{x \in X | d(x, \phi(x)) = D(\phi)\}$, which may be empty. Elliptic isometries have fixed points, i.e. $D(\phi) = 0$ and the min set is non-empty, hyperbolic isometries have $D(\phi) > 0$ and non-empty min sets, called axes, and parabolic isometries have empty min sets.

For each homotopy class α in $S_t - x_t$ there is a corresponding covering translation of the universal cover, which gives rise to an isometry of \bar{S}_t . These isometries may be elliptic or hyperbolic, but not

parabolic, as the the completion of a preimage of a fundamental domain for $S_t - x_t$ is compact in the universal cover. Let γ_t be the set of points which are moved the least distance by the isometry. These either form a geodesic which is the axis of a hyperbolic isometry, or consist of a single fixed point, if the isometry is elliptic. γ_t is the shortest length loop in the homotopy class. If γ_t consists of a single point, then we consider it to be the constant loop of length zero. γ_t is a piecewise geodesic associated to α , and is homotopic to α if it is disjoint from x_t . Otherwise is it is homotopic to α by an arbitrarily small perturbation off x_t . By abuse of notation we will call γ_t the geodesic representative of α .

The surfaces S_t are made from negatively curved triangles, which change continuously with t , so the geodesic representatives of homotopy classes change continuously with t as well. We now give a detailed proof of this. The basic point is that for surfaces which are close together in time, there is a quasi-isometry between them which is close to an isometry. A geodesic in one surface gets mapped to a quasigeodesic in the other surface, and it is well known [3] that a quasigeodesic lies in a K -neighbourhood of a geodesic for some K . It suffices to show that by choosing the times to be sufficiently close together, we may choose K to be as small as we like.

Lemma 3.4. *Let γ be a simple closed curve in $S - x$. The geodesic representatives γ_t of γ vary continuously with t .*

Proof. It suffices to consider S_{t_1} and S_{t_2} with t_1 and t_2 close together. So we may assume there is at most one change of triangulation between t_1 and t_2 . If there is a change of triangulation at time t_3 , we can break the interval up into two parts $[t_1, t_3]$ and $[t_3, t_2]$, where on each subinterval the surfaces have the same triangulation, with a possibly degenerate metric on some of the triangles. If things are continuous on both closed intervals, then they are continuous on the whole interval.

Let γ be a simple closed curve in $S_t - x_t$, and let γ_t be the geodesic representative of γ at time t .

We want to show that for all $\epsilon > 0$ there is a $\delta > 0$ such that if $|t_1 - t_2| < \delta$ then γ_{t_1} lies in $N_\epsilon(\gamma_{t_2})$.

There is a map $\phi : S_{t_1} \rightarrow S_{t_2}$ which is a $(l, 0)$ -quasi-isometry, with $l \rightarrow 0$ as $t_1 \rightarrow t_2$. So $\phi(\gamma_{t_1})$ is a $(l, 0)$ -quasigeodesic in S_{t_2} . It is a standard result of [3] that therefore $\phi(\gamma_{t_1})$ lies in a $N_K(\gamma_{t_2})$ for some K . We wish to show that by choosing t_1 close to t_2 , and hence l close to 1, we can make K as small as we like.

Suppose that $\phi(\gamma_{t_1})$ is not contained in a K -neighbourhood of γ_{t_2} . Then there must be a segment of $\phi(\gamma_{t_1})$ of length at least K lying outside a $K/2$ -neighbourhood of γ_{t_2} . Let x and y be the endpoints of this segment, so $d_\phi(x, y) \geq K$. Let $d_{t_2}(x, y)$ be the length of the geodesic segment between them in S_{t_2} . Let $d_N(x, y)$ be the distance between x and y along the boundary of the $K/2$ neighbourhood. This is illustrated in Figure 2 below.




Figure 2: A $K/2$ -neighbourhood of a geodesic.

A lower bound for the quasigeodesic constant is given by $K/2d(x, y)$, and the largest possible value of $d(x, y)$ occurs when the distance along $N_{K/2}(\gamma)$ is $K/2$. If $K > 0$, then the curvature of the boundary segment is also greater than 0, so $d(x, y) < K/2$. So $l > K/2d(x, y) > 1$, and $K/2d(x, y)$ tends to 1 as $K \rightarrow 0$. Therefore by choosing l to be close to 1, we can make $K/2d(x, y)$ as close to 1 as we like, forcing K to be as small as we like. \square

Definition 3.5. A simple closed curve in $S_t - x_t$ is **short** if the length of its geodesic representative is less than ϵ , the injectivity radius of M . \diamond

Let Γ_t be the collection of short simple closed curves. There are only finitely many short simple closed curves, and there is always at least one, as the loop around x_t has length zero.

We now define some special neighbourhoods of the short simple closed curves $\gamma_t \in \Gamma_t$ as follows. Consider a connected component $\tilde{\gamma}_t$ of the pre-image of γ_t in \bar{S}_t , and choose a transverse orientation for $\tilde{\gamma}_t$, so that the distance function from $\tilde{\gamma}_t$ is signed. If γ_t has length zero, then assume that distances are positive. For a connected interval $[a, b]$, let $\tilde{N}_{[a, b]}(\tilde{\gamma}_t)$ be the set of points q in \bar{S}_t with $a \leq d(q, \tilde{\gamma}_t) \leq b$. Let $N_{[a, b]}(\gamma_t)$ be the projection of $\tilde{N}_{[a, b]}(\tilde{\gamma}_t)$ into S_t . We will write $N_{[r]}$ if the interval consists of a single point r .

Definition 3.6. Annular and surgery neighbourhoods. Define an **annular neighbourhood** of $\gamma_t \in \Gamma_t$ as follows. Let $\mathcal{A}(\gamma_t)$ be the maximal $N_{[a, b]}(\gamma_t)$ such that each $N_{[r]}(\gamma_t)$ is embedded of length at most ϵ . $\mathcal{A}(\gamma_t)$ is non-empty as it contains $N_{[0]}(\gamma_t)$. The annular neighbourhoods $\mathcal{A}(\gamma_t)$ vary continuously with t .

Given an annular neighbourhood $\mathcal{A}(\gamma_t) = N_{[a, b]}(\gamma_t)$ of γ_t , define an annular **surgery neighbourhood** $\mathcal{E}(\gamma_t)$ of γ_t to be the subset of $\mathcal{A}(\gamma_t)$ of curves $N_{[r]}(\gamma_t)$ which are at least a distance $\epsilon/2$ from the boundary of $\mathcal{A}(\gamma_t)$. So $\mathcal{E}(\gamma_t) = N_{[a+\epsilon/2, b-\epsilon/2]}(\gamma_t)$, with the convention that this is the empty set if $b - a < \epsilon$. The surgery neighbourhood $\mathcal{E}(\gamma_t)$ varies continuously with t , but need not contain γ_t . \diamond




Figure 3: Annular and surgery neighbourhoods.

Lemma 3.7. *Let α_t and β_t be the geodesic representatives of distinct short curves in the sweepout surface S_t . Then their surgery neighbourhoods $\mathcal{E}(\alpha_t)$ and $\mathcal{E}(\beta_t)$ are disjoint.*

Proof. Suppose $\mathcal{E}(\alpha_t)$ and $\mathcal{E}(\beta_t)$ intersect. Then there is a boundary curve $N_{[r]}(\alpha_t)$ of $\mathcal{E}(\alpha_t)$ which intersects $\mathcal{E}(\beta_t)$. However the boundary curve $N_{[r]}(\alpha_t)$ has length at most ϵ , so it must therefore be contained in $\mathcal{A}(\beta_t)$. But this means $N_{[r]}(\alpha_t)$ is either null-homotopic, or homotopic to the other short loop β_t , a contradiction. \square

As the surgery neighbourhoods are disjoint, there are at most $2g - 1$ such neighbourhoods.

We now show how to create a new generalized sweepout \hat{S}_t from the original sweepout S_t by cutting off the annular surgery neighbourhoods and replace them with discs, using the modification described above.

First we describe how to remove a given surgery neighbourhood $\mathcal{E}(\gamma_t)$. Let $[a, b]$ be a maximal time interval on which $\mathcal{E}(\gamma_t)$ is non-empty, so $\mathcal{E}(\gamma_t)$ consists of a single curve $N_{[x]}(\gamma_a)$ at time a , and a single curve $N_{[y]}(\gamma_b)$ at time b . The union of the surgery annuli $\mathcal{E}(\gamma_{[a,b]}) = \{\mathcal{E}(\gamma_t) | t \in [a, b]\}$ is a solid torus in Σ . We now describe how to produce a new sweepout, identical to the first one outside $E(\gamma_{[a,b]})$, and which corresponds to doing $(0, 1)$ surgery on Σ . Choose a continuously varying family of basepoints for the boundary components of $\mathcal{E}(\gamma_t)$, which agree at times a and b .

First alter the sweepout by expanding times a and b to short intervals I_a and I_b , for which the map ϕ is constant, as illustrated in Figure 4 below.




Figure 4: Inserting short constant intervals.

In the interval I_a do a cut move which replaces the loop $\mathcal{E}(\gamma_a)$ by a pair of discs coned from the basepoint. The coned discs are ruled surfaces, so can be developed into \mathbb{H}^2 , so their area is bounded above by the area of a disc with perimeter the same length as the curve $\mathcal{E}_i(a)$, namely ϵ , so each disc has area at most $\epsilon/4$, and is contained in a ball of radius ϵ . In the interval (a, b) replace the annulus bounded by $\partial\mathcal{E}_i$ by the discs on each boundary component formed by coning each boundary component to its basepoint. Finally in the interval I_b paste the two coned discs back to the original surface.




Figure 5: Doing a cut move in I_a .

We can do this modification for each surgery annulus $\mathcal{E}(\gamma_t)$, and so remove all the surgery annuli from the sweepout, to produce a new sweepout which we shall call \hat{S}_t . As there are at most $2g - 1$ short annuli, the new sweepout surface \hat{S}_t also has area at most $4\pi g + 2(2g - 1)\epsilon/4$.

We now recall some useful facts about Voronoi decompositions of surfaces. Let x_i be a collection of points in a metric space (X, d) . Let $V_i = \{x \in X | d(x, x_i) \leq d(x, x_j) \text{ for all } j\}$. We say that V_i is the **Voronoi decomposition** of X with respect to x_i . If we have chosen a maximal collection of points x_i so that no two are closer together than ϵ , we say we have a **maximal ϵ -spaced Voronoi decomposition**.

Let S be a surface with a metric, which may be degenerate, so that the induced metric on the universal cover is CAT(-1). It will be useful to know that if we have chosen a maximal ϵ -spaced Voronoi decomposition for S , then every Voronoi region which is not a disc has boundary components which are essential in S .

Lemma 3.8. *Let S be a surface with a metric, which may be degenerate, so that the induced metric on the universal cover is CAT(-1), and let V_i be a maximal ϵ -spaced Voronoi decomposition for S . Then*

1. $B_{\epsilon/2}(x_i) \subset V_i \subset B_\epsilon(x_i)$.
2. A Voronoi region which is not a disc has boundary components which are essential in S .

Proof. For the first claim, note that any two points x_i and x_j are a distance at least ϵ apart, so if a point y is at most $\epsilon/2$ from x_i , then y must be at least as far from any other x_j , so $B_{\epsilon/2}(x_i)$ is contained in V_i . If any point in V_i is a distance greater than ϵ from x_i , then it must also be a distance at least ϵ from all of the other V_j , so we could add it to our collection x_i to produce a larger collection of points distance at least ϵ apart, contradicting maximality.

We now prove the second claim. Then $B_{\epsilon/2}(x_i) \subset V_i \subset B_\epsilon(x_i)$, and if a Voronoi region V_i is not a disc, then its boundary components are essential in S . This can be seen as follows. The boundary of V_i in S lies in $B_\epsilon(x_i) - B_{\epsilon/2}(x_i)$. If some boundary component of V_i is inessential, then it lifts to a simple closed curve in the universal cover, which lies in $B_\epsilon(\tilde{x}_i) - B_{\epsilon/2}(\tilde{x}_i)$, for some pre-image \tilde{x}_i of x_i . This closed curve bounds a disc, and as the distance function $d(x, \tilde{x}_i)$ is convex, it has maxima only on the boundary, so the entire disc is contained in $B_\epsilon(\tilde{x}_i)$. However the disc must contain at least one other pre-image of some x_j . This can't be a pre-image of x_i , as then there would be an essential loop in S intersecting the inessential simple closed once, so it must be a pre-image of some x_j with $j \neq i$. However this implies that there is an x_i and an x_j with $d(x_i, x_j) < \epsilon$, a contradiction. \square

Given a surface with a Voronoi decomposition, there is a one-vertex triangulation with edge lengths bounded in terms of the number and size of the Voronoi regions. The surface may have an arbitrary metric.

Lemma 3.9. *Let S be a closed surface, with a Voronoi decomposition V_i with N regions, such that each V_i is contained in a ball of radius ϵ . Then the surface has a one-vertex triangulation in which each edge has length at most $2\epsilon N$.*

Proof. Let \tilde{V}_i be the Voronoi decomposition of the universal cover \tilde{S} of S , corresponding to the pre-images of the points x_i in S . Let $\tilde{\Gamma}$ be the Delaunay graph dual to the Voronoi decomposition of the universal cover. We may choose $\tilde{\Gamma}$ to be equivariant, and to have edges of length at most 2ϵ . Let Γ be the projection of $\tilde{\Gamma}$ into S . The complementary regions of Γ are all discs, but Γ need not be a triangulation, as it may have complementary regions that have more than three sides.

Complementary regions with more than three sides correspond to points at which more than three distinct \tilde{V}_i meet. So any pair of pre-images \tilde{x}_i and \tilde{x}_j are distance at most 2ϵ apart, so we can add extra edges to triangulate each n -gon, to produce a triangulation in which every edge has length at most 2ϵ .




Figure 6: Add extra arcs to make a triangulation.

This produces a triangulation of S with the same number of vertices as the number of Voronoi regions, and in which every edge has length at most 2ϵ . So we can shrink a maximal tree to a point, and throw away extra edges, to produce a one vertex triangulation in which every edge has length at most $2\epsilon N$. \square

Lemma 3.10. *There is a constant $N(\epsilon, g)$, which only depends on the genus g of S_t and the injectivity radius ϵ of the manifold M , so that the surface \hat{S}_t has*

- ϵ -diameter at most N .
- A one-vertex triangulation in which each edge has length at most $3\epsilon N$.

We may take N to be $12g - 10 + 4g/\sinh^2(\epsilon/8)$.

Proof. Choose a maximal $\epsilon/2$ -spaced Voronoi decomposition for S_t , so that $x_1 = v_t$. We now show that there is an upper bound on the number of Voronoi regions. As before, we find bounds on the number of Voronoi regions which are discs, annuli, or have negative Euler characteristic.

If a Voronoi region V_i is a disc, with $i \neq 1$, then by 3.8.1 it contains a ball of radius $\epsilon/4$, so must have area at least $\pi \sinh^2(\epsilon/8)$, so the area of the surface is at most $4\pi g$, there may be at most $4g/\sinh^2(\epsilon/8) + 1$ Voronoi regions which are discs.

Voronoi regions with negative Euler characteristic have essential boundary components, so there may be at most $2g - 2$ of these.

If a Voronoi region is an annulus, then it contains a simple closed curve of length at most ϵ . So there are at most $2g - 1$ sets of annuli such that the simple closed curves in each set of annuli are parallel. The simple closed curves are at least $\epsilon/2$ apart, so if there are three annuli in a row, the center curve must be contained in \mathcal{E} , so there may be at most four annuli for each short curve not contained entirely in \mathcal{E} , so there are at most $8(g - 1)$ annular Voronoi regions intersecting $S_t - \mathcal{E}$.

Each Voronoi region is contained in a single ϵ -ball in M , and we add at most $2g - 1$ extra discs to make \hat{S}_t from $S_t - \mathcal{E}$, each of which is contained in a single ϵ -ball, so the ϵ -diameter of \hat{S}_t is at most $N = 4g/\sinh^2(\epsilon/8) + 1 + 2g - 2 + 8(g - 1) + 2g - 1$ as required.

We now show how to construct a triangulation with bounded edge length. Choose a Voronoi decomposition using the points x_i which lie in $S_t - \mathcal{E}$, the above argument shows that there may be

at most N such points. Each point in \hat{S}_t is distance at most $\epsilon/2$ from $S_t - \mathcal{E}$, and hence distance at most $3\epsilon/2$ from an x_i in $S_t - \mathcal{E}$. Therefore the resulting Voronoi decomposition of \hat{S}_t has each $\hat{V}_i \subset B_{3\epsilon/2}(x_i)$. The surface \hat{S}_t now has a triangulation of length $3\epsilon N$ by Lemma 3.9. \square

4 Heegaard gradient and virtual fibers.

We briefly summarize of the remainder of the argument to complete the proof of Theorem 1.1. If the volume of the covers grows, and χ_i/d_i tends to zero, but $c^+(M_i)$ remains bounded, then the volumes of the compression bodies in a Scharlemann-Thompson thin position of the manifold must become arbitrarily large. The compression bodies have bounded genus, and hence are swept out by generalised sweepout surfaces of bounded ϵ -diameter. As the sweepout surfaces have bounded ϵ -diameter, they can only intersect finitely many fundamental domains, so eventually there must be many surfaces intersecting the same pattern of fundamental domains in different parts of a cover. Each surface is made from compressing the higher genus boundary of the compression body along discs, and if they are disjoint and nested in the compression body, then they must be homotopic. We use Gabai's result that the singular norm on homology is equal to the Thurston norm to replace the immersed surfaces by embedded surfaces, and then show we can find many parallel embedded surfaces. As the surfaces have bounded diameter, then if there are enough surfaces we can find a pair which intersect matching sets of fundamental domains. We can therefore cut along the matching sets of fundamental domains and glue the ends of the resulting piece of the cover together to form a fibered manifold. We now fill in the details of this argument.

If the Heegaard gradient of the covers χ_i/d_i tends to zero, then the volume of the covers grows faster than the number of handles in the given Heegaard splittings of the covers. The number of compression bodies in the Scharlemann-Thompson thin position for the splitting is bounded by the number of handles, so there must be a sequence of compression bodies in the covers whose volumes become arbitrarily large.

The components of ∂_- of a compression body consists of a collection of incompressible surfaces, which may be empty, and the ∂_+ boundary component consists of a strongly irreducible Heegaard splitting for the union of the two adjacent compression bodies. We need to be able to homotop the boundary to a bounded ϵ -diameter surface close to the original surface, and we will use the fact that the boundary components of a compression body in a thin-position Heegaard untelescoping are closely related to minimal surfaces, which have bounded diameter.

By Freedman, Hass and Scott [6], the incompressible surfaces are either isotopic to minimal surfaces, or double cover one-sided minimal surfaces in M . In the latter case, this would mean that two adjacent compression bodies formed a strongly irreducible splitting of a twisted I -bundle over a non-orientable surface. Such a splitting would lift to a Heegaard splitting of the double cover which is an I -bundle over an orientable surface, and these splittings are classified by Scharlemann and Thompson [16]. Given their classification, it is easy to see that the original splitting could not be strongly irreducible. Furthermore, the incompressible surfaces in the untelescoping are all disjoint, so their minimal surface representatives are also all disjoint, as minimal surfaces may only have essential intersections.

A strongly irreducible Heegaard splitting is either homotopic to a minimal surface, or is homotopic to a minimal surface union some number of arcs, by Pitts and Rubinstein [14]. We may use the incompressible minimal surfaces as barrier surfaces for the sweepout, so we may assume that the

arcs lie in the union of the compression bodies adjacent to the Heegaard splitting surface. As the Heegaard surface is strongly irreducible, the arcs may lie on only one side. If the arcs lie outside the compression body, then we may compress them along discs disjoint from ∂_- to produce a new compression of lower genus, but similar volume. If the arcs lie on the inside of the compression body, we can remove small open neighbourhoods of the arc so that ∂_+ still has curvature at most -1 , without changing the volume of the compression body very much.

We wish to homotop the sweepout to a simplicial sweepout, and we would like to do this in a complete manifold of sectional curvature at most -1 . We can construct such a manifold by starting with the compression body, and for each boundary component S , gluing on a copy of $S \times [0, \infty)$ with a warped product metric. By the formula for sectional curvature in a warped product (e.g. see Bishop and O'Neill [2, page 26]), given a surface S of Gauss curvature at most -1 , there is a warped product metric on $S \times [0, \infty)$ so that the metric is complete, with sectional curvature at most -1 . This produces a manifold homeomorphic to the interior of the compression body, with a metric of varying negative curvature at most -1 . We can perturb this metric to make it smooth.

We will show a minimal surface has a triangulation of bounded edge length, where the bound depends on the injectivity radius of the manifold and the genus of the surface. This implies we can straighten the minimal surfaces to simplicial surfaces in a bounded neighbourhood of the original minimal surfaces. If the minimal surface corresponding to the Heegaard surface has extra arcs, then we can triangulate the surface so that any edge in the triangulation runs over the arc at most once, so that the part of the surface corresponding to the arcs collapses down to a geodesic arc, and the rest of the surface has bounded diameter.

At various points we will wish to replace an immersed surface by an embedded surface. An immersed surface S contains embedded surfaces in a regular neighbourhood of S in the same homology class. There is a fundamental result of Gabai [7] that the singular norm on homology is the same as the Thurston norm, so we may choose the embedded surface to have genus at most the genus of the immersed surface S .

We now show that we can homotop a minimal surfaces to a simplicial surface, which is not too far away in the manifold. This fact follows from the proof of Proposition 6.1 in Lackenby [10], however, for the convenience of the reader, we present a proof, relying only on the following bound from [10].

Lemma 4.1. *[10, proof of Proposition 6.1, claim 1] Let F be a minimal surface in a Riemannian manifold of curvature at most -1 , and let α be a simple closed curve homotopic to a geodesic γ . If α is distance at least $L/2\pi + 1/2$ from γ then the length of α is at least L .*

We will show that a minimal surface has bounded ϵ -diameter, and a triangulation of bounded edge length.

Lemma 4.2. *Let S be a minimal surface in a closed Riemannian manifold M of curvature at most -1 . Then there is a constant N which depends only on the genus g of S and the injectivity radius ϵ of M , so that S has*

- ϵ -diameter at most N .
- a one-vertex triangulation in which each edge has length at most $2\epsilon N$.

We may take N to be $-\chi(21/4 + 3/4\pi + 3/4\epsilon + 2/\sinh^2(\epsilon/4))$.

Proof. Choose a maximal collection of points x_i , such that no two points are closer together than ϵ , and let V_i be the Voronoi regions for the x_i , i.e. $V_i = \{x \in S \mid d(x, x_i) \leq d(x, x_j) \text{ for all } j \neq i\}$. Each point in a Voronoi region is distance at ϵ from the basepoint x_i for the Voronoi region, and each Voronoi region contains all the points distance at most $\epsilon/2$ from the basepoint. Each Voronoi region is either a disc, an annulus, or a surface of negative Euler characteristic, and we now consider each of these possibilities in turn.

Each Voronoi region which is not a disc has essential boundary in S , by Lemma 3.8.2. This means that every Voronoi region of negative Euler characteristic must have boundary components which are essential curves in the surface, and so there may be at most $-3/2\chi(S)$ such pieces.

If a Voronoi region is an annulus, then it contains an essential loop of length at most ϵ , which is homotopic to a geodesic at most a distance $\epsilon/2\pi + 1/2$ away, by Lemma 4.1. This means there may be at most $1/2\epsilon + 5/2 + 1/2\pi$ annuli parallel to any given curve, so there may be at most $-3/2\chi(1/2\epsilon + 5/2 + 1/2\pi)$ annuli in total.

If a Voronoi region is a disc, then it contains a ball of radius $\epsilon/2$ in the surface. By the monotonicity formula for minimal surfaces in 3-manifolds of negative curvature at most -1 , for example, see Choe [4], the area of this disc must be at least $\pi \sinh^2(\epsilon/4)$, so there are at most $-2\chi(S)/\sinh^2(\epsilon/4)$ discs.

So $N = -\chi(21/4 + 3/4\pi + 3/4\epsilon + 2/\sinh^2(\epsilon/4))$ is an upper bound for the total number of Voronoi regions, and hence an upper bound for the ϵ -diameter of the surface. Then by Lemma 3.9, the surface has a triangulation in which each edge has length at most $2\epsilon N$. \square

Given simplicial surfaces homotopic to the components of ∂_- , we can join them together with arcs to form a spine for the compression body. We next show that we can find a simplicial surface homotopic to this spine, by a homotopy which does not sweep out very much volume.

Lemma 4.3. *Let S_1, \dots, S_n be a collection of simplicial surfaces, with basepoints x_i , and of total genus g in a complete Riemannian 3-manifold of curvature at most -1 . Connect the basepoint x_1 to each of the other basepoints by a geodesic arc to form a 2-complex Σ . Then there is a simplicial surface of genus g which is homotopic to Σ , by a homotopy which sweepouts out a volume of at most $3(2g + 2)$ times the volume of an ideal hyperbolic tetrahedron.*

Proof. Use the construction from Lemma 2.5 to homotop the basepoint of each S_i to x_1 , for $i \neq 1$. Each geodesic triangle sweeps out a geodesic prism, which can be triangulated with three tetrahedra. So each triangle sweeps out a volume at most 3 times the volume of an ideal tetrahedron, so the total homotopy sweeps out a volume of at most $3(2g + 2)\Delta$, as there are at most $2g + 2$ triangles. The surfaces all now share a common basepoint, so we can map a simplicial surface of genus g with a one vertex triangulation onto them, so that the vertex goes to the basepoint. \square

These two lemmas show that if we have a compression body in a complete Riemannian manifold of curvature at most -1 , we can construct a sweepout of the compression body, which is simplicial, except possibly in a bounded neighbourhood of the boundary, where the bound depends only on the genus of the boundary.

Definition 4.4. We say a collection of surfaces in a compression body is **nested** if for any pair of surfaces, one of them separates the other from the higher genus boundary of the compression body. \diamond

As the volume of a compression body C becomes large, the diameter of C must also become large. We use this to show there is a cover which contains a compression body with many disjoint

nested sweepout surfaces.

Lemma 4.5. *If the volume of a compression body C in one of the covers M_i is at least $\text{vol}(B_{Kn}(x))$, then there are at least n disjoint nested surfaces, each one of which is formed by compressing the higher genus boundary of the compression body. The constant K depends on ϵ and the genus of the compression body.*

Proof. There is a generalised sweepout S_t of the compression body such that every sweepout surface has ϵ -diameter at most $K(g, \epsilon)$, where g is the genus of the handlebody, and ϵ is the injectivity radius of M . Let A be an upper bound for the volume of a $K\epsilon$ -neighbourhood of a sweepout surface S_t . Let $A = K\text{vol}(B_{K\epsilon+\epsilon})$.

If the diameter of C is d , then there is a point x of distance at least $d/2 - K$ from ∂_+ . Let γ be the geodesic arc from x to ∂_+ . Let γ have length L , which is at least $d/2 - K$, and let $\gamma(l)$ be an arc-length parameterisation of γ , so that $\gamma(0)$ is in ∂_+ .

The sweepout surfaces S_t sweep out the compression body C . More precisely, we say that a point $x \in C - S_t$ is separated from ∂_+ by S_t if any path from x to ∂_+ has algebraic intersection number $+1$ with S_t . Let C_t be the subset of C swept out by S_t , i.e. C_t is the closure of $\{x \in C \mid x \text{ is separated from } \partial_+ \text{ by } S_t\}$. Then C_t starts out as the empty set, and ends up as all of C . Furthermore, C_t varies continuously with t , as S_t varies continuously with t . In particular, any subset of C , such as γ , is also swept out by the sweepout S_t , i.e. $\gamma \cap C_t$ varies continuously with S_t , and $\gamma \cap C_0$ is the empty set, and $\gamma \cap C_1$ is all of γ .

There is a sweepout surface S_t which divides γ into two parts of equal length, i.e. the length of $C_t \cap \gamma$ is half the length of γ . The sweepout surface S_t may have at most g components. Each component of S_t has diameter at most K , so if γ hits some component of S_t , and then travels a distance greater than K past that component of S_t , then it can never return and hit the same component of S_t again. As S_t splits γ in two, there must be some component of S_t which intersects $\{\gamma(l) \mid l \in ([L/4, 3L/4])\}$ which separates γ into two pieces each of length at least $L/4$. Choose this component of S_t to be T_1 . Note that T_1 separates x from ∂_+ , is disjoint from both of them, and is distance at least $d/8 - K$ from x . Let l_1 be the smallest value of l such that $\gamma(l) \in T_1$.

So we can now apply the same argument again, this time using $\{\gamma(l) \mid l \in [0, l_1]\}$ instead of γ . We can continue in this way finding a sequence of n surfaces T_i , until $L/2^{n+1}$ is comparable to K . \square

Therefore, as the volume of a compression body becomes large, we can find many disjoint nested surfaces inside it, formed from compressing the higher genus boundary, so in particular we can find many nested connected surfaces of the same genus.

We next show that if two connected sweepout surfaces are nested in the compression body, then the compressing discs for the outer one may be chosen to be a subset of the compressing discs for the inner one. In particular, if the surfaces have the same genus, then the compressing discs are the same, so they are homotopic.

Lemma 4.6. *Let S_1 and S_2 be two disjoint immersed sweepout surfaces each of which is homotopic to a surface obtained from compressing ∂_+ along collections of discs Δ_1 and Δ_2 respectively. Furthermore suppose that S_2 separates S_1 from ∂_+ . Then we may choose compressing discs Δ'_1 and Δ'_2 for S_1 and S_2 so that the discs Δ'_2 are a subset of the discs Δ'_1 .*

Proof. The sweepout gives a homotopy from S_1 to $\partial_+ \cup \Delta_1$. We may homotop the compressing discs D_1 so they are disjoint and embedded in the compression body C . Similarly we can homotop the

compressing discs Δ_2 so they are embedded and disjoint. Note that two different sets of compressing discs Δ_1 and Δ'_1 can give rise to isotopic surfaces.

We say a properly embedded meridional disc D in C is a compressing disc for S_i if there is a collection of disjoint compressing discs Δ'_i containing D , so that ∂_+ cut along Δ'_i is isotopic to ∂_+ cut along Δ_i .

We first prove the lemma for non-separating discs. Let D be a non-separating meridional disc in C , which is *not* a compressing disc for S_1 . Then there is a curve α in $\partial_+ - \Delta_1$, which has algebraic intersection number one with D , and this is preserved under homotopy, so S_1 must intersect D . As S_2 separates S_1 from ∂_+ , S_2 must also intersect D , and in fact the intersection of S_2 with D must separate the intersections of S_1 in D from ∂D . Furthermore, there must be an essential intersection of D with S_2 , as otherwise discs in S_2 would have to intersect the image of α in S_1 , a contradiction. This means that D may not be a compressing disc for S_2 either.

If D is a separating compressing disc for C , which is *not* a compressing disc for S_1 , then it still has homotopically essential intersection with S_1 . Explicitly, Δ separates ∂_+ into two components, neither of which is planar. Choose curves α and β in each component which are essential in C , and connect them with an arc γ which intersects D once. The resulting trivalent graph, or eyeglass, cannot be homotoped to be disjoint from D . This can be seen by considering a connected pre-image θ in the universal cover, which intersects a particular lift \tilde{D} of D . If the pre-image of the graph could be homotoped off \tilde{D} , then this homotopy would lift to a homotopy of the pre-image in the universal cover, but θ has unbounded components on both sides of \tilde{D} , and the homotopy can only move each point a finite distance, a contradiction. So S_1 intersects D , as must S_2 , as S_2 separates S_1 from ∂_+ . There must be essential intersections of S_2 with D , as if all intersections were inessential then in the universal cover an unbounded part of θ would be separated from the pre-image of ∂_+ by a bounded part of S_2 , a contradiction. So D may not be a compressing disc for S_2 .

So we have shown that a compressing disc for S_2 is also a compressing disc for S_1 , so we may choose a set of compressing discs for S_2 , and then extend this to a set of compressing discs for S_1 \square

We have shown that if the volume of a compression body is sufficiently large, there must be many disjoint nested homotopic surfaces inside it. However, we require the further property that the homotopy from S_n to S_i is disjoint from S_j for $j < i$. We now show that we can find a collection of surfaces with this additional property.

Lemma 4.7. *Let S_1, \dots, S_n be a collection of disjoint, nested, homotopic surfaces in a compression body C . Then there is a collection of surfaces $S'_1, \dots, S'_{n-1}, S_n$ which are disjoint, nested and homotopic, and furthermore the homotopy from S_n to S'_i is disjoint from S'_j for $j < i$.*

Proof. The surfaces S_i have triangulations of bounded length, so they are homotopic to simplicial surfaces of bounded ϵ -diameter, a bounded distance away in M . Therefore there is a homotopy which is a simplicial sweepout turned into a bounded diameter sweepout S'_t from S_n to S_1 . Let S'_i be the first S'_t which hits S_i , so the sweepout from S_n to S'_i is disjoint from S'_j for $j < i$ by construction. Suppose the genus of some S'_i is less than the genus of S_i . As the Thurston norm is equal to the singular norm, there is therefore an embedded surface T'_i of genus lower than S_i separating S_1 from S_n , contradicting Lemma 4.8. So in fact no compressions take place, and the S'_i are all homotopic, and the homotopy from S_n to S'_i is disjoint from S'_j for $j < i$. The S'_i are disjoint as they have bounded diameter, and are far apart. Finally, we show that $S'_1, \dots, S'_{n-1}, S_n$ are nested. Let C_n be the subset of C separated from ∂_+ by S_n . First note that S'_{n-1} intersects S_{n-1} , which is far from S_n , and hence far from ∂C_n , so $S'_{n-1} \subset C_n$, so S'_{n-1} is nested with respect to S_n .

Similarly, S'_{n-k} is contained within C_n , so for any point x in C_{n-k} there is a path γ from x to ∂_+ with algebraic intersection number 1 with S_n . Now S'_{n-k+1} is homotopic to S_n by a homotopy which is disjoint from C'_{n-k} , so this does not change the algebraic intersection number, so $S'_{n-k} \subset C'_{n-k+1}$, as required. \square

As the singular norm on homology is equal to the Thurston norm on homology, we can replace an immersed surfaces S_i with an embedded surface contained in a regular neighbourhood of S_i , in the same homology class, and with genus at most the genus of S_i . We now show that for S_i with $i > 1$, the new embedded surfaces must in fact have the same genus as the original immersed surfaces.

Lemma 4.8. *Let S be a connected immersed surface in a compression body, which is homotopic to ∂_+ compressed along some collection of discs. Let T be a least genus connected embedded surface which separates S from ∂_+ . Then T is incompressible in $C - S$, and the genus of T is at least as big as the genus of S .*

Proof. As T is an embedded surface in a compression body, it is separating, and it can be compressed to a surface parallel to components of ∂_- by a sequence of compressions along embedded discs. If any compressing disc is disjoint from S then we can compress along the disc to reduce the genus of T , so T is incompressible in $C - S$.

The surface S is homotopic to ∂_+ compressed along some collection of discs D say, so $C - D$ is a compression body C' , and we can choose a spine Γ for C' which can be homotoped into S . The map on first homology induced by inclusion $H_1(\Gamma) \rightarrow H_1(C)$ is injective. As T bounds discs on one side only, T bounds a compression body C'' in C . Consider the maps induced by inclusion $H_1(\Gamma) \rightarrow H_1(C'') \rightarrow H_1(C)$. This composition must be injective, so the rank of $H_1(C'')$ is at least as big as the rank of $H_1(\Gamma)$, so by Poincaré-Lefschetz duality, the rank of $H_1(T)$ is at least as big as the rank of $H_1(S)$, which implies the genus of T is at least as big as the genus of S , as required. \square

If we choose S to be S_1 , and T to be an embedded surface corresponding to S_i , with $i > 1$, then the lemma above shows that the genus of the embedded surface is the same as genus of the immersed surface S_i , and furthermore, that the embedded surface is incompressible in the complement of S_1 .

Any two embedded surfaces T_i and T_j , with $1 < i < j < n$ bound a submanifold in the compression body C , and we now show that any homotopy between S_1 and S_n sweeps out this submanifold in a degree one manner.

Lemma 4.9. *Let S_1, \dots, S_n be a collection of disjoint connected nested immersed homotopic surfaces in a compression body C , and let $T_i \subset N(S_i)$ be embedded surfaces in the same homology class. Let X be the part of C between T_{n-1} and T_2 . Then any homotopy $\phi : S \times I \rightarrow C$ from S_1 to S_n is degree one onto X , i.e. $\phi_* : H_3(S \times I, \partial) \rightarrow H_3(X, \partial X)$ is an isomorphism.*

Proof. The map $\phi : S \times I \rightarrow C$ gives a map of pairs $\phi : (S \times I, \partial) \rightarrow (C, C - X)$, and by excision, $H_3(C, C - X_2) \cong H_3(X, \partial X)$. Continuous maps between pairs induce natural homomorphisms between the homology long exact sequences of pairs, so the following diagram commutes.

$$\begin{array}{ccc} H_3(S \times I, \partial) & \xrightarrow{\partial} & H_2(S \times 0 \cup S \times 1) \\ \phi_* \downarrow & & \phi_* \downarrow \\ H_3(C, C - X) & \xrightarrow{\partial} & H_2(C - X) \end{array}$$

The group $H_3(S \times I, \partial)$ is generated by $[S \times I]$, which is mapped to $[S \times 1] - [S \times 0]$ by the boundary map, and ϕ_* maps this to $[S_n] - [S_1]$, which is homologous to $[T_{n-1}] - [T_2]$. If there is no component of ∂_- contained in the component of $C - X$ containing T_2 , then $[T_2]$ will be zero. However, $[T_{n-2}]$ is not zero, as it lies in the same component of $C - X$ as ∂_+ .

The group $H_3(C, C - X) \cong H_3(X, \partial X)$ is generated by $[X]$, which gets mapped to $[T_{n-1}] - [T_2]$ by the boundary map, so $[S \times I]$ must get mapped to $[X]$, as required. \square

We next show that the embedded surfaces are in fact homotopic to the original immersed surfaces.

Lemma 4.10. *Let S_1 and S_n be disjoint immersed surfaces in a compression body C , which are homotopic, by a homotopy $\phi : S \times I \rightarrow C$. Let T be an embedded surface which separates S_1 from S_n , which is incompressible in $C - (S_1 \cup S_2)$, and for which the map $\phi_* : H_3(S \times I, \partial) \rightarrow H_3(N(T), \partial)$ is degree one, where $N(T)$ is a regular neighbourhood of T . Then T is homotopic to S_t .*

Proof. Let $\phi : S \times I \rightarrow M$ be a smooth homotopy from S_1 to S_n . The embedded surface T is contained in the image of ϕ , and T separates S_1 from S_n , so $\phi^{-1}(T)$ is an embedded surface in $S \times I$ which separates $S \times 0$ from $S \times 1$. Furthermore the map ϕ is degree one from $\phi^{-1}(T)$ to T .

Suppose the surface $\phi^{-1}(T)$ is compressible in $S \times I$. By the loop theorem, we may assume the compressing discs are embedded. Let D be a maximal collection of disjoint embedded compressing for $\phi^{-1}(T)$. Construct a map of a surface $\psi : F \rightarrow \phi^{-1}(T) \cup D$ as follows. Take the closure of $\phi^{-1}(T) - D$ with the induced path metric to produce a surface F' with $2|D|$ boundary components. Now glue in a disc along each boundary component, and let F be the resulting surface. We now define a map from F to $\phi^{-1}(T) \cup D$. Send each point in F' to the corresponding point in $\phi^{-1}(T)$. Send each disc in $F - F'$ to the corresponding compressing disc in $S \times I$. This is illustrated below in Figure 7.



Figure 7: A map from F to $\phi^{-1}(T) \cup D$.

The surface T is incompressible in the complement of S_1 , so the images of the discs D in M are all inessential, and so they are homotopic into T , as M is irreducible. So the map $\phi \circ \psi : F \rightarrow M$ is homotopic onto the image of T , and is still degree one, as if it were degree zero, then there would be a path in C from S_1 to S_n with zero algebraic intersection number with $\phi \circ \psi(F)$, which would pull back to a path from $S \times 0$ to $S \times 1$ with zero intersection number with $\psi(F)$, but $\psi(F)$ separates $S \times 0$ from $S \times 1$, a contradiction. This means that F is in fact homotopic to T in C . Furthermore, $\psi(F)$ is an incompressible surface in $S \times I$, so is homotopic to $S \times I$, hence $\phi \circ \psi(F)$ is homotopic to S_t in C , as required. \square

We now show that the embedded surfaces T_i are not just homotopic, but actually parallel in the compression body C .

Lemma 4.11. *Let T_{n-1} and T_3 be embedded incompressible surfaces contained in the image of a homotopy ϕ from S_n to S'_2 , such that the image of ϕ is disjoint from S'_1 . Then T_{n-1} and T_3 are parallel.*

Proof. Let X be the region bounded by T_{n-1} and T_3 . Consider $\phi^{-1}(X)$ in $S \times I$.




Figure 8: $\phi^{-1}(X)$

Suppose some component of $\phi^{-1}(X)$ has compressible boundary. Then we can change the homotopy ϕ in a neighbourhood of the compressing disc to compress the boundary. If any component of $\phi^{-1}(X)$ is a 3-ball, we can change the homotopy in a neighbourhood of the 3-ball to remove it. The remaining components of $\phi^{-1}(X)$ must all have boundary components isotopic to $S \times t$. We can change the product structure on $S \times I$ so that these boundary components are horizontal, and then there is a subset of I which gives a homotopy from T_{n-1} to T_3 entirely contained in X . So now X is a 3-manifold with two incompressible boundary components, which are homotopic inside X , so X has the same fundamental group as a surface, so X is a product $S \times I$, by Waldhausen [17]. \square

Lemma 4.12. *If there are enough disjoint parallel surfaces then the manifold is virtually fibered.*

Proof. Each parallel surface has bounded diameter, so is contained in finitely many fundamental domains, and we can choose parallel surfaces so that the fundamental domains they hit are disjoint. There are only finitely many ways of gluing finitely many fundamental domains together, so if there are enough parallel surfaces, there must be at least three which hit the same pattern of fundamental domains, so we cut the manifold along a pair of fundamental domains with compatible orientation, and glue back together to get a fibered cover of the original manifold. \square

This completes the proof of Theorem 1.1.

References

- [1] D. Bachman, D. Cooper, M. White, Large embedded balls and Heegaard genus in negative curvature, Preprint, math.GT/0305290.

- [2] R. L. Bishop, B. O’Neill, Manifolds of negative curvature, *Trans. Amer. Math. Soc.* **145** (1969) 1–49.
- [3] M. Bridson, A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften, 319, Springer-Verlag, Berlin, 1999.
- [4] J. Choe, The isoperimetric inequality for minimal surfaces in a Riemannian manifold, *J. reine angewandte Mathematik* **506** (1999), 495–503.
- [5] T. H. Colding, C. De Lellis, The min–max construction of minimal surfaces, Preprint, math.AP/0303305.
- [6] M. Freedman, J. Hass, P. Scott, Least area incompressible surfaces in 3-manifolds, *Invent. Math.* **71** (1983), no. 3, 609–642.
- [7] D. Gabai, Foliations and the topology of 3-manifolds, *J. Differential Geom.* **18** (1983), no. 3, 445–503.
- [8] A. Hatcher, On triangulations of surfaces, *Top. Appl.* **40** (1991), no. 2, 189–194.
- [9] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [10] M. Lackenby, Heegaard splittings, the virtually Haken conjecture and Property (τ) , Preprint, math.GT/0210316.
- [11] M. Lackenby, The asymptotic behaviour of Heegaard genus, Preprint, math.GT/0205327.
- [12] A. Lubotzky, *Discrete groups, expanding graphs and invariant measures*, Progress in Mathematics 125, Birkhauser, 1994.
- [13] J. Maher, J. H. Rubinstein, Period three actions on the three sphere, *Geometry and Topology*, **Vol. 7** (2003), 329–397.
- [14] J. Pitts, J. H. Rubinstein, Existence of minimal surfaces of bounded topological type in three-manifolds, *Miniconference on geometry and partial differential equations (Canberra, 1985)*, 163–176.
- [15] M. Scharlemann and A. Thompson, Thin position for 3-manifolds, *Geometric topology (Haifa, 1992)*, 231–238, *Contemp. Math.*, 164.
- [16] M. Scharlemann and A. Thompson, Heegaard splittings of $(\text{surface}) \times I$ are standard, *Math. Ann.* **295** (1993), no. 3, 549–564.
- [17] F. Waldhausen, On irreducible 3-manifolds which are sufficiently large, *Ann. of Math.* **87** (1968), 56–88.