

QUANTITATIVE NULL-COBORDISM

GREGORY R. CHAMBERS, DOMINIC DOTTERRER, FEDOR MANIN,
AND SHMUEL WEINBERGER,
WITH AN APPENDIX BY FEDOR MANIN AND SHMUEL WEINBERGER

1. INTRODUCTION

This paper is about two intimately related problems. One of them is quantitative algebraic topology: using powerful algebraic methods, we frequently know a lot about the homotopy classes of maps from one space to another, but these methods are extremely indirect, and it is hard to understand much about what these maps look like or how the homotopies come to be. The other is the analogous problem in geometric topology. The paradigm of this subject since immersion theory, cobordism, surgery, etc., has been to take geometric problems and relate them to problems in homotopy theory and, sometimes, algebraic K-theory and L-theory, and to solve those algebraic problems by whatever tools are available. As a result, we can solve many geometric problems without understanding at all what the solutions look like.

A beautiful example of this paradoxical state of affairs is the result of Nabutovsky that, despite the result of Smale (proved *inter alia* in the proof of the high-dimensional Poincaré conjecture) that every smooth codimension 1 sphere in the unit n -disk ($n > 4$) can be isotoped to the boundary, the minimum complexity of the embeddings required in the course of such an isotopy (measured by how soon normal exponentials to the embedding intersect) cannot be bounded by any recursive function of the original complexity of the embedding. Effectively, an easy isotopy would give such a sphere a certificate of its own simple connectivity, which is known to be impossible.

In other situations, such as those governed by an h -principle, a hard logical aspect of this sort does not arise. In this paper we introduce some tools of quantitative algebraic topology which we hope can be applied to showing that various geometric problems have solutions of low complexity.

As a first and, we hope, typical example, we study the problem, emphasized by Gromov, of trying to understand the work of Thom¹ on cobordism. Given a closed smooth (perhaps oriented) manifold, the cobordism question is whether it

Received by the editors December 21, 2016, and, in revised form, July 24, 2017, January 21, 2018, and May 18, 2018.

2010 *Mathematics Subject Classification*. Primary 53C23; Secondary 57R75.

The first author was partially supported by NSERC Postdoctoral Fellowship PDF-487617-2016.

The fourth author was partially supported by NSF grant DMS-1510178.

¹Thom solved the unoriented version of this exactly, and he only solved the rational version of the oriented question. However, later work of Milnor and Wall did the more difficult homotopy theory necessary for the oriented case.

bounds a compact (oriented) manifold. The answer to this is quite checkable: it is determined by whether the cycle represented by the manifold in the relevant (i.e., \mathbb{Z} or $\mathbb{Z}/2\mathbb{Z}$) homology of a Grassmannian (where the manifold is mapped in via the Gauss map classifying the manifold's stable normal bundle) is trivial.

This raises two questions: First, how is the geometry of a manifold reflected in the algebraic topological problem? Second, how difficult is it to find the null-homotopy predicted by the algebraic topology? As a test of this combined problem, Gromov suggested the following question: Given a manifold, assume away small scale problems by giving it a Riemannian metric whose injectivity radius is at least 1 and whose sectional curvature is everywhere between -1 and 1 . These properties can be achieved through a rescaling. A manifold possessing these properties will be said to have *bounded local geometry*. The geometric complexity of such a manifold can be measured by its volume.

If M is a smooth compact manifold, without a specified metric, we measure its (differential-) topological complexity by the infimum of the geometric complexity over all metrics with bounded local geometry. (If M is not closed, we require it to look like a collar $\partial M \times [0, 1]$ within distance 1 of the boundary.) This is a reasonable complexity measure: there are only finitely many diffeomorphism classes of manifolds with a given bound on complexity; see [Che70], [Pet84], [Gro98, §8D].

The central question is as follows. Given a smooth (oriented) manifold M^n of complexity V which is null-cobordant, what is the least complexity of a null-cobordism? That is, if W is an (oriented) compact Riemannian $(n + 1)$ -manifold of bounded local geometry, which bounds a manifold diffeomorphic to M , how small can the volume of W be? Gromov has observed [Gro96, §5 $\frac{5}{7}$ II] that tracing through the relevant mathematics would give a tower of exponentials of V (of size around the dimension of the manifold minus 2), but he has suggested [Gro99] that the truth might be linear.

The linearity problem, if it has an affirmative solution, would require very new geometric ideas and seemingly a solution to the cobordism problem essentially different from Thom's. We build on Thom's work to obtain the following:

Theorem A. *If M is an (oriented) closed smooth null-cobordant manifold of complexity V , then it has a null-cobordism of complexity at most*

$$c_1(n)V^{c_2(n)}.$$

The degree of this polynomial, obtained by tracing through our arguments, grows exponentially with dimension. In the Appendix, we improve this result to give an only slightly superlinear bound on the size of the null-cobordism. F. Costantino and D. Thurston have already shown that for 3-manifolds, one does not need worse than quadratic growth for the complexity of the null-cobordism² [CT08].

Our proof follows the ideas of Thom quite closely and is based on making those steps quantitative (if suboptimally) and then getting an a priori estimate on the size of the most efficient null-homotopy of a Thom map when the homological condition holds.

²Though they use a PL measure of complexity, the number of simplices in a triangulation.

Thom's work starts by embedding M into a sphere (or equivalently Euclidean space). This is already an act of violence: one knows that this will automatically introduce distortion. This is one source of growth that we do not know how to avoid.³

For manifolds embedded in the sphere, the Lipschitz constant of the Thom map is closely related to the complexity of the submanifold⁴ and the thickness of a tubular neighborhood. Conversely, if we know something about the Lipschitz constant of a null-homotopy of the Thom map, we can extract a geometrically bounded transverse inverse image.

Zooming in, we see three issues that need to be taken care of.

- (1) We need to bound the Lipschitz constants of the maps at time t in a null-homotopy (its "thickness"). Gromov has suggested [Gro99] that these frequently have a linear bound for maps of finite complexes into finite simply connected complexes.⁵
- (2) Bounding the worst Lipschitz constant arising in a null-homotopy does not quite suffice. One needs to bound the width⁶ of the null-homotopy as well. This is a nontrivial issue: a null-homotopy of thickness L can in general be replaced by one of width $\exp(L^d)$ where d is the dimension of the domain, but this is the best "automatic" bound.
- (3) Even provided such bounds, a transverse inverse image may be very large compared to the original manifold.

We deal with (1) and (2) simultaneously; this is the homotopy-theoretic result mentioned earlier. The real loss in our theorem comes from (3). In order to find a quantitative embedding of our manifold into S^N , we are forced to take N to be very large, and the embedded submanifold has small support in the resulting sphere. However, the support of a null-homotopy may still be quite large. This problem of the increase in the support is also one we have made no progress on and which seems important in a context broader than just cobordism theory.

1.1. Building Lipschitz homotopies. The main technical result of the paper is the following:

Theorem B. *Let X be an n -dimensional finite complex, and let Y be a finite complex which is rationally equivalent to a product of simply connected Eilenberg–MacLane spaces through dimension n . If $f, g : X \rightarrow Y$ are L -Lipschitz homotopic maps, then there is a homotopy between them which is $C(X, Y)L$ -Lipschitz as a map from $X \times [0, 1]$ to Y .*

The simplest settings in which this theorem applies are those in which Y is an odd-dimensional sphere or in which Y is a $2k$ -sphere and $n \leq 4k - 2$. More generally, Y may be any Lie group or, even more generally, H-space. Given that the targets in many topological problems are H-spaces, we are optimistic that this partial result

³A proof of the nonoriented cobordism theorem was given by [BH81] without using embedding. However, at a key moment there is a "squaring trick" in the proof, which also ends up giving, as a result of an induction, a polynomial estimate with an $\exp(n^2)$ -degree polynomial.

⁴Thom produces the null-cobordism from a null-homotopy by taking a transverse inverse image.

⁵If the domain is a circle and the target is a 2-complex, then for manifolds with an unsolvable word problem, there can be no computable upper bound for the worst Lipschitz constant in a null-homotopy. But for many groups with small Dehn function, it is possible to do this with only a linear increase. In particular, simple connectivity is an extremely natural requirement.

⁶The Lipschitz constant in the time direction.

regarding the linearity of homotopies will have more general application. (We give an example below showing that this theorem cannot be extended to arbitrary simply connected complexes in place of Y .)

One antecedent to this result is given in [FW13], where maps with target possessing finite homotopy groups are studied. In that setting, the width of a null-homotopy is actually bounded universally, independent of X . On the other hand, that paper shows that for any space with infinite homotopy groups there cannot be too uniform of an estimate of a linear upper bound on null-homotopies.

The obstruction in [FW13] has to do ultimately with homological filling functions. Isoperimetry, likewise, comes up in our result and is best appreciated by considering the following very concrete setting:

Lemma. *If $f : S^2 \rightarrow S^2$ is a degree 0 map with Lipschitz constant L , then there is a CL -Lipschitz null-homotopy for some C .*

This can be proved following the classical idea of Brouwer of cancelling point inverses with opposite local degree, but in a careful layered way so as to be able to control the Lipschitz constants. We will give a careful explanation of this as it provides the main intuition for the proof of Theorem B0.

1.2. Obstruction theory. Let $f : S^2 \rightarrow S^2$ be a null-homotopic L -Lipschitz map. We assume this has a very particular structure; later we will see that such a structure can be obtained with only small penalties on constants. The domain sphere X is a subdivision of a tetrahedron into grid isometric subsimplices, L to a side. The map f maps its 1-skeleton to the basepoint; for every 2-simplex either it also maps it to the basepoint or it maps a ball in the simplex homeomorphically to S^2 minus the basepoint, with degree ± 1 .

To construct a null-homotopy of f , we need to connect the positive and negative preimages with tubes in $X \times [0, 1]$. Care must be taken to route these tubes in such a way that there are not too many clustered in any given spot, as in Figure 1. To






FIGURE 1. Connecting preimages of opposite orientations with tubes: the global picture. Note that the Lipschitz constant of a null-homotopy depends only on the thickness of the tubes; therefore, inefficiencies in routing only matter insofar as they force many tubes to bunch up in the same region.



(a) Degrees on 2-cells of prisms.



(b) Connecting homeomorphic preimages of $S^2 \setminus *$ with tubes.

FIGURE 2. Constructing a null-homotopy: the local picture.

do this, we decide beforehand how many tubes need to go through any particular part of $X \times [0, 1]$ and then connect them up in any available way.

To make this precise, assume that the tubes miss $X^{(0)} \times [0, 1]$. Then we can count the number of tubes going through $p \times [0, 1]$ for each 1-simplex p of X . Every tube that goes into $q \times [0, 1]$, for any 2-simplex q , must either come out through another edge or come back to 0. In other words, if $\alpha \in C^1(X; \mathbb{Z})$ is the cochain which indicates the number of tubes (with sign!) going through $p \times [0, 1]$, then $\omega = \delta\alpha$ gives the degree of f on 2-simplices of X . In the language of obstruction theory, ω is the obstruction to null-homotoping f , and the existence of α demonstrates that the obstruction can be resolved.

To ensure that it can be resolved efficiently, we need to pick a relatively small α . The best we can do is to choose an α which takes values $\leq CL$. By considering a situation with degree $O(L^2)$ on one side of X canceling out degree $-O(L^2)$ on the other side, we see that we can do no better. That this is also the worst possible situation follows from the classical isoperimetric inequality for spheres; this is discussed in much greater generality in section 3.

In effect, once we have set α , deciding how many tubes must go through a given point, we can connect them up in an entirely local way. We give $X \times [0, 1]$ a cellulation by prisms of length $1/CL$ and base the 2-simplices of X . We then construct the map F by skeleta on this cellulation:

- (1) First, map the 1-skeleton to the basepoint.
- (2) Next, we can map the 2-cells via maps of degree between -3 and 3 in such a way that the map on the boundary of each prism has total degree 0, as in Figure 2(a). (It is here that we “layer” the null-homotopy.)

- (3) Finally, we choose a way to connect pairs of preimages on each prism via tubes, as in Figure 2(b). Since the number of tubes in each prism is bounded, we can do this with bounded Lipschitz constant.

For the second step, we need to use our α . If we ensure that for each 1-simplex p of X , the degree of F on $p \times [0, 1]$ is $\langle \alpha, p \rangle$, then F will have degree 0 on the boundary of each “long prism” $q \times [0, 1]$, where q is a 2-simplex of X .

It remains to make sure that the degree is 0 on the “short prisms”. To do this, we spread $\langle \alpha, p \rangle$ as evenly as possible along the unit interval: for every integer $1 \leq t \leq CL$, the degree of α on $p \times [0, t/CL]$ is $\lfloor \frac{t}{CL} \langle \alpha, p \rangle \rfloor$. This then also determines the required degree on $q \times \{t/CL\}$ for every 2-simplex q and time t to make the total degree on the boundary of each prism 0. It is easy to check that the resulting degrees on all 2-cells are at most 3.

1.3. Outline of proof of Theorem B. We now describe how the proof of the above Lemma leads to the proof of Theorem B. The motto is the same: if we can kill the obstruction to finding a homotopy, then we can do the killing in a bounded way.

The first step is to reduce to a case where obstruction theory applies. For this, we simplicially approximate our map in a quantitative way. That is, given a map $X \rightarrow Y$ between metric simplicial complexes, the fineness of the subdivision of X must be inversely proportional to the Lipschitz constant of the map.

From here the general strategy is to build a homotopy by induction on the skeleta of $X \times I$ with a product cell structure. This homotopy will not in general be simplicial, but it will have the property that restrictions to each cell form a fixed finite set depending only on X and Y . Every time we run into a null-cohomologous obstruction cocycle, we use a cochain that it bounds to modify the map on the previous skeleton. We ensure that these modifications are chosen from a fixed finite set of maps, leaving us with a fixed finite set of maps on the boundaries of cells one dimension higher. Then we can fill each such map in a fixed way, preserving the desired property.

When the obstructions are torsion, the main issue is the well-known one that killing obstructions “blindly” will sometimes lead to a dead end even when a homotopy exists. On the other hand, since there is a finite number of choices of torsion values for a cochain to take, we may avoid this by following a “road map” given by a known, but potentially uncontrolled, null-homotopy of f . This is the content of Lemma 4.1.

On one hand, when we get integral obstructions, our choice of rational homotopy structure ensures that such issues do not come up. On the other hand, we do need to worry about isoperimetry. This is covered by Theorem 4.2, which generalizes the argument above.

2. PRELIMINARIES

In this section, we discuss how to subdivide a metric simplicial complex so that the edges all have length approximately $1/L$ for a specified L . We also show that, for any simplicial map $f : X \rightarrow Y$ and any L , we can subdivide X as above to form X_L and homotope f through a short homotopy to $\tilde{f} : X_L \rightarrow Y$.

2.1. Regular subdivision of simplices.

Definition. Define a *simplicial subdivision scheme* to be a family, for every n and L , of metric simplicial complexes $\Delta^n(L)$ isometric to the standard Δ^n with length 1 edges, such that $\Delta^n(L)$ restricts to $\Delta^{n-1}(L)$ on all faces. A subdivision scheme is *regular* if for each n there is a constant A_n such that $\Delta^n(L)$ has at most A_n isometry classes of simplices and a constant r_n such that all 1-simplices of $\Delta^n(L)$ have length in $[r_n^{-1}L^{-1}, r_nL^{-1}]$.

Given a regular subdivision scheme, we can define the *L-regular subdivision* of any metric simplicial complex, where each simplex is replaced by an appropriately scaled copy of $\Delta^n(L)$.

Note that L times barycentric subdivision is *not* regular. On the other hand, there are at least two known examples of regular subdivision. One is the edgewise subdivision of Edelsbrunner and Grayson [EG00], which has the advantages that the L -regular subdivision of $\Delta^n(M)$ is $\Delta^n(LM)$ and that the lengths of edges vary by a factor of only $\sqrt{2}$. Roughly, the method is to cut the simplex into small polyhedra by planes parallel to the $(n - 1)$ -dimensional faces, then partition each such polyhedron into simplices in a standard way. The other is described by Ferry and Weinberger [FW13]: the trick is to subdivide Δ^n into $n + 1$ identical cubes, then subdivide these in the obvious way into L^n cubes, and finally subdivide these in a canonical way into simplices. This method has the advantage of being easy to describe.

None of the listed advantages is crucial for our continued discussion, so we may remain agnostic as to how precisely we subdivide our simplices.

2.2. Simplicial approximation.

Proposition 2.1 (Quantitative simplicial approximation theorem). *For finite simplicial complexes X and Y with piecewise linear metrics, there are constants C and C' such that any L -Lipschitz map $f : X \rightarrow Y$ has a CL -Lipschitz simplicial approximation via a $(CL + C')$ -Lipschitz homotopy.*

Proof. We trace constants through the usual proof of the simplicial approximation theorem, as given in [Hat01].

Denote the open star of a vertex v by $\text{st } v$. Let c be a Lebesgue number for the open cover $\{\text{st } w \mid w \text{ is a vertex of } Y\}$ of Y , that is, a number such that every c -ball in Y is contained in one of the sets in the cover. Then c/L is a Lebesgue number for the open cover $\{f^{-1}(\text{st } w)\}$ of X . Take a regular subdivision X_L of X so that for some $0 < d(X) < 1/2$ each simplex of X_L has diameter between dc/L and $c/2L$. Hence f maps the closed star of each vertex v of X_L to the open star of some vertex $g(v)$ of Y . This gives us a map $g : X_L^{(0)} \rightarrow Y^{(0)}$ which takes adjacent vertices of X_L to adjacent vertices of Y , and hence if ℓ is the maximum edge length of Y , g is $\ell L/dc$ -Lipschitz.

By a standard argument, this map g extends linearly to a map $g : X_L \rightarrow Y$ with the same Lipschitz constant. The linear homotopy from f to g has Lipschitz constant $\max\{\ell L/dc, \ell\}$. □

Remark. Suppose that Y and X are n -dimensional and made up of standard simplices of edge length 1. Then $c = \frac{1}{\sqrt{2n(n+1)}}$ is the inradius of a standard

simplex, and by using the edgewise subdivision, we can make sure that $d > 1/2\sqrt{2}$. Thus the Lipschitz constant of the map increases by a factor of at most

$$C \leq 4\sqrt{n(n+1)}.$$

Furthermore, if X is two-dimensional, then all of the edge lengths of the subdivision are equal. Therefore, in this case, $C \leq 4\sqrt{3}$ and, in fact, it approaches $2\sqrt{3}$ for large L , since we can choose a subdivision parameter very close to L and, thus, d very close to 1.

We will use simplicial approximation mainly as a way of ensuring that our maps have a uniformly finite number of possible restrictions to simplices. Almost all instances of “simplicial” in this paper can be replaced with “such that the restrictions to simplices are chosen from a finite set associated with the target space”. This formulation makes sense even when the target space is not a simplicial complex. In particular, it is preserved by postcomposition with any map, for example one collapsing certain simplices.

3. ISOPERIMETRY FOR INTEGRAL COCHAINS

The goal of this section is to prove the following (co)isoperimetric inequality.

Lemma 3.1 (ℓ^∞ coisoperimetry). *Let X be a finite simplicial complex equipped with the standard metric, and let X_L be the cubical or edgewise L -regular subdivision of X , and let $k \geq 1$. Then there is a constant $C_{\text{IP}} = C_{\text{IP}}(X, k)$ such that for any simplicial coboundary $\omega \in C^k(X_L; \mathbb{Z})$ there is an $\alpha \in C^{k-1}(X_L; \mathbb{Z})$ with $d\alpha = \omega$ such that $\|\alpha\|_\infty \leq C_{\text{IP}}L\|\omega\|_\infty$.*

We will start by proving the much easier version over a field; in the rest of the section \mathbb{F} will denote \mathbb{Q} or \mathbb{R} . Then we will demonstrate how to find an integral-filling cochain near a rational or real one.

Lemma 3.2. *Let X be a finite simplicial complex equipped with the standard metric, and let X_L be an L -regular subdivision of X . Then for any k , there is a constant $K = K(X, k)$ such that for any simplicial coboundary $\omega \in C^k(X_L; \mathbb{F})$, there is an $\alpha \in C^{k-1}(X_L; \mathbb{F})$ with $d\alpha = \omega$ such that $\|\alpha\|_\infty \leq KL\|\omega\|_\infty$.*

Proof. We first show a similar isoperimetric inequality and then demonstrate that it is equivalent to the coisoperimetric version.

Lemma 3.3. *There is a $K = K(X, k)$ such that boundaries $b \in C_{k-1}(X_L; \mathbb{F})$ of simplicial volume V bound chains of simplicial volume at most KL .*

Proof. There are two ways we can measure the volume of a simplicial i -chain in X_L . The first, *simplicial volume*, is given by assigning every simplex volume 1, i.e.,

$$\text{vol} \left(\sum \alpha_i p_i \right) = \sum |\alpha_i|.$$

Alternatively, we can measure the *i -mass* of chains: the mass of a simplex p is its Riemannian i -volume, and in general

$$\text{mass} \left(\sum \alpha_i p_i \right) = \sum |\alpha_i| \text{vol}_i(p_i).$$

Thus there are constants K_i and K'_i , depending on the choice of subdivision scheme, such that for every i -chain c ,

$$K_i L^i \text{mass } c \leq \text{vol } c \leq K'_i L^i.$$

Therefore to prove the lemma it suffices to show that a boundary whose $(k-1)$ -mass in X is V bounds a chain whose k -mass is at most KV .

Our main tool here is the Federer–Fleming deformation theorem, a powerful result in geometric measure theory which allows very general chains to be deformed to simplicial ones in a controlled way. One proves this result by shining a light from the right spot inside each simplex so that the resulting shadow on the boundary of the simplex is not too large. By iterating this procedure on simplices of each dimension between n and $k+1$, we eventually end up with a shadow in the k -skeleton, which is the desired simplicial chain. Federer and Fleming’s original version [FF60, Thm. 5.5] was based on deformation to the standard cubical lattice in \mathbb{R}^n . However, everything in their proof, except for the precise constants, translates to simplicial complexes. (See [EPC⁺92, Thm. 10.3.3] for a proof of a slightly narrower analogue in the case of triangulated manifolds, which however also applies to any simplicial complex.)

Federer and Fleming’s theorem works for normal currents. To avoid this rather technical concept, we state the result for Lipschitz chains, that is, singular chains whose simplices are Lipschitz.

Theorem (Federer–Fleming deformation theorem). *Let W be an n -dimensional simplicial complex with the standard metric on each simplex. There is a constant $\rho(k, n)$ such that the following holds. Let T be a Lipschitz k -chain in W with coefficients in \mathbb{F} . Then we can write $T = P + Q + \partial S$, where*

- (1) $\text{mass}(P) \leq \rho(k, n)(\text{mass}(T) + \text{mass}(\partial T))$;
- (2) $\text{mass}(Q) \leq \rho(k, n) \text{mass}(\partial T)$;
- (3) $\text{mass}(S) \leq \rho(k, n) \text{mass}(T)$;
- (4) P can be expressed as an \mathbb{F} -linear combination of k -simplices of W ;
- (5) if ∂T can already be expressed as a combination of $(k-1)$ -simplices of W (for example, if T is a cycle), then $Q = 0$ and

$$\text{mass}(P) \leq \rho(k, n)(\text{mass}(T)).$$

Now suppose that W is given a metric d_W whose simplices are not standard, but such that the identity map $\iota : (W, d_{\text{std}}) \rightarrow (W, d_W)$ satisfies

$$\lambda_1 d(x, y) \leq d(\iota(x), \iota(y)) \leq \lambda_2 d(x, y)$$

for all $x, y \in W$. When mass is measured with respect to d_W , the bounds in the theorem become

- (1) $\text{mass}(P) \leq \rho(k, n) \left(\frac{\lambda_2^k}{\lambda_1^k} \text{mass}(T) + \frac{\lambda_2^{k-1}}{\lambda_1^k} \text{mass}(\partial T) \right)$;
- (2) $\text{mass}(Q) \leq \frac{\lambda_2^{k-1}}{\lambda_1^k} \rho(k, n) \text{mass}(\partial T)$;
- (3) $\text{mass}(S) \leq \frac{\lambda_2^{k+1}}{\lambda_1^k} \rho(k, n) \text{mass}(T)$.

We apply the theorem twice. First, we apply it to b as a Lipschitz cycle in X to show that it is homologous to a $(k-1)$ -cycle $P \in C_{k-1}(X; \mathbb{F})$ of volume $\leq C(k)V$ via a Lipschitz k -chain S of volume $\leq C(k)V$. Next, we apply it to S as a k -chain in X_L . Notice that the ratio λ_2/λ_1 is bounded independent of L for a regular subdivision; therefore, S deforms rel boundary to a chain in $C_k(X_L; \mathbb{F})$ of volume $\leq C(k)C'(k)V$, where C' depends on the subdivision scheme. Finally, P bounds a

chain in X of volume $\leq C^\dagger C(k)V$, where C^\dagger depends only on the geometry of X . Thus we can set $K = (C^\dagger + C')C$. \square

Lemma 3.4. *Let X be a finite simplicial complex. Then the following are equivalent for any constant C :*

- (1) *any boundary $\sigma \in C_{k-1}(X; \mathbb{F})$ has a filling $\tau \in C_k(X; \mathbb{F})$ with*

$$\text{vol } \tau \leq C \text{ vol } \sigma;$$
- (2) *any coboundary $\omega \in C^k(X; \mathbb{F})$ is the coboundary of some $\alpha \in C^{k-1}(X; \mathbb{F})$ with $\|\alpha\|_\infty \leq C\|\omega\|_\infty$.*

The authors would like to thank Alexander Nabutovsky and Vitali Kapovitch for pointing out this simplified proof.

Proof. The cochain complex is dual to the chain complex, and the L_∞ -norm on cochains is dual to the volume norm on chains. So consider the general situation of a linear transformation between two normed vector spaces $T : (V, \|\cdot\|_V) \rightarrow (W, \|\cdot\|_W)$, and let $C(T)$ be the operator norm of the transformation

$$\bar{T}^{-1} : (\text{im } T, \|\cdot\|_W) \rightarrow (V/\ker T, \|\cdot\|_{\bar{V}}),$$

where the norm of an equivalence class $\bar{v} \in V/\ker T$ is given by $\|\bar{v}\|_{\bar{V}} = \min_{v \in \bar{v}} \|v\|_V$. When T is the boundary operator on $C_k(X; \mathbb{F})$, $C(T)$ is exactly the minimal constant C in condition (1). Hence this is also the operator norm of the dual transformation $(\bar{T}^{-1})^* : \text{im } T^* \rightarrow W^*/\ker T^*$. It remains to investigate the dual norms on these spaces.

By the Hahn–Banach theorem, any operator on $\text{im } T$ extends to an operator of the same norm on all of W . Hence the dual norm of $\|\cdot\|_W|_{\text{im } T}$ is exactly the norm $\|\bar{\varphi}\|_{\bar{W}^*} = \min_{\varphi \in \bar{\varphi}} \|\varphi\|_{W^*}$ on $W^*/\ker T^*$, and similarly the dual norm of $\|\cdot\|_{\bar{V}}$ is $\|\cdot\|_{V^*}|_{\text{im } T^*}$. Therefore, the operator norm of $(\bar{T}^{-1})^*$ is the minimal constant of condition (2). \square

Combining Lemmas 3.3 and 3.4, we complete the proof of the rational and real versions of the coisoperimetry lemma. \square

Now we introduce the ingredients for proving the integral version.

Definition. A k -spanning tree of a simplicial complex X is a k -dimensional subcomplex T which contains $X^{(k-1)}$, such that the induced map

$$H_{k-1}(T; \mathbb{Q}) \rightarrow H_{k-1}(X; \mathbb{Q})$$

is an isomorphism and $H_k(T; \mathbb{Q}) = 0$. A k -wrapping tree of X is a k -dimensional subcomplex U which contains $X^{(k-1)}$ and such that the induced maps

$$H_{k-1}(U; \mathbb{Q}) \rightarrow H_{k-1}(X; \mathbb{Q}) \quad \text{and} \quad H_k(U; \mathbb{Q}) \rightarrow H_k(X; \mathbb{Q})$$

are both isomorphisms.

Lemma 3.5. *Every simplicial complex X has a k -spanning tree and a k -wrapping tree.*

Proof. A k -spanning tree for any X can be constructed greedily starting from $X^{(k-1)}$. At each step, we find a k -simplex c in X such that ∂c represents a nonzero class in $H_{k-1}(T; \mathbb{Q})$ and add it to T . Once there are no such simplices left, $H_{k-1}(T; \mathbb{Q}) \rightarrow H_{k-1}(X; \mathbb{Q})$ is an isomorphism. By construction, T has no rational k -cycles.

Notice that every k -simplex of X outside T is a cycle in $C_k(X, T; \mathbb{Q})$. To build a k -wrapping tree from a k -spanning tree, we may choose a basis for $H_k(X, T; \mathbb{Q})$ from among the simplices and add it to the tree. \square

Informally speaking, a k -spanning tree should be thought of as the least subcomplex T so that every k -simplex outside T is a cycle mod T ; a k -wrapping tree is the least subcomplex U so that every k -simplex outside U is a boundary mod U . In both cases the minimality means that there is a unique “completion” for a k -simplex q , i.e., a chain c supported in T (resp., U) so that $c + q$ is a cycle (resp., boundary).

Such spanning trees have been previously studied by Kalai [Kal83] and Duval, Klivans, and Martin [DKM09] and [DKM11] in the case where k is the dimension of the complex. In that case the k -simplices not contained in a spanning tree T form a basis for $H_k(X, T; \mathbb{Q})$ (and a k -wrapping tree is simply the whole complex). When X contains simplices in dimension $k + 1$, however, there may be relations between the simplices when viewed as cycles in X modulo T . The next definition attempts to quantify the extent to which such relations constrain the behavior of cocycles in the pair (X, T) .

Definition. Let T be a k -spanning tree of X . Consider the set \mathcal{A} of vectors in $H_k(X, T; \mathbb{Q})$ which are images of k -simplices of X . We define the *gnarledness*

$$G(T) = \min \left\{ \max_{a \in \mathcal{A}} \|a\|_1 : \text{bases } \mathcal{B} \text{ for } H_k(X, T; \mathbb{Q}) \text{ such that } \mathcal{A} \subset \mathbb{Z}\mathcal{B} \right\}.$$

We say that T is $G(T)$ -*gnarled*; we say a basis is *optimal* if $\max_{a \in \mathcal{A}} \|a\|_1$ is minimal in it.

The gnarledness measures the extent to which certain simplices are homologically “larger” than others. For example, consider a two-dimensional simplicial complex which is homeomorphic to the mapping telescope of a degree 2 self-map of S^1 ,

$$X = S^1 \times [0, 1] / (x, 1) \sim (-x, 1).$$

Let us say we take a one-dimensional spanning tree T which includes all but one of the simplices of both $S^1 \times \{0\}$ and $S^1 \times \{1\}$; let e_0 and e_1 , respectively, be the relevant 1-simplices in $X \setminus T$. Then in $H_1(X, T; \mathbb{Q}) \cong \mathbb{Q}$, $[e_0] = 2$ and $[e_1] = 1$. For any basis for $H_1(X, T; \mathbb{Q})$ in which e_1 is a lattice point, $\|e_0\|_1 \geq 2$, so the tree T is at least 2-gnarled. Indeed, the same will happen for any spanning tree of this complex.

Lemma 3.6. *The cubical and edgewise L -regular subdivisions of X both admit k -spanning trees which are at most $C(X)$ -gnarled; the gnarledness is bounded independent of L .*

We will actually show this for grids in a cube complex. It is routine to modify this proof to work for the cubical subdivision of a simplicial complex; a similar construction works for the edgewise subdivision, since it consists of a grid of subspaces parallel to the faces which is then subdivided in a fixed way depending on dimension.




FIGURE 3. An illustration of the subcomplex T for $n = 2$, $k = 1$ and $n = 3$, $k = 2$.

We first show that the subdivision of a cube has a “ k -spanning tree rel boundary” with good geometric properties. To be precise:

Lemma 3.7. *Let $K = I^n$ be cubulated by a grid of side length $1/r$, and let $k \leq n$. We refer to*

- cells, *i.e.*, faces of the cubulation;
- faces, *i.e.*, subcomplexes corresponding to faces of the unit cube; and
- boxes, *i.e.*, subcomplexes which are products of subintervals.

Then there is a k -subcomplex $T \supset K^{(k-1)}$ of K with the following properties:

- (1) $T \cap \partial K = (\partial K)^{(k-1)}$.
- (2) T deformation retracts to $(\partial K)^{(k-1)}$.
- (3) Every k -cell of $K \setminus T$ is homologous rel T to a chain in ∂K whose intersection with each $(n-1)$ -face is a box.
- (4) More generally, every k -dimensional box in K is homologous rel T to a chain in ∂K whose intersection with each $(n-1)$ -face is a box.

This subcomplex is illustrated in low dimensions in Figure 3.

Suppose now that we equip every face of K in dimensions $k \leq i \leq n$ with subcomplexes satisfying these properties, and we let T' be the union of all these subcomplexes. Then by induction using property (4), any k -cell of $K \setminus T'$ is homologous rel T' to a union of at most $2^{n-k}n!/k!$ boxes in the k -faces of K . In turn, by property (2), each of these boxes has at most one cell outside T' . Therefore any k -cell is homologous rel T' to a sum of at most $2^{n-k}n!/k!$ cells in the k -faces. This is the property that we use to prove Lemma 3.6.

Proof. We construct $T = T_{n,k}$ by induction on n and k . For $k = 0$, we can set $T_{n,0} = \emptyset$. Similarly, for $k = n$, we can take $T_{n,n}$ to be K less the interior of any one cell—for concreteness, let that be the cell that includes the origin.

Now we construct $T_{n,k}$ for $n > k > 0$ by induction on n . Write $K = K' \times I$; then for every $0 < i < r$, let $T|_{K' \times \{i/r\}} = T_{n-1,k} \times \{i/r\}$, and for every $0 \leq i < r$, let

$$T|_{K' \times (i/r, i+1/r)} = T_{n-1, k-1} \times (i/r, i+1/r).$$

Finally, we throw in $K^{k-1} \cap K' \times \{0, 1\}$. It remains to show that the resulting complex $T = T_{n,k}$ satisfies the lemma.

It is clear that T contains $K^{(k-1)}$ and that condition (1) holds. Moreover, T deformation retracts first to $T|_{\partial K \cup K' \times \{i/r: i=1,2,\dots,r-1\}}$ using a retraction of $T_{n-1,k-1}$, and thence to $(\partial K)^{(k-1)}$ via a retraction of each layer individually. This demonstrates (2). It remains to show that (3) and (4) hold.

In order to do this more easily, we present an alternate rule for determining whether a k -cell c is contained in T . Showing that it is indeed equivalent to the previous definition is tedious but straightforward. Let c be a k -cell of $K \setminus \partial K$ and let $\mathcal{I}(c) \subset \{1, \dots, n\}$ be the set of directions in which it has positive width. Write π_i for the projection of K onto its i th interval factor, and write $\ell(c)$ for the greatest integer such that $\{1, \dots, \ell\} \subset \mathcal{I}(c)$. Then c is in T if and only if $\pi_i c \neq [0, 1/r]$ for some $i \leq \ell$. In particular, if $\ell = 0$, then $c \notin T$.

Now let c be a k -cell of $K \setminus T$. If $c \in \partial K$, then it already fits the bill, so suppose it is in $K \setminus (T \cup \partial K)$. We will argue that c is bounded rel $T \cup \partial K$ by a box B with positive width in directions $\mathcal{I}(c) \cup \{\ell(c) + 1\}$. Specifically, the projections of B onto each interval factor of K are as follows:

$$\begin{aligned} \pi_i B &= I, & \text{if } 1 \leq i \leq \ell(c); \\ \pi_i B &= [x, 1], \text{ where } \pi_i c = \{x\}, & \text{if } i = \ell(c) + 1; \\ \pi_i B &= \pi_i c, & \text{otherwise.} \end{aligned}$$

By the criterion above, $\partial B \setminus \partial K$ contains only one cell which is not in T , namely c . Thus $\partial B \cap \partial K$ is the chain desired for (3).

More generally, given a k -dimensional box in K , one can take the union of the B 's constructed for each cell in the box. This gives a solution for (4). □

Proof of Lemma 3.6 for cubulations. Let X_L be the complex obtained by dividing X into grids at scale $1/L$. We begin by choosing a k -spanning tree T for X , then use it to build a k -spanning tree T_L for X_L . We include all cells of X_L contained in T ; for every cell of X not contained in T , we include a complex as in Lemma 3.7. The resulting subcomplex includes $X_L^{(k-1)}$ and, by induction on $n - k$, deformation retracts to T . Therefore it is a k -spanning tree for X_L .

Now let \mathcal{B} be an optimal basis for $H_k(X, T; \mathbb{Q}) \cong H_k(X_L, T_L; \mathbb{Q})$. By the argument above, any k -cell of X_L is homologous rel T_L to a sum of at most $2^{n-k}n!/k!$ cells in the k -faces, where n is the dimension of X . In turn, any k -cell of $X_L \setminus T_L$ which is contained in a k -face represents the same homology class modulo T_L as that face does modulo T , and therefore can be represented as a sum of at most $G(T)$ elements of \mathcal{B} . Therefore, $G(T_L) \leq G(T) \cdot 2^{n-k}n!/k!$. □

We now have the tools we need to prove Lemma 3.1. We will do this by way of two auxiliary lemmas. The first states that any cochain with coefficients in \mathbb{F}/\mathbb{Z} which can be lifted to \mathbb{F} can be lifted to a cochain which is not too big.

Lemma 3.8 (Bounded lifting). *Let X be a finite simplicity complex. Fix a cocycle $z \in Z^k(X; \mathbb{F}/\mathbb{Z})$ and a k -spanning tree T of X . Then if z lifts to a cocycle $\tilde{z} \in Z^k(X; \mathbb{F})$, we can find such a lift \tilde{z} with $\|\tilde{z}\|_\infty \leq k + 1 + G(T)$.*

Proof. Fix U , a $(k - 1)$ -wrapping tree of X . Then for every $(k - 1)$ -simplex p of $X \setminus U$, there is a unique k -chain $F(p)$ supported in T which fills $p \bmod U$. Moreover,

$$\mathcal{F} = \{F(p) : p \text{ is a } (k - 1)\text{-simplex of } X \setminus U\}$$

is a basis for $C_k(T)$: they are linearly independent since their boundaries are linearly independent in $C_{k-1}(X)$, and any k -simplex q in T can be expressed as an integral linear combination $\sum_{p \in \partial q} F(p)$. We can therefore extend F by linearity to an isomorphism $F : C_{k-1}(X; \mathbb{F}) \rightarrow C_k(T; \mathbb{F})$.

Now let \mathcal{B} be an optimal basis for $H_k(X, T; \mathbb{F})$ which demonstrates that T is $G(T)$ -gnarled. For every $b \in \mathcal{B}$, choose a $\hat{b} \in C_k(X, T; \mathbb{F})$ representing it, and let

$$\tilde{\mathcal{B}} = \{\hat{b} - F(\partial \hat{b}) : b \in \mathcal{B}\}.$$

These are cycles and form a basis for $H_k(X; \mathbb{F})$.

Now for any cocycle $w \in C^k(T; \mathbb{F})$ and any k -simplex q of X , we can write

$$\langle w, q \rangle = \left\langle w, \sum_{p \in \partial q} F(p) \right\rangle + \left\langle w, q - \sum_{p \in \partial q} F(p) \right\rangle.$$

The chain $q - \sum_{p \in \partial q} F(p)$ is a cycle, and hence homologous to the sum of at most $G(T)$ elements of $\tilde{\mathcal{B}}$ (with signs). Thus w is determined by its values on $\mathcal{F} \cup \tilde{\mathcal{B}}$. Conversely, any function $\mathcal{F} \cup \tilde{\mathcal{B}} \rightarrow \mathbb{F}$ extends to a k -cocycle on X : the values on \mathcal{F} determine its values on simplices of T , while the values on $\tilde{\mathcal{B}}$ determine its values on cycles. Since there are no cycles in T , these are independent.

Now let \tilde{z}_0 be any lift of z to a cocycle in $C^k(X; \mathbb{F})$. If we change \tilde{z}_0 by changing the values on $\mathcal{F} \cup \tilde{\mathcal{B}}$ by integers, we get a new cocycle; in particular, we can do this to get a new \tilde{z} such that its values on $\mathcal{F} \cup \tilde{\mathcal{B}}$ are in $[0, 1)$. Now, for every k -simplex q , $\langle \tilde{z}, q \rangle = \sum \pm \langle \tilde{z}, c \rangle$ where the sum is over $k + 1 + G(T)$ elements $c \in \mathcal{F} \cup \tilde{\mathcal{B}}$. Therefore, \tilde{z} is still a lift of z and has $\|\tilde{z}\|_\infty \leq k + 1 + G(T)$. \square

Now we show that if a chain has a filling with \mathbb{Z} coefficients, we can find such a filling near any filling with \mathbb{F} coefficients.

Lemma 3.9. *Let X be a finite simplicial complex equipped with the standard metric, let X_L be the cubical or edgewise L -regular subdivision of X , and let $k \geq 0$. Then there is a constant $C(X, k)$ such that for any $\alpha \in C^k(X_L; \mathbb{F})$, such that $\delta\alpha$ takes integer values and is a coboundary over \mathbb{Z} , there is an $\tilde{\alpha} \in C^k(X_L; \mathbb{Z})$ such that $\delta\alpha = \delta\tilde{\alpha}$ and $\|\tilde{\alpha}\|_\infty \leq \|\alpha\|_\infty + C$.*

Proof. By Lemma 3.6, X_L admits a spanning tree whose gnarledness is bounded by a constant $C_0(X, k)$. Then by Lemma 3.8, the cocycle $\alpha \bmod \mathbb{Z} \in C^k(X_L; \mathbb{F}/\mathbb{Z})$ has a lift $\Delta\alpha \in C^k(X_L; \mathbb{F})$ with $\|\Delta\alpha\|_\infty \leq k + 1 + C_0$. Then we can set $\tilde{\alpha} = \alpha - \Delta\alpha$ and $C = k + 1 + C_0$. \square

Proof of Lemma 3.1. If $\omega = 0$, we can take $\alpha = 0$, so suppose $\omega \neq 0$.

By Lemma 3.2, we can find an $\alpha \in C^{k-1}(X_L; \mathbb{Q})$ which satisfies $d\alpha = \omega$ and $\|\alpha\|_\infty \leq KL\|\omega\|_\infty$. Then by Lemma 3.9 we can find an $\tilde{\alpha} \in C^{k-1}(X_L; \mathbb{Z})$ such that $d\tilde{\alpha} = \omega$ and

$$\|\tilde{\alpha}\|_\infty \leq KL\|\omega\|_\infty + k + 1 + C_0 \leq (KL + k + 1 + C_0)\|\omega\|_\infty.$$

This gives us an estimate for the isoperimetric constant $C_{IP}(X, k)$. \square

4. BUILDING LINEAR HOMOTOPIES

In this section we prove Theorem B. The proof is based on two lemmas: one to take care of obstructions posed by finite homotopy groups, and the other for infinite obstructions.

We start with a fairly general result for finite homotopy groups. It shows that if a map $X \rightarrow Z$ can be retracted to a subspace $Y \subset Z$ with finite relative homotopy groups, then one can force this retraction to be geometrically bounded. The special case in which Y is a point is proven in [FW13, Theorem 1].

Lemma 4.1. *Let $Y \subset Z$ be a pair of finite simplicial complexes such that $\pi_k(Z, Y)$ is finite for $k \leq n + 1$. Then there is a constant $C(n, Y, Z)$ with the following property. Let X be an n -dimensional simplicial complex, and let $f : X \rightarrow Z$ be a simplicial map which is homotopic to a map $g : X \rightarrow Y$. Then there is a short homotopy of f to a map g' which is homotopic to g in Y , that is, a homotopy which is C -Lipschitz under the standard metric on the product cell structure on $X \times [0, 1]$.*

Note that the constant C does not depend on X and in particular on the choice of a subdivision of X . Thus if we consider Lipschitz and not just simplicial maps from X to Y , the width of the homotopy remains constant, rather than linear, in the Lipschitz constant as is the case with some of our later results.

We will actually use the following relative version: if $f : (X, A) \rightarrow (Z, Y)$ homotopes into Y rel A , then there is a corresponding short homotopy rel A . The proof below works just as well for this variant; one merely has to check that at every stage $f|_A$ remains invariant.

Proof. Let $H : X \times [0, 1] \rightarrow Z$ be a homotopy with $H_0 = f$ and $H_1 = g$; we have no control over this homotopy, only over f . Our strategy will be to push both f and the homotopy into Y via a second-order homotopy. Let Δ^2 be the 2-simplex with edges e_0, e_1 , and e_2 opposite vertices 0, 1, and 2. At the end of the construction, we will obtain a map $F : X \times \Delta^2 \rightarrow Z$ such that $F|_{e_2} = H, F|_{e_0}$ lands in Y , and $F|_{e_1} : X \times [0, 1] \rightarrow Z$ is the short homotopy we are looking for (see also Figure 4(a)).

We will construct this map one skeleton of X at a time. At each step we ensure that the restrictions $F|_{q \times e_1}$ for simplices q of X are chosen from a finite set of Lipschitz maps depending only on Y and Z . In this way we get a universal bound on the Lipschitz constant. We start by setting $B_0 = X \times e_2$ and

$$F^{(0)} = H : B_0 \rightarrow Z.$$

In general, for $k \geq 0$, let

$$B_k = (X^{(k-1)} \times \Delta^2) \cup (X \times e_2)$$

and $A_k = B_k \cap X \times e_0$. Then suppose by induction we have a map

$$F^{(k)} : (B_k, A_k) \rightarrow (Z, Y)$$

such that the restrictions $F|_{q \times e_1}$ for $(k-1)$ -simplices q of X are contained in a finite set $\mathcal{F}_k(Z, Y)$. We would now like to extend this (over cells of the form $q \times \Delta^2$, for every k -simplex q of X ; see Figure 4(b)) to a map $F^{(k+1)} : (B_{k+1}, A_{k+1}) \rightarrow (Z, Y)$.

To avoid doing an extra ad hoc step we will use the convention $S^{-1} = \emptyset$. Let $k \geq 0$. Given a k -simplex q , let

$$\cap q = (q \times \{0\} \cup \partial q \times e_1, \partial q \times \{2\}) \subset (B_k, A_k).$$




FIGURE 4. Illustrations for the proof of Lemma 4.1

We think of this as a map $(D^k, S^{k-1}) \rightarrow (B_k, A_k)$. It can be homotoped into A_k rel boundary via a null-homotopy which is constant on the X -coordinate and sends e_1 to e_0 , keeping the vertex $\{2\}$ constant. Therefore the map

$$F^{(k)} \circ \cap q : (D^k, S^{k-1}) \rightarrow (Z, Y)$$

homotopes rel boundary into Y . Moreover, the set of homotopy classes of maps homotoping $F^{(k)} \circ \cap q$ into Y (more precisely, of maps

$$(q \times e_1, q \times \{2\}) \rightarrow (Z, Y)$$

which restrict to $F^{(k)} \circ \cap q$ on $\cap q$) is in (noncanonical) bijection with $\pi_{k+1}(Z, Y)$. One such bijection u_q is obtained by sending a map $\varphi : (q \times e_1, q \times \{2\}) \rightarrow (Z, Y)$ to the map

$$u_q(\varphi) : \mathbb{m}q := (q \times (e_1 \cup e_2) \cup \partial q \times \Delta^2, q \times \{1, 2\} \cup \partial q \times e_0) \rightarrow (Z, Y)$$

which restricts to φ on $(q \times e_1, q \times \{2\})$ and to $F^{(k)}$ everywhere else.

Now, by our inductive assumption, the number of different possibilities for the map $F^{(k)} \circ \cap q$ is bounded above by

$$|\mathcal{F}_k(Z, Y)|^{k+1} \cdot \#\{(k-1)\text{-simplices of } Z\}.$$

Let $\mathcal{F}_{k+1}(Z, Y)$ contain one Lipschitz map

$$(\Delta^k \times e_1, \Delta^k \times \{2\}) \rightarrow (Z, Y)$$

for each possible value of $F^{(k)} \circ \cap q$ and each homotopy class of null-homotopy; thus there are at most

$$|\mathcal{F}_k(Z, Y)|^{k+1} \cdot \#\{(k-1)\text{-simplices of } Z\} \cdot |\pi_{k+1}(Z, Y)|$$

such maps. We then set $F^{(k+1)}|_{q \times e_2}$ to be the element φ of $\mathcal{F}_{k+1}(Z, Y)$ for which $u_q([\varphi]) = 0$. With this choice, the map can be extended in some way to $q \times \Delta^2$. Since this part of the map does not need to be controlled, we can do this in an arbitrary way.

At the end of the induction, we have our map F : the Lipschitz constant of $F|_{X \times e_1}$ is at most $\max\{\text{Lip } \varphi : \varphi \in \mathcal{F}_{n+1}(Z, Y)\}$. \square

Now we prove Theorem B in the case where the target space is an Eilenberg–MacLane space. This will also be incorporated into the proof of the general case.

Theorem 4.2. *Let X be a finite n -dimensional simplicial complex, and let Y be a finite simplicial complex which has an $(n + 1)$ -connected map $Y \rightarrow K(\mathbb{Z}, m)$, for some $m \geq 2$. Then there are constants $C_1(n, Y)$ and $C_{\text{IP}}(X, m)$ such that any two homotopic L -Lipschitz maps $f, g : X \rightarrow Y$ are $C_1 C_{\text{IP}}(L + 1)$ -Lipschitz homotopic through $C_1(L + 1)$ -Lipschitz maps.*

This theorem is the main geometric input into the proof of Theorem B and is by itself enough to prove certain important cases. For example, it shows directly that any L -Lipschitz map $f : S^3 \rightarrow \mathbb{C}P^2$ is CL -null-homotopic, as is any null-homotopic L -Lipschitz map $X \rightarrow S^n$ for any n -dimensional X . The general proof strategy is that described in §1.2.

Proof. Y is homotopy equivalent to the CW complex obtained from it by contracting an m -spanning tree. In order to create maps that we can homotope combinatorially, we simplicially approximate f and g on an L -regular subdivision of X and then compose with this contraction. After the homotopy is constructed, we can compose with the homotopy equivalence going back to get to the original Y . This increases constants multiplicatively and adds short homotopies to the ends; both of these can be absorbed into C_1 .

For the rest of the proof we assume that Y is the contracted complex and that f and g are compositions of simplicial maps with the contraction.

We construct the homotopy by induction on skeleta of $X \times I$. In particular $f(X^{(m-1)}) = \{*\}$. Let $C_{\text{IP}} = C_{\text{IP}}(X, m)$ be the isoperimetric constant from Lemma 3.1, and let Z be the polyhedral complex given by the product cell structure on $X \times I$, where I is split into $C_{\text{IP}}L$ subintervals $[i/C_{\text{IP}}L, (i + 1)/C_{\text{IP}}L]$. We define

$$F_{m-1} : X \times \{0, 1\} \cup Z^{(m-1)} \rightarrow Y$$

by letting $F_{m-1}|_{X \times \{0\}} = f|_{X^{(m)}}$, $F_{m-1}|_{X \times \{1\}} = g|_{X^{(m)}}$, and sending the rest to $*$.

Now define a simplicial cocycle $\omega \in C^m(X; \pi_m(Y))$ by setting

$$\langle \omega, q \rangle = [f|_{(q, \partial q)}] - [g|_{(q, \partial q)}] \in \pi_m(Y)$$

for m -simplices q of X . Since Y has a finite number of cells, there is a finite number of possible values of ω on simplices. In particular, $\|\omega\|_\infty \leq C$ for some $C = C(Y)$.

By assumption, since $f \simeq g$, ω is a coboundary. By Lemma 3.1, $\omega = d\alpha$ for some cochain $\alpha \in C^{m-1}(X; \pi_m(Y))$ with $\|\alpha_i\|_\infty \leq C_{\text{IP}}CL$. We will use α to construct a cochain $\beta \in C^m(Z; \pi_m(Y))$ which we will use to extend F_{m-1} to $Z^{(m)}$.

Define an extension $\hat{\alpha} \in C^{m-1}(Z; \pi_m(Y))$ of α by

$$\begin{aligned} \left\langle \hat{\alpha}, p \times \left\{ \frac{i}{C_{\text{IP}}L} \right\} \right\rangle &= \left[\left(1 - \frac{i}{C_{\text{IP}}L} \right) \langle \alpha, p \rangle \right] && \text{for } (m - 1)\text{-simplices } p \text{ of } X, \\ & && 0 \leq i \leq C_{\text{IP}}L; \\ \left\langle \hat{\alpha}, s \times \left[\frac{i}{C_{\text{IP}}L}, \frac{i + 1}{C_{\text{IP}}L} \right] \right\rangle &= 0 && \text{for } (m - 2)\text{-simplices } s \text{ of } X, \\ & && 0 \leq i < C_{\text{IP}}L. \end{aligned}$$

Clearly, $\beta = d\hat{\alpha}$ is a cocycle. Moreover, since $\left| \sum_{p \in \partial q} \langle \alpha, p \rangle \right| = |\langle \omega, q \rangle| \leq C$, one can see that

$$\|\beta\|_\infty \leq C + m + 1.$$

In particular the bound depends only on Y .

For each possible value of β on cells, choose representatives

$$(\Delta_m, \partial\Delta_m) \rightarrow (Y, *) \quad \text{and} \quad (\Delta_{m-1} \times I, \partial(\Delta_{m-1} \times I)) \rightarrow (Y, *)$$

and extend F_{m-1} to each m -cell of Z using the appropriate representative to get $F|_{X \times \{0,1\} \cup Z^{(m)}}$. By construction, for each $(m+1)$ -cell c of Z , $F|_{\partial c}$ is null-homotopic.

Now suppose we have constructed $F|_{Z^{(k)}}$ for some $m \leq k \leq n$. By induction, there is a finite number, depending only on k and Y , of possible restrictions $F|_{\partial c}$, where c is a $(k+1)$ -cell of Z . Moreover, if $k \geq m+1$, $F|_{\partial c}$ is null-homotopic since $\pi_k(Y) \cong 0$. Thus for each possible restriction $F|_{\partial c}$, we can choose an extension to c . Extending F to $X \times \{0,1\} \cup Z^{(k+1)}$ in this way gives us a finite set, depending on $k+1$ and Y , of possible restrictions to $(k+2)$ -cells.

At the conclusion of the induction, we obtain a map F which is the desired null-homotopy. \square

In general, the constant C_1 increases by a multiplicative factor in each dimension, depending on the topology of Y . It is worth attempting to analyze C_1 and C_{IP} in simple cases, for example for maps $S^2 \rightarrow S^2$. Here, simplicial approximation multiplies the Lipschitz constant by slightly more than $2\sqrt{3}$. The induction has one step, and if ω satisfies $\|\omega\|_\infty = 1$, then β satisfies $\|\beta\|_\infty \leq 4$. With a bit of care in plumbing as we connect preimages of $S^2 \setminus *$ on the surface of our 3-cells, we can build the null-homotopy by increasing the Lipschitz constant by a factor of 3. This gives a total multiplicative factor of $C_1 = 6\sqrt{3} + \varepsilon \approx 10.4$ when L is large. The isoperimetric constant C_{IP} depends on the exact geometric model for the preimage sphere; in the case of the tetrahedron, it is 1.

Putting together Lemma 4.1 and Theorem 4.2, we can now prove Theorem B. We recall this result below:

Theorem. *Let X be an n -dimensional finite complex. If Y is a finite simply connected complex which is rationally equivalent through dimension n to a product of Eilenberg–MacLane spaces, then there are constants $C_1(n, Y)$ and $C_2(X)$ such that homotopic L -Lipschitz maps from X to Y are $C_1 C_2(L+1)$ -Lipschitz homotopic through $C_1(L + C_2)$ -Lipschitz maps.*

A corollary for highly connected Y follows from the rational Hurewicz theorem.

Corollary 4.3. *Let Y be a rationally $(k-1)$ -connected finite complex, and let X be an n -dimensional finite complex. Then if $n \leq 2k-2$, then there are constants $C_1(n, Y)$ and $C_2(X)$ such that homotopic L -Lipschitz maps from X to Y are $C_1 C_2(L+1)$ -Lipschitz homotopic through $C_1(L + C_2)$ -Lipschitz maps.*

Before giving the proofs of the corollary and the theorem, we recall some facts about maps to Eilenberg–MacLane spaces which derive from properties of the obstruction-theoretic isomorphism

$$[(X, A), (K(G, n), *)] \cong H^n(X, A; G)$$

induced by cell-wise degrees on cellular maps. See for example [Spa81, Chapter 8] for details. Let X be any CW complex, let $n \geq 2$, and let G be an abelian group, and consider a CW model of $K(G, n)$ whose $(n - 1)$ -skeleton is a point $*$. Then:

- $K(G, n)$ is an H-space: the element in $\text{Hom}(G^2, G)$ sending $(a, b) \mapsto ab$ induces a multiplication map $\text{mult} : K(G, n) \times K(G, n) \rightarrow K(G, n)$. This has identity $*$, i.e., it sends

$$(K(G, n) \times *) \cup (* \times K(G, n)) \mapsto *,$$

and is associative and commutative up to homotopy. It can also be assumed cellular.

- Let $f : X \rightarrow K(G, n)$ be a map. Then the group

$$\pi_1(\text{Map}(X, K(G, n)), f) \cong [X \times [0, 1], K(G, n)]_f$$

of self-homotopies of f is naturally isomorphic to $H^{n-1}(X; G)$.

- Denote the map that sends X to $*$ in $K(G, n)$ also by $*$. Then

$$\pi_1(\text{Map}(X, K(G, n)), *) \cong [SX, K(G, n)]$$

acts freely and transitively on $\pi_1(\text{Map}(X, K(G, n)), f)$ via the multiplication map; the above isomorphism takes this to the action of $H^{n-1}(X; G)$ on itself via multiplication.

Proof of Corollary 4.3. The rational Hurewicz theorem (see, e.g., [KK04]) states that if X is a simply connected space such that $\pi_i(X) \otimes \mathbb{Q} = 0$ for $i \leq k - 1$, then the Hurewicz map

$$\pi_i(X) \otimes \mathbb{Q} \rightarrow H_i(X; \mathbb{Q})$$

induces an isomorphism for $i \leq 2k - 2$. Therefore, for $i \leq 2k - 2$,

$$[X, K(\pi_i(X) \otimes \mathbb{Q}, i)] \cong H^i(X; \pi_i(X) \otimes \mathbb{Q}) \cong \text{Hom}(\pi_i(X) \otimes \mathbb{Q}, \pi_i(X) \otimes \mathbb{Q}).$$

In particular, we can find a map $\varphi_i : X \rightarrow K(\pi_i(X) \otimes \mathbb{Q}, i)$ which induces the identity on π_i . Then the map

$$(\varphi_2, \varphi_3, \dots, \varphi_{2k-2}) : X \rightarrow \prod_{i=1}^{2k-2} K(\pi_i(X) \otimes \mathbb{Q}, i)$$

is rationally $(2k - 1)$ -connected. This allows us to apply Theorem B. □

Proof of Theorem B. Suppose that Y is rationally homotopy equivalent through dimension n to $\prod_{i=1}^r K(\mathbb{Z}, n_i)$. This gives us a map $Q : Y \rightarrow \prod_{i=1}^r K(\mathbb{Q}, n_i)$ inducing an isomorphism on $H^*(-; \mathbb{Q})$. For each i , let $\alpha_i \in H^{n_i}(Y; \mathbb{Z})$ be in the preimage of the copy of \mathbb{Q} corresponding to $H^{n_i}(K(\mathbb{Q}, n_i))$; this induces a map $\varphi_i : Y \rightarrow K(\mathbb{Z}, n_i)$. Then

$$\varphi = (\varphi_1, \dots, \varphi_r) : Y \rightarrow Z = \prod_{i=1}^r Z_i$$

is again a rational homology isomorphism, and so by the rational Hurewicz theorem, (Z, Y) is a pair with $\pi_k(Z, Y)$ finite for $k \leq n + 1$.

Let $f, g : X \rightarrow Y$ be homotopic L -Lipschitz maps, and let $C_{2,k} = C_{\text{IP}}(X, k)$. Then by Theorem 4.2, for each i , there is a $C_{1,i}(Y)$ such that $\varphi_i \circ f$ and $\varphi_i \circ g$ are

$C_{1,i}C_{2,n_i}L$ -Lipschitz null-homotopic through $C_{1,i}L$ -Lipschitz maps via homotopies $F_i : X \times [0, 1] \rightarrow Z_i$. Then

$$F := (F_1, \dots, F_r) : X \times [0, 1] \rightarrow Z$$

is a $\sum_{i=1}^r C_{1,i}C_2L$ -Lipschitz homotopy. Suppose first that we can homotope F to an uncontrolled homotopy of f and g in Y . Then by the relative version of Lemma 4.1 applied to the pair $(X \times [0, 1], X \times \{0, 1\})$, there is a $C_1(n, Y, Z)$ such that f and g are C_1C_2L -Lipschitz homotopic in Y through C_1L -Lipschitz maps.

Note that such a homotopy may not exist a priori; we will need to modify F so that it does. For this we use an algebraic construction. We know that there is some homotopy $G : X \times [0, 1] \rightarrow Y$ between f and g . So we can concatenate the homotopies F and $\varphi \circ G$ to give a map $H : X \times S^1 \rightarrow Z$ defined by


$$H(x, t) = \begin{cases} F(x, 2t), & 0 \leq t \leq 1/2, \\ \varphi \circ G(x, 2(1-t)), & 1/2 \leq t \leq 1 \end{cases}$$

(where we think of S^1 as \mathbb{R}/\mathbb{Z}) representing an element of $\pi_1(\text{Map}(X, Z), \varphi \circ f)$. Since each factor Z_i is a high-dimensional skeleton of an H-space, there is a multiplication map $\text{mult} : Z^{(M)} \times Z^{(M)} \rightarrow Z$ for some large enough M . This induces a free transitive action of $[SX, Z]$ on each $\pi_1(\text{Map}(X, Z), \varphi \circ f)$.

We now analyze the cokernel of the group homomorphism

$$\pi_1(\text{Map}(X, Y), f) \rightarrow \pi_1(\text{Map}(X, Z), \varphi \circ f).$$

Consider the relative Postnikov tower



of the inclusion $\varphi : Y \hookrightarrow Z$. Here, P_k is a space such that $\pi_i(P_k, Y) = 0$ for $i \leq k$ and $\pi_i(Z, P_k) = 0$ for $i > k$. The map p_k therefore only has one nonzero relative homotopy group, $\pi_k(Z, Y)$. In this setting there is an obstruction theory long exact sequence ([Bau77, §2.5]; cf. also [GM81, Prop. 14.3] and [Sul74, Lemma 2.7]) of groups

$$\begin{aligned} \dots \rightarrow H^{k-1}(X; \pi_k(Z, Y)) \rightarrow \pi_1(\text{Map}(X, P_k), \varphi_k \circ f) \\ \rightarrow \pi_1(\text{Map}(X, P_{k-1}), \varphi_{k-1} \circ f) \rightarrow H^k(X; \pi_k(Z, Y)) \rightarrow \dots \end{aligned}$$

In particular, an element of $|\pi_k(Z, Y)|\pi_1(\text{Map}(X, P_{k-1}), \varphi_{k-1} \circ f)$ is the image of some loop of maps to P_k based at $\varphi_k \circ f$. Hence, independently of $\varphi \circ f$,

$$R\pi_1(\text{Map}(X, Z), \varphi \circ f), \quad \text{where } R := \prod_{k=2}^n |\pi_k(Z, Y)|$$

always lifts to $\pi_1(\text{Map}(X, Y), f)$. Let \mathcal{H} be the (finite!) collection of linear combinations with coefficients between 0 and $R - 1$ of some finite generating set for

$[SX, Z]$. Then for any $f : X \rightarrow Y$, the finite set

$$\{\text{mult}(a, \text{id}_{\pi_1(\text{Map}(X,Z), \varphi \circ f)}) : a \in \mathcal{H}\}$$

surjects onto the cokernel we are interested in.

We can then choose $a \in \mathcal{H}$ so that $\text{mult}([a], [H])$ can be homotoped into Y . Now define a map $\tilde{H} : X \times S^1 \rightarrow Z$ by

$$\tilde{H}(x, t) = \begin{cases} \text{mult}(F(x, 2t), a(x, 2t)), & 0 \leq t \leq 1/2, \\ \varphi \circ G(x, 2(1-t)), & 1/2 \leq t \leq 1. \end{cases}$$

Then \tilde{H} is in the same homotopy class as $\text{mult}([a], [H])$. This means that the map $\tilde{F} : X \times [0, 1] \rightarrow Z$ given by $\tilde{F}(x, t) = \tilde{H}(x, t/2)$ is a homotopy between f and g which homotopes into Y and whose Lipschitz constant is bounded by

$$\text{Lip}(\text{mult}) \cdot \left(\text{Lip } F + \max_{a \in \mathcal{H}} \text{Lip } a \right).$$

This is linear in $\max\{\text{Lip } f, \text{Lip } g\}$, and except for $\text{Lip } a$, the coefficients depend only on n, Y , and Z , so \tilde{F} can be plugged into the argument above. \square

Remark. Note that in this proof, the dependence of $\max_{a \in \mathcal{H}} \text{Lip } a$ on X lies only in the choice of generating set for $[SX, Z]$. In certain special cases, this constant can be independent of X . For example, suppose that we know that X is an n -sphere (or even just an n -dimensional PL homology sphere). Then $[SX, Z] = \pi_{n+1}(Z)$ is generated by maps whose degree on simplices is at most 1—regardless of the geometry of X . This means that for such homology spheres X , L -Lipschitz maps $f, g : X \rightarrow Y$ can be homotoped through maps of Lipschitz constant $C(Y)L$, though the width of the homotopy required may depend on the geometry. This may have applications such as finding skinny metric tubes between “comparable” metrics on the sphere. In contrast, results of Nabutovsky and Weinberger imply that without this comparability condition, such tubes may have to be extremely (uncomputably) thick.

5. A COUNTEREXAMPLE

One may ask whether the linear bound of Theorem B holds for any simply connected target space, not just products of Eilenberg–MacLane spaces. The answer is emphatically no. Here we give, for each $n \geq 4$, a space Y and a sequence of null-homotopic maps $S^n \rightarrow Y$ such that volume of any Lipschitz null-homotopy grows faster than the $(n + 1)$ -st power of the Lipschitz constant of the maps. This forces the Lipschitz constant of the null-homotopy to grow superlinearly.

To make this precise: by the volume of a map $F : S^n \times [0, 1] \rightarrow Y$, we mean

$$\text{vol } F = \int_{S^n \times [0,1]} |\text{Jac } F(x)| d \text{vol}$$

(recall that by Rademacher’s theorem the derivative of a Lipschitz map is defined almost everywhere). By this definition,

$$\text{vol } F \leq \text{vol}(S^n \times [0, 1]) \sup_{x \in S^n \times [0,1]} |\text{Jac } F(x)| \leq \text{vol}(S^n \times [0, 1]) (\text{Lip } F)^{n+1}.$$

To construct the space Y , we take $S^2 \vee S^2$ and attach $(n + 1)$ -cells via attaching maps which form a basis for $\pi_n(S^2 \vee S^2) \otimes \mathbb{Q}$. Note that by rational homotopy theory, $\pi_{*+1}(S^2 \vee S^2) \otimes \mathbb{Q}$ is a free graded Lie algebra on two generators of degree 1 whose

Lie bracket is the Whitehead product (see [GM81, Exercise 44] or [FHT12, §24(f), Example 1]). In particular, if f and g are the identity maps on the two copies of S^2 , the iterated Whitehead product

$$h_1 = [f, [f, \dots [f, g] \dots]] : S^n \rightarrow S^2 \vee S^2,$$

with f repeated $n-2$ times, represents a nonzero element of $\pi_n(S^2 \vee S^2)$. Moreover, the map

$$h_L = [L^2 f, [L^2 f, \dots [L^2 f, L^2 g] \dots]] : S^n \rightarrow S^2 \vee S^2$$

is an $O(L)$ -Lipschitz representative of $L^{2n-2}[h_1]$. Thus in Y we can define a null-homotopy H of h_L by first homotoping it inside $S^2 \vee S^2$ to $h_1 \circ \varphi_{2n-2}$ for some map $\varphi_{2n-2} : S^n \rightarrow S^n$ of degree L^{2n-2} , and then null-homotoping each copy of h_1 via a standard null-homotopy.

Since h_1 is not null-homotopic in $S^2 \vee S^2$, this standard null-homotopy must have degree $C \neq 0$ on at least one of the $(n+1)$ -cells, giving a closed $(n+1)$ -form ω on Y such that $\int_{S^n \times I} \omega^* H = L^{2n-2} C$. Now, suppose H' is some other null-homotopy of h_L . Then gluing H and H' along the copies of $S^n \times \{0\}$ gives a map $p : S^{n+1} \rightarrow Y$. Note that if any map $(D^{n+1}, S^n) \rightarrow (Y, S^2 \vee S^2)$ had nonzero degree on cells, then the map $S^n \rightarrow S^2 \vee S^2$ on the boundary would be homotopically nontrivial. This shows that p must have total degree 0 on cells, in other words, that $\int_{S^n \times I} \omega^* H' = L^{2n-2} C$. Thus the volume of a null-homotopy of h_L grows at least as L^{2n-2} .

In the sequel to this paper [Cha18], we show that for $n = 4$, this estimate is sharp, in the sense that we can always produce a null-homotopy whose Lipschitz constant is quadratic in the time coordinate and linear in the others.

6. QUANTITATIVE COBORDISM THEORY

The goal of the rest of the article is to prove Theorem A, which we recall below.

Theorem. *If M is an oriented closed smooth null-cobordant manifold which admits a metric of bounded local geometry and volume V , then it has a null-cobordism which admits a metric of bounded local geometry and volume*

$$\leq c_1(n) V^{c_2(n)}.$$

Moreover, $c_2(n)$ can be chosen to be $O(\exp(n))$.

As described in the introduction, we will prove this theorem by executing the following steps. We begin by choosing a metric g on M such that (M, g) has bounded local geometry and such that the volume V of (M, g) is bounded by twice the complexity of M . We then proceed as follows:

- (1) We embed M into \mathbb{R}^{n+k} for an appropriately large k (depending on n) so that the embedding has bounded curvature, bounded volume, and has a large tubular neighborhood. We will use this map to embed the manifold into the standard round sphere \mathbb{S}^{n+k} while maintaining bounds on its geometry.
- (2) We show that the Pontryagin–Thom map from this sphere to the Thom space of the universal bundle of oriented k -planes in \mathbb{R}^{n+k} (relative to the embedded manifold and its tubular neighborhood) has Lipschitz constant bounded as a function of n and the volume of M .

- (3) We analyze the rational homotopy type of the Thom space and determine that, up to dimension $n + k + 1$, it is rationally equivalent to a product of Eilenberg–MacLane spaces. Since M is null-cobordant, this map is null-homotopic and so, as a result, we can apply Theorem B to conclude that there is a null-homotopy which has Lipschitz constant bounded as a function of n and the volume of M . This translates to a map from the ball with boundary \mathbb{S}^{n+k} to the Thom space with the same bound on the Lipschitz constant.
- (4) The proof is completed by simplicially approximating this map from the ball, then using PL transversality theory to obtain an $(n + 1)$ -dimensional manifold, embedded in this ball, which fills M and satisfies the conclusions of the theorem.

Throughout this section, we use the following notation. We write $x \lesssim y$ to mean that there is a constant $c(n) > 0$, depending only on n , such that $x \leq c(n)y$. Similarly, we write $x \lesssim A^y$ to imply that there are constants $c_1(n) > 0$ and $c_2(n) > 0$, again depending only on n , such that $x \leq c_1(n)A^{c_2(n)y}$. We define the same expression with \gtrsim analogously. Throughout this section we will also use V to denote the volume of M . Lastly, we will write $\text{Gr}(n + k, n)$ to denote the Grassmannian of oriented n -dimensional planes in \mathbb{R}^{n+k} and $\text{Th}(n + k, n)$ to denote the Thom space of the universal bundle over this Grassmannian. $\text{Gr}(n + k, n)$ is given the standard metric, which induces a metric on $\text{Th}(n + k, n)$. Furthermore, we denote by p^* the basepoint of the Thom space $\text{Th}(n + k, n)$.

We begin by explicitly defining what “bounded local geometry” means in Theorem A.

Definition. Suppose that (M, g) is a closed Riemannian manifold of dimension n . Following [CG85], we say that M has *bounded local geometry* $\text{geo}(M) \leq \beta$ if it has the following properties:

- (B1) M has injectivity radius at least $1/\beta$.
- (B2) All elements of the curvature tensor are bounded below by $-\beta^2$ and above by β^2 .

The manifold (M, g) satisfies $\widetilde{\text{geo}}(M) \leq \beta$ if in addition it satisfies the following condition:

- (B3) The k th covariant derivatives of the curvature tensor are bounded by constants $C(n, k)\beta^{k+2}$. (The $C(n, k)$ are defined once and for all, but we will not specify them.)

Conditions (B1)–(B3) taken together agree with the standard definition used by Riemannian geometers, except that we require explicit quantitative bounds. A theorem of Cheeger and Gromov [CG85, Thm. 2.5] states that for any given $\varepsilon > 0$, a metric g on M with $\text{geo}(M, g) \leq 1$ can be ε -perturbed to g_ε with $\text{geo}(M, g_\varepsilon) \lesssim 1$ which satisfies (B3). In particular, $\text{vol}(M, g_\varepsilon) \leq (1 + \varepsilon)^n \text{vol}(M, g)$. By rescaling, we get a metric \hat{g} with $\text{geo}(M, \hat{g}) \leq 1$ and

$$\text{vol}(M, \hat{g}) \lesssim (1 + \varepsilon)^n \text{vol}(M, g).$$

Therefore, for the rest of the proof we can assume that (B3) holds, with a constant multiplicative penalty on the volume of our manifold.

Finally, if M has boundary, we say, following [Sch01], that it satisfies $\text{geo}(M) \leq \beta$ if (B1) holds at distance at least β from the boundary, (B2) holds everywhere, and in addition the neighborhood of ∂M of width 1 is isometric to a collar $\partial M \times [0, \beta]$. In particular, this implies that $\text{geo}(\partial M) \leq \beta$.

6.1. Embedding M into \mathbb{R}^{n+k} . To begin constructing the embedding described in the first step, we first choose a suitable atlas of M . A similar set of properties defines uniformly regular Riemannian manifolds, a notion due to H. Amann (see, for example, [DSS16, p. 4]). However, we require our quantitative bounds on the geometry of the maps to be much more uniform, depending only on the dimension; we also require that the charts can be partitioned into a uniform number of subsets consisting of pairwise disjoint charts.

Lemma 6.1. *Suppose that M is a compact orientable n -dimensional manifold with $\widetilde{\text{geo}}(M) \leq 1$. There exists a finite atlas \mathfrak{U} with the following properties, expressed in terms of constants $\mu \leq 3/25$, c , and q depending only on n , as well as a natural number $10 \leq m \leq \kappa \exp(\kappa n)$ for some universal constant $\kappa > 0$.*

- (1) *Every map in \mathfrak{U} is the exponential map from the Euclidean n -ball of radius μ to M which agrees with the orientation of M . Since the injectivity radius of M is at least 1 and $\mu < 1$, this is well-defined. We write*

$$\mathfrak{U} = \{\phi_i : B_\mu \rightarrow M_\mu\}.$$

Here, M_μ is a geodesic ball of M of radius μ , and B_μ is the Euclidean ball of radius μ in \mathbb{R}^n .

- (2) *\mathfrak{U} can be written as the disjoint union of sets $\mathfrak{U}_1, \dots, \mathfrak{U}_m$ of charts such that any pair of charts from the same \mathfrak{U}_j have disjoint image.*
- (3) *When we restrict all the maps in \mathfrak{U} to $B_{\mu/4}$, they still cover M .*
- (4) *The pullback of the metric with respect to every $\phi \in \mathfrak{U}$ is comparable to the Euclidean metric, that is,*

$$\frac{1}{q}(\rho \cdot \rho) \leq \phi^*g(x)(\rho, \rho) \leq q(\rho \cdot \rho)$$

*for every $\rho \in \mathbb{R}^n$, for every $x \in T_x M$, and where $\phi^*g(x)$ is the pullback of g at x .*

- (5) *The first and second derivatives of all transition maps are bounded by c .*

Proof. As mentioned above, this list of properties is closely related to one used in the definition of a uniformly regular Riemannian manifold. Every compact manifold is uniformly regular, and it is known that a (potentially noncompact) orientable manifold M with $\text{geo}(M) \leq \beta$ for some β is uniformly regular; this is shown in [Ama15]. This guarantees an atlas with properties similar, though not identical, to the above. We use a similar set of arguments to those compiled by Amann.

To begin, we cover M by balls of radius $\frac{\mu}{12} \leq \frac{1}{100}$. Since M is compact, we require only finitely many balls to cover M . Furthermore, by the Vitali covering lemma, we can choose a finite subset B_1, \dots, B_k of these balls such that $3B_1, \dots, 3B_k$ also cover M , and such that B_1, \dots, B_k are disjoint. We also have that the balls $12B_1, \dots, 12B_k$ cover M , and that these balls have radius μ .

Fix a ball $12B_i$ for some i . We would like to count how many other balls in $12B_1, \dots, 12B_k$ intersect $12B_i$. Call these balls $12B_{j_1}, \dots, 12B_{j_m}$. Then B_i and B_{j_1}, \dots, B_{j_m} all lie inside $50B_i$, and all are disjoint. Since M has bounded local geometry, the volume of $50B_i$ is bounded above in terms of n , and the volumes of

B_i and B_{j_1}, \dots, B_{j_p} are bounded below in terms of n . This yields an exponential bound on m in terms of n . As a result, the balls $12B_1, \dots, 12B_k$ can be partitioned into m sets of pairwise disjoint balls. We define these sets as $\mathcal{B}_1, \dots, \mathcal{B}_m$. This proof is analogous to a standard proof of the Besicovitch covering lemma in \mathbb{R}^n .

For every j with $1 \leq j \leq m$, \mathcal{U}_j is defined as follows. For every ball $B \in \mathcal{B}_i$, the exponential map goes from the Euclidean ball of radius μ to B ; furthermore, it can be chosen so that it agrees with the orientation of M . These are exactly the charts that comprise \mathcal{U}_j . The first three properties that we desire are now satisfied.

Property (4) is part of [HKW77, Lemma 1]. Indeed, all k th derivatives of the metric tensor are also bounded by a constant depending only on n and k ; see [Eic91], [Sch01]. This allows us to also bound the derivatives of the pullback of the Euclidean metric along transition functions between the charts. Property (5) follows immediately from this. \square

We will also need the following simple observation.

Lemma 6.2. *There is a C^∞ function ζ from $[0, \mu]$ to $[0, 1]$ such that:*

- (1) ζ is monotonically increasing with $\zeta(0) = 0$ and $\zeta(\mu) = 1$.
- (2) $\zeta(t) = t/\mu$ for all $t \in [0, \mu/2]$.
- (3) $\zeta^{(k)}(\mu) = 0$ for all $k \in \mathbb{Z}_{>0}$.
- (4) For every $k \in \mathbb{Z}_{\geq 0}$, there is some $c(k) \in \mathbb{R}$ such that

$$|\zeta^{(k)}(t)| \leq c(k)$$

for all $t \in [0, \mu]$.

We will now embed M into \mathbb{R}^{n+k} so that we have control over its geometry. In particular, we will prove the following proposition.

Proposition 6.3. *Suppose that (M^n, g) is a compact orientable n -dimensional Riemannian manifold with volume V and bounded local geometry. Then there is some $k \leq \kappa(n + 1) \exp(\kappa n)$ (in particular k depends only on n) such that M is diffeomorphic to a submanifold $M' \subset \mathbb{R}^{n+k}$ with the following properties:*

- (1) M' lies in a ball of radius $\lesssim 1$.
- (2) The smooth map $F : M' \rightarrow \text{Gr}(n + k, n)$ sending $x \in M'$ to $T_x M' \subset \mathbb{R}^{n+k}$, the oriented tangent space of M' at x , has Lipschitz constant $\lesssim 1$.
- (3) M' has a normal tubular neighborhood of size $\gtrsim 1/V$.

Proof. We will use the chart \mathcal{U} constructed in Lemma 6.1 to define an embedding of (M, g) into \mathbb{R}^N , with $N > 2n + 3$ depending only on n . By property (2), \mathcal{U} can be written as a disjoint union $\bigsqcup_{j=1}^m \mathcal{U}_j$ of sets of charts with disjoint images. The number of elements in each \mathcal{U}_j is $\lesssim V$ since these disjoint images have volume $\gtrsim 1$. Let $R = \max_{1 \leq j \leq m} \#\mathcal{U}_j$. We define n -dimensional spheres $\mathbb{S}_1, \dots, \mathbb{S}_R$ in \mathbb{R}^{n+1} by the following properties:

- (1) \mathbb{S}_i has radius $1 + i/R$;
- (2) every \mathbb{S}_i passes through the origin;
- (3) the center of every \mathbb{S}_i lies on the ray from the origin in the direction $(1, 0, \dots, 0)$.

The radii of the spheres are between 1 and 2, and the difference between any two of the radii is $\gtrsim 1/V$. An example of such a sequence of spheres is shown in Figure 5. We will refer to the antipode of the origin on each sphere as its “north pole”.




FIGURE 5. A sequence of 1-spheres in \mathbb{R}^2 , with north poles spaced at distance $1/V$.

Define $N = m(n + 1)$ and $k = N - n$. Fix a point $x \in M$. Our embedding $E : M \rightarrow M' \subset (\mathbb{R}^{n+1})^m$ will map x to $(\vec{y}_1, \dots, \vec{y}_m)$, where each $\vec{y}_j \in \mathbb{R}^{n+1}$, as follows. For every j with $1 \leq j \leq m$, if x is not in the image of any chart of \mathfrak{U}_j , then we set $\vec{y}_j = \vec{0}$. If not, then x is in the image of exactly one chart $\phi_i : B_\mu \rightarrow M$ in \mathfrak{U}_j , $1 \leq i \leq R$. In this case, we set \vec{y}_j to be the point on \mathbb{S}_i given by composing $\phi_i^{-1}(x)$ with a map $\kappa_i : B_\mu \rightarrow \mathbb{S}_i$ which is defined as follows: take the origin to the north pole of \mathbb{S}_i , and then map the geodesic sphere of radius r in B_μ homothetically to the geodesic sphere around the north pole in \mathbb{S}_i of radius $\zeta(r)D_i$. Here ζ is defined as in Lemma 6.2 and D_i is the intrinsic diameter of \mathbb{S}_i .

Define a map $\widehat{\phi}_i : M \rightarrow \mathbb{R}^{n+1}$ by

$$\widehat{\phi}_i = \begin{cases} \kappa_i \circ \phi_i^{-1}(x), & x \in \phi_i(B_\mu), \\ \vec{0}, & \text{otherwise.} \end{cases}$$

Since ζ is smooth and all its derivatives go to 0 at μ , this is a smooth map whose derivative has rank n on B_μ . If the original charts in \mathfrak{U}_j are $(\phi_1^j, \dots, \phi_{q_j}^j)$, then we can write

$$E(x) = \left(\sum_{i=1}^{q_1} \widehat{\phi}_i^1(x), \dots, \sum_{i=1}^{q_m} \widehat{\phi}_i^m(x) \right).$$

Since \mathfrak{U} is an atlas, for any x , some $\widehat{\phi}_i^j(x)$ is nonzero. On the other hand, at most one of $\widehat{\phi}_1^j(x), \dots, \widehat{\phi}_{q_j}^j(x)$ is nonzero. This shows that E is an immersion. Moreover, if $E(x_1) = E(x_2)$, then for some chart x_1 and x_2 are in the image of that chart, and in fact $x_1 = x_2$. This shows that E is injective. Since every \mathbb{S}_i is contained in a ball of radius 2 around the origin, every point in M is mapped to a point in \mathbb{R}^{n+k} of norm $\lesssim 1$.

We have a natural set \mathfrak{U}' of oriented charts for the embedded manifold M' given by $E \circ \phi_i^j$ for each $\phi_i^j \in \mathfrak{U}$. Since the first and second derivatives of all of the transition maps are bounded $\lesssim 1$, since ζ has bounded derivatives, and since the radii of the balls are all bounded below by 1 and above by 2, the first and second derivatives of all charts are $\lesssim 1$. Moreover, since every point of M is contained in $\phi_i^j(B_{\mu/4})$ for some i and j , and $\frac{d\zeta}{dt} = 1$ for $t \leq \mu/4$, the first derivative of each chart is $\gtrsim 1$.

Combined with the property that the pullback of the metric of M using each chart ϕ_i^j is comparable to the Euclidean metric, this shows that the map from M to M' with its intrinsic Riemannian metric is bi-Lipschitz with constant $\lesssim 1$.

Let us now consider the map F as defined in the statement of Proposition 6.3. Fix a point $x' \in M'$, choose one of the above charts ϕ' which covers x' , and define $x \in B_\mu$ to be the unique point with $\phi'(x) = x'$. Choose unit vectors v_1, \dots, v_n in \mathbb{R}^n such that

$$\frac{D_{v_1}\phi'(x)}{|D_{v_1}\phi'(x)|}, \dots, \frac{D_{v_n}\phi'(x)}{|D_{v_n}\phi'(x)|}$$

is an orthonormal set of vectors that spans the tangent plane of M' at x' . For any unit vector $w \in \mathbb{R}^n$, consider

$$\left| D_w \frac{D_{v_1}\phi'(x)}{|D_{v_1}\phi'(x)|} \right|, \dots, \left| D_w \frac{D_{v_n}\phi'(x)}{|D_{v_n}\phi'(x)|} \right|.$$

Since all first and second derivatives of ϕ' are bounded above by $\lesssim 1$, and since the first derivatives of ϕ' are bounded from below by $\gtrsim 1$, all of these values are bounded by $\lesssim 1$. Since the original vectors are orthonormal, for ϵ sufficiently small the distance in $\text{Gr}(N, n)$ between the tangent plane at $\phi'(x)$ and the tangent plane at $\phi'(x + \epsilon w)$ is $\lesssim \epsilon$. Since ϕ' is $\lesssim 1$ -bi-Lipschitz, this completes the proof that F is $\lesssim 1$ -Lipschitz.

Lastly, we want to show that M' has a normal tubular neighborhood of width $\gtrsim 1/V$. Suppose that x' and y' are two points on M' , and suppose $v_{x'}$ and $v_{y'}$ are normal vectors at x' and y' , respectively, such that $x' + v_{x'} = y' + v_{y'}$. We would like to show that $\max(|v_{x'}|, |v_{y'}|) \gtrsim 1/V$.

Let θ be the angle between $v_{x'}$ and $v_{y'}$. Consider a minimal-length geodesic γ , parametrized by arclength, between x' and y' ; $v_{x'}$ and $v_{y'}$ lie in the orthogonal $(N - 1)$ -planes to this geodesic at x' and y' , respectively. The above arguments imply that the tautological embedding $M' \rightarrow \mathbb{R}^N$ has second derivatives $\lesssim 1$. Therefore, the second derivative of γ is $\lesssim 1$.

Proposition. *Let $\ell = \text{length}(\gamma)$. Then $\ell \gtrsim \theta$.*

Proof. Let V be the plane spanned by $v_{x'}$ and $v_{y'}$, and let π_V and π_{V^\perp} be orthogonal projections to V and V^\perp . Then:

- the average over $[0, \ell]$ of $\pi_{V^\perp} \frac{d\gamma}{dt}$ is 0;
- $\frac{d\gamma}{dt}(0) \cdot v_{x'} = 0$ and $\frac{d\gamma}{dt}(\ell) \cdot v_{y'} = 0$.

The bounds on the second derivative then imply that for every t ,

$$\pi_{V^\perp} \frac{d\gamma}{dt} \lesssim \ell \quad \text{and} \quad \pi_V \frac{d\gamma}{dt} \lesssim \frac{\ell}{\sin(\theta/2)} \lesssim \frac{\ell}{\theta}.$$

Therefore,

$$\ell = \int_0^\ell \sqrt{\left| \pi_V \frac{d\gamma}{dt} \right|^2 + \left| \pi_{V^\perp} \frac{d\gamma}{dt} \right|^2} dt \lesssim \ell^2 \sqrt{\theta^{-2} + 1},$$

and therefore $\ell \gtrsim \frac{\theta}{\sqrt{1+\theta^2}} \gtrsim \theta$. □

Now let ϕ be a chart in some \mathfrak{U}_j such that $x' \in E \circ \phi(B_{\mu/4})$. Suppose first that $y' \in E \circ \phi(B_{\mu/2})$. Then the properties of any κ_i imply that $|x' - y'| \gtrsim \text{length}(\gamma)$; in particular, $|x' - y'| \gtrsim \theta$ and so $\max(|v_{x'}|, |v_{y'}|) \gtrsim 1$.

On the other hand, suppose that y' is not in $E \circ \phi(B_{\mu/2})$. Suppose first that it is in $E \circ \phi(B_{\mu})$ but not $E \circ \phi(B_{\mu/2})$. Here again the properties of any κ_i imply that $x' - y' \gtrsim 1$. The same is true if y' is not in the image of any $\phi' \in \mathfrak{U}_j$. Finally, if y' is in $\phi' \in \mathfrak{U}_j$ for some $\phi' \neq \phi$, then the properties of the κ_i imply that $x' - y' \gtrsim 1/V$. In all these cases it must be the case that

$$\max(|v_{x'}|, |v_{y'}|) \geq \frac{x' - y'}{2} \gtrsim 1/V.$$

This completes the proof that M' has a large tubular neighborhood. \square

Finally, we prove a lemma which allows us to embed M' into a round sphere.

Lemma 6.4. *Suppose that M' is an embedded submanifold of \mathbb{R}^{n+k} satisfying the conclusions of Proposition 6.3. Then there is an embedding $\tilde{E} : M' \rightarrow \tilde{M} \subset \mathbb{S}^{n+k}$ into the round unit sphere such that*

- (1) \tilde{M} has a tubular neighborhood of width $\gtrsim 1/V$. Additionally, \tilde{E} can be extended to a $\lesssim 1$ -Lipschitz diffeomorphism from this tubular neighborhood to a neighborhood of width $\gtrsim 1/V$ of M' .
- (2) The map $\tilde{F} : \tilde{M} \rightarrow \text{Gr}(n+k, n)$ given by $F \circ \tilde{E}^{-1}$ has Lipschitz constant $\lesssim 1$. Here, F is the map from M' to $\text{Gr}(n+k, n)$ from Proposition 6.3.

Proof. M' is contained in a ball of radius $\lesssim 1$, and without loss of generality we may assume that this ball is centered at the origin. If we restrict the stereographic projection to M' , we obtain an embedded manifold of \mathbb{S}^{n+k} which satisfies all of the above properties. \square

6.2. Proof of Theorem A. To complete the proof of Theorem A, we use the embedding of M in \mathbb{S}^{n+k} produced by combining Proposition 6.3 with Lemma 6.4. We begin by describing the Pontryagin–Thom map and by computing its Lipschitz constant.

We map \mathbb{S}^{n+k} into $Y = \text{Th}(n+k, n)$, the Thom space of the universal bundle of oriented n -dimensional planes in \mathbb{R}^{n+k} , via a map $G : \mathbb{S}^{n+k} \rightarrow Y$ defined as follows. Let $z \in \mathbb{S}^{n+k}$. If z is outside of the tubular neighborhood of \tilde{M} of width $c_1(n)/V$ (here the constant depending on n is the same as that in Lemma 6.4), then it is mapped to p^* (the basepoint of $\text{Th}(n+k, n)$). If not, then applying \tilde{E}^{-1} to z produces a point in the tubular neighborhood of M' of width $c_2(n)/V$ (this constant depending on n is the same as that in Proposition 6.3). Hence, $\tilde{E}^{-1}(z) = x + y$, where $x \in M'$ and y is a point in the oriented normal plane \mathcal{N} of M' at x , and y has length $< c_2(n)/V$. Both x and y are unique. We then take

$$G(z) = \left(\mathcal{N}, \frac{V}{c_2(n)} y \right) \in \text{Th}(n+k, n).$$

Since the map \tilde{F} from Lemma 6.4 is Lipschitz with Lipschitz constant $\lesssim 1$, the map from $x \in \tilde{M}$ to the oriented normal plane of M' at $\tilde{E}^{-1}(x)$ is also Lipschitz with Lipschitz constant $\lesssim 1$. If we assume that $c_2(n)/V$ is at most half the critical radius of the tubular neighborhood, then the projection $z \mapsto x$ has Lipschitz constant ≤ 2 . Furthermore, the tubular neighborhood of M' is dilated by a factor of $\lesssim V$ when it is mapped to $\text{Th}(n+k, n)$ and the map \tilde{E}^{-1} has Lipschitz constant $\lesssim 1$ on the tubular neighborhood of width $c_1(n)V$ of \tilde{M} . Hence, the Lipschitz constant of G is $\lesssim V$.

By [MS74, Theorem of Thom, p. 215], the map G is null-homotopic, since M (and so M' and \widetilde{M} with the orientation induced by the charts ϕ' as in the proof of Proposition 6.3 and the stereographic projection from Lemma 6.4) is null-cobordant. $\text{Th}(n+k, n)$ is $(k-1)$ -connected by [MS74, Lemma 18.1]. We can assume, perhaps by adding extra “empty” dimensions, that $k > n + 3$ and so $2(k-1) > n+k+1$.

By Corollary 4.3, since $\text{Th}(n+k, k)$ is a metric CW complex, there is a null-homotopy of G with Lipschitz constant $\lesssim C_{\mathbb{S}^{n+k}, \text{Th}(n+k, n)} V$. This constant depends only on n , and so there is a null-homotopy H of G of Lipschitz constant $\lesssim V$. This extends to a map from a ball B of radius 1 in \mathbb{R}^{n+k+1} to $\text{Th}(n+k, n)$ with Lipschitz constant $\lesssim V$.

We now observe that we can consider both B and $Y = \text{Th}(n+k, n)$ as finite simplicial complexes in the following sense. Since the result follows from standard arguments, we omit the proof.

Lemma 6.5. *There is a finite simplicial complex \widetilde{Y} and a scale $L_1(n)$ such that if we give each simplex the metric of the standard simplex of side length $L_1(n)$, then there is a 2-bi-Lipschitz function f_Y from Y to \widetilde{Y} . Moreover, the image of the 0 section of Y under this map is a subcomplex (and a simplicial submanifold) of \widetilde{Y} .*

Similarly, there is a finite simplicial complex \widetilde{B} and a scale $L_2(n)$ such that if every simplex is given the metric of the standard simplex of side length $L_2(n)$, then there is a 2-bi-Lipschitz function f_B from B to \widetilde{B} . We can also choose f_B so that $f_B : \partial B \rightarrow \partial \widetilde{B}$ is a homeomorphism from ∂B to $\partial \widetilde{B}$.

Both $L_1(n)$ and $L_2(n)$ depend only on n .

We can now consider the map $\widetilde{H} : \widetilde{B} \rightarrow \widetilde{Y}$ given by $f_Y \circ G \circ f_B^{-1}$. Since the maps are 2-bi-Lipschitz, \widetilde{H} is still $\lesssim V$ bi-Lipschitz. With a slight abuse of notation, we will refer to \widetilde{Y} by Y , \widetilde{H} by H , and \widetilde{B} by B . By using Proposition 2.1, we can subdivide the simplices of B to form B' such that H can be homotoped to a simplicial map from B' to Y with Lipschitz constant $\lesssim V$. We also know that the side lengths of the simplices in B' are $\gtrsim 1/V$. We will define Z to be the simplicial submanifold formed by applying f_Y on the 0-bundle of $\text{Th}(n+k, n)$.

Clearly, $H^{-1}(Z) \cap \partial B$ is a PL manifold which is homeomorphic to M . This is because the map f_B was assumed to be a homeomorphism from the boundary of the ball to the boundary of the simplicial approximation of the ball. We will begin by perturbing Z to Z' , a PL manifold embedded in Y . We want Z' to have the following properties:

- (1) Z' is an n -dimensional PL manifold.
- (2) $G^{-1}(Z') \cap \partial B$ is homeomorphic to M .
- (3) For every open k -simplex c of Y , Z' is transverse to c .
- (4) Z' depends only on n .

We can find such a PL manifold by perturbing Z using PL transversality theory. There are several standard references for this; see for example [RS72, Theorem 5.3]. This theorem does not yield this result directly but can be adapted to do so.

We will use the transverse inverse image of Z' to construct our filling. We know that $H^{-1}(Z') \cap \partial B$ is homeomorphic to M from property (2). Furthermore, the fact that the map is simplicial combined with properties (1) and (3) implies that $H^{-1}(Z')$ is an $(n+1)$ -dimensional PL manifold with boundary, and its boundary is

$H^{-1}(Z') \cap \partial B$. Furthermore, since the sphere, the ball, the simplicial approximations to them, and the embedded manifold \widetilde{M} are all orientable, from the discussion in [MS74, p. 210] we see that we also have that this manifold is orientable, and agrees with the orientation of its boundary (which is homeomorphic to M).

We now estimate the volume of $H^{-1}(Z')$. Since B only depends on n , the number of simplices of B' is $\lesssim V^{n+k+1}$. Since H is a simplicial map, the intersection of $H^{-1}(Z')$ with a given simplex belongs to a finite set of subsets which depends only on n ; since the simplices are at scale $\sim 1/V$, the $(n+1)$ -dimensional volume of this intersection is $\lesssim V^{-(n+1)}$. Therefore, the volume of $H^{-1}(Z')$ is $\lesssim V^k$, where k is $O(\exp(n))$.

To build our manifold, we smooth out $W = H^{-1}(Z') \cap c$ and ∂W . We can do this so that the volumes do not increase very much and so that ∂W , after smoothing, is diffeomorphic to M . As above, since Z' and Y depend only on n , since Y is a finite complex, and since the side lengths of the simplices in B are $\gtrsim 1/V$, this smoothing can be done so that the result has $\text{geo} \lesssim V$ (including on the boundary). After dilating the smoothed version of W by a factor which is $\lesssim V$, we have a compact oriented manifold \widetilde{W} with $\text{geo}(\widetilde{W}) \leq 1$ whose boundary is (orientation preserving) diffeomorphic to M . The dilation increases the volume of the resulting manifold by a factor of $\lesssim V^{n+1}$, and so the result still has volume bounded by $\lesssim V^k$.

In particular, after the dilation has been performed, we obtain a manifold with bounded local geometry with volume bounded by $\lesssim V^k$, and which bounds a manifold diffeomorphic to M with locally bounded geometry. Thus the complexity of the null-cobordism of M is $\lesssim V^k$. Since V is within a factor of 2 of the complexity of M , this completes the proof of the theorem.

APPENDIX A. THE GROMOV–GUTH–WHITNEY EMBEDDING THEOREM

1. SUMMARY

By using a different method of embedding manifolds in Euclidean space, the bound of Theorem A can be improved to achieve one tantalizingly close to Gromov's linearity conjecture:

Theorem A'. *Every closed smooth null-cobordant manifold of complexity V has a filling of complexity at most $\varphi(V)$, where $\varphi(V) = o(V^{1+\varepsilon})$ for every $\varepsilon > 0$.*

As with the original Theorem A, this holds for both unoriented and oriented cobordisms.

Recall that the polynomial bound on the complexity of a null-cobordism follows from a quantitative examination of the method of Thom:

- (1) One embeds the manifold M in S^N , with some control over the shape of a tubular neighborhood.
- (2) This induces a geometrically controlled map from S^N to the Thom space of a Grassmannian; one constructs a controlled extension of this map to D^{N+1} .
- (3) Finally, from a simplicial approximation of this null-homotopy, one can extract a submanifold of D^{N+1} which fills M and whose volume is bounded by the number of simplices in the approximation.

Part (2) is the result of the quantitative algebraic topology done to control Lipschitz constants of null-homotopies. Abstracting away the method of embedding, we extract the following:

Theorem. *Let M^n be an oriented closed smooth null-cobordant manifold which embeds with thickness 1 in a ball in \mathbb{R}^N of radius R ; that is, there is an embedding whose exponential map on the unit ball normal bundle is also an embedding. Then M has a filling of complexity at most $C(n, N)R^{N+1}$. (For unoriented cobordism, $C(n, N)R^N$ is sufficient.)*

This is optimal in the sense that the asymptotics of the estimates in steps (2) and (3) cannot be improved. Then to prove Theorem A', we simply need the following estimate, which may also be of independent interest.

Theorem B'. *Let M be a closed Riemannian n -manifold of complexity V . Then for every $N \geq 2n + 1$, M has a smooth 1-thick embedding $g : M \rightarrow \mathbb{R}^N$ into a ball of radius*

$$R = C(n, N)V^{\frac{1}{N-n}}(\log V)^{2n+2}.$$

This then implies that for every N , M has a filling of complexity at most

$$C(n, N)V^{1+\frac{n+1}{N-n}}(\log V)^{(N+1)(2n+2)},$$

proving Theorem A'.

The embedding estimate is in turn derived from a similar estimate of Gromov and Guth [GG12] for piecewise linear embeddings of simplicial complexes. The combinatorial notion of thickness used in that paper does not immediately translate into a bound on the thickness of a smoothing. Rather, in order to prove our estimate, we first prove a version of Gromov and Guth's theorem, largely using their methods, with a stronger notion of thickness which controls what happens near every simplex. We then translate this into the smooth world using the following result.

Theorem C' (Corollary of [BDG17, Thm. 3]). *Every Riemannian n -manifold of bounded geometry and volume V is $C(n)$ -bi-Lipschitz to a simplicial complex with $C(n)V$ vertices with each vertex lying in at most $L(n)$ simplices. In particular, every smooth n -manifold of complexity V has a triangulation with $C(n)V$ vertices and each vertex lying in at most $L(n)$ simplices.*

The PL picture. In dimensions < 8 , all PL manifolds are smoothable. Therefore Theorems A' and C' together imply that for $n \leq 6$, every PL null-cobordant manifold with V vertices and at most L simplices meeting at a vertex admits a PL filling with $C(n, L)\varphi(V)$ vertices and at most L simplices meeting at a vertex, where the function φ satisfies $\varphi(V) = o(V^{1+\varepsilon})$ for every $\varepsilon > 0$. For $n = 3$, this complements the result of Costantino and D. Thurston [CT08] which gives bounded geometry fillings of quadratic volume with no restrictions on the local geometry of M .

On the other hand, in high dimensions the PL cobordism problem is still open, and poses interesting issues since, unlike in the smooth category, BPL is not an explicit compact classifying space for PL structures. We hope to return to this in a future paper.

So, is it linear? Gromov’s linearity conjecture appears even more interesting now that we know that it is so close to being true. On the other hand, at least in the oriented case, linearity cannot be achieved by Thom’s method. Suppose that one could always produce “optimally space-filling” embeddings $M \hookrightarrow S^N$, that is, 1-thick embeddings in a ball of radius $V^{1/N}$. Even in this case, an oriented filling would have volume $C(n, N)V^{1+1/N}$.

Moreover, recent results of Evra and Kaufman [EK16] on high-dimensional expanders imply that, at least for simplicial complexes, the Gromov–Guth embedding bound is near optimal and space-filling embeddings of this type cannot be found. While n -manifolds are quite far from being n -dimensional expanders, it is possible that a similar or weaker but still nontrivial lower bound can be found. This would show that Thom’s method is not sufficient for constructing linear-volume unoriented fillings, either.

On the other hand, at the moment we cannot reject the possibility that it is possible to find linear fillings for manifolds by some method radically different from Thom’s. In particular, it is completely unclear how to go about looking for a counterexample to Gromov’s conjecture, although we believe that ideas related to expanders may play an important role.

2. PL EMBEDDINGS WITH THICK LINKS

In [GG12] Gromov and Guth describe “thick” embeddings of k -dimensional simplicial complexes in unit n -balls, for $n \geq 2k + 1$. They define the thickness T of an embedding to be the maximum value such that disjoint simplices are mapped to sets at least distance T from each other. [GG12, Thm. 2.1] gives a nearly sharp upper bound on the optimal thickness of such an embedding in terms of the volume and bounds on the geometry.

This condition is insufficient to produce smooth embeddings of bounded geometry, because as thickness decreases, adjacent 1-simplices of length ~ 1 may make sharper and sharper angles. In this section we show that Gromov and Guth’s construction can be improved to obtain embeddings that also have large angles. Recall that the *link* $\text{lk } \sigma$ of a i -simplex σ inside a simplicial complex X is the simplicial complex obtained by taking the locus of points at any sufficiently small distance $\varepsilon > 0$ from any point of σ in all directions normal to σ . This complex contains an $(r - i)$ -simplex for every r -simplex of X incident to p . If X is linearly embedded in \mathbb{R}^n , there is an obvious induced embedding $\text{lk } \sigma \rightarrow S^{n-i-1}$. We show the following:

Theorem 2.1. *Suppose that X is a k -dimensional simplicial complex with V vertices and each vertex lying in at most L simplices. Suppose that $n \geq 2k + 1$. Then there are $C(n, L)$ and $\alpha(n, L) > 0$ and a subdivision X' of X which embeds linearly into the n -dimensional Euclidean ball of radius*

$$R \leq C(n, L)V^{\frac{1}{n-k}}(\log V)^{2k+2}$$

with Gromov–Guth thickness 1 and such that for any i -simplex σ of X' , the induced embedding $\text{lk } \sigma \rightarrow S^{n-i-1}$ is $\alpha(n, L)$ -thick.

Proof. The proof proceeds with the same major steps as in [GG12]. We first show that a random linear embedding which satisfies the condition that all links are thick, while not having the right thickness, is sparse in a weaker sense: most balls have few simplices crossing them. Gromov and Guth then show that the simplices can

be bent locally, at a smaller scale, in order to thicken the embedding; this produces a linear embedding of a finer complex. We note that if the scale is small enough, this finer, bent embedding also has thick links.

We write $A \lesssim B$ for $A \leq C(n, L)B$ and $A \sim B$ to mean $B \lesssim A \lesssim B$. Following Gromov–Guth, we actually embed X in a $V^{\frac{1}{n-k}}$ -ball with thickness $\sim (\log V)^{-(2k+2)}$; for simplicity, write $R = V^{\frac{1}{n-k}}$.

We start by choosing, uniformly at random, an assignment of the vertices of X to points of ∂B_R from those such that for some $\alpha_0(n, L) > 0$, the following hold:

- (1) Adjacent vertices are mapped to points at least distance $\alpha_0 R$ apart.
- (2) The linear extension to an embedding of X has α_0 -thick links.

We call the resulting linear embedding $I_0(X)$. We can choose α_0 so that this is possible since the thickness of the link of some vertex v (and of incident higher-dimensional simplices) only depends on the placement of vertices at most distance 2 away. Moreover, this implies the following:

- (*) The probability distribution of v conditional on some prior distribution on the other vertices is pointwise \lesssim the uniform distribution. This follows from the fact that this is true even when all vertices within distance 2 from v are fixed.

This implies that given a d -simplex σ , the probability distribution of σ (conditional on any distribution on the vertices outside σ) is likewise pointwise \lesssim the uniform distribution where every vertex is mapped independently.

- (†) If $d(v, w) \leq 2$, then v and w are mapped at least $c_0(n, L)R$ units apart. In particular, every embedded edge has length $\sim R$.

Lemma 2.2. *With high probability, each unit ball $B_1(p) \subset B_R$ meets $\lesssim \log V$ simplices of $I_0(X)$.*

Proof. By an argument of Gromov and Guth, the probability that a random $B_1(p)$ meets a fixed d -simplex σ is $\lesssim V^{-1}$.

Therefore, the expected number of simplices hitting $B_1(p)$ is $\lesssim 1$. If each simplex hitting $B_1(p)$ was an independent event, then the probability that S simplices meet $B_1(p)$ would be $\lesssim e^{-S}$; therefore, with high probability, for every p the number of simplices hitting $B_1(p)$ would be $\lesssim \log V$. Indeed, complete independence is not necessary for this; the condition (*) is sufficient.

This condition holds when the simplices have no common vertices. Therefore, we can finish with a coloring trick, as in Gromov–Guth. We color the simplices of X so that any two simplices that share a vertex are different colors. This can be done with $(k + 1)L$ colors. With high probability, the number of simplices of each color meeting $B_1(p)$ is $\lesssim \log V$. Since the number of colors is $\lesssim 1$, we are done. \square

Now we decompose each simplex into finer simplices, using the family of edgewise subdivisions due to Edelsbrunner and Grayson [EG00]. This is a family of subdivisions of the standard d -simplex with parameter L which has the following relevant properties:

- All links of interior vertices are isometric, and all links of boundary vertices are isometric to part of the interior link.
- The subdivided simplices fall into at most $\frac{d!}{2}$ isometry classes; in particular, all edges have length $\sim 1/L$.

When we apply the edgewise subdivision with parameter L , with the appropriate linear distortion, to $I_0(X)$, we get an embedding $I_0(X')$ of a subdivided complex X' such that all edges have length ~ 1 by (\dagger) and all links have thickness $\gtrsim \alpha_0$ and hence $\gtrsim 1$.

Now we use the following lemma of Gromov and Guth:

Lemma 2.3. *For every $0 < \tau \leq 1$, there is a way to move the vertices of X' by $\leq \tau$ such that the resulting embedding $I_\tau(X)$ is $\gtrsim \tau \cdot (\log V)^{-(2k+2)}$ -thick.*

If we choose $\tau(n, L)$ sufficiently small compared to the edge lengths of $I(X')$, then there is an $\alpha(n, L)$ such that however we move vertices by $\leq \tau$, the links will still be α -thick. Since these edge lengths are uniformly bounded below, this completes the proof. \square

3. THICK SMOOTH EMBEDDINGS

We now use Theorem 2.1 to build thick smooth embeddings of manifolds of bounded geometry.

Theorem 3.1. *Let M be a closed Riemannian m -manifold with $\text{geo}(M) \leq 1$ and volume V . Then for every $n \geq 2m + 1$, there is a smooth embedding $g : M \rightarrow \mathbb{R}^n$ such that*

- $g(M)$ is contained in a ball of radius $R = C(m, n)V^{\frac{1}{n-m}}(\log V)^{2m+2}$.
- For every unit vector $v \in TM$,

$$K_0(m, n)R \leq |Dg(v)| \leq K_1(m, n)R.$$

- The reach of g is greater than 1, that is, the extension of g to the exponential map on the normal bundle of vectors of length ≤ 1 is an embedding.

Proof. We prove this by reducing it to Theorem 2.1. That is, first we build a simplicial complex which is bi-Lipschitz to M , with a bi-Lipschitz constant depending only on m . We apply Theorem 2.1 to this complex to obtain a PL embedding and then smooth it out, using the fact that PL embeddings in the Whitney range are always smoothable. The quantitative bound on the smoothing follows from the fact that the local behavior of the PL embedding comes from a compact parameter space, allowing us to choose from a compact parameter space of local smoothings.

Throughout this proof we write $A \lesssim B$ to mean $A \leq C(m, n)B$. This is different from the usage in the section 2. The first step is achieved by the following result.

Theorem 3.2. *There is a simplicial complex X with at most $L = L(m)$ simplices meeting at each vertex and a homeomorphism $h : X \rightarrow M$ which is ℓ -bi-Lipschitz for some $\ell = \ell(m)$ when X is equipped with the standard simplex-wise metric.*

Proof. We start by constructing an ε -net x_1, \dots, x_V of points on M for an appropriate $\varepsilon = \varepsilon(m) > 0$. We do this greedily: once we have chosen x_1, \dots, x_t , we choose x_{t+1} so that it is outside $\bigcup_{i=1}^t B_\varepsilon(x_i)$. In the end we get a set of points such that the $\frac{\varepsilon}{2}$ -balls around them are disjoint and the ε -balls cover M .

Now, [BDG17, Theorem 3] in particular gives the following:

Lemma 3.3. *If $\varepsilon(m)$ is small enough, there is a perturbation of x_1, \dots, x_V to $x'_1, \dots, x'_V \in M$ and a simplicial complex X with a bi-Lipschitz homeomorphism*

$X \rightarrow M$ as well as the following properties:

- Its vertices are x'_1, \dots, x'_V .
- It is equipped with the piecewise linear metric determined by edge lengths $d(x'_i, x'_j)$ which are geodesic distances in M .
- Its simplices have “thickness” $\geq C(m)$; this is defined to be the ratio of the least altitude of a vertex above the opposite face to the longest edge length. In particular, since the edge lengths are $\sim \varepsilon$, this means that each simplex is $C(m)$ -bi-Lipschitz to a standard one.

This automatically gives a bi-Lipschitz map to X with the standard simplex-wise metric. Moreover, since M has sectional curvatures ≤ 1 , we immediately get a uniform bound on the local combinatorics of X . □

After applying this result to get $h : X \rightarrow M$, we apply Theorem 2.1, finding an embedding $X' \rightarrow \mathbb{R}^n$ of a subdivision X' of X which is 1-thick, lands in an R -ball for $R = C(m)V^{\frac{1}{n-k}}(\log V)^{2k+2}$, has $\alpha(m)$ -thick links, and expands all intrinsic distances by $\sim R$. In other words, we get a PL embedding $f : M \rightarrow \mathbb{R}^n$.

For the sake of uniformity, we expand the metric of M by a factor of R ; this makes the embedding f locally uniformly bi-Lipschitz. That is, for any $x, y \in M$ such that $d(x, y) \leq 1$,

$$d(f(x), f(y)) \sim d(x, y).$$

This is the property of f which we actually use to construct a smoothing.

As in the main part of the paper, we assume that M additionally has controlled k th covariant derivatives of its curvature tensor for every k . This allows us, as in Lemma 6.1, to fix an atlas $\mathfrak{U} = \{\phi_i : B_\mu \rightarrow M\}$ for M , with the following properties:

- (1) The $\phi_i(B_{\mu/2})$ also cover M .
- (2) \mathfrak{U} is the disjoint union of sets $\mathfrak{U}_1, \dots, \mathfrak{U}_r$ each consisting of pairwise disjoint charts.
- (3) The charts are uniformly bi-Lipschitz, and the k th derivatives of all transition maps between charts are uniformly bounded depending only on m and n .

Here μ and r both depend only on m and n . We construct our smoothing first on \mathfrak{U}_1 , then extend to \mathfrak{U}_2 , and so on by induction.

At each step of the induction, we use the following form of the weak Whitney embedding theorem [Hir76, §2.2, Thm. 2.13]: for $s \geq 2r + 1$, the set of smooth embeddings $D^r \rightarrow \mathbb{R}^s$ is C^0 -dense in the set of continuous maps. Moreover, the set of smooth maps which restrict to some specific smooth map on a closed codimension 0 submanifold is likewise dense in the set of such continuous maps [Hir76, §2.2, Ex. 4].

The strategy is as follows. Note that the space of L -bi-Lipschitz maps $B_\mu \rightarrow \mathbb{R}^n$ up to translation is compact by the Arzelà–Ascoli theorem. At every stage we also have a C^∞ -compact space of possible partial local smoothings. Then Whitney will allow us to choose an extension from a space of possibilities which is also C^∞ -compact.

We now give a detailed account of the inductive step. Suppose that we have defined a partial smooth embedding $g : K \rightarrow \mathbb{R}^n$, where K is a compact codimension 0 submanifold of M with

$$(1) \quad \bigcup_{\substack{\phi \in \mathfrak{U}_i \\ 1 \leq i < j}} \phi \left(B_{\mu \cdot \frac{2r-j}{2r}} \right) \subset K \subset \bigcup_{\substack{\phi \in \mathfrak{U}_i \\ 1 \leq i < j}} \phi(B_\mu).$$

Moreover, suppose that g is ρ_{j-1} -close to f for some sufficiently small ρ_{j-1} depending on m and n , and that for each $\phi \in \mathfrak{U}_i$, $i < j$, the partially defined function $g \circ \phi$ is an element of a C^∞ -compact moduli space \mathcal{L}_{j-1} of maps each from one of a finite set of subdomains of B_μ to D^n .

Fix a fine cubical mesh in B_μ ; it should be fine enough that any transition function sends a distance of $\mu/2r$ to at least four times the diagonal of the cubes. The purpose of this mesh is to provide a uniformly finite set of subsets on which maps may be defined. Then, again by Arzelà–Ascoli, for any set K which is a union of cubes in this mesh, the space of potential transition maps $K \rightarrow B_\mu$ satisfying the bounds on the covariant derivatives in all degrees is C^∞ -compact.

Fix $\phi \in \mathfrak{U}_j$. By the above, $g|_{K \cap \phi(B_\mu)} \circ \phi$, again restricted to the union of cubes on which it is fully defined (call this domain $\hat{K} \subset B_\mu$), is also chosen from a C^∞ -compact moduli space \mathcal{M}_j , whose elements are patched together from a bounded number of compositions of elements of \mathcal{L}_{j-1} with transition maps as above. Of course, \mathcal{M}_1 consists of the unique map from the empty set.

Let \mathcal{N}_j be the C^0 -compact set of L -bi-Lipschitz embeddings $B_{\mu(1-1/2r)} \rightarrow D^n$. Notice that the subset $\Delta \subset \mathcal{M}_j \times \mathcal{N}_j$ consisting of pairs whose C^0 distance is $\leq \rho_{j-1}$ is compact; this Δ contains the pair $(g|_{\hat{K}} \circ \phi, f \circ \phi)$.

Fix a smooth embedding $u : B_{\mu(1-1/2r)} \rightarrow D^n$. We say that $(\varphi, \psi) \in \Delta$ is ε -good for u , for some $\varepsilon > 0$, if:

- The C^0 distance between u and ψ is $< \rho_j$, where $\rho_j > \rho_{j-1}$ is fixed.
- The map interpolating between ϕ and u via a bump function, only depending on \hat{K} , whose transition lies within the layer of cubes touching the boundary of \hat{K} , has reach $> \varepsilon$. (Here, we simply delete all boundary cubes outside of $B_{\mu \cdot (1-1/2r)}$ from the domain. Thus at this step the domain of g actually recedes slightly; this is the motivation for the condition (1).)

For any fixed pair (u, ε) , these are both open conditions in Δ , so there is an open set $V_{u, \varepsilon} \subseteq \Delta$ of good pairs (φ, ψ) . Moreover, since (by Whitney) we can always choose a u which coincides on \hat{K} with a given element of \mathcal{M}_j , these sets cover Δ . Therefore, we can take a finite subcover corresponding to a set of pairs (u_i, ε_i) . Taking a cover by compact subsets subordinate to this, we get a compact set of allowable extensions of elements of \mathcal{M}_j to $B_{\mu \cdot (1-1/2r)}$; together with the modified sets of allowable maps on previous \mathfrak{U}_i 's (cut back so as to be defined on a domain of cubes) this makes \mathcal{L}_j .

We choose an extension of g from the set of allowable extensions above. Doing this for every $\phi \in \mathfrak{U}_j$ completes the induction step, giving some bound on the local geometry and reach by the compactness argument. Moreover, if we pick ρ_j small enough compared to $\mu/2r$, then the embedding outside $\phi(B_\mu)$ stays far enough away from the embedding inside. Nevertheless, all of these bounds become worse with every stage of the induction.

At the end of the induction, we have a smooth embedding of M . Every choice we made was from a compact set of local smoothings depending ultimately only on m and n , which in turn controlled various bi-Lipschitz and C^k bounds. Thus the resulting submanifold $\tilde{M} = g(M) \subset B_R$ has $\text{geo}(\tilde{M}) \lesssim 1$. For the same reason, g (as a map from M with its original metric) has all directional derivatives $\sim R$. Moreover, since we did not move very far from f , points from disjoint simplices cannot have gotten too close to each other. This, together with the local conditions, shows that \tilde{M} has an embedded normal bundle of radius $\gtrsim 1$. By expanding everything by some additional $C(m, n)$, we achieve the bounds desired in the statement of the theorem. \square

ACKNOWLEDGMENTS

This paper owes a lot to the ideas of Gromov, as the introduction makes amply clear. The last author would like to thank Steve Ferry for a collaboration that began this work. Essentially, the polynomial bound in the nonoriented case can be obtained by combining [FW13] with some of the embedding arguments in this paper. We also thank MSRI for its hospitality during a semester (long ago) when we began working toward the results reported here. The authors would also like to thank Alexander Nabutovsky and Vitali Kapovitch for pointing out simplifications to several proofs, and for many useful conversations. Finally, we would like to thank the anonymous referee for a large number of remarkably insightful suggestions and corrections.

The authors of the appendix would like to thank Larry Guth and Sasha Berdnikov for stimulating correspondence during the writing of this appendix, and two anonymous referees for helpful comments regarding the exposition.

REFERENCES

- [Ama15] H. Amann, *Uniformly regular and singular Riemannian manifolds*, Elliptic and parabolic equations, Springer Proc. Math. Stat., vol. 119, Springer, Cham, 2015, pp. 1–43, DOI 10.1007/978-3-319-12547-3_1. MR3375165
- [Bau77] H. J. Baues, *Obstruction theory on homotopy classification of maps*, Lecture Notes in Mathematics, Vol. 628, Springer-Verlag, Berlin-New York, 1977. MR0467748
- [BDG17] J.-D. Boissonnat, R. Dyer, and A. Ghosh, *Delaunay triangulation of manifolds*, Found. Comput. Math. **18** (2018), no. 2, 399–431, DOI 10.1007/s10208-017-9344-1. MR3777784
- [BH81] S. Buoncristiano and D. Hacon, *An elementary geometric proof of two theorems of Thom*, Topology **20** (1981), no. 1, 97–99, DOI 10.1016/0040-9383(81)90016-1. MR592572
- [Cha18] G. R. Chambers, F. Manin, and S. Weinberger, *Quantitative nullhomotopy and rational homotopy type*, Geom. Funct. Anal. **28** (2018), no. 3, 563–588, DOI 10.1007/s00039-018-0450-2. MR3816519
- [CG85] J. Cheeger and M. Gromov, *On the characteristic numbers of complete manifolds of bounded curvature and finite volume*, Differential geometry and complex analysis, Springer, Berlin, 1985, pp. 115–154. MR780040
- [Che70] J. Cheeger, *Finiteness theorems for Riemannian manifolds*, Amer. J. Math. **92** (1970), 61–74, DOI 10.2307/2373498. MR0263092
- [CT08] F. Costantino and D. Thurston, *3-manifolds efficiently bound 4-manifolds*, J. Topol. **1** (2008), no. 3, 703–745, DOI 10.1112/jtopol/jtn017. MR2417451
- [DKM09] A. M. Duval, C. J. Klivans, and J. L. Martin, *Simplicial matrix-tree theorems*, Trans. Amer. Math. Soc. **361** (2009), no. 11, 6073–6114, DOI 10.1090/S0002-9947-09-04898-3. MR2529925

- [DKM11] A. M. Duval, C. J. Klivans, and J. L. Martin, *Cellular spanning trees and Laplacians of cubical complexes*, Adv. in Appl. Math. **46** (2011), no. 1-4, 247–274, DOI 10.1016/j.aam.2010.05.005. MR2794024
- [DSS16] M. Disconzi, Y. Shao, and G. Simonett, *Some remarks on uniformly regular Riemannian manifolds*, Math. Nachr. **289** (2016), no. 2-3, 232–242, DOI 10.1002/mana.201400354. MR3458304
- [EG00] H. Edelsbrunner and D. R. Grayson, *Edgewise subdivision of a simplex*, Discrete Comput. Geom. **24** (2000), no. 4, 707–719, DOI 10.1145/304893.304897. ACM Symposium on Computational Geometry (Miami, FL, 1999). MR1799608
- [Eic91] J. Eichhorn, *The boundedness of connection coefficients and their derivatives*, Math. Nachr. **152** (1991), 145–158, DOI 10.1002/mana.19911520113. MR1121230
- [EK16] S. Evra and T. Kaufman, *Bounded degree cosystolic expanders of every dimension*, STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2016, pp. 36–48. MR3536553
- [EPC⁺92] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston, *Word processing in groups*, Jones and Bartlett Publishers, Boston, MA, 1992. MR1161694
- [FF60] H. Federer and W. H. Fleming, *Normal and integral currents*, Ann. of Math. (2) **72** (1960), 458–520, DOI 10.2307/1970227. MR0123260
- [FHT12] Y. Félix, S. Halperin, and J.-C. Thomas, *Rational homotopy theory*, Graduate Texts in Mathematics, vol. 205, Springer-Verlag, New York, 2001. MR1802847
- [FW13] S. Ferry and S. Weinberger, *Quantitative algebraic topology and Lipschitz homotopy*, Proc. Natl. Acad. Sci. USA **110** (2013), no. 48, 19246–19250, DOI 10.1073/pnas.1208041110. MR3153953
- [GG12] M. Gromov and L. Guth, *Generalizations of the Kolmogorov–Barzdin embedding estimates*, Duke Math. J. **161** (2012), no. 13, 2549–2603, DOI 10.1215/00127094-1812840. MR2988903
- [GM81] P. A. Griffiths and J. W. Morgan, *Rational homotopy theory and differential forms*, Progress in Mathematics, vol. 16, Birkhäuser, Boston, Mass., 1981. MR641551
- [Gro96] M. Gromov, *Positive curvature, macroscopic dimension, spectral gaps and higher signatures*, Functional analysis on the eve of the 21st century, Vol. II (New Brunswick, NJ, 1993), Progr. Math., vol. 132, Birkhäuser Boston, Boston, MA, 1996, pp. 1–213, DOI 10.1007/s10107-010-0354-x. MR1389019
- [Gro98] M. Gromov, *Metric structures for Riemannian and non-Riemannian spaces*, vol. 152, Birkhäuser Boston, 1998.
- [Gro99] M. Gromov, *Quantitative homotopy theory*, Invited Talks on the Occasion of the 250th Anniversary of Princeton University (H. Rossi, ed.), Prospects in Mathematics, 1999, pp. 45–49.
- [Hat01] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002. MR1867354
- [Hir76] M. W. Hirsch, *Differential topology*, Springer-Verlag, New York–Heidelberg, 1976. Graduate Texts in Mathematics, No. 33. MR0448362
- [HKW77] S. Hildebrandt, H. Kaul, and K.-O. Widman, *An existence theorem for harmonic mappings of Riemannian manifolds*, Acta Math. **138** (1977), no. 1-2, 1–16, DOI 10.1007/BF02392311. MR0433502
- [Kal83] G. Kalai, *Enumeration of \mathbf{Q} -acyclic simplicial complexes*, Israel J. Math. **45** (1983), no. 4, 337–351, DOI 10.1007/BF02804017. MR720308
- [KK04] S. Klaus and M. Kreck, *A quick proof of the rational Hurewicz theorem and a computation of the rational homotopy groups of spheres*, Math. Proc. Cambridge Philos. Soc. **136** (2004), no. 3, 617–623, DOI 10.1017/S0305004103007114. MR2055050
- [MS74] J. W. Milnor and J. D. Stasheff, *Characteristic classes*, Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1974. Annals of Mathematics Studies, No. 76. MR0440554
- [Pet84] S. Peters, *Cheeger's finiteness theorem for diffeomorphism classes of Riemannian manifolds*, J. Reine Angew. Math. **349** (1984), 77–82, DOI 10.1515/crll.1984.349.77. MR743966

- [RS72] C. P. Rourke and B. J. Sanderson, *Introduction to piecewise-linear topology*, Springer-Verlag, New York-Heidelberg, 1972. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 69. MR0350744
- [Sch01] T. Schick, *Manifolds with boundary and of bounded geometry*, Math. Nachr. **223** (2001), 103–120, DOI 10.1002/1522-2616(200103)223:1<103::AID-MANA103>3.3.CO;2-J. MR1817852
- [Spa81] E. H. Spanier, *Algebraic topology*, Springer-Verlag, New York-Berlin, 1981. Corrected reprint. MR666554
- [Sul74] D. Sullivan, *Genetics of homotopy theory and the Adams conjecture*, Ann. of Math. (2) **100** (1974), 1–79, DOI 10.2307/1970841. MR0442930

DEPARTMENT OF MATHEMATICS, RICE UNIVERSITY, HOUSTON, TEXAS 77005
Email address: `gchambers@rice.edu`

DEPARTMENT OF COMPUTER SCIENCE, STANFORD UNIVERSITY, STANFORD, CALIFORNIA 94305
Email address: `dominid@cs.stanford.edu`

DEPARTMENT OF MATHEMATICS, OHIO STATE UNIVERSITY, COLUMBUS, OHIO 43210
Email address: `manin@math.toronto.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, CHICAGO, ILLINOIS 60637
Email address: `shmuel@math.uchicago.edu`