MIT Libraries | DSpace@MIT

MIT Open Access Articles

Volumes of balls in Riemannian manifolds and Uryson width

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Guth, Larry. "Volumes of Balls in Riemannian Manifolds and Uryson Width." Journal of Topology and Analysis (February 22, 2016): 1–25.

As Published: http://dx.doi.org/10.1142/S1793525317500029

Publisher: World Scientific

Persistent URL: http://hdl.handle.net/1721.1/110192

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

Massachusetts Institute of Technology

arXiv:1504.07886v1 [math.DG] 29 Apr 2015

VOLUMES OF BALLS IN RIEMANNIAN MANIFOLDS AND URYSON WIDTH

LARRY GUTH

Abstract. If (M^n, g) is a closed Riemannian manifold where every unit ball has volume at most ϵ_n (a sufficiently small constant), then the $(n-1)$-dimensional Uryson width of (M^n, g) is at most 1.

If X is a compact metric space, then we say that X has Uryson q-width $\leq W$ if there is a q-dimensional simplicial complex Y and a continuous map $\pi : X \to Y$ so that every fiber $\pi^{-1}(y)$ has diameter $\leq W$ in X. In other words, if $x_1, x_2 \in X$ lie in the same fiber, $\pi^{-1}(y)$, then $d_X(x_1, x_2) \leq W$. We denote the Uryson q-width of X by $UW_q(X)$. If X is homeomorphic to a q-dimensional simplicial complex, then we can choose $Y = X$ and π to be the identity, and so $UW_q(X) = 0$. Roughly speaking $UW_q(X)$ is small if X is "close to being q-dimensional".

To get a first perspective, let $S^1(L)$ denote a circle of length L and let X be the product $S^1(W) \times S^1(L)$, where $W < L$. The Uryson 1-width of this space is $\sim W$. If we let π denote the projection from $S^1(W) \times S^1(L)$ to $S^1(L)$, then each fiber of π has diameter $W/2$. Therefore, $UW_1(X) \leq W/2$. On the other hand, the Uryson 1-width of X is $\gtrsim W$ because of the Lebesgue covering lemma, which we discuss more later in the introduction.

The Uryson width originally appeared in the early 1900's in connection with topological dimension theory. In [Gr1], Gromov began to investigate the Uryson widths of Riemannian manifolds in connection with systolic geometry. In [Gr3], he made a conjecture about Uryson widths and the volumes of balls, which we prove in this paper.

Theorem 0.1. *There exists $\epsilon_n > 0$ so that the following holds. If (M^n, g) is a closed Riemannian manifold, and if there is a radius R so that every ball of radius R in (M^n, g) has volume at most $\epsilon_n R^n$, then $UW_{n-1}(M^n, g) \leq R$.*

This theorem is a slightly stronger version of an estimate about filling radius from [Gu2]. The filling radius is another metric invariant introduced in [Gr1]. The filling radius of (M^n, g) has a similar flavor to $UW_{n-1}(M^n, g)$. For example, the filling radius of a product $S^1(W) \times S^1(L)$ is also around W (assuming $W < L$). It is not hard to show that $FillRad(M^n, g) \lesssim UW_{n-1}(M^n, g)$ – see Appendix 1 of [Gr1], (B)-(D). It is not known whether there is an inequality in the other direction. I

suspect that there is a closed manifold M^n and a sequence of metrics g_j where the ratio $UW_{n-1}(M^n, g_j)/FillRad(M^n, g_j)$ gets arbitrarily large, but I don't know of any examples. If so, a bound on Uryson (n-1)-width is slightly stronger than a bound on filling radius.

One of the main theorems of [Gr1] bounds the filling radius of a Riemannian manifold in terms of its volume.

Theorem 0.2. *(Gromov,* [Gr1]*, Filling radius inequality) If (M^n, g) is any closed Riemannian manifold, then $FillRad(M^n, g) \le C_n Vol(M^n, g)^{1/n}$.*

The paper [Gr1] raised the question whether $UW_{n-1}(M^n, g)$ is also bounded in terms of the volume of (M^n, g). As a corollary of Theorem 0.1, we get such an estimate.

Corollary 0.3. *If (M^n, g) is any closed Riemannian manifold, then $UW_{n-1}(M^n, g) \le C_n \operatorname{Vol}(M^n, g)^{1/n}$.*

Proof. We choose R so that $\epsilon_n R^n = \operatorname{Vol}(M^n, g)$. By Theorem 0.1, $UW_{n-1}(M^n, g) \le R = \epsilon_n^{-1/n} \operatorname{Vol}(M^n, g)^{1/n}$. □

In [Gr3], Gromov conjectured Theorem 0.1, and he also conjectured a slightly weaker theorem about the filling radius. The paper [Gu2] proved the filling radius version of this inequality.

Theorem 0.4. *(Local filling radius inequality) If (M^n, g) is a closed Riemannian manifold, $R > 0$, and every ball in (M^n, g) of radius R has volume $< \epsilon_n R^n$, then the filling radius of (M^n, g) is $\le R$.*

Theorem 0.1 is a slightly stronger version of the local filling radius inequality, Theorem 0.4, which in turn is a stronger version of the filling radius inequality, Theorem 0.2. For more context about metric geometry and systolic geometry, the reader can consult [Gr1] or [Gu3].

The proof of Theorem 0.1 closely follows the proof of Theorem 0.4 with one new ingredient, which we call a pushout lemma for small surfaces. This pushout lemma is the new observation in this paper.

0.1. **A push-out lemma for small surfaces.** In the late 1950's, Federer-Fleming proved (a close relative of) the following result.

Lemma 0.5. *(Federer-Fleming push-out lemma) Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. Suppose that $K \subset \mathbb{R}^N$ is a bounded, open convex set, and $\phi_0 : (X, \partial X) \to (K, \partial K)$ is a piecewise smooth map. If $n < N$, then ϕ_0 may be homotoped to a map $\phi_1 : X \to \partial K$ so that the following holds.*

1. *The map ϕ_1 agrees with ϕ_0 on ∂X.*
2. $\operatorname{Vol}_n \phi_1(X) \le C(K, N, n) \operatorname{Vol}_n \phi_0(X)$.

This Lemma plays an important role in metric geometry. See [Gu1] for a proof and for a description of the connection with the filling radius inequality. In the proof of Theorem 0.1, we need to construct a homotopy of this type, but we need a much more careful bound for the volume of $\phi_1(X)$. We need to find a homotopy where $\mathrm{Vol}_n \phi_1(X)$ is only very slightly larger than $\mathrm{Vol}_n \phi_0(X)$. We show that if $\mathrm{Vol}_n \phi_0(X)$ is small, then we can homotope ϕ_0 to ϕ_1 so that ϕ_1 lies near to ∂K and $\mathrm{Vol}_n \phi_1(X) \leq \mathrm{Vol}_n \phi_0(X)$. Here is the statement of our pushout lemma for small surfaces.

Lemma 0.6. *(Pushout lemma for small surfaces) For each dimension $n \geq 2$, there is a constant σ_n so that the following holds.*

Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. Suppose that $K \subset \mathbb{R}^N$ is a (bounded open) convex set, and $\phi_0 : (X, \partial X) \to (K, \partial K)$ is a piecewise-smooth map. Let $W = \sigma_n \mathrm{Vol}_n \phi_0(X)^{1/n}$. Then ϕ_0 may be homotoped to a map $\phi_1 : (X, \partial X) \to (K, \partial K)$ so that the following holds.

1. *The map ϕ_1 agrees with ϕ_0 on ∂X.*
2. $\mathrm{Vol}_n \phi_1(X) \leq \mathrm{Vol}_n \phi_0(X)$.
3. *The image $\phi_1(X)$ lies in the W-neighborhood of ∂K.*

When K is a ball, we have the following corollary.

Corollary 0.7. *For every n and every $\epsilon > 0$, we can find $\delta > 0$ so that the following holds. Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. Let $\phi_0 : (X, \partial X) \to (B^N, \partial B^N)$ be a map to the unit N-ball. If $\mathrm{Vol}_n \phi_0(X) < \delta$, then we can homotope ϕ_0 to $\phi_2 : X \to \partial B^N$ so that ϕ_2 agrees with ϕ_0 on ∂X and $\mathrm{Vol}_n \phi_2(X) < (1+\epsilon) \mathrm{Vol}_n \phi_0(X)$.*

Proof. By Lemma 0.6, we can homotope ϕ_0 rel ∂X to a map ϕ_1 with $\mathrm{Vol}_n \phi_1(X) \leq \mathrm{Vol}_n \phi_0(X)$, and with $\phi_1(X)$ lying within the W-neighborhood of ∂B^N for $W = \sigma_n \mathrm{Vol}_n \phi_0(X)^{1/n} \leq \sigma_n \delta^{1/n}$. Now we radially project ϕ_1 into ∂B^N, pushing away from the origin. The resulting map is ϕ_2. Since ϕ_1 already lies in the W-neighborhood of ∂B^N, the Lipschitz constant of this radial projection on $\phi_1(X)$ is $\leq (1-W)^{-1}$. Hence $\mathrm{Vol}_n \phi_2(X) \leq (1-W)^{-n} \mathrm{Vol}_n \phi_0(X)$. If we choose δ sufficiently small, we can arrange that $(1-W)^{-n} \leq (1+\epsilon)$. □

We remark that without the assumption that $\mathrm{Vol}_n \phi_0(X)$ is small, we cannot hope to push ϕ_0 into ∂B^N without stretching it significantly. For example, if $\phi_0(X)$ is a flat n-disk through the origin of B^N, and if $\phi_0(\partial X)$ is the boundary of the disk, then any homotopic map ϕ_2 will have to have volume at least equal to that of an n-dimensional unit hemisphere. The corollary says that for very small surfaces, we can push the surface into ∂B^N with only slight stretching.

The intuition for our push-out lemma comes from minimal surface theory. Suppose we consider all of the maps ϕ homotopic to ϕ_0 rel ∂X. We let V denote the infimal

volume of $\phi(X)$ among all these maps. Clearly $V \leq \mathrm{Vol}_n\, \phi_0(X)$. It seems plausible that V is actually realized by a map ϕ_1 whose image is some kind of stationary object, such as a stationary varifold, in the interior of K. For the sake of intuition let us suppose that such a ϕ_1 exists. Since the image is a stationary varifold in the interior of K, it obeys the monotonicity formula. Namely, for any regular point p in $\phi_1(X)$ and any $B(p,r) \subset K$, we have $\mathrm{Vol}_n\, \phi_1(X) \cap B(p,r) \geq \omega_n r^n$, where ω_n is the volume of the unit n-ball. We see that $\omega_n r^n \leq \mathrm{Vol}_n\, \phi_0(X)$. Hence $\phi_1(X)$ lies in the W-neighborhood of ∂K, for $W = \omega_n^{-1/n} \mathrm{Vol}_n\, \phi_0(X)^{1/n}$. This thought experiment provides intuition for our push-out lemma, and it also suggests what the sharp constant σ_n should be: $\omega_n^{-1/n}$.

The proof of Theorem 0.4 used some minimal surface theory and the monotonicity formula. We give a direct constructive proof of Lemma 0.6, so no minimal surface theory is required in the proof of Theorem 0.1. Now Theorem 0.1 implies Theorem 0.4, so in particular, we get a modified proof of the local filling radius inequality without any use of minimal surface theory. The proof of Lemma 0.6 is based on Wenger's recent proof of Gromov's filling volume inequality, [W].

0.2. **Background on Uryson width in topological dimension theory.** The Uryson width first appeared in topological dimension theory in the early 1900's. In this subsection, we review some topological dimension theory. We will see the role of Uryson width in topology, and we will compare our main theorem to some classical results of topological dimension theory. In particular, we will see that our Theorem 0.1 is a more quantitative version of a classical theorem of Szpilrajn comparing topological dimension with Hausdorff dimension. We will outline the proof of Szpilrajn's theorem, and we will see how "push out" homotopies, as in Lemma 0.6, come into the story. A number of other characters in the proof of Theorem 0.1 also come from topological dimension theory.

We first note that the Uryson width has an equivalent definition in terms of open covers. If $\{O_i\}$ is an open cover of X, we say that it has multiplicity $\leq m$ if each point of x lies in $\leq m$ different sets O_i. We say that the diameter of the open cover is $\sup_i \mathrm{Diam}(O_i)$.

Open covers and Uryson width are connected via the idea of the nerve of a covering. We recall the definition of the nerve of a cover to set the notation. Suppose that $\{O_i\}_{i=1,...,S}$ is an open cover of X. The nerve N of this cover is a simplicial complex with one vertex v_i for each set O_i. Given a $(k+1)$-tuple of vertices, $v_{i_0}, ..., v_{i_k}$, there is a simplex in the nerve with these vertices if and only if $\cap_{j=0}^k O_{i_j}$ is non-empty. The nerve N naturally sits in $[0,1]^S$, where the vertex v_i corresponds to the unit vector in the ith direction. In particular, any map $F : X \to N$ has coordinates $F_1, ..., F_S$. We say that the map F is subordinate to the cover $\{O_i\}$ if each F_i is supported in O_i.

Lemma 0.8. *A compact metric space X has $UW_q(X) < W$ if and only if there is an open cover of X with multiplicity $\leq q+1$ and diameter $< W$.*

We sketch the proof of Lemma 0.8. Suppose that X has an open cover with multiplicity $\leq q+1$ and diameter $< W$. Let N denote the nerve of the cover. Note that N is a simplicial complex of dimension q. Let ϕ_i be a partition of unity on X with ϕ_i supported in O_i. These ϕ_i give a continuous map ϕ from X to the nerve N. Note that ϕ is subordinate to the cover $\{O_i\}$. In particular, each fiber $\phi^{-1}(y)$ is contained in one of the sets O_i, and so it has diameter $< W$. Hence $UW_q(X) < W$.

Now suppose that $UW_q(X) < W$. Let $\pi : X \to Y$ be a map to a q-dimensional polyhedron so that each fiber has diameter $\leq UW_q(X)+\epsilon < W$. Since X is compact, the image of X is compact, and so we can assume that Y is compact. A q-dimensional polyhedron has an open cover O_i' with multiplicity at most $q+1$ where the diameters of the sets O_i' can be made arbitrarily small. Then, $O_i = \pi^{-1}(O_i')$ give an open cover of X with multiplicity at most $q+1$. If the diameter of O_i' is small enough, then each O_i lies in an arbitrarily small neighborhood of a fiber $\pi^{-1}(y)$, and so each O_i has diameter $< W$.

The most famous estimate about Uryson width is the Lebesgue covering lemma.

Lebesgue covering lemma. *Suppose that $\{O_i\}$ is an open cover of the unit n-cube $Q_n = [0,1]^n$ with multiplicity $\leq n$. Then the diameter of $\{O_i\}$ is ≥ 1. In other words, $UW_{n-1}(Q_n) = 1$.*

The Lebesgue covering lemma was stated by Lebesgue and first proven by Brouwer a little later. For a proof, see [HW]. The Lebesgue covering lemma has many applications in topological dimension theory. Since $UW_{n-1}(Q_n) > 0$, there is no injective continuous map from Q_n into an (n-1)-dimensional polyhedron. In particular, there is no injective continuous map from $\mathbb{R}^n$ into $\mathbb{R}^{n-1}$. This implies the famous topological invariance of dimension.

Topological invariance of dimension. *(Brouwer, 1909) If $m \neq n$, then $\mathbb{R}^m$ and $\mathbb{R}^n$ are not homeomorphic.*

The Lebesgue covering dimension of a compact metric space X is the smallest q so that $UW_q(X) = 0$. In particular, the Lebesgue covering dimension of Q_n is equal to n, which follows from the Lebesgue covering lemma.

If X_1 and X_2 are homeomorphic compact metric spaces, then the Lebesgue covering dimension of X_1 is equal to that of X_2. In other words, $UW_q(X_1) = 0$ if and only if $UW_q(X_2) = 0$. This follows fairly easily from the definition of Uryson width and the following observation. If $F : X_1 \to X_2$ is a homeomorphism of compact metric spaces, then for every ϵ there exists a δ so that if $A \subset X_2$ has diameter $< \delta$ then $F^{-1}(A) \subset X_1$ has diameter $< \epsilon$.

(The Lebesgue covering lemma is a precise quantitative estimate for $UW_{n-1}(Q_n)$. It easily implies non-sharp lower bounds on the Uryson (n-1)-width of other manifolds. For instance, it follows that the Uryson (n-1)-width of the unit n-sphere is $\gtrsim 1$ with an explicit lower bound. This was sharpened by Katz in [K], who proved the precise value of $UW_{n-1}(S^n)$.)

There are several notions of the "dimension" of a compact metric space X, and it is interesting to see how they relate. One important notion is the Hausdorff dimension. Szpilrajn proved that the Lebesgue covering dimension of any X is at most the Hausdorff dimension of X. In fact, he proved an even stronger theorem, which we now state.

Szpilrajn dimension inequality. *Let X be a compact metric space with n-dimensional Hausdorff measure equal to zero. Then $UW_{n-1}(X) = 0$. In other words, the Lebesgue covering dimension of X is $\leq n-1$.*

(Szpilrajn's inequality was pointed out to me by Anton Petrunin and Vitali Kapovitch.)

Our main theorem is a quantitative version of the Szpilrajn theorem. In particular, Szpilrajn's theorem implies that if X is a compact metric space and every unit ball in X has n-dimensional Hausdorff measure exactly equal to zero, then $UW_{n-1}(X) \leq 1$. In Theorem 0.1, we suppose that X is a closed n-dimensional Riemannian manifold. Instead of assuming that every unit ball in X has zero n-dimensional volume, we assume that every unit ball has n-dimensional volume at most a tiny constant ϵ_n. Under this weaker assumption, Theorem 0.1 says that $UW_{n-1}(X)$ is still at most 1, provided that ϵ_n is small enough.

We sketch here the proof of Szpilrajn's theorem. The proof of Theorem 0.1 follows the same outline as this proof, but it requires careful estimates at several steps.

Proof. Pick any $\epsilon > 0$. Cover X by finitely many open sets O_i with diameter $< \epsilon$. Let $\phi : X \to N$ be a map to the nerve coming from a partition of unity on X, as in the discussion of Lemma 0.8 above.

The idea of a mapping subordinate to the cover $\{O_i\}$ is central to the proof. Recall that a map $F : X \to N$ is subordinate to the cover $\{O_i\}$ if each component F_i is supported in O_i. The map ϕ coming from a partition of unity (subordinate to the cover $\{O_i\}$) is indeed subordinate to the cover. We can also arrange that the map ϕ is Lipschitz. Therefore, the image $\phi(X)$ has n-dimensional Hausdorff measure zero.

Let D denote the dimension of the nerve N. We will construct a sequence of homotopic maps from X to N, $\phi = \phi^{(D)} \sim \phi^{(D-1)} \sim ... \sim \phi^{(n)} \sim \phi^{(n-1)}$, where $\phi^{(k)}$ maps X into the k-skeleton of the nerve N. We denote the k-skeleton of N by $N^{(k)}$. All maps $\phi^{(k)}$ are Lipschitz and all maps subordinate to the cover.

For $n \leq k \leq D$, we homotope $\phi^{(k)}$ to $\phi^{(k-1)}$ as follows. Since $\phi^{(k)}$ is Lipschitz, $\phi^{(k)}(X)$ has n-dimensional Hausdorff measure zero. Since $k \geq n$, the image $\phi^{(k)}(X)$

does not contain any k-face of N. For each k-face $F_j \subset N$, we pick a point $y_j \in F_j$ with y_j not in the image of $\phi^{(k)}$. We let $\tilde{N}^{(k)}$ denote the k-skeleton of N with the points y_j removed.

We let R_k be a retraction from $\tilde{N}^{(k)}$ into $N^{(k-1)}$, taking $F_j \setminus \{y_j\}$ into ∂F_j. We define $\phi^{(k-1)}$ to be $R_k \circ \phi^{(k)}$. We can arrange that the map R_k is Lipschitz on any compact subset of $\tilde{N}^{(k)}$. Therefore, $\phi^{(k-1)}$ is Lipschitz.

Since R_k maps each face $F \subset N^{(k)}$ into its closure $\bar{F}$, and since $\phi^{(k)}$ is subordinate to the cover, it follows that $\phi^{(k-1)}$ is subordinate to the cover as well.

In particular, we get a map $\phi^{(n-1)}$ from X to the (n-1)-dimensional polyhedron $N^{(n-1)}$ subordinate to the cover. Because $\phi^{(n-1)}$ is subordinate to the cover, each fiber of the map lies in one of the sets O_i and has diameter $< \epsilon$. Hence $UW_{n-1}(X) < \epsilon$. □

The proof of Theorem 0.1 follows this outline, but it requires careful estimates. By assumption, each unit ball of X has small volume, and so each set O_i in the cover can be taken to have small volume. The map $\phi = \phi^{(D)}$ is qualitatively Lipschitz, but we need to control the Lipschitz constant of ϕ in order to bound the volume of $\phi(O_i)$. This part of the argument is handled by ideas from [Gu2]. Suppose now that $\phi(O_i) = \phi^{(D)}(O_i)$ has very small volume for each i. We will construct a sequence of homotopies, $\phi^{(D)} \sim \phi^{(D-1)} \sim ... \sim \phi^{(n)}$. At each step, we need to prove that $\phi^{(k)}(O_i)$ remains small. If $\mathrm{Vol}_n\, \phi^{(k)}(O_i) = 2\,\mathrm{Vol}_n\, \phi^{(k+1)}(O_i)$, then it turns out that $\phi^{(n)}(O_i)$ will be far too large. So we need to homotope $\phi^{(k+1)}$ to $\phi^{(k)}$ while being very careful about how much volumes stretch. This step is accomplished with Lemma 0.6, the pushout lemma for small surfaces.

For more information about topological dimension theory, one can read the book [HW]. For a more geometric point of view about Uryson width see the survey paper [Gr2].

To end this discussion, we remark that trying to transform a qualititative theorem from topological dimension theory into a quantitative estimate in Riemannian geometry doesn't always work. For example we consider the following theorem on the locality of topological dimension.

Locality of Lebesgue covering dimension. *Let X be a compact metric space. Suppose that X is covered by open sets O_i and suppose that $UW_q(O_i) = 0$ for each i. Then $UW_q(X) = 0$.*

This locality theorem follow from the central results of topological dimension theory in [HW]. The main theme of the book [HW] is that various ways of defining the 'topological dimension' of a compact metric space are equivalent to each other. The central definition in the book is the topological dimension (or Menger topological dimension) defined inductively as follows. A compact metric space has topological

dimension -1 if it is empty. A compact metric space X has topological dimension $\leq q$ at a point $x \in X$ if x has arbitrarily small open neighborhoods U such that ∂U has topological dimension $\leq q-1$. A compact metric space X has topological dimension $\leq q$ if it has topological dimension $\leq q$ at each point $x \in X$. Theorem V8 of [HW] says that for a compact metric space X, the topological dimension and the Lebesgue covering dimension are the same. But the definition of topological dimension is clearly local, because it only involves arbitrarily small neighborhoods of every point. If X is covered by open sets O_i, and each set O_i has topological dimension at most q, then it follows immediately that X has topological dimension at most q. Since the Lebesgue covering dimension is equivalent to the topological dimension, the Lebesgue covering dimension is also local. Finally, X has Lebesgue covering dimension at most q if and only if $UW_q(X) = 0$.

We consider a more quantitative version of this result for Riemannian manifolds:

Question 0.9. *Is there a constant $\epsilon(q,n) > 0$ so that the following holds: if (M^n, g) is a closed Riemannian manifold, and every unit ball in (M, g) has Uryson q-width at most $\epsilon(q,n)$, then (M^n, g) has Uryson q-width at most 1?*

This question is related to the locality of dimension theorem in exactly the same way that Theorem 0.1 is related to Szpilrajn's theorem. But the answers turn out to be different. At least when $n = 3$ and $q = 2$, the answer to question 0.9 is no because of the following counterexample:

Proposition 0.10. *For any $\epsilon > 0$, there is a metric g_ϵ on S^3 so that every unit ball in (S^3, g_ϵ) has Uryson 2-width $< \epsilon$ and yet the whole manifold (S^3, g_ϵ) has Uryson 2-width at least 1.*

The main idea of this counterexample comes from [Gr2], Example (H_1'').

So from a quantitative point of view, the Szpilrajn theorem is more robust than the locality theorem.

0.3. **Outline of the paper.** In Section 1 of the paper, we prove Theorem 0.1 using the pushout lemma for small surfaces. The proof closely follows the arguments from [Gu2], and we use some lemmas from [Gu2]. In Sections 2 and 3, we prove the pushout lemma for small surfaces. Section 2 introduces a key tool, a variation on the isoperimetric inequality which we call an isoperimetric extension lemma. Section 3 uses the extension lemma to prove Lemma 0.6. In Section 4, we prove Proposition 0.10. In Section 5 we discuss a few open questions, and in Section 6, we discuss complete manifolds and manifolds with boundary.

1. The proof of the main theorem

The proof of the main theorem is closely based on the argument in [Gu2]. The new ingredient is the pushout lemma for small surfaces, which we will prove in Section 3 below.

Suppose (X^n, g) is a closed Riemannian manifold of dimension n. We assume that each unit ball in (X^n, g) has volume $< \epsilon(n)$, a sufficiently small number that we can choose later. We have to construct a continuous map $\pi : X \to Y^{n-1}$ to a polyhedron of dimension $n-1$ so that each fiber $\pi^{-1}(y)$ has diameter at most 1.

Suppose that B_i in X, $i = 1, ..., D$, so that $\frac{1}{2}B_i$ cover X and so that each ball B_i has radius $< (1/100)$. We will choose a particular set of balls B_i later, obeying some additional geometric estimates. Next we build a map to the "rectangular nerve" of this cover.

The rectangular nerve of the covering $\{B_i\}$ is defined as follows. We begin with the rectangle $\prod_{i=1}^{D}[0, r_i]$ where r_i is the radius of the ball B_i. We let ϕ_i, $i = 1, ..., D$ be the coordinate functions on this rectangle. The rectangular nerve is a closed subcomplex of this product. An open face F of the rectangle is determined by dividing the dimensions $1, ..., D$ into three sets: I_0, I_1, and $I_{(0,1)}$. The open face F is defined by the equations $\phi_i = 0$ if $i \in I_0$, $\phi_i = r_i$ if $i \in I_1$, and $0 < \phi_i < r_i$ if $i \in I_{(0,1)}$. We let I_+ be the union of I_1 and $I_{(0,1)}$. Now an open face F is contained in the rectangular nerve N iff the set $I_1(F)$ is not empty, and the intersection $\cap_{i \in I_+(F)} B_i$ is not empty.

As with the regular nerve, we can construct a naturally defined map from X to the rectanglar nerve. We call the map $\phi : X \to N$. To define ϕ, we let ϕ_i be a real-valued function supported on B_i with $\phi_i(x) = r_i$ for $x \in \frac{1}{2}B_i$ and $\phi_i(x)$ decreasing to zero as x approaches ∂B_i. We can choose ϕ_i to be piecewise smooth with values in $[0, r_i]$, and with Lipschitz constant less than 3. Taking ϕ_i as coordinates, we get a map $\phi : X \to \prod_{i=1}^{D}[0, r_i]$, and the image of ϕ lies in the rectangular nerve N. To see this, suppose that $\phi(x)$ lies in an open face F. Because $\frac{1}{2}B_i$ cover M, $\phi_i(x) = r_i$ for some i, and so $I_1(F)$ is not empty. Since ϕ_i is supported in B_i, x lies in $\cap_{i \in I_+(F)} B_i$, which must be non-empty.

The polyhedron Y will be the (n-1)-skeleton of the rectangular nerve N. To construct $\pi : X \to Y$, we will homotope ϕ until its image lands in Y.

We say that a map $\psi : X \to N$ is subordinate to our covering if the ith component, ψ_i, is supported in B_i for each $i = 1, ..., D$. (Since $N \subset \prod_{i=1}^{D}[0, r_i]$, we can speak of the ith coordinate of ψ.) By construction our map ϕ is subordinate to the cover. The map $\pi : X \to Y \subset N$ will also be subordinate to the covering. If $\psi : X \to N$ is subordinate to the covering, then any inverse image $\psi^{-1}(y)$ is contained in some ball B_i of radius $< 1/100$. So any fiber of ψ has diameter $\leq 1/50 < 1$.

The strategy for constructing π is based on the proof of Szpilrajn's theorem. We let $\phi^{(D)}$ denote our original map ϕ, and we construct a sequence of homotopies

$\phi^{(D)} \sim \phi^{(D-1)} \sim ... \sim \phi^{(n)} \sim \phi^{(n-1)}$. Each map $\phi^{(k)}$ maps X to the k-skeleton of N. And each of these maps is subordinate to the covering. The map $\phi^{(n-1)}$ is our desired map π, and to prove our theorem, we only need to construct a map $\phi^{(n-1)}$ from X to the $(n-1)$-skeleton of N which is subordinate to the cover.

Since we will need to check that various maps are subordinate to the cover, the following simple observation is useful.

Lemma 1.1. *Suppose that $\Phi_1 : X \to N$ is subordinate to the cover, and that $\Phi_2 : X \to N$ is another map. Suppose that for each $x \in X$, if $\Phi_1(x)$ lies in an open face $F \subset N$, then $\Phi_2(x)$ lies in its closure $\bar{F}$. Then Φ_2 is also subordinate to the cover.*

Proof. Pick an index i. We want to show that i^{th} coordinate of Φ_2 is supported in B_i. Pick a point $x \in X$ with $\Phi_{2,i}(x) \neq 0$. The image $\Phi_2(x)$ lies in an open face F with $i \in I_+(F)$. Let us say that $\Phi_1(x)$ lies in the open face F_1. By our hypothesis $F \subset \bar{F}_1$. Hence $I_+(F) \subset I_+(F_1)$. In particular, $i \in I_+(F_1)$, and so $\Phi_{1,i}(x) > 0$. Since Φ_1 is subordinate to the cover, $x \in B_i$. □

This lemma applies to the push-out type construction that we used in Section 1 to prove Szpilrajn's theorem. In particular, we get the following lemma.

Lemma 1.2. *Suppose that $\phi^{(k)} : X \to N^{(k)}$ is a map from X to the k-skeleton of N subordinate to the cover. Suppose that for each k-face $F \subset N$, the image $\phi^{(k)}(X) \cap F$ is not the entire k-face. Then we can homotope $\phi^{(k)}$ to a map $\phi^{(k-1)} : X \to N^{(k-1)}$ subordinate to the cover.*

Proof. For each k-face, $F_j \subset N$, pick a point y_j which lies in F_j but does not lie in $\phi^{(k)}(X)$. We let N' denote the k-skeleton of N take away the points y_j we just chose. So $\phi^{(k)}$ maps X into N'. Now we let Ψ be a retraction from N' to $N^{(k-1)}$, constructed by radially pushing $F_j \setminus \{y_j\}$ into ∂F_j. We define $\phi^{(k-1)}$ to be $\Psi \circ \phi^{(k)} : X \to N^{(k-1)}$. Since Ψ maps each face F_j into $\bar{F}_j$, Lemma 1.1 implies that $\phi^{(k-1)}$ is subordinate to the cover. □

Using this lemma, we can construct a sequence of homotopies subordinate to the cover, $\phi^{(D)} \sim ... \sim \phi^{(n)}$. As long as $k > n$, any piecewise smooth map $\phi^{(k)} : X \to N^{(k)}$ will not contain any k-face in its image, and so we can apply Lemma 1.2. The real obstruction comes when we try to homotope $\phi^{(n)}$ to $\phi^{(n-1)}$.

In the case of the Spzilrajn theorem, we knew that $\phi^{(n)}(X)$ has n-dimensional Hausdorff measure 0. Hence $\phi^{(n)}(X)$ does not cover any n-dimensional face $F \subset N$, and we can Lemma 1.2. In our case, we will prove the following volume estimate.

$$\text{For each n-face } F \subset N, \mathrm{Vol}_n\, \phi^{(n)}(X) \cap F < \mathrm{Vol}_n\, F. \qquad (*)$$

(The rectangular nerve has a metric given by the restriction the Euclidean metric on the rectangle $\prod_{i=1}^{D}[0, r_i]$.)

Given this key estimate $(*)$, we can apply Lemma 1.2 again to homotope $\phi^{(n)}$ to $\phi^{(n-1)}$ subordinate to the cover. This will prove the main theorem.

Establishing the volume estimate $(*)$ takes some care. Since $\phi^{(n)}$ is subordinate to our covering, the preimage of any face F lies one of our balls B_i. The ball B_i has volume at most ϵ, which we can choose very small. But the initial map $\phi = \phi^{(D)}$ may stretch this volume. Then each homotopy from $\phi^{(k)}$ to $\phi^{(k-1)}$ may stretch it further. After all this potential stretching, we need a bound for the volume of $\phi^{(n)}(X) \cap F$. To get such a bound, we need to choose the covering carefully, in order to control the volume of $\phi^{(D)}(X) \cap F$. Then we need to choose our homotopies from $\phi^{(k)}$ to $\phi^{(k-1)}$ carefully, in order to inductively control the volume of each $\phi^{(k-1)}(X) \cap F$.

We need a little bit of vocabulary in order to state our estimates. Given an open face $F \subset N$, we let $\mathrm{Star}(F)$ denote the union of F and all open faces F' such that $F \subset \bar{F}'$. If F has dimension k, then each open face in $F' \subset \mathrm{Star}(F)$ has dimension $\geq k$, and the only k-face in $\mathrm{Star}(F)$ is F. We let $r_1(F)$ denote the shortest length of any of the sides of F.

The paper [Gu2] constructs a good covering $\{B_i\}$ and proves the estimates that we need about $\phi = \phi^{(D)}$. In particular, we use Lemma 5 from [Gu2]:

Lemma 1.3. *(Lemma 5 in* [Gu2]*.) There are constants $C(n)$ and $\beta(n) > 0$ depending only on n so that the following estimate holds. Let $\epsilon > 0$ be any number. Suppose that (X^n, g) is a closed Riemannian manifold of dimension n, and that each unit ball in (X^n, g) has volume $< \epsilon$. Then there is a covering B_i as above and a map $\phi^{(D)} : X \to N$ subordinate to the cover so that that following volume estimate holds:*

For any face $F \subset N$ of dimension $d(F)$,

$$\mathrm{Vol}_n[\phi^{(D)}(X) \cap \mathrm{Star}(F)] \leq C\epsilon r_1(F)^{n+1} e^{-\beta d(F)}. \tag{1}$$

We will construct a sequence of homotopic maps $\phi = \phi^{(D)} \sim \phi^{(D-1)} \sim ... \sim \phi^{(n)}$ subordinate to the cover. Moreover, every map $\phi^{(k)}$ will obey the following estimate, slightly weaker than the estimate that ϕ obeys.

$$\mathrm{Vol}_n[\phi^{(k)}(X) \cap \mathrm{Star}(F)] < 2C\epsilon r_1(F)^n e^{-\beta d(F)}. \tag{$**$}$$

In particular, for each n-face F, $\phi^{(n)}(X) \cap F$ obeys the following estimate:

$$\mathrm{Vol}_n[\phi^{(n)}(X) \cap F] < 2C\epsilon r_1(F)^n.$$

We will choose ϵ less than $(2C)^{-1}$, so we conclude that $\mathrm{Vol}_n[\phi^{(n)}(X) \cap F] < \mathrm{Vol}_n(F)$, proving $(*)$. Then by Lemma 1.2, $\phi^{(n)}$ is homotopic to a map $\phi^{(n-1)}$ to the $(n-1)$-skeleton of N, subordinate to the cover, and this proves the main theorem. So it suffices to construct the maps $\phi^{(k)}$ for each $k = D, D-1, ..., n$, subordinate to the cover and obeying the key estimate $(**)$.

The homotopy from $\phi^{(k)}$ to $\phi^{(k-1)}$ is based on our pushout lemma, Lemma 0.6. We recall the statement here.

Push-out lemma for small surfaces. *For each dimension $n \geq 2$, there is a constant σ_n so that the following holds.*

Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. Suppose that $K \subset \mathbb{R}^N$ is a convex set, and $\phi : (X, \partial X) \to (K, \partial K)$ is a piecewise-smooth map. Let $W = \sigma_n \operatorname{Vol}_n \phi_0(X)^{1/n}$. Then ϕ may be homotoped to a map $\tilde{\phi} : (X, \partial X) \to (K, \partial K)$ so that the following holds.

1. *The map $\tilde{\phi}$ agrees with ϕ on ∂X.*
2. $\operatorname{Vol}_n \tilde{\phi}(X) \leq \operatorname{Vol}_n \phi(X)$.
3. *The image $\tilde{\phi}(X)$ lies in the W-neighborhood of ∂K.*

We will apply this lemma on each k-face $F_j \subset N^{(k)}$. We let $K_j \subset F_j$ be a closed convex subset of F_j, containing almost all of F_j but in general position. We let X_j be the preimage of K_j under $\phi^{(k)}$. So we have a map $\phi^{(k)} : (X_j, \partial X_j) \to (K_j, \partial K_j)$, and we can apply the pushout lemma for small surfaces to this map. The lemma gives us a new map $\tilde{\phi}^{(k)} : (X_j, \partial X_j) \to (K_j, \partial K_j)$ agreeing with $\phi^{(k)}$ on ∂X_j.

We have just defined $\tilde{\phi}^{(k)}$ on $\cup_j X_j \subset X$. We extend $\tilde{\phi}^{(k)}$ to all of X by letting $\tilde{\phi}^{(k)} = \phi^{(k)}$ on $X \setminus \cup_j X_j$. The boundary between $X \setminus \cup_j X_j$ and $\cup X_j$ is $\cup \partial X_j$. Since $\tilde{\phi}^{(k)}$ agrees with $\phi^{(k)}$ on each ∂X_j, this definition gives a piecewise smooth map from X to $N^{(k)} \subset N$.

By induction, we know that

$$\operatorname{Vol}_n \phi^{(k)}(X_j) \leq \operatorname{Vol}_n \phi^{(k)}(X) \cap F_j \leq 2C\epsilon r_1(F_j)^n e^{-\beta k}.$$

The pushout lemma for small surfaces tells us that $\tilde{\phi}^{(k)}(X_j)$ lies in the w_j-neighborhood of ∂K_j where $\sigma_n w_j^n = 2C\epsilon r_1(F_j)^n e^{-\beta k}$. Since we can choose K_j as close as we like to the whole face F_j, we see that $\tilde{\phi}^{(k)}(X) \cap F_j$ lies in the W_j-neighborhood of ∂F_j for $\sigma_n W_j^n = 3C\epsilon r_1(F_j)^n e^{-\beta k}$. Rearranging this formula, we get the following inequality for W_j.

$$W_j / r_1(F_j) \leq [3C\sigma_n^{-1}\epsilon e^{-\beta k}]^{1/n}. \tag{1}$$

If $\phi^{(k)}(x)$ lies in a face F, then $\tilde{\phi}^{(k)}(x)$ lies in the same face F, and so $\tilde{\phi}^{(k)}$ is subordinate to the cover.

We define a map that pulls a small neighborhood of the (k-1)-skeleton of N into the (k-1)-skeleton. Our map will be called R_δ, and it depends on a number δ in the range $0 < \delta < 1/2$. The basic map is a map from an interval $[0, r]$ to itself, which takes the set $[0, \delta r]$ to 0, and the set $[r - \delta r, r]$ to r, and linearly stretches the set $[\delta r, r - \delta r]$ to cover $[0, r]$. The Lipschitz constant of this map is $(1 - 2\delta)^{-1}$. Now

we apply this map separately to each coordinate ϕ_i of the rectangle $\prod_{i=1}^{D}[0, r_i]$. The resulting map is R_δ.

Our map $\phi^{(k-1)}$ will be $R_{\delta(k)} \circ \tilde{\phi}^{(k)}$ for a well-chosen $\delta(k)$. Notice that for any δ, the map R_δ sends each open face F into $\bar{F}$. By Lemma 1.1, $R_\delta \circ \tilde{\phi}^{(k)}$ is subordinate to the cover.

Also, for sufficiently big δ, $R_\delta \circ \tilde{\phi}^{(k)}$ maps X into $N^{(k-1)}$. In particular, if $\delta \geq [3C\sigma_n^{-1}\epsilon e^{-\beta k}]^{1/n}$, then inequality (1) guarantees that for each k-face F_j, $W_j/r_1(F_j) \leq \delta$. Since $\tilde{\phi}^{(k)} \cap F_j$ lies in the W_j-neighborhood of ∂F_j, $R_\delta \circ \tilde{\phi}^{(k)}$ lies in the (k-1)-skeleton of N. We define $\delta(k) := [3C\sigma_n^{-1}\epsilon e^{-\beta k}]^{1/n}$, and then we define $\phi^{(k-1)} = R_{\delta(k)} \circ \tilde{\phi}^{(k)}$.

To close the induction, we just need to check the estimate $(**)$ for $\phi^{(k-1)}$. In other words, we need to prove that

$$\mathrm{Vol}_n[\phi^{(k-1)}(X) \cap \mathrm{Star}(F)] < 2C\epsilon r_1(F)^n e^{-\beta d(F)}.$$

By Lemma 1.3, we know that

$$\mathrm{Vol}_n[\phi^{(D)}(X) \cap \mathrm{Star}(F)] \leq C\epsilon r_1(F)^n e^{-\beta d(F)}.$$

We know that the pushout construction does not increase any volumes. In particular, for any $k \leq l \leq D$, we know that in every face $F \subset N$, $\mathrm{Vol}_n[\tilde{\phi}^{(l)}(X) \cap F] \leq \mathrm{Vol}_n[\phi^{(l)}(X) \cap F]$. Hence we know that

$$\mathrm{Vol}_n[\tilde{\phi}^{(l)}(X) \cap \mathrm{Star}(F)] \leq \mathrm{Vol}_n[\phi^{(l)}(X) \cap \mathrm{Star}(F)].$$

On the other hand, the map R_δ does not increase volumes very much if δ is small. The Lipschitz constant of R_δ is $[1 - 2\delta]^{-1}$. Therefore, for any n-dimensional surface, Σ,

$$\mathrm{Vol}_n[R_\delta(\Sigma)] \leq [1 - 2\delta]^{-n} \mathrm{Vol}_n[\Sigma].$$

Also, we have seen that R_δ maps each face F into $\bar{F}$. Hence $R_\delta^{-1}(\mathrm{Star}(F)) \subset \mathrm{Star}(F)$. Therefore we get the following inequality.

$$\mathrm{Vol}_n[\phi^{(l-1)}(X) \cap \mathrm{Star}(F)] \leq [1 - 2\delta(l)]^{-n} \mathrm{Vol}_n[\tilde{\phi}^{(l)}(X) \cap \mathrm{Star}(F)] \leq$$

$$\leq [1 - 2\delta(l)]^{-n} \mathrm{Vol}_n[\phi^{(l)}(X) \cap \mathrm{Star}(F)].$$

By combining these inequalities for all $k \leq l \leq D$, and using inequality 1 for $\phi^{(D)}$, we get the following inequality.

$$\mathrm{Vol}_n[\phi^{(k-1)}(X) \cap \mathrm{Star}(F)] \leq \prod_{l=k}^{D} [1 - 2\delta(l)]^{-n} C\epsilon r_1(F)^n e^{-\beta d(F)}.$$

Therefore, it suffices to prove that $\prod_{l=k}^{D}[1-2\delta(l)]^{-n} \leq \prod_{l=n}^{\infty}[1-2\delta(l)]^{-n} < 2$. Recalling the definition of $\delta(l)$, it suffices to prove that

$$\prod_{l=n}^{\infty}\left[1-2(3C\sigma_n^{-1}\epsilon e^{-\beta l})^{1/n}\right]^{-n} < 2.$$

This formula is a little messy, but C, σ, and β are just dimensional constants. Because of the exponential decay $e^{-\beta l}$, the infinite product converges. And we can choose ϵ sufficiently small so that the product is less than 2.

This finishes the proof of the main theorem, except for the proof of the pushout lemma for small surfaces. We prove the pushout lemma in Section 3, following preliminary work in Section 2.

2. An isoperimetric extension lemma

In order to prove the pushout lemma, we need the following version of the isoperimetric inequality. Suppose that X is a piecewise smooth n-dimensional manifold and $F : X \to \mathbb{R}^N$ is a piecewise smooth map. We write $\mathrm{Vol}_n F(X)$ for the volume of X with the induced Riemannian metric given by pullback with F. Equivalently, we can think of $\mathrm{Vol}_n F(X)$ as the volume of the image of F counted with multiplicity.

Lemma 2.1. *(Extension lemma) For each dimension $n \geq 2$, there exists an "isoperimetric constant" $I(n)$ so that the following holds.*

Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. If F_0 is a map from ∂X to $\mathbb{R}^N$, then there is an extension $F : X \to \mathbb{R}^N$ so that

$$\mathrm{Vol}_n F(X) \leq I(n)[\mathrm{Vol}_{n-1} F_0(\partial X)]^{\frac{n}{n-1}}.$$

This lemma can be thought of as a (minor) generalization of the isoperimetric inequality. For any integral cycle $y^{n-1} \subset \mathbb{R}^N$, the Michael-Simon isoperimetric inequality says that there is a chain x^n with $\partial x = y$ and $Vol_n(x) \leq C_n Vol_{n-1}(y)^{\frac{n}{n-1}}$ [MS]. Our extension lemma is a version of this inequality for maps instead of chains.

Recently, Wenger gave a short constructive proof of the Michael-Simon isoperimetric inequality [W]. Our proof of the extension lemma closely follows Wenger's proof. (Wenger's proof works more generally in Banach spaces, and our extension lemma also generalizes to Banach spaces, but we don't pursue it here.)

Proof of extension lemma. We write $A \lesssim B$ if $A \leq C(n)B$. The constants do not depend on the ambient dimension N.

We begin with a cone-type inequality which allows us to construct extensions when the diameter of $F_0(\partial X)$ is not too big.

Lemma 2.2. *(Cone inequality) Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. If F_0 is a map from ∂X to ball of radius R, $B^N(R) \subset \mathbb{R}^N$, then there is an extension $F : X \to B^N(R)$ so that*

$$\mathrm{Vol}_n F(X) \leq C_n R[\mathrm{Vol}_{n-1} F_0(\partial X)].$$

Proof. There is a tubular neighborhood E of ∂X in X which is diffeomorphic to $\partial X \times (-1, 1]$. We choose coordinates (x, t) on this neighborhood, where $x \in \partial X$ and $t \in [0, 1)$. The boundary ∂X is given by the equation $t = 0$. Let $\rho(t)$ be a smooth non-negative function on $[0, 1)$ with $\rho(0) = 1$ and $\rho(t) = 0$ for $t \geq 1/2$. Then define F on this tubular neighborhood by $F(x, t) = \rho(t)F_0(x)$. Note that F maps the subset of E where $1/2 \leq t < 1$ to the origin. Therefore, we can smoothly extend F to all of X by mapping the complement of E to the origin.

The image of F is the cone over the image of F_0. By standard calculations in Riemannian geometry, we get

$$\mathrm{Vol}_n F(X) \leq (1/n)R\,\mathrm{Vol}_{n-1} F_0(\partial X).$$

(If we were working in a Banach space, then the constant $1/n$ would have to be replaced by a larger constant C_n.) □

(Notice that cone inequality proves the extension lemma in the special case that $Vol_{n-1}F_0(\partial X) = 0$. So in the rest of the proof, we may assume that $Vol_{n-1}F_0(\partial X) > 0$.)

Our argument will be by induction on n. The base case is $n = 2$, which we now discuss. In this case, ∂X is a 1-dimensional manifold. It consists of finitely many connected components, which we call $\partial X_1, \partial X_2$, etc. For each connected component, the diameter of $F_0(\partial X_i)$ is at most the length of $F_0(\partial X_i)$, which we denote by L_i. For each i, we choose a point y_i in the image of $F_0(\partial X_i)$ in $\mathbb{R}^N$. We observe that $F_0(\partial X_i)$ is contained in $B^N(y_i, L_i)$.

We let E_i be a neighborhood of ∂X_i in X, diffeomorphic to $\partial X_i \times [0, 1)$. By the same construction as the cone inequality, we can extend F_0 to a map $F : E_i \to B^N(y_i, L_i)$ so that F maps all but a compact subset of E_i to y_i and so that $\mathrm{Vol}_2 F(E_i) \leq CL_i^2$. Now consider $X' = X \setminus \cup_i E_i$. We have to define F on X' so that F maps ∂E_i to y_i. We can choose F so that the image of F is a tree, and $\mathrm{Vol}_2 F(X') = 0$. We have now defined an extension F on X with

$$\mathrm{Vol}_2 F(X) \leq C \sum_i L_i^2 \leq C(\sum_i L_i)^2 = C(\mathrm{Vol}_1 F_0(\partial X))^2.$$

This finishes the proof of the extension lemma in the base case $n = 2$.

Now we begin the proof of the inductive step. We assume that the extension lemma holds in dimension $n-1$. We construct the extension F by repeatedly using the following partial extension lemma.

Partial extension lemma. *Suppose that X is a compact piecewise-smooth n-dimensional manifold with boundary. If F_0 is a map from ∂X to $B^N(R) \subset \mathbb{R}^N$, then we can decompose X as a union of two n-dimensional submanifolds with ∂, $X = X_1 \cup G_1$, and we can extend F_0 to a map $F_1 : G_1 \to B^N(R)$ so that the following estimates hold.*

$$\mathrm{Vol}_n F_1(G_1) \lesssim \mathrm{Vol}_{n-1} F_0(\partial X)^{\frac{n}{n-1}}.$$

Next, we note that $\partial X_1 \subset \partial G_1 \cup \partial X$. Therefore, F_1 is defined on ∂X_1.

$$\mathrm{Vol}_{n-1} F_1(\partial X_1) \leq (1-\delta_n) Vol_{n-1} F_0(\partial X).$$

Using the partial extension lemma and the cone inequality, we can quickly finish the proof of the extension lemma. Suppose that $F_0 : \partial X \to \mathbb{R}^N$. Choose a large radius R so that $F_0(\partial X) \subset B^N(R)$. We use the partial extension lemma to define X_1, G_1, F_1. Then we consider $F_1 : \partial X_1 \to B^N(R)$, and we apply the extension lemma to it. We use the partial extension lemma J times, where J is a large number that we will choose below. We get a sequence of subsets $X_J \subset X_{J-1} \subset ... \subset X_1 \subset X_0 = X$, and a sequence of maps $F_j : X_{j-1} \setminus X_j \to B^N(R)$. The maps F_j fit together to define a single piecewise smooth map $F : X \setminus X_J \to B^N(R)$, extending $F_0 : \partial X \to B^N(R)$.

By the second estimate in the partial extension lemma, we know that

$$\mathrm{Vol}_{n-1} F_j(\partial X_j) \leq (1-\delta_n) \mathrm{Vol}_{n-1} F_{j-1}(\partial X_{j-1}).$$

Therefore,

$$\mathrm{Vol}_{n-1} F_j(\partial X_j) \leq (1-\delta_n)^j \mathrm{Vol}_{n-1} F_0(\partial X).$$

Next we can bound the volume of $F_j(X_{j-1} \setminus X_j)$, using the first estimate in the partial extension lemma.

$$\mathrm{Vol}_n F_j(X_{j-1} \setminus X_j) \lesssim (\mathrm{Vol}_{n-1} F_{j-1}(\partial X_{j-1}))^{\frac{n}{n-1}} \lesssim (1-\delta_n)^{\frac{n}{n-1} j} \mathrm{Vol}_{n-1} F_0(\partial X)^{\frac{n}{n-1}}.$$

Summing the exponential sum, we see that

$$\mathrm{Vol}_n F(X \setminus X_J) \lesssim \mathrm{Vol}_{n-1} F_0(\partial X)^{\frac{n}{n-1}}.$$

This estimate holds uniformly in the choice of J. Finally, we extend F to X_J using the cone inequality. We get

$$\mathrm{Vol}_n F(X_J) \lesssim R \, \mathrm{Vol}_{n-1} F_J(\partial X_J) \lesssim R(1-\delta_n)^J \mathrm{Vol}_{n-1} F_0(\partial X).$$

We now choose J sufficiently large in terms of R and $\mathrm{Vol}_{n-1} F_0(\partial X)$ so that this final term is dominated by the previous term.

This finishes the proof of the extension lemma from the partial extension lemma. Now we turn to the proof of the partial extension lemma.

Proof of partial extension lemma. At this point, it is convenient to know that our mapping is an embedding. To accomplish this, we add extra dimensions to the target space $\mathbb{R}^N$. We let $F_0^+ : \partial X \to \mathbb{R}^N \times \mathbb{R}^E$ be given by F_0 in the first factor and by a nice embedding in the second factor. By scaling the second factor, we can assume that $\mathrm{Vol}_{n-1} F_0^+(\partial X) \leq (1+\epsilon)\,\mathrm{Vol}_{n-1} F_0(\partial X)$ and we can assume that $F_0^+(\partial X) \subset B^{N+E}(R)$. Now F_0^+ is an embedding. Our construction will give a partial extension $F_1^+ : G_1 \to B^{N+E}(R)$, obeying good estimates. Finally, we define F_1 to be F_1^+ composed with the projection from $B^{N+E}(R)$ to $B^N(R)$. This projection can only decrease volumes, so the resulting partial extension $F_1 : G_1 \to B^N(R)$ obeys the same estimates as F_1^+. In summary, it suffices to consider the case that F_0 is an embedding.

The proof of the partial extension lemma is based on a ball-covering argument, the extension lemma in dimension $n-1$, and the cone inequality. It closely follows the argument in [W].

We consider the volumes $\mathrm{Vol}_{n-1}[F_0(\partial X) \cap B^N(p,r)]$ for various balls $B^N(p,r) \subset \mathbb{R}^N$. Because F_0 is an embedding, we know that for every regular point p in the image $F_0(\partial X)$, the volume $\mathrm{Vol}_{n-1}[F_0(\partial X) \cap B(p,r)] \geq c_n r^{n-1}$ for all sufficiently small r. (We may take c_n to be one half the volume of the unit $(n-1)$-ball.) Since F_0 is a piecewise smooth embedding, almost every point of the image is regular.

Fix any regular point p. Let $V(r)$ denote $\mathrm{Vol}_{n-1}[F_0(\partial X) \cap B^N(p,r)]$. Let $\epsilon_n > 0$ be a small constant that we will choose below. We let r_0 denote the largest radius r so that $V(r) \geq \epsilon_n r^{n-1}$. Since $V(r) \leq Vol F_0(\partial) < \infty$, such an r exists, and we have $V(r_0) = \epsilon_n r_0^{n-1}$. Since p is a regular point, $r_0 > 0$. By the definition of r_0, we see that $V(5r_0) < \epsilon_n (5r_0)^{n-1}$.

Now we consider the intersections $F_0(\partial X) \cap \partial B(p,r)$. Since F_0 is in general position, for almost every r, F_0 is transverse to $\partial B(p,r)$. By the coarea inequality, $\int_{r_0}^{5r_0} \mathrm{Vol}_{n-2}[F_0(\partial X) \cap \partial B(p,r)] \leq V(5r_0)$. Therefore, we can choose a generic value $r \in (r_0, 5r_0)$ so that $\mathrm{Vol}_{n-2} F_0(\partial X) \cap \partial B(p,r) \lesssim \epsilon_n r_0^{n-2} \leq \epsilon_n r^{n-2}$. We call this radius r a good radius, and we call $B(p,r)$ a good ball.

The good ball $B(p,r)$ obeys three important geometric estimates.

$$1.\ \mathrm{Vol}_{n-1}[F_0(\partial X) \cap B(p,r)] \gtrsim \epsilon_n r^{n-1}.$$

$$2.\ \mathrm{Vol}_{n-1}[F_0(\partial X) \cap B(p,5r)] \lesssim \epsilon_n r^{n-1}.$$

$$3. \operatorname{Vol}_{n-2}[F_0(\partial X) \cap \partial B(p,r)] \lesssim \epsilon_n r^{n-2}.$$

Using the Vitali covering lemma, we choose a finite collection of disjoint good balls $B_i = B(p_i, r_i)$ so that $\cup B(p_i, 5r_i)$ covers most of $F_0(\partial X)$. More precisely:

$$\operatorname{Vol}_{n-1}[F_0(\partial X) \cap (\cup B(p_i, 5r_i))] \geq \frac{1}{2} \operatorname{Vol}_{n-1} F_0(\partial X).$$

By properties 1 and 2, it follows that

$$\operatorname{Vol}_{n-1}[F_0(\partial X) \cap (\cup_i B_i)] \gtrsim \operatorname{Vol}_{n-1} F_0(\partial X). \tag{2}$$

Now we are ready to define the sets G_1 and X_1. We first extend F_0 to X in an arbitrary (generic) way. We denote this extension by F_0. We define $G_1 \subset X$ to be $F_0^{-1}(\cup_i B_i)$, and we define $X_1 \subset X$ to be the complement of G_1. Since we chose the radii r_i generically, we can assume that F_0 is transverse to ∂B_i, and so G_1 and X_1 are (piecewise smooth) manifolds with boundary.

We still have to define the partial extension $F_1 : G_1 \to \mathbb{R}^N$, extending F_0. We let $G_1(i) := F_0^{-1}(B_i)$. The boundary of $G_1(i)$ has two parts. One part of $\partial G_1(i)$ lies in ∂X. We let $Y_1(i)$ denote the rest of the boundary:

$$Y_1(i) := \partial G_1(i) \setminus \partial X.$$

We can think of $Y_1(i)$ as the set of x in the interior of X so that $F_0(x) \in \partial B_i$. Now $Y_1(i)$ is an $(n-1)$-dimensional piecewise smooth manifold with boundary. Its bounday lies in ∂X, and $F_0 : \partial Y_1(i) \to \partial B_i$. By Property 3 above, we know that

$$\operatorname{Vol}_{n-2} F_0(\partial Y_1(i)) = \operatorname{Vol}_{n-2}[F_0(\partial X) \cap \partial B(p_i, r_i)] \lesssim \epsilon_n r_i^{n-2}.$$

Now we can use the extension lemma for dimension $n-1$ to define a good map F_1 on $Y_1(i)$ extending $F_0 : \partial Y_1(i) \to B(p_i, r_i)$. The map F_1 will obey the following estimate:

$$\operatorname{Vol}_{n-1} F_1(Y_1(i)) \lesssim \operatorname{Vol}_{n-2} F_0(\partial Y_1(i))^{\frac{n-1}{n-2}} \lesssim I(n-1)\epsilon_n^{\frac{n-1}{n-2}} r_i^{n-1}.$$

(Here $I(n-1)$ denotes the constant in the extension lemma in dimension $n-1$.)

We have now defined F_1 on all of ∂X_1, and we are ready to prove the crucial volume estimate $\operatorname{Vol}_{n-1} F_1(\partial X_1) \leq (1-\delta_n) Vol_{n-1} F_0(\partial X)$. To do this, let $\partial X(i) := \partial X \cap F_0^{-1}(B_i)$. Now ∂X_1 is formed from ∂X by deleting each $\partial X(i)$ and adding in each $Y_1(i)$. Therefore,

$$\operatorname{Vol}_{n-1} F_1(\partial X_1) = \operatorname{Vol}_{n-1} F_0(\partial X) - \sum_i \operatorname{Vol}_{n-1} F_0(\partial X) \cap B_i + \sum_i \operatorname{Vol}_{n-1} F_1(Y_1(i)).$$

By the first property of a good ball, we know that $\mathrm{Vol}_{n-1} F_0(\partial X) \cap B_i \gtrsim \epsilon_n r_i^{n-1}$. On the other hand, $\mathrm{Vol}_{n-1} F_1(Y_1(i)) \lesssim I(n-1)\epsilon_n^{\frac{n-1}{n-2}} r_i^{n-1}$. The key observation is that we have a better power of ϵ_n in the volume bound for $F_1(Y_1(i))$. At this point we choose ϵ_n sufficiently small compared to the other dimensional constants to guarantee that

$$\mathrm{Vol}_{n-1} F_1(Y_1(i)) \leq \frac{1}{2} \mathrm{Vol}_{n-1} F_0(\partial X) \cap B_i. \tag{3}$$

Plugging this estimate in, we see that

$$\mathrm{Vol}_{n-1} F_1(\partial X_1) \leq \mathrm{Vol}_{n-1} F_0(\partial X) - \frac{1}{2} \sum_i \mathrm{Vol}_{n-1} F_0(\partial X) \cap B_i.$$

But as we noted above in Equation 2, the first two properties of a good ball imply that imply that $\sum_i \mathrm{Vol}_{n-1} F_0(\partial X) \cap B_i \gtrsim \mathrm{Vol}_{n-1} F_0(\partial X)$. Therefore, we conclude

$$\mathrm{Vol}_{n-1} F_1(\partial X_1) \leq (1 - \delta_n) \mathrm{Vol}_{n-1} F_0(\partial X).$$

This is one of the two estimates in the conclusion of the partial extension lemma.

Recall that G_1 is defined to be $X \setminus X_1$. In other words, $G_1 = F_0^{-1}(\cup_i B_i)$. We let $G_1(i)$ be $F_0^{-1}(\bar{B}_i)$. Next we have to define F_1 on G_1 and bound the volume $\mathrm{Vol}_n F_1(G_1)$. We will use the cone inequality to define F_1 on each $G_1(i)$.

Note that $\partial G_1(i) \subset Y_1(i) \cup \partial X$. We have already defined F_1 on $Y_1(i)$, and $F_1 = F_0$ on ∂X, and so we've already defined F_1 on $\partial G_1(i)$. We also have estimates about the volume of $F_1(Y_1(i))$ and $F_1(\partial X \cap G_1(i))$.

Using Property 2 of good balls, we see that

$$\mathrm{Vol}_{n-1} F_1(\partial X \cap G_1(i)) = \mathrm{Vol}_{n-1} F_0(\partial X) \cap B(p_i, r_i) \lesssim r_i^{n-1}.$$

And by Equation 3, we know that

$$\mathrm{Vol}_{n-1} F_1(Y_1(i)) \leq \frac{1}{2} \mathrm{Vol}_{n-1} F_0(\partial X) \cap B(p_i, r_i) \lesssim r_i^{n-1}.$$

Altogether, we see that $\mathrm{Vol}_{n-1} F_1(\partial G_1(i)) \lesssim r_i^{n-1}$. By the cone inequality, we can extend F_1 to all of $G_1(i)$ so that

$$\mathrm{Vol}_n F_1(G_1(i)) \lesssim r_i \mathrm{Vol}_{n-1} F_1(\partial G_1(i)) \lesssim r_i^n \sim [\mathrm{Vol}_{n-1} F_0(\partial X) \cap B_i]^{\frac{n}{n-1}}.$$

Adding the contributions from different balls B_i, we see that

$$\mathrm{Vol}_n F_1(G_1) \lesssim \sum_i [\mathrm{Vol}_{n-1} F_0(\partial X) \cap B_i]^{\frac{n}{n-1}} \lesssim Vol_{n-1} F_0(\partial X)^{\frac{n}{n-1}}.$$

We have now defined the partial extension $F_1 : G_1 \to \mathbb{R}^N$, and we have proven both estimates about F_1.

(We would like the map F_1 to send G_1 to $B^N(R)$. The construction above may not map G_1 to $B^N(R)$, but we can fix that problem in a simple way. Let $\pi : \mathbb{R}^N \to B^N(R)$ be the closest point map. (The map π is the identity on $B^N(R)$, and it maps each point outside of $B^N(R)$ to the closest point on the boundary of $B^N(R)$.) Since balls are convex, π is distance-decreasing, and so π decreases k-dimensional volumes for all k. The map $\pi \circ F_1$ maps G_1 to $B^N(R)$, and it obeys all the same estimates as F_1.)

□

This finishes the proof of the extension lemma.

□

3. Proof of the push-out lemma for small surfaces

In this section, we prove Lemma 0.6.

We explained in the introduction that there is some intuition for the pushout lemma coming from the monotonicity formula in minimal surface theory. We will prove the push-out lemma without using minimal surface theory, but the proof is based on some version of the monotonicity idea. We intersect $\phi_0(X)$ with various balls $B(p,r) \subset K$. If the volume of the intersection does not obey a monotonicity-type estimate, we can homotope ϕ_0 to a new map with smaller volume. This homotopy is accomplished using the extension lemma from the last section. We continue performing this type of homotopy, decreasing the total volume each time, until the amount of volume outside the W-neighborhood of ∂K gets as tiny as we like. Finally, we push this tiny volume into the W-neighborhood of ∂K using the Federer-Fleming push-out lemma.

Now we turn to the details. We homotope ϕ_0 to ϕ_1 in a sequence of small steps, and each step is given by the following lemma.

For any $W > 0$, we let $K(W)$ denote the set $\{x \in K | \operatorname{Dist}(x, \partial K) > W\}$.

Lemma 3.1. *Let $\delta > 0$ be any number. Suppose that $\Phi : (X, \partial X) \to (K, \partial K)$. Let $V = Vol_n \Phi(X)$. Define W so that $\sigma_n W^n = V$. Suppose that the volume of $\Phi(X) \cap K(W)$ is at least δ. Then we can homotope Φ rel ∂X to a new map Φ' so that its total volume is decreased by a definite amount:*

$$\operatorname{Vol}_n[\Phi'(X)] \leq \operatorname{Vol}_n[\Phi(X)] - c(\delta, n, K).$$

Proof. The constant σ_n in the formula above is the same constant as in our push-out lemma for small surfaces. We haven't chosen the constant yet. We will choose a sufficiently small $\sigma_n > 0$ below.

We find a point $p \in K(W)$ so that

$$\mathrm{Vol}_n[\Phi(X) \cap B(p, W/2)] \geq c(\delta, n, K).$$

We can find p by averaging over all points in $K(W)$. The total volume of $\Phi(X) \cap K(W)$ is assumed to be at least δ, and so we can find a ball that contains a certain definite amount of volume $c(\delta, n, K)$.

We let $V(R)$ denote the volume of $\Phi(X) \cap B(p, R)$. We let $A(R)$ denote the area of $\Phi(X) \cap \partial B(p, R)$.

Our argument will be based on the extension lemma from the last section. We recall that $I(n)$ is the isoperimetric constant in the extension lemma.

We call a radius R good if $W/2 < R < W$ and if $V(R)$ obeys the inequality

$$V(R) \geq 2I(n)A(R)^{\frac{n}{n-1}}.$$

For generic R, $\Phi^{-1}[\partial B(p, R)] \subset X$ is a piecewise smooth (n-1)-manifold, and so $\Phi^{-1}[B(p, R)] = X' \subset X$ is a piecewise smooth manifold with boundary (of the same dimension n as X). Note that $\mathrm{Vol}_{n-1} \Phi(\partial X') = A(R)$. Also, since $p \in K(W)$ and $R < W$, the ball $B(p, R)$ is contained in the interior of K, and so X' does not intersect ∂X.

If there is a generic good radius R, then we let $B := B(p, R)$, and we apply the extension lemma to the map $\Phi : (X', \partial X') \to (B, \partial B)$. The extension lemma gives us a new map $\Phi' : (X', \partial X') \to (B, \partial B)$ so that Φ' agrees with Φ on $\partial X'$ and we get the following volume estimate:

$$\mathrm{Vol}_n \Phi'(X') \leq I(n)A(R)^{\frac{n}{n-1}} \leq \frac{1}{2}V(R).$$

Now we extend Φ' to all of X by letting Φ' agree with Φ on $X \setminus X'$. Since K is convex, we can easily homotope Φ to Φ' rel ∂X. Most importantly, the volume $\Phi'(X)$ is smaller than the volume of $\Phi(X)$ by a definite amount:

$$\mathrm{Vol}_n \Phi(X) - \mathrm{Vol}_n \Phi'(X) = \mathrm{Vol}_n \Phi(X') - \mathrm{Vol}_n \Phi'(X') \geq \frac{1}{2}V(R) \geq 1/2\, \mathrm{Vol}_n[\Phi(X) \cap B(p, W/2)] \geq c(\delta, n, K).$$

This estimate suffices to prove the Lemma, provided that we can find a generic good radius R. Suppose that almost every radius R in the range $W/2 < R < W$ is not good. In other words, we have the following inequality.

$$V(R) \leq 2I(n)A(R)^{\frac{n}{n-1}}, \text{ for almost every } R \in [W/2, W].$$

But $V'(R) \geq A(R)$. Hence we have

$$V'(R) \geq [2I(n)]^{-\frac{n-1}{n}} V(R)^{\frac{n-1}{n}} \text{ for almost every } R \in [W/2, W].$$

Equivalently,

$$\frac{d}{dR}[V(R)^{1/n}] = (1/n)V'(R)V(R)^{-\frac{n-1}{n}} \geq (1/n)[2I(n)]^{-\frac{n-1}{n}}.$$

Integrating this inequality from $W/2$ to W, we see that

$$V(W) \geq C(n)W^n,$$

where $C(n) = n^{-n}2^{-(n-1)}I(n)^{-(n-1)}$. But $V(W) \leq V = \sigma_n W^n$. If we choose $\sigma_n > 0$ sufficiently small, we get a contradiction. Hence there exists a generic good radius R, and the lemma is proved. □

With this lemma, we can prove the push-out lemma for small surfaces.

Proof of the push-out lemma for small surfaces. We fix a $\delta > 0$ and we use Lemma 3.1 repeatedly. We get a sequence of maps ψ_j homotopic to ϕ_0 rel ∂X, and with $\mathrm{Vol}_n\, \psi_j(X) \cap K(W)$ decreasing steadily until $\mathrm{Vol}_n\, \psi_j(X) \cap K(W) < \delta$. We label this last map $\phi_{1/2}$. We know that $\phi_{1/2}$ is homotopic to ϕ_0 rel ∂X and that $\mathrm{Vol}_n\, \phi_{1/2}(X) \leq \mathrm{Vol}_n\, \phi_0(X)$ and the volume of $\phi_{1/2}(X) \cap K(W)$ is less than δ.

Next, we use the Federer-Fleming push-out lemma to remove the tiny volume from $K(W)$. Note that $K(W)$ is itself a convex set. We define $X' = \phi_{1/2}^{-1}(K(W))$. We note that $\phi_{1/2} : (X', \partial X') \to K(W)$. Using Lemma 0.5, we homotope $\phi_{1/2}$ rel $\partial X'$ to a new map $\phi_1 : X' \to \partial K(W)$ so that $\mathrm{Vol}_n\, \phi_1(X') \leq C(n, N, K)\delta$. Now we get a new map ϕ_1 from X to K, with image lying in the W-neighborhood of ∂K, and with $\mathrm{Vol}_n\, \phi_1(X) \leq \mathrm{Vol}_n\, \phi_0(X) + C(n, N, K(W))\delta$. Since we can choose δ as small as we like, we can arrange that $\mathrm{Vol}_n\, \phi_1(X)$ is as close as we like to $\mathrm{Vol}_n\, \phi_0(X)$.

In fact, with a little more work, we can arrange that $\mathrm{Vol}_n\, \phi_1(X) \leq \mathrm{Vol}_n\, \phi_0(X)$. This is not an important point, but it is convenient for keeping the notation simple when we apply the pushout lemma. We consider two cases. If $\mathrm{Vol}_n\, \phi_0(X) \cap K(W) = 0$, then we just apply the Federer-Fleming pushout as in the last paragraph to homotope ϕ_0 to ϕ_1 with image in the W-neighborhood of ∂K. If $\mathrm{Vol}_n\, \phi_0(X) \cap K(W) = V_0 > 0$, then we apply Lemma 3.1. The first application gives a map $\phi_{1/4}$ where $\mathrm{Vol}_n\, \phi_{1/4}(X) \leq \mathrm{Vol}_n(\phi_0(X)) - c(n, K, V_0)$. After noting the constant $c(n, K, V_0) > 0$, we now choose some $\delta > 0$, and we continue to apply Lemma 3.1 until we arrive at a map $\phi_{1/2} : X \to K$ so that $\mathrm{Vol}_n\, \phi_{1/2}(X) \cap K(W) < \delta$. At this point, we use the Federer-Fleming pushout argument to homotope $\phi_{1/2}$ to a map ϕ_1 with image in the W-neighborhood of K. The volume of $\phi_1(X)$ is at most

$$\mathrm{Vol}_n\, \phi_1(X) \leq \mathrm{Vol}_n(\phi_0(X)) - c(n, K, V_0) + c(n, N, K(W))\delta.$$

After noting $c(n, K, V_0)$, we choose δ sufficiently small so that $\mathrm{Vol}_n\, \phi_1(X) \leq \mathrm{Vol}_n\, \phi_0(X)$.

□

4. An example related to Uryson width

In this section, we give an example of a Riemannian manifold where each unit ball has small Uryson width, and yet the whole space has large Uryson width.

Proposition 4.1. *For any $\epsilon > 0$, there is a metric g_ϵ on S^3 so that the following holds. Every unit ball in (S^3, g_ϵ) has Uryson 2-width $< \epsilon$. The whole manifold (S^3, g_ϵ) has Uryson 2-width at least 1.*

The main idea of the proof comes from Example (H_1'') in [Gr2].

Proof. Let g_0 denote the standard unit-sphere metric on S^3. Let T denote a fine triangulation of (S^3, g_0). We choose T sufficiently fine so that the lengths of edges are at most δ, for a small constant $\delta > 0$ we may choose later. We let K_1 denote the 1-skeleton of the triangulation, and we let K_2 denote the dual 1-skeleton. There is one vertex of K_2 in each 3-face of T, and there is an edge connecting two vertices of K_2 if the corresponding 3-faces share a common 2-face in their boundaries.

Now we let U_1 denote the δ_1-neighborhood of K_1 for a small constant $\delta_1 << \delta$. We let U_2 denote $S^3 \setminus \bar{U}_1$ and we let Σ denote ∂U_1. By choosing δ_1 sufficiently small, we can arrange that $K_2 \subset U_2$.

Next we choose retractions $r_1 : \bar{U}_1 \to K_1$, and $r_2 : \bar{U}_2 \to K_2$. We can choose r_1 and r_2 to obey the following diameter estimates:

- For each $y \in K_i$, the diameter of $r_i^{-1}(y)$ within $\bar{U}_i$ is $\lesssim \delta$.
- For each $y \in K_i$, the diameter of $r_i^{-1}(y) \cap \Sigma$ within Σ is $\lesssim \delta$.

We define the space X to be $U_1 \cup (\Sigma \times [0, 10]) \cup U_2$, where the three pieces are glued together as follows. We glue $\partial U_1 = \Sigma$ to $\Sigma \times \{0\}$ using the identity map, and we glue $\partial U_2 = \Sigma$ to $\Sigma \times \{10\}$ using the identity map. We put a metric on X, where U_1, U_2, and Σ have the metric inherited from (S^3, g_0), and $\Sigma \times [0, 10]$ has the product metric. The space X is homeomorphic to S^3, and it is $(1 + \delta)$ bilipschitz to a Riemannian metric (S^3, g).

The Riemannian metric (S^3, g) is formed by taking a small tubular neighborhood of $\Sigma \subset S^3$, say $N = \Sigma \times (-\delta_2, \delta_2)$, and stretching it so that the $(-\delta_2, \delta_2)$ direction becomes long. As such $UW_2(S^3, g) \geq (1 + \delta)^{-1} UW_2(S^3, g_0) \gtrsim 1$ by the Lebesgue covering lemma.

(The exact value of $UW_2(S^3, g_0)$ is known by work of Katz, see [K].)

On the other hand, if $B_1 \subset (S^3, g)$ denotes any unit ball, then $UW_2(B_1) \leq (1 + \delta) UW_2(B)$ where $B \subset X$ is a ball of radius $(1 + \delta)$. Any such ball $B \subset X$ lies either in $X \setminus U_1$ or $X \setminus U_2$. Therefore, the result follows from the following lemma.

Lemma 4.2. *If $\delta > 0$ is sufficiently small, the Uryson width $UW_2(X \setminus U_i) \lesssim \delta$ for $i = 1, 2$.*

Proof of Lemma. The situation is the same for $i = 1, 2$. We write down the proof for $i = 1$.

We have to find a continuous map π from $X \setminus U_1$ to a 2-dimensional complex Y with small fibers. Our target Y is $K_2 \times [0, 10]$. The domain $X \setminus U_1$ is equal to $\Sigma \times [0, 10] \cup U_2$. We let (x, t) be coordinates on $\Sigma \times [0, 10]$. We define π on $\Sigma \times [0, 10]$ by $\pi(x, t) = (r_2(x), t) \in K_2 \times [0, 10]$. We let x denote the coordinate on $\bar{U}_2$. We define π on $\bar{U}_2$ by $\pi(x) = (r_2(x), 10) \in K_2 \times [0, 10]$. Since $\Sigma \times [0, 10]$ and U_2 are glued together by identifying $\Sigma \times \{10\}$ with $\partial U_2 = \Sigma$, the map π is a continuous map from $X \setminus U_1$ to $Y = K_2 \times [0, 10]$.

Next, we estimate the size of the fiber $\pi^{-1}(k, t)$, where $k \in K_2$ and $t \in [0, 10]$. The estimate has two cases. If $t < 10$, then the fiber has the form $(r_2^{-1}(k) \cap \Sigma) \times \{t\} \subset \Sigma \times [0, 10]$. So the diameter of $\pi^{-1}(k, t)$ is at most the diameter of $r_2^{-1}(k) \cap \Sigma$ within Σ, which is $\lesssim \delta$.

If $t = 10$, then the fiber $\pi^{-1}(k, t)$ has the form $r_2^{-1}(k) \subset \bar{U}_2$. So the fiber has diameter at most the diameter of $r_2^{-1}(k)$ in $\bar{U}_2$, which is also $\lesssim \delta$. □

So our metric (S^3, g) has Uryson 2-width $\gtrsim 1$ and yet every ball of radius 1 in (S^3, g) has Uryson width $\lesssim \delta$. By taking δ small and rescaling the metric a little, we get a metric (S^3, g_ϵ) as desired. □

Remark: We also note that this metric (S^3, g_ϵ) has volume ~ 1 and diameter ~ 1.

5. Open problems

The Szpilrajn theorem holds very generally for all compact metric spaces. Our theorem is only proven for Riemannian manifolds. But I don't know any counterexample to prevent Theorem 0.1 from generalizing to compact metric spaces.

Question 5.1. *Suppose that X is a compact metric space. Suppose that each unit ball of X has n-dimensional Hausdorff measure $< \epsilon_n$. If ϵ_n is chosen sufficiently small, does this imply that $UW_{n-1}(X) \leq 1$?*

In fact, something even more general based on the Hausdorff content looks very plausible. Recall that the n-dimensional Hausdorff content of a subset S in a metric space X is the infimum of $\sum_i r_i^n$ among all coverings of S by countably many balls $B(x_i, r_i)$.

Question 5.2. *Suppose that X is a compact metric space, and that each unit ball in X has n-dimensional Hausdorff content $< \epsilon_n$. If ϵ_n is chosen sufficiently small, does this imply that $UW_{n-1}(X) \leq 1$?*

This result would be stronger than our theorem even for Riemannian manifolds. In particular, it may apply to a Riemannian manifold (X^d, g) with dimension $d > n$. When $d > n$, the n-dimensional Hausdorff *measure* of a unit ball in X^d is always

infinite. But the n-dimensional Hausdorff content of a unit ball in X^d is always finite, and for some X it can be very small.

There are also open questions related to the funny example in Proposition 0.10.

Question 5.3. *Suppose that (M^n, g) is a Riemannian manifold so that each unit ball $B \subset M$ has $UW_q(B) < \epsilon$. If ϵ is sufficiently small, does this inequality imply anything about $UW_{q'}(M)$ for some $q' \geq q$?*

6. Appendix: non-compact manifolds and manifolds with boundary

Our main theorem also holds for compact Riemannian manifolds with boundary.

Corollary 6.1. *There exists $\epsilon_n > 0$ so that the following holds. If X is a compact Riemannian manifold with boundary, and if there is a radius R so that every ball of radius R in X has volume at most $\epsilon_n R^n$, then $UW_{n-1}(X) \leq R$.*

Proof. Let DX denote the double of X, which is a closed Riemannian manifold. If $x \in X$, then the ball in DX centered at x of radius R is contained in the double of the ball in X centered at x of radius R. Therefore, every ball of radius R in DX has volume at most $2\epsilon_n R^n$. If ϵ_n is sufficiently small, then Theorem 0.1 implies that $UW_{n-1}(DX) \leq R$. In other words, there is an (n-1)-dimensional polyhedron Y and a map $\pi : DX \to Y$ so that each fiber $\pi^{-1}(y)$ has diameter at most R in DX. We restrict π to a map $X \to Y$. Finally, we note that for two points $x_1, x_2 \in X$, the distance from x_1 to x_2 in X is equal to the distance in DX. Therefore, the diameter of each fiber $\pi^{-1}(y)$ in X is at most R, and we conclude that $UW_{n-1}(X) \leq R$. □

Our main theorem also extends to complete Riemannian manifolds in the following sense:

Theorem 6.2. *There exists $\epsilon_n > 0$ so that the following holds. If (M^n, g) is a complete Riemannian manifold, and if there is a radius R so that every ball of radius R in (M^n, g) has volume at most $\epsilon_n R^n$, then there is a continuous map $\pi : M \to Y$ for an infinite $(n-1)$-dimensional complex Y so that each fiber $\pi^{-1}(y)$ has diameter at most R in (M^n, g).*

The proof is essentially the same as the proof for closed manifolds. The main tricky issue is that the rectangular nerve N is not finite-dimensional but only locally finite dimensional. The cover by balls B_i will be locally finite, but the multiplicity of the cover may go to infinity at a sequence of points going to infinity in M. (See Section 6 of [Gu2] for an explanation of how to choose the balls B_i on a complete Riemannian manifold.) Therefore, each face of N is finite dimensional, and each face of N is adjacent to only finitely many other faces of N, but the dimension of the faces may be unbounded.

In the proof of Theorem 0.1, we began with a map $\phi = \phi^{(D)} : M \to N$, and we built a sequence of homotopic maps $\phi = \phi_D \sim \phi_{D-1} \sim ... \sim \phi_{n-1}$, where D was the dimension of N and ϕ_k maps M to the k-skeleton of N. In general, the dimension of N is not finite, but is only locally finite, and we must proceed a little differently. Instead, we construct an infinite sequence of maps $\phi_k : M \to N$ where each ϕ_k maps M to the k-skeleton of N and each ϕ_k is subordinate to our cover.

In a region of N where the dimension is less than k, we define ϕ_k to be the infinite composition $R_{\delta(k+1)} \circ R_{\delta(k+2)} \circ ...$ applied to ϕ. (This infinite composition is defined to be the limit of the maps $R_{\delta(k+1)} \circ ... \circ R_{\delta(N)}$ as N goes to infinity. The sequence of maps converges uniformly on compact sets.) In a region where the dimension of N is at least k, we define ϕ_k from ϕ_{k+1} as in the proof of Theorem 0.1. All the maps ϕ_k are subordinate to the cover, and the volume bounds work in the same way as in the proof of Theorem 0.1.

References

[Gr1] Gromov, M., Filling Riemannian manifolds, J. Differential Geom. 18 (1983) no. 1, 1-147.

[Gr2] Gromov, M., Width and related invariants of Riemannian manifolds. Astrisque No. 163-164 (1988), 6, 93-109, 282 (1989).

[Gr3] Gromov, M., Large Riemannian manifolds, in *Curvature and topology of Riemannian manifolds* (Katata, 1985), 108-121, Lecture Notes in Math. 1201, Springer, Berlin, 1986.

[Gr4] Gromov, M., *Metric Structures for Riemannian and Non-Riemannian Spaces*, Progress in Math. 152, Birkhauser Boston Inc., Boston, MA, 1999.

[Gu1] Guth, L., Notes on Gromov's systolic estimate. Geom. Dedicata 123 (2006), 113-129.

[Gu2] Guth, L., Guth, Volumes of balls in large Riemannian manifolds. Ann. of Math. (2) 173 (2011), no. 1, 51-76.

[Gu3] Guth, L. Metaphors in systolic geometry. Proceedings of the International Congress of Mathematicians. Volume II, 745-768, Hindustan Book Agency, New Delhi, 2010.

[HW] Hurewicz, W. and Wallman, H., *Dimension Theory*, Princeton Mathematical Series, v. 4. Princeton University Press, Princeton, N. J., 1941. vii+165 pp.

[K] Katz, M., Katz, M., The filling radius of two-point homogeneous spaces, J. Differential Geom. 18 (1983) no. 3, 505-511.

[MS] Michael, J.; Simon, L., Sobolev and mean-value inequalities on generalized submanifolds of $\mathbb{R}^n$, Comm. Pure Appl. Math 26 (1973), 361-379.

[W] Wenger, S., A short proof of Gromov's filling inequality. Proc. Amer. Math. Soc. 136 (2008), no. 8, 2937-2941.