

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

A hardness of approximation result in metric geometry

Permalink

<https://escholarship.org/uc/item/4td3p12p>

Journal

Selecta Mathematica, 26(4)

ISSN

1022-1824

Authors

Brady, Zarathustra

Guth, Larry

Manin, Fedor

Publication Date

2020-09-01

DOI

10.1007/s00029-020-00585-3

Peer reviewed

A hardness of approximation result in metric geometry

Zarathustra Brady, Larry Guth, and Fedor Manin*

June 30, 2020

Abstract

We show that it is NP-hard to approximate the hyperspherical radius of a triangulated manifold up to an almost-polynomial factor.

1 Introduction

This paper is concerned with the problem of estimating the smallest Lipschitz constant of a map with a given degree. We consider this problem from the point of view of computational complexity.

Suppose that Σ is a closed oriented triangulated manifold of dimension n . We equip Σ with a metric by making each simplex a unit equilateral Euclidean simplex. We let $L_{\neq 0}(\Sigma)$ and $L_1(\Sigma)$ denote the smallest Lipschitz constant of a map from Σ to the unit n -sphere with non-zero degree and with degree 1, respectively. When Σ is a triangulated 2-sphere, the paper [14] gives a polynomial time algorithm to estimate $L_{\neq 0}(\Sigma)$ up to a constant factor. In contrast, if Σ is a triangulated surface of arbitrary degree, we show that it is NP-hard to approximate $L_{\neq 0}(\Sigma)$ up to a constant factor, or indeed to within a far wider range. Similarly, if $n \geq 3$ and Σ is a triangulation of S^n , then we show that it is NP-hard to estimate $L_{\neq 0}(\Sigma)$. Here is a more precise and general statement.

Theorem 1.1. *Let $n \geq 2$. For Σ a triangulation of S^n (if $n \geq 3$) or a triangulated surface (if $n = 2$), both $L_{\neq 0}(\Sigma)$ and $L_1(\Sigma)$ are NP-hard to approximate to within $N^{c/\log \log N}$, where $N = \text{vol } \Sigma$ and $c > 0$ depends on n .*

Compare the total range of possible values for these quantities. The “roundest” possible shape for Σ is approximately a round n -sphere of radius $N^{1/n}$, and in the worst case one can contract all but a single simplex. Therefore,

$$N^{-1/n} \lesssim L_{\neq 0}(\Sigma) \lesssim L_1(\Sigma) \lesssim 1.$$

*MSC 2010 classes: 53C23, 68Q17. Keywords: hyperspherical radius, hardness of approximation.

Background. In the 1970s Gromov began to investigate the smallest Lipschitz constant of a map in a given homotopy class [11]. For instance, he showed that the smallest Lipschitz constant of a degree D map from the unit n -sphere to itself is $\sim D^{1/n}$. In [13], Gromov and Lawson defined the hyperspherical radius of a (closed oriented) Riemannian n -manifold (Σ^n, g) as the largest radius R so that (Σ^n, g) admits a 1-Lipschitz map to the round n -sphere of radius R with non-zero degree. The hyperspherical radius of Σ is $1/L_{\neq 0}(\Sigma)$. Gromov and Lawson proved that if the scalar curvature of g is at least 1, and Σ is spin, then the hyperspherical radius of (Σ^n, g) is at most a constant $C(n)$.

The most well-known open problem about hyperspherical radius involves the universal covers of aspherical manifolds. Suppose that Σ is a closed manifold. If we choose a metric g on Σ , then we can consider the hyperspherical radius of the universal cover $(\tilde{\Sigma}, \tilde{g})$. Different metrics g may lead to different hyperspherical radii, but whether the hyperspherical radius is infinite does not depend on the metric. For instance, on $S^1 \times S^2$, the hyperspherical radius of the universal cover is always finite, but on T^3 the hyperspherical radius of the universal cover is always infinite. It is an open question whether every closed aspherical manifold has universal cover with infinite hyperspherical radius. This open question is connected to scalar curvature (cf. [13]) and to the Novikov conjecture (cf. [6]).

This problem is difficult partly because the hyperspherical radius is harder to estimate than it may sound at first. Our main theorem shows from a different perspective that hyperspherical radius is hard to estimate. That said, our result is the best possible for example in the sense that in exponential time, hyperspherical radius can be estimated to within a constant; we make this precise in an appendix. This contrasts, for example, with the related problem of determining the least Lipschitz constant of a degree nonzero map between two simplicial manifolds, both given as input. In that case, the existence of such a map is undecidable, and even if it exists, its minimal Lipschitz constant may grow faster than any computable function of the volumes of the two manifolds. These observations are rooted in the work of Nabutovsky and Weinberger, cf. [19].

Idea of proof. The proof of Theorem 1.1 builds on recent progress in quantitative topology due to Chambers, Dotterrer, Ferry, Weinberger, and the last author—cf. [3], [4], and [18]. This work answered a number of questions from quantitative topology that were raised by Gromov in the 1990s. For instance, they prove that if f is a contractible map from the unit m -sphere to the unit n -sphere with Lipschitz constant L , then f extends to a nullhomotopy $F : B^{m+1} \rightarrow S^n$ with Lipschitz constant at most $C(m, n)(L + 1)$ (see [3] for the case when n is odd and [2] and [1] for the case when n is even). This progress involves new ideas about how to construct maps with controlled Lipschitz constant and we build on those ideas in our proof.

Let us briefly describe the examples of high genus surfaces Σ where $L_{\neq 0}(\Sigma)$ is hard to estimate. We begin with a complicated simplicial complex X —in par-

ticular the dimension of $H_2(X; \mathbb{Z})$ is quite large. The surface Σ sits inside of X . Topologically, it represents a nonzero homology class $[\Sigma] \in H_2(X, \mathbb{Z})$. As a metric space, Σ is very similar to X : it comes within a distance ϵ of every point of X , and moreover, for any two points in Σ , the distance between them in Σ is almost the same as the distance between them in X (using the path metric on both spaces). This last feature requires Σ to have high genus—we can achieve it by adding lots of handles to provide short routes between various points.

Since Σ and X are so similar as metric spaces, if f is a map from Σ to the unit 2-sphere with Lipschitz constant not too large, then f extends to a map $g : X \rightarrow S^2$ with a similar Lipschitz constant. Each map $g : X \rightarrow S^2$ induces a cohomology class $\alpha = g^*([S^2]^*) \in H^2(X, \mathbb{Z})$. For each choice of $\alpha \in H^2(X, \mathbb{Z})$, let $\text{Lip}(\alpha)$ be the smallest Lipschitz constant of a map $g : X \rightarrow S^2$ with $g^*([S^2]^*) = \alpha$. Using the recent ideas in quantitative topology we described above, we are able to accurately estimate $\text{Lip}(\alpha)$ for each $\alpha \in H^2(X, \mathbb{Z})$. But $L_{\neq 0}(\Sigma)$ is approximately the minimum of $\text{Lip}(\alpha)$ over all $\alpha \in H^2(X, \mathbb{Z})$ with $\alpha([\Sigma]) \neq 0$. This is a minimization problem over $H^2(X, \mathbb{Z})$, which is isomorphic to \mathbb{Z}^D for a large D . It turns out to be closely related to the shortest vector problem for lattices in L^∞ . Dinur [5] has a hardness of approximation result for the shortest vector problem in L^∞ , and using her paper we get our hardness of approximation result for $L_{\neq 0}$.

Further questions. There are many quantities studied in metric geometry which it is not obvious how to estimate, and it would be interesting to understand the computational complexity of estimating them to various degrees of accuracy. For instance, it is unknown how hard it is to approximate Uryson widths (cf. [12]) or minimax volumes (cf. [15]). As another example, given a simplicial complex X and a homology class $h \in H_k(X, \mathbb{Z})$, the mass of h is defined to be the smallest mass of an integral k -cycle in the class h (and the mass of a k -cycle is its k -dimensional volume counted with multiplicity). How hard is it to approximate the mass of a homology class?

Organization. The paper is organized as follows. In Section 2 we introduce the comass and explain how it is related to Lipschitz constants of maps. In Section 3, we efficiently approximate the comass using linear programming. In Section 4, we use these tools to prove a hardness of approximation result for homologically non-trivial maps from a complicated complex to the unit sphere. In Section 5, we prove our main theorem for the case $n = 2$. In Section 6, we prove our main theorem for triangulations of S^n with $n \geq 3$. Finally, we have also included an appendix explaining some upper bounds on the complexity of computing hyperspherical radius.

Acknowledgements. We would like to thank an anonymous referee for correcting a number of typos and offering other useful suggestions. F. Manin was partially supported by NSF individual grant DMS-2001042.

2 Comass and the Lipschitz constant

Let X be a finite m -dimensional simplicial complex endowed with the standard simplexwise linear metric with edge length 1. In this section, we discuss the geometric relationship between maps $X \rightarrow S^n$ and the cohomology of X . The discussion applies when n is odd, or when n is even and $m \leq 2n - 2$.

For each homotopy class $\alpha \in [X, S^n]$, we associate $\alpha^*[S^n]$. This association gives a map from $[X, S^n]$ to $H^n(X; \mathbb{Z})$. Obstruction theory shows that this map is uniformly finite-to-one: the cardinality of a point preimage is bounded by

$$\prod_{k=n+1}^{\dim X} \#\{i\text{-simplices of } X\} \cdot \#\pi_k(S^n).$$

Also, the image of this map contains a finite-index subgroup of $H^n(X; \mathbb{Z})$, since there is a map $K(\mathbb{Z}, n)^{(m)} \rightarrow S^n$ of positive degree.

In this section, we study the relationship between the geometric properties of a homotopy class $\alpha \in [X, S^n]$ and the geometric properties of the cohomology class $\alpha^*[S^n]$. We will study the Lipschitz norm $\text{Lip}(\alpha)$ of a class $\alpha \in [X, S^n]$ and compare it to the comass norm of the cohomology class $\alpha^*[S^n]$.

First we recall the definition of the comass norm. The comass norm is a norm on $H^n(X; \mathbb{R})$. Recall that a (real) Lipschitz n -chain is a formal sum $\sum_i a_i \phi_i$ where $a_i \in \mathbb{R}$ and $\phi_i : \Delta^n \rightarrow X$ is a Lipschitz map. Lipschitz chains can be used to define homology just as well as singular (i.e. continuous) chains. One advantage of Lipschitz chains is that there is a volume associated to each ϕ_i . The mass of a chain is then given by

$$\text{mass} \left(\sum_i a_i \phi_i \right) = \sum_i |a_i| \text{vol}(\phi_i).$$

The mass descends to a norm on $H_n(X, \mathbb{R})$, written $\|\cdot\|_{\text{mass}}$. The dual norm on $H^n(X; \mathbb{R})$ is called the comass—in other words,

$$\|w\|_{\text{comass}} := \sup_{\substack{h \in H_n(X, \mathbb{R}) \\ \|h\|_{\text{mass}}=1}} |w(h)|.$$

Theorem 2.1. *Let $n \geq 2$ and let X be a simplicial complex. If n is odd or $\dim X \leq 2n - 2$, then for every homotopy class of maps $\alpha \in [X, S^n]$*

$$(\text{vol}(S^n) \|\alpha^*[S^n]\|_{\text{comass}})^{1/n} \leq \text{Lip}(\alpha) \leq C(\dim X, n) (\|\alpha^*[S^n]\|_{\text{comass}}^{1/n} + 1).$$

When n is even and $\dim X = 2n - 1$, the first inequality holds, but the second holds only in a weaker sense: for every element $a \in H^n(X; \mathbb{R})$ which is in the image of $[X, S^n]$, there is some $\alpha \in [X, S^n]$ such that $\alpha^[S^n] = a$ and*

$$\text{Lip}(\alpha) \leq C(\dim X, n) (\|a\|_{\text{comass}}^{1/n} + 1).$$

Notice that if the comass is greater than 1, this theorem allows us to compute the Lipschitz norm up to a multiplicative constant; on the other hand, if the comass is tiny, it gives almost no information. If Σ is a triangulated n -manifold, and $\alpha \in [\Sigma, S^n]$ is the class of degree 1 maps, then the bounds given by this theorem are not very strong. In this case, $\|\alpha^*[S^n]\|_{\text{comass}} = \|\Sigma\|_{\text{comass}} = 1/\text{vol}(\Sigma)$. Theorem 2.1 then gives $\text{vol}(\Sigma)^{-1/n} \lesssim \text{Lip}(\alpha) \lesssim 1$. In Section 4 we will see some homotopy classes where Theorem 2.1 gives sharp bounds.

Theorem 2.1 is a corollary of the more general shadowing principle in quantitative homotopy theory, [18, Thm. 4.1], given by the last author. However, it is simpler than the general shadowing principle, and all the tools are already present in the second author's survey paper [16]. We give a full proof here which may also be of use as an introduction to [18].

Proof. Let $f : X \rightarrow S^n$ be any representative of the homotopy class α .

We first show the left inequality. For any $h \in H_n(X; \mathbb{R})$, we have

$$|f_*h| \leq \frac{(\text{Lip } f)^n \|h\|_{\text{mass}}}{\text{vol}(S^n)};$$

By definition, this gives

$$(\text{vol}(S^n) \|\alpha^*[S^n]\|_{\text{comass}})^{1/n} \leq \text{Lip}(\alpha).$$

Note that this does not depend on the dimension of X .

To show the second inequality when n is odd or $\dim X \leq 2n - 2$, we will homotope f to a map with the appropriate Lipschitz bound. Define

$$R = \lceil \|\alpha^*[S^n]\|_{\text{comass}}^{1/n} \rceil,$$

and let X_R be a subdivision of X at scale $\sim 1/R$: specifically, we ensure that the volume of each n -simplex of X_R is $\leq \|\alpha^*[S^n]\|_{\text{comass}}^{-1}$, and that all simplices are $C(\dim X, n)$ -bilipschitz to the standard linear simplex with edge lengths $1/R$. There are many ways to do this [3, §2].

If the comass is very small, the best we can do is no subdivision at all, hence the theorem does not tell us very much in that instance.

We will homotope $f : X_R \rightarrow S^n$ to a map which, for every k , restricts to on each k -simplex of X_R to one of a finite set \mathcal{F}_k of maps. This automatically gives the Lipschitz bound we are looking for. Specifically:

- For $k \leq n - 1$, \mathcal{F}_k consists of a single map, the constant map to the basepoint of S^n .
- For $k = n$, \mathcal{F}_n contains a map $(\Delta^n, \partial\Delta^n) \rightarrow S^n$ of each degree between $-(n + 2)$ and $(n + 2)$.
- For $k > n$, \mathcal{F}_k contains a map for each map $\partial\Delta^k \rightarrow S^n$ which restricts to an element of \mathcal{F}_{k-1} on each $(k - 1)$ -simplex and relative homotopy class of extensions to Δ^k . The dimension restriction means that this is a finite set.

We will build a sequence of maps f_k homotopic to f such that the restrictions of f_k to k -simplices of X_R are chosen from \mathcal{F}_k .

First we homotope f to a map $f_{n-1} : X_R \rightarrow S^n$ which takes $X_R^{(n-1)}$ to the basepoint of S^n .

We now homotope this f_{n-1} to a map f_n which has degree 0 or ± 1 on every n -simplex of X_R . This is done as follows. Note that f_{n-1} has a well-defined degree on each n -simplex of X_R . This defines a simplicial cocycle $z \in C^n(X_R; \mathbb{Z})$.

Now notice that the comass of $\alpha^*[S^n]$ is at least as large as its maximum value on *simplicial* n -cycles of mass ≤ 1 in $Z_n(X_R; \mathbb{R})$. In particular, this includes all simplicial cycles whose ℓ^1 -norm with respect to the obvious basis is $\leq \|\alpha^*[S^n]\|_{\text{comass}}$. By the Hahn–Banach theorem, this means that the ℓ^∞ -minimal simplicial representative $\hat{z} \in C^n(X_R; \mathbb{R})$ of $\alpha^*[S^n]$ satisfies $\|\hat{z}\|_\infty \leq 1$.

Now we can find a chain $b \in C^{n-1}(X; \mathbb{R})$ such that $db = z - \hat{z}$. Let $\hat{b} = [b] \in C^{n-1}(X; \mathbb{Z})$ be the nearest integer chain to b . Then the chain $z - d\hat{b}$ satisfies

$$\|z - d\hat{b}\|_\infty \leq n + 2.$$

We now define a homotopy $h_n : X_R \times [0, 1] \rightarrow S^n$ between f_{n-1} and a new map f_n on cells of the product cell structure on $X_R \times [0, 1]$, with $[0, 1]$ thought of as the 1-simplex:

- Send $(X_R \times [0, 1])^{(n-1)}$ to the basepoint of S^n .
- Map n -cells of the form $q \times [0, 1]$, where q is an $(n-1)$ -simplex of X_R , to S^n via a map of degree $\langle \hat{b}, q \rangle$
- This forces the degree on $p \times \{1\}$, where p is an n -simplex of X_R , to be

$$\langle z - d\hat{b}, p \rangle.$$

We let $f_n|_p$ be the corresponding element of \mathcal{F}_n .

- Finally, we extend arbitrarily to the rest of $X_R \times [0, 1]$.

Now, suppose we have constructed $f_k : X_R \rightarrow S^n$, $k \geq n$. We homotope to f_{k+1} by homotoping, on each $(k+1)$ -simplex p , to the element of \mathcal{F}_{k+1} corresponding to the relative homotopy class of $f_k|_p$. Then we extend the homotopy arbitrarily to higher skeleta.

The map $f_{\dim X}$ is the map we want.

Finally, suppose that $\dim X = 2n - 1$ and let $f : X \rightarrow S^n$ be some map. Letting R be as before, we can proceed with the same construction of f_{2n-2} on the $(2n-2)$ -skeleton of X_R . Finally, in the last step there is no obstruction to filling in the map on the $(2n-1)$ -cells. Therefore, for each map $\partial\Delta^k \rightarrow S^n$ which restricts to an element of \mathcal{F}_{2n-2} on each $(2n-2)$ -simplex, we can fix a filling, and then use those fillings to extend f_{2n-2} to the $(2n-1)$ -cells of X_R with bounded Lipschitz constant. The resulting map is not homotopic to f , but induces the same class in $H^n(X; \mathbb{R})$. \square

3 Computing the comass

A good feature of the comass is that it can be approximated efficiently using linear programming.

Lemma 3.1. *There is a polynomial-time algorithm which, given an m -dimensional simplicial complex X , an integer n , and an element $\beta \in H^n(X; \mathbb{R})$, computes a constant $\text{comass}_\Delta(\beta)$ such that for some constant $C(m, n)$,*

$$\text{comass}_\Delta(\beta) \leq \|\beta\|_{\text{comass}} \leq C(m, n) \text{comass}_\Delta(\beta).$$

Proof. The quantity $\text{comass}_\Delta(\beta)$ is the *simplicial comass*, that is, the supremum of $\langle \beta, z \rangle$ over simplicial cycles z of mass 1.

Clearly, the simplicial comass is less than or equal to the usual comass. Conversely, the Federer–Fleming inequality allows us to push any cycle to a simplicial one while only increasing the mass by a constant factor $C(m, n)$. This gives the second inequality.

It remains to show that the simplicial comass is polynomially computable. In fact, it is a linear programming problem: given a simplicial representative $b \in C^n(X; \mathbb{R})$ of β ,

$$\text{maximize } a(z) \quad \text{subject to} \quad \text{mass}(z) \leq 1, \quad \partial z = 0.$$

We now give a bit more detail as to the implementation. We introduce a variable $z_p \geq 0$ for each *oriented* n -simplex p of X (that is, two for each unoriented simplex). The p -coefficient of the chain z is then given by $z_p - z_{p_-}$, where p_- is the same simplex with opposite orientation.

Then the constraint $\text{mass}(z) \leq 1$ can be written as

$$\text{vol}(\Delta^n) \sum_p z_p \leq 1,$$

since every chain has a representative where for every n -simplex p either z_p or z_{p_-} is zero. The constraint $\partial z = 0$ can be written with an equation for each $(n-1)$ -simplex of X .

The total number of variables and constraints is linear in $|X|$; hence the computation is polynomial in $|X|$. \square

Suppose that $\alpha \in [X, S^n]$. In polynomial time, we can approximate the comass of $\alpha^*[S^n]$ up to a constant factor. Suppose that n is odd or $\dim X \leq 2n-2$, so that we can apply Theorem 2.1. If in addition $\|\alpha^*[S^n]\|_{\text{comass}} \geq 1$, then Theorem 2.1 guarantees that $\|\alpha^*[S^n]\|_{\text{comass}}^{1/n}$ agrees with $\text{Lip}(\alpha)$ up to a constant factor $C(\dim X, n)$. This gives a significant class of examples in which we can approximate $\text{Lip}(\alpha)$ up to a constant factor in polynomial time.

4 Hardness of approximation

In the last two sections, we have given some conditions when it is easy to estimate $\text{Lip}(\alpha)$ for a given $\alpha \in [X, S^n]$. We will use that work as a tool to eventually give an example where it is hard to compute $\text{Lip}(\alpha)$. In this section, we prove our first hardness of approximation result, for a slightly different quantity.

For a simplicial complex X , we define $L_{\text{HNT}}(X)$ to be the smallest Lipschitz constant of a map f from X to the unit n -sphere so that $f^*[S^n]$ is non-zero in $H^n(X; \mathbb{R})$. (HNT stands for ‘‘homologically non-trivial’’.)

Theorem 4.1. *Let $n \geq 2$. Let X denote a simplicial complex of dimension $(n+1)$. It is NP-hard to approximate $L_{\text{HNT}}(X)$ to within a factor of $N^{C/\log \log N}$, where N is the number of simplices in X .*

Let us say something about the proof before we begin the details. In the family of examples X we construct,

$$L_{\text{HNT}}(X) \sim \min\{\|\beta\|_{\text{comass}}^{1/n} \mid \text{all non-torsion } \beta \in H^n(X; \mathbb{Z})\}. \quad (*)$$

Estimating this minimum is a cousin of the shortest vector problem. Dinur [5] proved that a version of the shortest vector problem in L^∞ is NP-hard to approximate. Using her theorem, we will show that L_{HNT} is also NP-hard to approximate.

Actually to get (*), we only need that every non-torsion $\beta \in H^n(X; \mathbb{Z})$ has comass at least some constant (in our case $1/(n+1)$.) Then Theorem 2.1 tells us that for individual $\alpha \in [X, S^n]$, $\text{Lip}(\alpha) \sim \|\alpha^*[S^n]\|_{\text{comass}}$; moreover, since X is $(n+1)$ -dimensional, every $\beta \in H^n(X; \mathbb{Z})$ is represented by some $\alpha \in [X, S^n]$.

Eventually, we will prove that it is hard to approximate $L_{\neq 0}(\Sigma)$ and $L_1(\Sigma)$, where Σ is a triangulated manifold. As a bridge to get to that result, we will also prove that two cousins of $L_{\text{HNT}}(X)$ are hard to approximate. Suppose that $h \in H_n(X, \mathbb{Z})$. We define $L_{\neq 0}(X, h)$ to be the smallest Lipschitz constant of a map $f : X \rightarrow S^n$ so that $f^*[S^n](h) \neq 0$. Similarly, we define we define $L_1(X, h)$ to be the smallest Lipschitz constant of a map $f : X \rightarrow S^n$ so that $f^*[S^n](h) = 1$.

Theorem 4.2. *Let $n \geq 2$. Let X denote a simplicial complex of dimension $(n+1)$ and h denote a homology class in $H_n(X, \mathbb{Z})$. It is NP-hard to approximate $L_{\neq 0}(X, h)$ or $L_1(X, h)$ to within a factor of $N^{C/\log \log N}$, where N is the number of simplices in X . In fact, it is NP-hard to distinguish the case where $L_{\neq 0}(X, h) \sim 1$ from the case where $L_{\neq 0}(X, h) \geq N^{C/\log \log N}$, and similarly for L_1 .*

Now we begin the proofs. All of the results are based on a hardness of approximation result for shortest vectors in lattices in L^∞ .

Theorem 4.3. *(Dinur, [5]) For a subgroup $\Gamma \subset \mathbb{Z}^N$ specified by generators of size polynomial in N , it is NP-hard to approximate the smallest length of a non-zero vector in Γ to within $O(N^{C/\log \log N})$.*

This would be enough to prove Theorem 4.1, but we need a small refinement to prove Theorem 4.2. This refinement was also proven in [5]: it follows from inspecting the proof of Theorem 4.3, and in particular condition (a) appears in the proof of Proposition 22 of the paper.

Theorem 4.4. *Consider a lattice $\Gamma \subset \mathbb{Z}^N$ with basis $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_M$, with $\|\mathbf{u}_j\|_\infty$ at most polynomial in N . It is NP-hard to distinguish the following two cases:*

- (a) *There is a vector $\mathbf{v} \in \Gamma$ with $\|\mathbf{v}\|_\infty = 1$ of the form $\mathbf{v} = \mathbf{u}_0 + \sum_{j=1}^M a_j \mathbf{u}_j$, where $a_j \in \{0, 1\}$.*
- (b) *There is no nonzero vector \mathbf{v} in the subgroup with $\|\mathbf{v}\|_\infty \leq N^{C/\log \log N}$.*

Now we begin the proof of Theorems 4.1 and 4.2. We start with a lattice $\Gamma \in \mathbb{Z}^N$ with basis $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_M$ as in Theorem 4.4. We use this lattice to build an $(n+1)$ -dimensional complex X with an isomorphism $\gamma : \Gamma \rightarrow H^n(X; \mathbb{Z})$, and so that the l^∞ norm on Γ is closely related to the simplicial comass norm on $H^n(X; \mathbb{Z})$. More precisely, we will show that

- (a) If $\mathbf{v} \in \Gamma$ has the form $\sum_{j=0}^M a_j \mathbf{u}_j$ with $a_j \in \{0, 1\}$, then

$$\|\mathbf{v}\|_\infty = \text{comass}_\Delta(\gamma(\mathbf{v})).$$

- (b) For any $\mathbf{v} \in \Gamma$, $\text{comass}_\Delta(\gamma(\mathbf{v})) \geq \|\mathbf{v}\|_\infty$.

Here is the construction of the complex X . We start with a wedge of n -spheres $\Sigma_0, \Sigma_1, \dots, \Sigma_M$, each triangulated as $\partial\Delta^{n+1}$. For each coordinate of \mathbb{Z}^N , $i = 1, \dots, N$, we attach a n -sphere S_i , also triangulated as $\partial\Delta^{n+1}$, together with a triangulated mapping cylinder between S_i and a map $S^n \rightarrow \bigvee_{j=0}^M \Sigma_j$ in the homotopy class

$$\sum_{j=0}^M (n+1)u_{ji} \text{id}_{\Sigma_j}, \quad \text{where } \mathbf{u}_j = (u_{j0}, \dots, u_{jN}) \in \mathbb{Z}^N.$$

The complex X deformation retracts to $\Sigma_0 \vee \dots \vee \Sigma_M$, so in particular $H^n(X; \mathbb{Z})$ is the free abelian group generated by $[\Sigma_j]^*$. We define γ so that $\gamma(\mathbf{u}_j) = [\Sigma_j]^*$.

Next we have to estimate the comass norm. Let $\beta \in H^n(X; \mathbb{Z})$ satisfy $\beta(\Sigma_j) = \beta_j$; in other words, $\beta = \gamma(\sum_{j=0}^M \beta_j \mathbf{u}_j)$. Then

$$\beta([S_i]/(n+1)) = \sum_{j=1}^M \beta_j u_{ji}$$

and therefore

$$\text{comass}_\Delta(\beta) \geq \sup_{1 \leq i \leq N} \left| \sum_{j=1}^M \beta_j u_{ji} \right| = \left\| \sum_{j=1}^M \beta_j \mathbf{u}_j \right\|_\infty.$$

Conversely, if $z \in C_n(X; \mathbb{R})$ is a simplicial cycle of mass 1, then it is homologous to a combination of S_i 's and Σ_j 's with coefficients summing up to at most $\frac{1}{n+1}$. Then

$$\text{comass}_\Delta(\beta) = \max \left\{ \sup_j \frac{\beta_j}{n+1}, \sup_i \left| \sum_{j=1}^M \beta_j u_{ji} \right| \right\},$$

and if the β_j are 0 or 1, the term on the right is larger.

Using Theorem 4.4 and our estimate connecting the comass of $\gamma(\mathbf{v})$ with the l^∞ norm of \mathbf{v} in \mathbb{Z}^n , we see that it is NP-hard to distinguish the following two cases:

- (a) There is a non-zero class $\beta \in H^n(X; \mathbb{Z})$ with $\text{comass}_\Delta(\beta) = 1$ and β of the form $[\Sigma_0]^* + \sum_{j=1}^M a_j [\Sigma_j]^*$ with $a_j \in \{0, 1\}$.
- (b) For every non-zero $\beta \in H^n(X, \mathbb{Z})$, $\text{comass}_\Delta(\beta) \geq N^{C/\log \log N}$.

Since X is homotopy equivalent to a wedge of n -spheres, $[X, S^n] = H^n(X; \mathbb{Z})$. For each $\alpha \in [X, S^n]$, there is a unique $\mathbf{v}_\alpha \in \Gamma$ so that $\gamma(\mathbf{v}_\alpha) = \alpha^*[S^n]$. We apply Theorem 2.1 to estimate $\text{Lip}(\alpha)$. If $\alpha \neq 0$, then we know that

$$\|\alpha^*[S^n]\|_{\text{comass}} \geq \|\mathbf{v}_\alpha\|_\infty \geq 1.$$

Theorem 2.1 tells us that $\text{Lip}(\alpha) \sim \|\alpha^*[S^n]\|_{\text{comass}}^{1/n}$. Therefore, we see that

$$L_{\text{HNT}}(X)^n \sim \min_{0 \neq \beta \in H^n(X; \mathbb{Z})} \|\beta\|_{\text{comass}}.$$

But it is NP-hard to distinguish the case when the right-hand side is ~ 1 from the case when the right-hand side is $\geq N^{C/\log \log N}$. This shows Theorem 4.1.

To prove Theorem 4.2, let X be the same complex and let $h = [\Sigma_0]$. Then

$$L_{\neq 0}(X, h)^n \sim \min_{\beta = \sum_{j=0}^M a_j [\Sigma_j]^*, a_0 \neq 0} \|\beta\|_{\text{comass}},$$

and

$$L_1(X, h)^n \sim \min_{\beta = \sum_{j=0}^M a_j [\Sigma_j]^*, a_0 = 1} \|\beta\|_{\text{comass}}.$$

In both cases, we have seen that it is NP-hard to distinguish the case when the right-hand side is ~ 1 from the case when the right-hand side is $\geq N^{C/\log \log N}$. This shows Theorem 4.2.

4.1 On the proof of Dinur's result

For the sake of completeness, we give an idea of Dinur's proof of Theorem 4.3. This is a reduction which is performed in two steps. The first step reduces SAT to an intermediate optimization problem \mathcal{SSAT}_∞ (which is somewhat complicated to define) and the second reduces \mathcal{SSAT}_∞ to SVP_∞ , the problem of deciding the

length of the ℓ^∞ -shortest nonzero vector in a lattice. The first reduction creates a gap, that is, any satisfiable instance of SAT produces an instance of \mathcal{SSAT}_∞ whose smallest solution has norm 1, and any unsatisfiable instance produces an instance of \mathcal{SSAT}_∞ whose smallest solution has norm $\geq g \sim N^{c/\log \log N}$; and the second reduction maintains this gap.

The second reduction, whose particulars were used to state Theorem 4.4, is essentially an act of tidying up: as we'll see, the solution set to an instance of \mathcal{SSAT}_∞ already looks like a lattice, and the reduction adds an extra degree of freedom and massages the vectors so that the norm in \mathcal{SSAT}_∞ (a complicated combination of ℓ^1 and ℓ^∞) is translated into the ℓ^∞ -norm.

The first reduction is much more complicated, but the core idea is that of *linear relaxation*. That is, think of an instance of SAT as a constraint satisfaction problem, where each clause is a constraint: is there a choice of assignments for each clause such that all the variables have the same value each time? What Dinur now does is *relax* this problem to get a new problem \mathcal{SSAT}_∞ . An instance of \mathcal{SSAT}_∞ has the same input as an instance of SAT, but an allowable assignment for each clause is a formal \mathbb{Z} -linear combination of satisfying assignments. The value of a variable in the clause is defined as the projection of this vector to a formal linear combination of variable assignments. For example, if there are two variables x and y and two clauses, then

$$\text{Clause 1: } \{x = T, y = T\} + \{x = F, y = F\}$$

$$\text{Clause 2: } \{x = T, y = F\} + \{x = F, y = T\}$$

is a consistent assignment, because the value of each variable in each clause is $T+F$. Despite this extra flexibility, however, a consistent assignment in which each clause vector has ℓ^1 -norm 1 corresponds to a genuine solution to the corresponding instance of SAT. Therefore, the natural norm on allowable assignments is the maximum over all clauses of the ℓ^1 -norm of the clause vector, and the set of all consistent assignments forms a lattice in the vector space spanned by all satisfying assignments of every clause.

The remaining and most complicated ingredient is introducing the gap: how to make it so that the “fake” solutions have very large norm even when the original SAT instance is very close to being satisfiable. This is done using something analogous to an error-correcting code, via applications of algebra on finite fields.

5 The proof of the main theorem for $n = 2$

In this section, we show that it is hard to approximate $L_{\neq 0}(\Sigma)$ or $L_1(\Sigma)$ when Σ is a triangulated surface (possibly with high genus). We connect this problem to the hardness of approximation results in the last section by adapting a method given by the second author in [14, §5]. Essentially, we show that given a complex X as in the previous section and a homology class $h \in H_2(X, \mathbb{Z})$, we can construct

a triangulated surface $\Sigma \subset X$ in the homology class h , so that the metric on Σ closely imitates the metric on X . Then $L_{\neq 0}(\Sigma)$ will be closely related to $L_{\neq 0}(X, h)$, which is NP-hard to approximate.

By ‘‘closely imitates’’, we mean the following. Given $\varepsilon > 0$, we say that the ε -girth of a map $p : \Sigma \rightarrow X$ is δ if its image is ε -dense in X and the preimages of 2ε -balls have diameter at most δ .¹ It turns out that this condition is sufficient for being able to extend maps from Σ to X :

Lemma 5.1. *Suppose that X has a cover by ε -balls with Lebesgue number $\varepsilon/2$ and multiplicity μ . If a map $p : \Sigma \rightarrow X$ has ε -girth δ and $f : \Sigma \rightarrow S^n$ is a δ^{-1} -Lipschitz map to the unit sphere, then there is a map $g : X \rightarrow S^n$ such that $f \simeq g \circ p$ and $\text{Lip } g \leq C(\mu, n)\varepsilon^{-1}$.*

Before proving the lemma, we give an example of a map of small girth. Suppose that X is a 3-dimensional simplicial complex with equilateral simplices of sidelength 1. We build a high genus surface Σ and a 1-Lipschitz map $p : \Sigma \rightarrow X$ as follows. Let $s > 0$ be a small number, and let $\{x_i\}$ be an s -net in X . For each i , let Σ_i be a 2-sphere of radius $s/10$ and let p project it to x_i . If the distance between Σ_i and Σ_j is at most $10s$, then join Σ_i and Σ_j by a tube of length at most $10s$ and radius $\sim |X|^{-1/2}s$,² and let p project the tube to the geodesic between x_i and x_j . This gives us the surface Σ and the map σ . We can triangulate Σ with simplices of sidelength $\sim |X|^{-1/2}s$, and the total number of simplices is $\lesssim s^{-3}|X|^2$.

Clearly p is ε -dense for any $\varepsilon > 10s$. Suppose $\varepsilon > 10s$; then any two points x_1, x_2 in $p^{-1}B(x, 2\varepsilon)$ can be joined by a path in Σ of length at most 100ε as follows. If the points x_1, x_2 are in thin tubes, we move them to points x'_1, x'_2 contained in the spheres $\Sigma_{i_1}, \Sigma_{i_2}$ via paths of length $\lesssim 10s \leq \varepsilon$. Then connect Σ_{i_1} and Σ_{i_2} through a sequence of thin tubes of total length \lesssim the distance from Σ_{i_1} to Σ_{i_2} , which in turn is $\leq 4\varepsilon$. Therefore, the map p has ε -girth δ whenever $\varepsilon > 10s$ and $\delta > 100\varepsilon$.

Now we turn to the proof of Lemma 5.1.

Proof. Let $\{U_i\} = \{B_\varepsilon(x_i)\}$ be the assumed cover of X . We first pick a partition of unity $\{\phi_i\}$ subordinate to $\{U_i\}$. This defines a map ϕ from X to the nerve N of $\{U_i\}$ via barycentric coordinates. Choose

$$\hat{\phi}_i(x) = \begin{cases} 1 & 0 \leq d(x_i, x) \leq \varepsilon/2 \\ 2 - 2\varepsilon^{-1}d(x_i, x) & \varepsilon/2 \leq d(x_i, x) \leq \varepsilon \\ 0 & \varepsilon \leq d(x_i, x) \end{cases}$$

and let $\phi_i(x) = \left(\sum_i \hat{\phi}_i(x)\right)^{-1} \hat{\phi}_i(x)$. It is not hard to see that with these choices, $\text{Lip}(\phi) = O(\mu^{1/2})\varepsilon^{-1}$.

Then we can define a map $\tilde{g} : N \rightarrow S^n$ as follows:

¹This is slightly different from the definition in [14].

²Here $|X|$ denotes the number of simplices of X ; since we put no restrictions on the combinatorics of X , the number of facets incident to a vertex may be $\sim |X|$.

- For each U_i , choose a point of $f(p^{-1}U_i)$ and send the corresponding vertex of N to that point.
- Now for any simplex of N , the images of all vertices are contained in a 1-ball in S^n . Thus we can project to S^n from the affine simplex spanned by the vertices; such maps have uniformly bounded Lipschitz constant depending only on n .

Defined this way, we can set $g = \tilde{g} \circ \phi$. Then for every point $x \in \Sigma$, $g \circ p(x)$ is at most 2 units from $f(x)$. Therefore there is a linear homotopy from $g \circ p$ to f . \square

Remark. This lemma is closely based on Lemma 5.2 in [14], which, however, works for a much larger class of target manifolds than just the sphere. The argument above can probably be generalized, for example by replacing the “simplex flattening” strategy by barycentric coordinates coming from Karcher mean, as in [7].

Proof of Theorem 1.1 for $n = 2$. We deal with $L_{\neq 0}$. The proof for L_1 is similar.

Start with (X, h) as in Theorem 4.2, where X is a 3-dimensional simplicial complex and $h \in H_2(X, \mathbb{Z})$. Let N be the number of simplices in X . We know that it is NP-hard to distinguish the case that $L_{\neq 0}(X, h) \lesssim 1$ from the case that $L_{\neq 0}(X, h) \geq N^{c/\log \log N}$.

Choose $s = N^{-c/\log \log N}$. Now given (X, h) , we construct a triangulated surface Σ_s with a map $p : \Sigma_s \rightarrow X$ so that $p([\Sigma]) = h$ and so that p has ε -girth δ for all $\varepsilon > 10s$ and $\delta > 100\varepsilon$. The construction is basically the same as the one in the paragraph just after the statement of Lemma 5.1. The only problem is that in that description, $[\Sigma] = 0$. To fix that problem, we let Σ_h be a cycle in the class h . Then we take Σ_h together with a dense set of spheres Σ_i and join them with thin tubes as above.

If $g : X \rightarrow S^n$ with $g^*[S^n](h) \neq 0$, then $g \circ p : \Sigma_s \rightarrow S^n$ has nonzero degree and $\text{Lip}(g \circ p) \leq \text{Lip}(g) \text{Lip}(p) \leq \text{Lip}(g)$. Therefore, $L_{\neq 0}(\Sigma_s) \leq L_{\neq 0}(X, h)$. In particular, if $L_{\neq 0}(X, h) \lesssim 1$, then $L_{\neq 0}(\Sigma_s) \lesssim 1$ as well.

On the other hand, suppose that $f : \Sigma_s \rightarrow S^n$ has nonzero degree. If $\text{Lip}(f) \leq (1/1000)s^{-1}$, then Lemma 5.1 tells us that f is homotopic to a composition $g \circ p$, where $\text{Lip}(g) \lesssim \text{Lip}(f)$. Since f is homotopic to $g \circ p$, $g^*[S^n](h) \neq 0$. Therefore, $L_{\neq 0}(\Sigma_s) \gtrsim \min(s^{-1}, L_{\neq 0}(X, h))$. Recall that we chose s so that $s^{-1} = N^{c/\log \log N}$. So if $L_{\neq 0}(X, h) \geq N^{c/\log \log N}$, then $L_{\neq 0}(\Sigma_s) \gtrsim N^{c/\log \log N}$ as well.

Theorem 4.2 tells us that it is NP-hard to distinguish the case $L_{\neq 0}(X, h) \sim 1$ from the case $L_{\neq 0}(X, h) \geq N^{c/\log \log N}$; therefore by the last two paragraphs it is also NP-hard to distinguish the case $L_{\neq 0}(\Sigma_s) \sim 1$ from the case $L_{\neq 0}(\Sigma_s) \geq N^{c/\log \log N}$. Therefore, it is NP-hard to approximate $L_{\neq 0}(\Sigma)$ to within a factor of $N^{c/\log \log N}$.

Recall that N was the number of simplices in X , but the number of simplices in Σ is $N_\Sigma \leq s^{-3}N^2$ and so eventually $\leq N^{2.1}$, so the same result holds with N_Σ in place of N . \square

6 Approximating simplicial complexes by spheres

If $n \geq 3$, the proof of Theorem 1.1 is similar, except that the step of constructing Σ and the map $p : \Sigma \rightarrow X$ is more complicated. If we just wanted to build triangulated n -manifolds Σ so that $L_{\neq 0}(\Sigma)$ is hard to approximate, we could use the same construction, replacing the 2-spheres by n -spheres, and replacing the annuli $S^1 \times [0, 1]$ by $S^{n-1} \times [0, 1]$. We will show that if $n \geq 3$, then $L_{\neq 0}(\Sigma)$ is hard to approximate even for triangulated spheres. To do that, we have to construct maps of small girth from a triangulated sphere Σ to X . This construction is based on work of Ferry and Okun [9]. Because the construction is a little complicated we explain it in two stages. First, we construct maps of small girth from S^n to X . This illustrates the main geometric idea of the construction. Next, we construct a triangulated sphere Σ with a controlled number of simplices and a simplicial map of small girth from Σ to a refinement of X . For this simplicial map, we have to check that the number of simplices is not too large and that the map can be constructed in polynomial time from the data of (X, h) .

Lemma 6.1. *Let X be a simply connected triangulated m -complex, $n \geq 3$, and $\varepsilon > 0$ a constant. Let $h : S^n \rightarrow X$ be any map. Then there is a piecewise Riemannian metric g and a 1-Lipschitz map $f : (S^n, g) \rightarrow X$ homotopic to h whose ε -girth is at most $C(m, n)\varepsilon$.*

Proof. This is a quantitative version of an asymptotic result due to Ferry and Okun [9]. We use a technique adapted from their result which is briefly described in [14, p. 1089], deforming h in several steps. We start by deforming h to a nearby embedding in a thickening of X and giving S^n the metric g_0 induced by that map. In the rest of the construction, we build up an $(n + 1)$ -dimensional stratified manifold, equipped with a Riemannian metric and a map to X , by iteratively attaching $(n + 1)$ -disks along n -disks on their boundary. At every step, the space remains homotopy equivalent to S^n , and we can think of it as a topologically trivial bordism from (S^n, g_0, h) to some (S^n, g, f) which, at the end of the process, will have the properties we desire.

The disks we glue in will take the form of “balls”, “sticks”, and “slabs”. That is, we think of them as $D^i \times D^{n+1-i}$, for $i = 0, 1$ and 2 respectively, equipped with a map to X which projects to a point, immersed curve, or immersed disk, and with the metric induced by the immersion in the D^i direction and a round metric of radius $\ll \varepsilon$ in the D^{n+1-i} direction.

In detail, the process is as follows:

1. Fix a subdivision X_ε of X at scale ε .
2. Connect all vertices of X_ε in the ε -neighborhood of $h(S^n)$ to $h(S^n)$ via geodesic segments. Call this set of vertices V_h and the set of segments E_h .
3. Choose a spanning tree T for the graph $X_\varepsilon^{(1)}/V_h$. We attach sticks along the edges of T and E_h and balls at the vertices. We may need to homotope h slightly so that it is constant on the n -disks where we attach the balls.

At this stage, the map $f_1 : (S^n, g_1) \rightarrow X$ given by restricting to the end of the bordism has ε -dense image in X , but points in $f_1(X)$ which are close in X may be far from each other in the induced metric on $f_1(X)$ (and hence in g_1).

4. Now let x and y be two vertices of X_ε which are adjacent via an edge not in T . In the next approximation $f_2 : S^n \rightarrow X$, the portions of $f_2(S^n)$ near x and y will be close to each other. We achieve this as follows.

Let x' and y' be preimages of x and y , respectively, in (S^n, g_1) . There is a path γ in (S^n, g_1) between x' and y' ; since X is simply connected, $f_1(\gamma)$ is homotopic to the edge between x and y via some immersed disk. We attach a slab which maps to this disk to our sphere. Again we may need to homotope h slightly near the attaching n -disk.

Similarly, there may be points in the original sphere whose images are close to each other or to vertices of X_ε to which they are not attached. So for some ε -net on the original sphere, consider all pairs of points x and y such that $d(f_1(x), f_1(y)) < \varepsilon$ but $d(x, y) > 2\varepsilon$. For each such pair we likewise attach a slab corresponding to a disk homotoping the image of a path in (S^n, g_1) to the geodesic path.

5. Now, f_2 still doesn't have small girth, since for a point p deep in the interior of a slab, $f_2(p)$ may be close to images of other, distant points. It is enough to create a shortcut from any such p to the ball mapping to a nearby vertex v of X_ε , since then any two such points will have a short path between each other through v .

So in each slab D , choose an ε -net V_D and a forest T_D connecting all the points of V_D to the boundary of D . This T_D can be deformed to lie in $X_\varepsilon^{(1)}$ via a homotopy which moves every point by at most 2ε . Thus we can glue in sticks and slabs spanning each such homotopy; after this surgery, each vertex $v \in V_D$ is 2ε -close in the induced metric to a point which approximates the nearest vertex of X_ε , and any point in the interior of one of the new slabs is 5ε -close to one of these vertices. In other words, because the new slabs are "narrow", they do not again have the problem of the slabs we added in the previous step. Thus the resulting map f at the end of the bordism has the desired small girth. \square

In order to perform the reduction in the proof of Theorem 1.1, it is not enough to show that such approximating spheres exist; we must also make sure that they can be represented by simplicial complexes of polynomial size and computed in polynomial time. For this argument, it's important to note that the complexes constructed in §4 are not only simply connected, but effectively so. In other words, we can construct a contraction of the 1-skeleton whose size is polynomial in the volume of the complex.

Since these complexes are constructed as wedges of spheres with attached mapping cylinders, we can contract the 1-skeleton by first collapsing the mapping cylinders onto the wedge of spheres and then puncturing each sphere in some simplex and contracting away from those points.

Call the relevant complex X . During each of the two stages, loops are expanded by a factor of $O(|X|)$ and each point travels a distance $O(|X|)$. In other words the overall contraction $X^{(1)} \times [0, 1] \rightarrow X$ can be made simplicial on a subdivision at scale $\sim 1/|X|^2$.

To complete the argument, we give an algorithmic, simplicial version of the construction in Lemma 6.1.

Lemma 6.2. *Suppose X is a simplicial m -complex equipped with an explicit contraction $c : X^{(1)} \times [0, 1] \rightarrow X$, expressed as a simplicial map on a subdivision with V facets. Let Σ_0 be a triangulation of the n -sphere and $f_0 : \Sigma_0 \rightarrow X$ a simplicial map. Let $\varepsilon > 0$. Then there is a subdivision X_ε of X and a triangulation Σ of the n -sphere and a simplicial map $f : \Sigma \rightarrow X_\varepsilon$ such that*

- Σ has $\text{poly}_n(|X|, |\Sigma_0|, V, \varepsilon^{-m})$ simplices;
- The ε -girth of f is at most $C(m, n)\varepsilon^{-1} \cdot \text{scale}(X_\varepsilon)$, when Σ is given the standard simplexwise linear metric;
- Σ and f can be computed from X , f_0 , and c in polynomial time for fixed n .

Proof. We proceed with a version of the construction in Lemma 6.1. At each step, we further subdivide both the domain and the codomain complexes.

Recall that we can subdivide an n -simplex into cubes by adding the midpoint of each subsimplex as a vertex. We can then subdivide those cubes as finely as necessary; if we prefer to work with simplicial maps, we can subdivide the cubes again to form simplices. This is somewhat less uniform than the edgewise subdivision of [8], but easier to work with.

In the beginning, we are given a simplicial map $f_0 : \Sigma_0 \rightarrow X$. We start by subdividing both the domain and codomain with scale parameter $E \sim 1/\varepsilon$ to form cubical complexes Σ_ε and X_ε so that the vertices of X_ε form an ε -net and the edge lengths range between about ε and $(m+1)\varepsilon$.

In order to give ourselves more space to cut and paste, we subdivide Σ_ε with another three cubes to a side to produce Σ_1 . We then precompose the map f_0 with a map $\Sigma_1 \rightarrow \Sigma_\varepsilon$ which sends the middle cube of each cube of Σ_ε to the corresponding full cube.

Now we build a relative spanning tree T for the pair of graphs $(X_\varepsilon^{(1)}, f_1(\Sigma_1)^{(1)})$. This has degree bounded by $C(m) \deg(X)$, and so certainly by $C(m)|X|$. Now fix a cubulation of the n -ball with enough boundary facets that we can shade $\deg T$ of them and have a path along the boundary between any two shaded facets which does not pass through any other shaded facets, and which has diameter $C(n)$.³ We

³For example a slight subdivision of the simplicial double cone on a sufficiently fine triangulation of D^n .

assemble these, together with $(n+1)$ -cubes for the sticks, into a cubical thickening of T , equipped with maps sending the balls to the vertices of X_ε and the sticks to the edges. Finally, for each vertex in $f_1(\Sigma_1)$, we arbitrarily choose an n -cube in Σ_1 mapping to the appropriate vertex and glue it to some boundary facet of the corresponding ball. We let Σ_2 be the n -dimensional “outer boundary” of the resulting complex, and $f_2 : \Sigma_2 \rightarrow X_\varepsilon$ the restriction of the associated map.

Now fix a set of n -cells V_ε of Σ_2 which are mapped to vertices of X_ε and form a net of a certain constant density on Σ_2 ; in particular, we include a boundary n -cell of the balls that we added in as well as a single n -cell adjacent to each vertex of Σ_ε .

The next step is to glue in slabs to make sure that the cells of V_ε that map to nearby vertices of X_ε are actually nearby. In order to have enough space to do this, we subdivide the domain and codomain again, this time with scale parameter

$$M = 10(n+1)|V_\varepsilon|^2.$$

Clearly this is polynomial in the various parameters. We call the resulting complexes Σ_M and X_M .

Now we find and totally order all pairs $v_i, v_j \in V_\varepsilon$ so that $f_2(v_i)$ and $f_2(v_j)$ are equal or adjacent as vertices of X_ε , but $d_{\Sigma_2}(v_i, v_j) > C(n)$, where $C(n)$ is large enough to exclude pairs of adjacent vertices in T or Σ_ε . For each such pair, we build a path p_{ij} from v_i to v_j of n -cells in Σ_2 which map to $X_\varepsilon^{(1)}$; a shortest such path can be found in polynomial time after building the graph of such n -cells. We can in addition assume that each p_{ij} goes “straight through” every cell which maps outside $X_\varepsilon^{(0)}$, so that the preimage of each open edge is a disjoint union of open cells each of which projects onto that edge by forgetting some coordinates.

We now build pairwise disjoint paths p'_{ij} of n -cubes in Σ_M , where each p'_{ij} is a refinement of p_{ij} . To make sure that they don't intersect, we assign each pair (i, j) a distinct distance from the middle of cubes of Σ_ε as well as an “entry/exit point” on each face, and have the entry/exit stretches and internal stretches map to different layers.⁴ The paths p'_{ij} consist of mostly straight chains of cubes with the occasional kink.

For each i and j , $f_2(p'_{ij})$ is a path in $X_\varepsilon^{(1)}$, and (together with the edge from $f_2(v_i)$ to $f_2(v_j)$ if they are not the same) it forms a loop $q_{ij} : S^1 \rightarrow X_\varepsilon^{(1)}$. We would like to glue in a slab equipped with a map to a contraction of this loop. This slab consists of several pieces:

1. A thickened annulus which straightens out the kinks in the chain of cubes; both sides map to the same path in $X_\varepsilon^{(1)}$.

Since the kinks are isolated and come in $\text{const}(n)$ types, this can be done with $\text{const}(n) \cdot \text{length}(p'_{ij})$ vertices, perhaps after subdividing both ends a

⁴See [17] for more details of a similar construction in three dimensions, used to fit many maximally separated paths into a cube.

bounded number of times (in which case we also must subdivide Σ_M and X_M .)

The remainder of the slab will have the product cubical structure of a disk with $[0, 1]^{n-1}$.

2. Another thickened annulus which maps to a composition of the path with a retraction of $X_\varepsilon^{(1)}$ to $X^{(1)}$.

We can perform such a retraction by expanding outward in each direction from one central ε -cube. This map has Lipschitz constant $\sim \varepsilon^{-1}$ and so the inner side of this annulus will consist of

$$\text{const}(n) \cdot \varepsilon^{-1} \cdot \text{length}(p'_{ij})$$

cubes.

3. A thickened annulus which maps to the composition of this with the provided contraction of $X^{(1)}$. This consists of

$$\text{const}(n) \cdot V\varepsilon^{-1} \cdot \text{length}(p'_{ij})$$

cubes.

4. A thickened disk which maps to the basepoint to which $X^{(1)}$ is contracted. This is a thickening of a cone on the inside of the previous annulus.

After all the slabs are glued on, we let Σ_3 be the n -dimensional “outer boundary” of the resulting complex, equipped with a map $f_3 : \Sigma_3 \rightarrow X_M$.

We now need to algorithmically glue in thin slabs so that the new disks are brought close to other parts of the complex which are mapped to nearby points of X_M . We first choose an ε -net of vertices in each disk and the shortest path via n -cubes from each point in the net to the boundary of the disk. (This means retaining the data of whether a given portion of Σ_3 is inherited from Σ_M .) Call the set of all the vertices in these nets V'_ε

Now, for each of these paths, we choose a path in Σ_M which maps to a nearby path in X_M . Specifically, we make the choice as follows. Choose a sequence of “signposts” x_i at intervals of length $\varepsilon/3m$ along the path. For each signpost x_i , choose the nearest vertex v_i of X_ε . Then the path between x_i and x_{i+1} is contained within the star of the vertex or edge spanned by v_i and v_{i+1} , and we can build a cubical approximation to the homotopy which retracts that path linearly to the vertex or edge.

The v_i give us a corresponding path in X_ε , and we can greedily construct a path of n -cubes in Σ_M which maps to that path, ending at the same point on the boundary of the disk. After subdividing Σ_3 and X_M yet again, this time with scale parameter

$$M' = 10(n+1)|V'_\varepsilon|,$$

we use the same technique as before to create non-intersecting sub-paths. Then for each pair of sub-paths, we glue in a slab whose cells map to X_M via the piecewise linear homotopy described above. The number of cubes in this slab is

$$O(\text{length of the path} \cdot M \cdot (M')^2).$$

This completes the construction of the space Σ and the map f . □

A NP algorithms for the Lipschitz constant

In this appendix, we show that $L_{\neq 0}(\Sigma)$ and $L_1(\Sigma)$ can be approximated up to a constant factor by an NP algorithm. In fact, NP algorithms can give constant factor approximations to $\text{Lip}(\alpha)$ in the more general setting of Theorem 2.1. We need a little background to state the result precisely.

An NP *optimization problem* is an optimization problem for which feasible solutions can be verified in polynomial time. In other words, it is a numerical function f from some input set A to \mathbb{N} (or some other discrete totally ordered set) such that the statement “ $f(x) \leq r$ ”, if true, has a polynomial-time certificate. Then the true optimum can be found by following all possible certificates in nondeterministic fashion.

Theorem A.1. *If n is odd or $m \leq 2n - 2$, there is an NP optimization problem which, given an m -dimensional simplicial complex X and a class $\alpha \in [X, S^n]$, outputs a value within a multiplicative constant $C(m, n)$ of $\text{Lip}(\alpha)$.*

Informally, we can say that computing the Lipschitz norm up to $C(m, n)$ is an NP optimization problem. In addition, if $\text{Lip}([f]) \geq 1$, then it can even be computed in polynomial time up to a multiplicative constant using the algorithm in §3.

The first step of the proof is the following absolute lower bound for $\text{Lip}(\alpha)$.

Lemma A.2. *If $\alpha \in [X, S^n]$ is not the zero class, then*

$$\text{Lip}(\alpha) \geq c(m, n)(\# \text{ of simplices of } X)^{-1/n}.$$

Proof. First, a non-nullhomotopic map to S^n always covers S^n : otherwise, it is a map to the disk and hence nullhomotopic.

Secondly, at the cost of a multiplicative constant, we can restrict to simplicial maps from a subdivision of X to a subdivision of Y , because of the following result:

Lemma A.3 (Quantitative simplicial approximation theorem [3]). *Let X and Y be simplicial complexes with the standard simplexwise linear metric. Then there is a constant $C_{\text{QSAT}} = C_{\text{QSAT}}(\dim X, \dim Y)$ such that every L -Lipschitz map $f : X \rightarrow Y$ can be approximated by a $C_{\text{QSAT}}(L + 1)$ -Lipschitz simplicial map on a subdivision of X at scale $\sim 1/L$.*

We can choose the subdivision on both sides to be the edgewise subdivision due to Edelsbrunner and Grayson [8] with scale parameters s and t , respectively, indicating the degree of subdivision of each edge. When we give the simplices of the subdivision of X the standard linear metric, the resulting metric space differs from an s -times dilation of X by a multiplicative constant $C(\dim X)$; likewise for Y and t . We call the resulting metric spaces X_s and Y_t .

In order to cover S^n , a simplicial map must hit every n -simplex of S^n . Thus to admit a surjective simplicial map $X_s \rightarrow S_t^n$ without subdividing further, X_s must have at least as many n -simplices as S_t^n . Now, X_s has

$$\leq (\# \text{ of simplices of } X) \cdot C(m) \cdot s^m$$

n -simplices. In particular, we can choose $t = C(m, n)(\# \text{ of simplices of } X)^{1/n}$ such that a surjective simplicial map $X_s \rightarrow S_t^n$ can only exist when $s \geq 2C_{\text{QSAT}}(m, n)$. By the lemma, this means that any surjective map $X \rightarrow S_t^n$ must have Lipschitz constant at least 1, and so any non-nullhomotopic map $X \rightarrow S^n$ must have Lipschitz constant at least $1/t$. \square

Thus, suppose we are given a map $f : X \rightarrow S^n$ which we know is not nullhomotopic. Then:

- A homotopic simplicial map $X_s \rightarrow S_t^n$, where $t = O((\# \text{ of simplices of } X)^{1/n})$ and $s \leq 2C_{\text{QSAT}}(m, n)$, is a certificate that $\text{Lip}([f]) \leq s/t$.
- Moreover, if there is *no* such simplicial map, then $\text{Lip}([f]) \geq 1/t$.

In our situation, the (polynomial-size) data of a simplicial map $X_s \rightarrow S_t^n$ in the right homotopy class is such a certificate, and its validity can be verified in polynomial time. In particular, if the homotopy class is specified using a representative map, then the fact that our map is homotopic to it can be computed in polynomial time for fixed dimension of X [10]. This shows Theorem A.1.

References

- [1] A. Berdnikov and F. Manin, *Scalable spaces*, arXiv preprint arXiv:1912.00590 (2019).
- [2] Aleksandr Berdnikov, *Lipschitz null-homotopy of mappings $S^3 \rightarrow S^2$* , arXiv preprint arXiv:1811.02606 (2018).
- [3] G. R. Chambers, D. Dotterrer, F. Manin, and S. Weinberger, *Quantitative null-cobordism*, J. Amer. Math. Soc. **31** (2018), no. 4, 1165–1203.
- [4] G. R. Chambers, F. Manin, and S. Weinberger, *Quantitative nullhomotopy and rational homotopy type*, Geometric and Functional Analysis (GAFA) **28** (2018), no. 3, 563–588.

- [5] I. Dinur, *Approximating SVP_∞ to within almost-polynomial factors is NP-hard*, Theoret. Comput. Sci. **285** (2002), no. 1, 55–71, Algorithms and complexity (Rome, 2000). MR 1925787
- [6] A. N. Dranishnikov, Steven C. Ferry, and Shmuel Weinberger, *Large Riemannian manifolds which are flexible*, Ann. of Math. (2) **157** (2003), no. 3, 919–938. MR 1983785
- [7] Ramsay Dyer, Gert Vegter, and Mathijs Wintraecken, *Riemannian simplices and triangulations*, Geom. Dedicata **179** (2015), 91–138. MR 3424659
- [8] H. Edelsbrunner and D. R. Grayson, *Edgewise subdivision of a simplex*, Discrete & Computational Geometry **24** (2000), no. 4, 707–719.
- [9] S. Ferry and B. Okun, *Approximating topological metrics by Riemannian metrics*, Proceedings of the American Mathematical Society **123** (1995), no. 6, 1865–1872.
- [10] M. Filakovský and L. Vokřínek, *Are two given maps homotopic? An algorithmic viewpoint*, Foundations of Computational Mathematics **20** (2020), no. 2, 311–330.
- [11] M. Gromov, *Homotopical effects of dilatation*, Journal of Differential Geometry **13** (1978), no. 3, 303–310.
- [12] ———, *Width and related invariants of Riemannian manifolds*, Asterisque **163-164** (1988), no. 6, 93–109.
- [13] Mikhael Gromov and H. Blaine Lawson, Jr., *Spin and scalar curvature in the presence of a fundamental group. I*, Ann. of Math. (2) **111** (1980), no. 2, 209–230. MR 569070
- [14] L. Guth, *Lipshitz maps from surfaces*, Geometric & Functional Analysis (GAFA) **15** (2005), no. 5, 1052–1090.
- [15] ———, *Minimax problems related to cup powers and Steenrod squares*, Geometric & Functional Analysis (GAFA) **18** (2009), no. 6, 1917–1987.
- [16] ———, *Recent progress in quantitative topology*, Surveys in Differential Geometry **22** (2017), 191–216.
- [17] A.N. Kolmogorov and Ya.M. Barzdin, *On the realization of nets in 3-dimensional space*, Problems in Cybernetics **8** (1967), no. 261-268, 259–260.
- [18] F. Manin, *Plato’s cave and differential forms*, Geometry & Topology **23** (2019), no. 6, 3141–3202.
- [19] Shmuel Weinberger, *Computers, rigidity, and moduli*, M. B. Porter Lectures, Princeton University Press, Princeton, NJ, 2005. MR 2109177