RECENT PROGRESS IN QUANTITATIVE TOPOLOGY

LARRY GUTH

Abstract. We discuss recent work of Chambers, Dotterrer, Ferry, Manin, and Weinberger, which resolved a fundamental question in quantitative topology: if $f : S^m \to S^n$ is a contractible map with Lipschitz constant L, what can we say about the Lipschitz constant of a null-homotopy of f?

In the mid-90s, Gromov wrote an article on quantitative topology [G], raising a number of interesting questions. We will focus on the following question.

Question 0.1. *Equip S^m and S^n with the unit sphere metrics, and suppose that $f : S^m \to S^n$ is a contractible map with Lipschitz constant L. What is the best Lipschitz constant of a null-homotopy $H : S^m \times [0,1] \to S^n$?*

For a long time, little was known about this question. In some simple cases, there were good estimates. For general m and n, Gromov sketched a construction giving a homotopy H with Lipschitz constant at most

$$\exp(\exp(...(\exp L)...)),$$

where the height of the tower of exponentials depends on the dimensions m and n [G]. On the other hand, he suggested that there may always be a homotopy with Lipschitz constant at most $C(m,n)L$. I thought about the problem a number of times but couldn't see any way to get started. Recently, the problem was almost completely solved by work of Chambers, Dotterrer, Ferry, Manin, and Weinberger. Here are their results.

Theorem 0.2. *([CDMW]) Suppose that n is odd and $f : S^m \to S^n$ is a contractible map with Lipschitz constant L. Then there is a null-homotopy $H : S^m \times [0,1] \to S^n$ with Lipschitz constant at most $C(m,n)L$.*

Theorem 0.3. *([CMW]) Suppose that n is even and $f : S^m \to S^n$ is a contractible map with Lipschitz constant L. Then there is a null-homotopy $H : S^m \times [0,1] \to S^n$ with Lipschitz constant at most $C(m,n)L^2$.*

More precisely, there is a null-homotopy H which is $C(m,n)L$-Lipschitz in the S^m-direction and $C(m,n)L^2$-Lipschitz in the $[0,1]$-direction. In other words:

$$\mathrm{dist}_{S^n}(H(x_1,t_1), H(x_2,t_2)) \le C(m,n) \left(L \, \mathrm{dist}_{S^m}(x_1,x_2) + L^2 |t_1 - t_2| \right).$$

When n is even, it is still an open question whether there is a null-homotopy H with Lipschitz constant at most $C(m,n)L$. Nevertheless, these results are a huge improvement over what was known before. They give a near complete solution to a fundamental problem.

Question 0.1 lies on the border between algebraic topology and metric geometry. It is an inventive variation on the theme of the isoperimetric inequality in the domain of homotopy theory. The proof involves ideas from both areas. On the homotopy theory side, it uses Serre's classification of the rational homotopy groups of spheres, as well as ideas from obstruction theory. On the metric geometry side, it uses estimates related to the isoperimetric inequality. One of the key observations in the proof is that the traditional isoperimetric inequality is connected to this more unusual homotopy-theoretic variation on the isoperimetric inequality.

The theorems in [CDMW] and [CMW] are more general than what we stated here. Using the more general theorem, Chambers, Dotterrer, Manin, and Weinberger give a nice application to a quantitative version of cobordism theory. But in this article I want to focus on the simplest statements. Our main goal will be to give a detailed sketch of the proof of Theorem 0.2. By focusing on this simplest case, we will be able to avoid some of the technical issues that appear in the more general theorem, while still explaining some of the key new ideas.

In the first section, I will describe some previous approaches to the problem, and try to give a sense of why it seemed difficult to me. In the second section, I will explain the new ideas from [CDMW] and give a detailed sketch of the proof of Theorem 0.2. In the last section, I will speculate on further directions and open problems.

Notation. We willl write $A \lesssim B$ to mean that there is a constant $C(m,n)$ so that $A \leq C(m,n)B$.

Acknowledgements. I would like to thank Sasha Berdnikov who read an earlier draft and made a number of helpful suggestions.

1. Background and previous approaches

In this section we will discuss some previous approaches to the quantitative null-homotopy problem. We will see some simple cases where good bounds had been proven already, and try to build up some intuition for the problem. One of the themes is taking ideas from topology and making them more quantitative. In some cases, when we quantify a classical argument from topology, it leads to good bounds. But in other cases, it leads to bounds that seem far from optimal. In particular, we will indicate where the tower of exponentials in Gromov's bound comes from.

1.1. Maps S^m to S^n when $m < n$.

The first case we consider is the case when the dimension of the domain S^m is smaller than the dimension of the range S^n. In this case, every map $f : S^m \to S^n$ is contractible. We start by recalling the proof and then we develop a quantitative version of the argument.

Proposition 1.1. *If $m < n$ and $f : S^m \to S^n$ is a continuous map, then f is contractible.*

Proof. The complement of a point in S^n is contractible. Therefore, if any non-surjective map to S^n is contractible. Since $m < n$, we may hope that $f : S^m \to S^n$ is non-surjective. However, we have to be careful here. Using Peano's construction of a space filling curve, one can construct continuous surjective maps $f : S^m \to S^n$ as long as $m \geq 1$. We first homotope f to a smooth or piecewise linear map $\tilde{f}$. The map $\tilde{f} : S^m \to S^n$ is non-surjective, and so it is contractible. □

Now suppose that $f : S^m \to S^n$ has Lipschitz constant L. We would like to adapt the argument above to produce a null-homotopy with controlled Lipschitz constant. Because f is Lipschitz, it is not surjective. The standard proof that f is not surjective gives a quantitative estimate which we record here as a lemma.

Lemma 1.2. *If $m < n$ and $f : S^m \to S^n$ has Lipschitz constant L, then the image of f misses a ball of radius r for $r \gtrsim L^{-\frac{m}{n-m}}$.*

Proof. For any $\rho > 0$, S^m can be covered by $\sim \rho^{-m}$ balls of radius ρ. The image of each such ball is contained in a ball of radius $L\rho$. Therefore, the image of f is contained in $\lesssim \rho^{-m}$ balls of radius

$L\rho$. We set $r = L\rho$. We can also say that the r-neighborhood of the image of f is contained in $\lesssim \rho^{-m}$ balls of radius $2r = 2L\rho$. We have to check that these balls cannot cover all of S^n. Their total volume is $\lesssim \rho^{-m}(L\rho)^n = L^n\rho^{n-m}$. We choose ρ so that this number is smaller than 1. So we get

$$\rho \sim L^{-\frac{n}{n-m}},$$

$$r \sim L\rho \sim L^{-\frac{m}{n-m}}.$$

□

We recalled above that the complement of a point in S^n is contractible. If we remove a ball from S^n, the leftover part can be contracted in a Lipschitz way.

Lemma 1.3. *For each radius r, there is a contraction $G : S^n \setminus B_r \times [0,1] \to S^n \setminus B_r$ with Lipschitz constant at most $1/r$. More precisely, G has Lipschitz constant $\lesssim r^{-1}$ in the S^n direction and $\lesssim 1$ in the $[0,1]$ direction.*

Proof. (proof sketch) We describe the sphere S^n using polar coordinates around a point p. The coordinates are $(\rho, \theta) \in [0, \pi] \times S^{n-1}$, and the standard spherical metric is

$$ds^2 = d\rho^2 + (\sin\rho)^2 d\theta^2.$$

We choose the point p to be the antipode of the center of B_r, so that in our coordinates, the domain $S^n \setminus B_r$ is given by $\rho \in [0, \pi - r]$. Now we write down our map

$$G(\rho, \theta, t) = ((1-t)\rho, \theta).$$

Clearly $G(\rho, \theta, 0) = (\rho, \theta)$, and $G(\rho, \theta, 1) = p$.

To check the Lipschitz constant of G, we compute the derivative of G and then compute $\sup |dG|$. It's not hard to check that this supremum is achieved when the boundary of B_r passes over the equator: i.e. when $\rho = \pi - r$ and t solves

$$(1-t)\rho = \pi/2.$$

At this point,

$$|dG| = \frac{\sin(\pi/2)}{\sin(\pi - r)} = (\sin r)^{-1} \sim r^{-1}.$$

Computing the derivative of G in the t-direction, we see that $|\partial_t G| \leq \rho \leq pi$, and so G has Lipschitz constant $\lesssim 1$ in the t direction.

□

Remark. It is a good exercise to prove that any null-homotopy of $S^n \setminus B_r$ has Lipschitz constant $\gtrsim 1/r$.

We can now build a null-homotopy of f. We know that the image of f lies in $S^n \setminus B_r$ for $r \sim L^{-\frac{m}{n-m}}$. We use the map G to contract this region. So our final homotopy H is the composition of G with $f \times id$. The Lipschitz constant of H is at most

$$\operatorname{Lip} H \leq \operatorname{Lip} G \cdot \operatorname{Lip} f \lesssim L^{\frac{m}{n-m}} L = L^{\frac{n}{n-m}}.$$

This construction, however, is not optimal. In fact the following better estimate holds.

Proposition 1.4. *If $m < n$ and $f : S^m \to S^n$ has Lipschitz constant L, then there is a null-homotopy with Lipschitz constant $\lesssim L$. In fact the null-homotopy has Lipschitz constant $\lesssim L$ in the S^m directions and $\lesssim 1$ in the $[0,1]$ direction.*

Looking back at the proof that $f : S^m \to S^n$ is contractible, there was a step where we approximate the map f with a smooth or piecewise linear map. We can get our better quantitative estimate by really incorporating that idea. We do so using simplicial approximation. This simplicial approximation idea will be useful all through the paper.

Let Tri_L be a triangulation of S^m into simplices which are bilipschitz to equilateral simplices of side length $\sim \frac{1}{L}$, with bilipschitz constant ~ 1. We will refer to this in the future as a triangulation into standard simplices of side length $\sim \frac{1}{L}$. We will have to be a bit more precise about the side length: let the side length be at most $\frac{c(m,n)}{L}$ for a small constant $c(m,n)$ to be chosen later. Let Tri_{S^n} be the triangulation of S^n as the boundary of the $(n+1)$-simplex. Simplicial approximation gives a controlled homotopy from (S^m, Tri_L) to $(S^n, \mathrm{Tri}_{S^n})$.

Lemma 1.5. *(Simplicial approximation) For any dimensions m and n, suppose that $f : S^m \to S^n$ has Lipschitz constant L. Let Tri_L be as in the previous paragraph. Then there is a simplicial map f_{simp} from (S^m, Tri_L) to $(S^n, \mathrm{Tri}_{S^n})$ and a homotopy $H_{simp} : S^m \times [0,1] \to S^n$ so that H is $\lesssim L$-Lipschitz in the S^m direction and $\lesssim 1$-Lipschitz in the $[0,1]$ direction.*

Remark. We use (S^m, Tri_L) to denote the simplicial complex given by the triangulation Tri_L. It is not a pair in the sense of relative homology, etc. By a simplicial map from (S^m, Tri_L) to $(S^n, \mathrm{Tri}_{S^n})$, we just mean a simplicial map from one simplicial complex to the other.

Proof. (proof sketch) We recall the construction of the simplicial approximation, following Hatcher [H], Section 2C, page 177-178. If v is a vertex of Tri_L, we let $\mathrm{St}\, v$ denote the closed star of v, the union of all closed simplices of Tri_L containing v. We see that $\mathrm{St}\, v$ is contained in a $\frac{c(m,n)}{L}$ ball in S^m and so $f(\mathrm{St}\, v)$ lies in a $c(m,n)$-ball in S^n. By choosing $c(m,n)$ small enough, we can guarantee that any such ball lies in the open star of some vertex $f_{simp}(v)$ in Tri_{S^n}. Then [H] explains that f_{simp} extends to a simplicial map from (S^m, Tri_L) to $(S^n, \mathrm{Tri}_{S^n})$ and that for each point $x \in S^m$, $f(x)$ and $f_{simp}(x)$ lie in a common simplex of $(S^n, \mathrm{Tri}_{S^n})$. We can define H_{simp} by taking the straight-line homotopy from $f(x)$ to $f_{simp}(x)$ inside this simplex. The Lipschitz constant of f_{simp} is at most $C(m,n)L$, and so the Lipschitz constant of this straight-line homotopy is at most $C(m,n)L$ in the S^m direction and at most $C(m,n)$ in the $[0,1]$ direction. □

To get a more efficient contraction of $f : S^m \to S^n$, we first simplicially approximate f using Lemma 1.5. The map f_{simp} has image in the m-skeleton of $(S^n, \mathrm{Tri}_{S^n})$. If $m < n$, f_{simp} misses a ball of radius ~ 1 in S^n. Applying our previous strategy to f_{simp}, we get an efficient null-homotopy of f, proving Proposition 1.4.

Remark. Simplicial approximation is a key tool in the proof of Theorem 0.2. The rest of the topics in our background discussion will not be needed later, so the reader who is interested in getting to the proof of Theorem 0.2 can jump from this point to Section 2.

1.2. **Maps between spheres of the same dimension.** Next we consider maps from S^n to S^n. Brouwer defined the degree of a map $f : S^n \to S^n$ and he proved that f is contractible if and only if the degree of f is zero. When the degree of f is zero, Brouwer constructed a null-homotopy of f. By using Brouwer's construction and adding quantitative bounds, Gromov was able to prove the following result (personal communication).

Proposition 1.6. *If $f : S^n \to S^n$ has degree zero and Lipschitz constant L, then f extends to a map $H : B^{n+1} \to S^n$ with Lipschitz constant $\lesssim L$.*

By the simplicial approximation argument in Lemma 1.5, we can assume without loss of generality that f is a simplicial map from (S^n, Tri_L) to $(S^n, \mathrm{Tri}_{S^n})$. Recall that Tri_L is a triangulation into standard simplices of side length $\sim 1/L$ and Tri_{S^n} is the triangulation of S^n as the boundary of the $(n+1)$-simplex.

Let us quickly recall the definition of the degree. Let y be a point in the center of one of the top-dimensional faces of the target $(S^n, \mathrm{Tri}_{S^n})$. Since f is simplicial, the preimage $f^{-1}(y)$ is a finite subset of the domain (S^n, Tri_L), consisting of at most one point in each top dimensional simplex, and no points in lower dimensional simplices. Near each point $x_i \in f^{-1}(y)$, the map f is either orientation-preserving or orientation-reversing, and we assign each point a multiplicity ± 1 based on the orientation. The degree of f is defined to be the sum of these multiplicities.

Suppose f has degree zero. The first step of Brouwer's construction is to build a collection of closed arcs γ_i in B^{m+1} each going from a positive point of $f^{-1}(y)$ to a negative point of $f^{-1}(y)$, in such a way that each point of $f^{-1}(y)$ lies in exactly one arc. To give a quantitative version of Brouwer's construction, the main new ingredients is to choose these arcs so that they don't come too close to each other.

Lemma 1.7. *Suppose that $\{x_{i,+}\}_{i=1,\ldots,I}$ and $\{x_{i,-}\}_{i=1,\ldots,I}$ are finite sets in S^n, with at most one point in any ball of radius $\sim 1/L$. Then there are curves $\gamma_i \subset B^{n+1}$, so that γ_i connects $x_{i,+}$ to $x_{i,-}$ for each i, and the distance between any two curves γ_i is $\gtrsim 1/L$.*

This lemma was proven by Kolmogorov and Barzdin [KB] in the case $n = 2$. It is a special case of their construction of thick embeddings of graphs into the ball. The case $n = 2$ is the hardest case, and their proof works in higher dimensions as well. The case $n = 1$ is different but easier. We will come back to discuss the proof of the lemma, but first sketch how to use it.

By straightening out the curves γ_i, we can arrange that each γ_i is a sequence of geodesic segments of length $\sim 1/L$, and the distance between any two non-consecutive segments is $\gtrsim 1/L$. The next step is to thicken each curve γ_i to a thin tube T_i. The tube T_i intersects the boundary S^n in disks around $x_{i,+}$ and $x_{i,-}$ of radius $\gtrsim 1/L$. We can assume that f maps each of these disks onto the ball $B_{\frac{1}{10}}(y)$ in the target S^n. The whole tube T_i has thickness $1/L$ and it is $C(n)$-bilipschitz to a cylinder $B^n(1/L) \times [0, \ell_i]$ where ℓ_i is the length of γ_i. Next we define H on each tube T_i so that H maps T_i to $B_{1/10}(y)$ and maps $\partial B^n \times [0, 1]$ to the boundary of that ball. So far, the Lipschitz constant of H is $\lesssim L$.

Let X denote $B^{n+1} \setminus \cup T_i$. We have defined H on the boundary of X and it remains to define H on the interior of X. Note that H maps the boundary of X to $S^n \setminus B_{1/10}(y)$, which is contractible, and so H can be extended to X. In fact $S^n \setminus B_{1/10}(y)$ is bilipschitz equivalent to the standard unit n-ball B^n and we identify them for convenience.

Lemma 1.8. *Suppose that X is a compact (piecewise smooth) Riemannian manifold with boundary. Suppose that $g : \partial X \to \partial B^n$ is given, and we want to extend g to $G : X \to B^n$. Suppose that*

$$\psi : N_{1/L}(\partial X) \to \partial X$$

is a retraction. Then there exists an extension G with

$$\mathrm{Lip}\, G \leq L + \mathrm{Lip}\, g \, \mathrm{Lip}\, \psi.$$

Proof. Let $d(x)$ denote the distance to the boundary of X. Our extension G will map points $x \in X$ with $d(x) \geq 1/L$ to the origin in B^n. The map G is

$$G(x) = (1 - Ld(x))_+ g(\psi(x)).$$

By the Leibniz rule and the chain rule,

$$|dG| \leq L + \operatorname{Lip} g \operatorname{Lip} \psi.$$

□

In our case, g is the map H we have constructed on ∂X, and it has Lipschitz constant $\lesssim L$. Since the distances between tubes T_i are $\gtrsim 1/L$, it is not hard to construct a retraction $\psi : N_{1/L}(\partial X) \to \partial X$ with Lipschitz constant $\lesssim 1$, and so our final homotopy has Lipschitz constant $\lesssim L$ as desired.

The proof of Lemma 1.7 is a little tricky. I learned about it in Arnold's reminiscences about Kolmogorov [A], page 94. Arnold writes that "when I mentioned it in a paper in Physics Today dedicated to Kolmogorov (1989), I received a sudden deluge of letters from engineers who were apparently working in miniaturization of computers, with requests for a precise reference to his work."

The spirit of the argument is as follows. Given two points, $x_{i,+}$ and $x_{i,-}$, we construct many different curves $\gamma_{i,j}$ from $x_{i,+}$ to $x_{i,-}$. We select γ_1 from among the curves $\gamma_{1,j}$. Next we select γ_2 from among the curves $\gamma_{2,j}$, making sure it does not pass too close to γ_1. And so on. At step i, we have to see that one of the curves $\gamma_{i,j}$ stays far enough away from the previous curves $\gamma_1, ..., \gamma_{i-1}$. In fact, we will show that at step i, if we choose j randomly, the probability that $\gamma_{i,j}$ comes too close to one of the previous curves is less than one half.

We do our construction in two stages. For simplicity, suppose that the points $x_{i,+}$ and $x_{i,-}$ all lie on the bottom face of unit $(n+1)$-cube. The curve $\gamma_{i,j}$ has the following form. Starting at $x_{i,+}$, we first draw a segment in the x_{n+1} direction to a height of the form $j_0/100nL$, where j_0 is an integer between 0 and $100nL$. Next we draw a segment in the x_1-direction a distance of the form $j_1/100nL$. Then we draw a segment in the x_2 direction a distance of the form $j_2/100nL$. We continue in this way up to a segment in the x_{n-1} direction of length $j_{n-1}/100nL$. We have about $(100nL)^n$ choices for $j_0, ..., j_{n-1}$. Then we draw a segment in the x_n direction which ends at the x_n coordinate of $x_{i,-}$. Then we draw a segment in the x_{n-1} direction which ends at the x_{n-1} coordinate of our target $x_{i,-}$, etc. Finally, we draw a segment in the x_{n+1}-direction which ends at $x_{i,-}$.

We claim that we can choose $j_0, j_1, ..., j_{n-1}$ so that $\gamma_{i,j}$ has only perpendicular intersections with the previously selected γ_i. This is the first stage of the construction. Since each curve is made of segments that point in the coordinate directions, we just have to check that none of the segments intersects a segment of a previous curve going in the same direction. Call a segment bad if it intersects a segment from a previous curve going in the same direction.

The initial vertical segment of $\gamma_{i,j}$ cannot intersect any of the vertical segments of previous γ_i just because $x_{i,+}$ is distinct from the other points. Consider the first segment in the x_1 direction. On this segment, the $2...n$ coordinates are fixed equal to those of $x_{i,+}$. This segment can intersect an x_1 segment of a previous curve $\gamma_{i'}$ only if $x_{i,+}$ has the same $2...n$ coordinates as $x_{i',+}$ or $x_{i',-}$. This leaves at most $2L$ worrisome values of i'. But on this segment of $\gamma_{i,j}$, the $(n+1)$ coordinate is fixed equal to $j_0/100nL$. This segment is bad only if it has the same x_{n+1} coordinate as an x_1 segment from one of the $2L$ worrisome values of i'. But there are more than $100nL$ choices of j_0. So the probability that this first segment is bad is at most $\frac{1}{50n}$.

A similar argument holds for the the second segment. This x_2 segment can intersect a previous curve $\gamma_{i'}$ only if $x_{i,+}$ has the same $3...n$ coordinates as $x_{i',+}$ or $x_{i',-}$. This leaves only $2L^2$ worrisome values of i'. But there are more than $(100nL)^2$ choices for (j_0, j_1). So the probability that the second segment is bad is again at most $\frac{1}{50n}$.

The same reasoning applies for the first n segments. And in fact the same reasoning applies for the following n segments as well. For instance, consider the second (and last) segment in the x_1 direction. Over the course of this segment, the $2...n$ coordinates are equal to those of $x_{i,-}$, and so this segment can intersect an x_1-segment of a previous curve $\gamma_{i'}$ only if $x_{i,-}$ has the same $2...n$ coordinates as $x_{i',+}$ or $x_{i',-}$. This leaves at most $2L$ worrisome values of i'. But there are more than $100nL$ choices of j_0, and so the probabilty that this segment is bad is at most $\frac{1}{50n}$.

In summary, there are $2n$ segments, and each has probabily at most $\frac{1}{50n}$ of being bad. So more than half the time, all the segments are good, and this gives us a curve γ_i which only intersects previous curves perpendicularly.

In the second stage, we get rid of these perpendicular intersections by adding extra wiggles at a small scale $\sim 1/L$. Consider an intersection point of the curves γ_i. Since the curves intersect perpendicularly, at most $(n+1)$ of them intersect at the given point. Moreover, there is a ball around the intersection point of radius $\gtrsim 1/L$ that intersects at most $n+1$ of the curves γ_i. We can get rid of the intersection point by adding small extra wiggles to these curves inside the ball in question. Repeating this process for all the intersection points finishes the construction.

This finishes our sketch of the proof of Proposition 1.6. Turning Brouwer's initial construction into a quantitative bound leads to an interesting new problem of embedding thick tubes, and it ultimately leads to a sharp estimate for our quantitative problem.

Incidentally, we did not yet cover the case of maps $S^1 \to S^1$. This case is easier because we can use the universal cover. We include the argument here in a bit more generality for future reference.

Proposition 1.9. *If $f : S^m \to S^1$ is null-homotopic with Lipschitz constant L, then f extends to a null-homotopy $h : S^m \times [0,1] \to S^1$ with Lipschitz constant $\le L$.*

Proof. Since f is contractible, we can lift f to the universal cover to get a map $\tilde{f} : S^m \to \mathbb{R}$. The Lipschitz constant of $\tilde{f}$ is also L. Let y_0 be a point in the image of $\tilde{f}$. Then we define the null-homotopy $h_1 : S^m \times [0,1] \to \mathbb{R}$ by

$$h_1(x,t) = ty_0 + (1-t)\tilde{f}(x).$$

We denote $\pi : \mathbb{R} \to S^1$ the map from the universal cover of S^1 to S^1, and we define $h : S^m \to S^1$ by $\pi \circ h_1$. □

1.3. Maps from S^3 to S^2.

The most interesting homotopy theory of spheres is about maps S^m to S^n with $m > n$, and our quantitative problem is also most difficult and interesting in this setting. The simplest example is maps from S^3 to S^2. The long exact sequence of the Hopf fibration gives one way to compute $\pi_3(S^2)$. Recall that the Hopf fibration is a fiber bundle $h : S^3 \to S^2$ where each fiber is a great circle. So the long exact sequence gives

$$\pi_3(S^1) \to \pi_3(S^3) \to \pi_3(S^2) \to \pi_2(S^1) \to ...$$

Since $\pi_3(S^1) = \pi_2(S^1) = 0$, we see that $h_* : \pi_3(S^3) \to \pi_3(S^2)$ is an isomorphism. By Brouwer's degree theory, $\pi_3(S^3)$ is isomorphic to $\mathbb{Z}$ and so $\pi_3(S^2)$ is also isomorphic to $\mathbb{Z}$. We will use the Hopf fibration to give quantitative bounds for null-homotopies, proving the following special case of Theorem 0.3:

Proposition 1.10. *Suppose that $f : S^3 \to S^2$ is a null-homotopic map with Lipschitz constant L. Then f extends to a homotopy $H : B^4 \to S^2$ with Lipschitz constant $\lesssim L^2$.*

The construction is based on the homotopy lifting property of the Hopf fibration. The homotopy lifting property says that given a map $g : K \times [0, T] \to S^2$, and a lift of g at time 0, $g^+ : K \to S^3$ with $h(g^+(x)) = g(x, 0)$, then we can extend g^+ to $K \times [0, T]$ with $h(g^+(x, t)) = g(x, t)$ for all x, t. We can construct this lift geometrically by using parallel transport with respect to the standard connection on the Hopf fibration. For our quantitative application, we need to know a quantitative estimate for this lifting.

Lemma 1.11. *If $g : K \times [0, 1] \to S^2$ has Lipschitz constant L_K in the K direction and $L_T \geq 1$ in the t direction, and $g^+ : K \to S^3$ has Lipschitz constant L_K, then g^+ extends to $K \times [0, 1]$ with Lipschitz constant $\lesssim L_K L_T$ in the K-direction and Lipschitz constant $\lesssim L_T$ in the t direction.*

We sketch the proof. The extension is given by parallel transport. To understand the Lipschitz constant in the t direction, we may as well consider K to be a point. Suppose $g : [0, 1] \to S^2$ is a map and $g^+ : [0, 1] \to S^3$ is a lift given by parallel translation. Suppose that $g'(t)$ is a vector v in $T_y S^2$, where $y = g(t)$. Let $y^+ = g^+(t)$. Then $(g^+)'(t)$ is the horizontal lift of v to $T_{y^+} S^3$, which we label v^+. It is straightforward to check that $|v^+| \lesssim |v|$, and this gives the desired bound.

To understand the Lipschitz constant in the K direction, we may as well consider K to be an interval $[0, 1]$. We need to show that the distance from $g^+(1.0)$ to $g^+(1, 1)$ is $\lesssim L_K L_T$.

To get a first intution, consider the special case that g maps $[0, 1] \times \{0\}$ and $[0, 1] \times \{1\}$ to a single point $q_0 \in S^2$ and g^+ maps $[0, 1] \times \{0\}$ to a single point p_0 in $h^{-1}(q_0) \subset S^3$. We know that g maps $\{0\} \times [0, 1]$ to a closed curve γ_0 from q_0 to itself, and that g maps $\{1\} \times [0, 1]$ to a closed curve γ_1 from q_0 to itself. We want to understand the distance between the points $g^+(0, 1)$ and $g^+(1, 1)$. These points are in the same fiber: $h^{-1}(q_0)$. This distance is the difference between the parallel transport of p_0 around γ_0 and the parallel transport of p_0 around γ_1. Because of our special assumption, this distance is also the parallel transport around the curve $g(\partial(K \times [0, 1]))$. We can express this parallel transport in terms of the curvature of the standard connection on the Hopf fibration. We let ω denote the curvature form of the connection. Then the holonomy of the connection around the loop $g(\partial K \times [0, 1])$ is equal to $\int_{K \times [0,1]} g^* \omega$. Because g has Lipschitz constant L_K in the K direction and L_T in the t direction, $|g^* \omega| \leq L_K L_T |\omega| \lesssim L_K L_T$. Therefore, the distance from $g^+(0, 1)$ to $g^+(1, 1)$ is

$$\leq \int_{[0,1] \times [0,1]} |g^* \omega| \lesssim L_K L_T.$$

Now let us sketch how to remove our special assumption that g^+ maps $[0, 1] \times \{0\}$ to p_0 and that g maps $[0, 1] \times \{1\}$ to q_0. In this case, we can bound the distance from $g^+(0, 1)$ to $g^+(1, 1)$ by the parallel transport around $g(\partial(K \times [0, 1]))$ plus a term related to $g(K \times \{0\})$ and a term related to $g(K \times \{1\})$. At time 0, we have to deal with the fact that the given lift $g^+ : K \times \{0\} \to S^3$ may not be parallel. The given map has Lipschitz constant L_K, and the parallel lift of g has Lipschitz constant $\lesssim L_K$, and so the error introduced here is $\lesssim L_K$. Similarly, at time 1, we have to subtract off the parallel lift of g over $K \times \{1\}$. This parallel lift has Lipschitz constant $\lesssim L_K$, so it introduces another error of size L_K. All together the distance from $g^+(0, 1)$ to $g^+(1, 1)$ is $\lesssim L_K L_T + L_K$. Since we assumed $L_T \gtrsim 1$, this bound is acceptable. This finishes the sketch of the proof of Lemma 1.11.

We can now give the proof of Proposition 1.10.

Proof. The plan is to lift the map $f : S^3 \to S^2$ to a map $f^+ : S^3 \to S^3$, and then contract f^+ using the Brouwer type argument from the last section.

We think of S^3 as $S^2 \times [0,1]$ with two 3-balls attached. Without loss of generality, we reduce to the case that $f : S^3 \to S^2$ maps each of these 3-balls to the basepoint of S^2. So $f : S^2 \times [0,1] \to S^2$ sends the boundary to the base point. We pick a base point of S^3 lying above the basepoint of S^2, and we define $f^+(x,0)$ to be this basepoint. Now we use homotopy lifting to lift f to a map $f^+ : S^2 \times [0,1] \to S^3$. By Lemma 1.11, the Lipschitz constant of f^+ is $\lesssim L^2$. Because f^+ is a lift, f^+ maps $S^2 \times \{1\}$ to the fiber above the base point, which is a great circle in S^3.

We would like to lift $f : S^3 \to S^2$ to $f^+ : S^3 \to S^3$. So far, we have defined f^+ on $S^2 \times [0,1]$. Since f^+ maps $S^2 \times \{0\}$ to the base point, we can easily extend f^+ to the 3-ball bounding $S^2 \times \{0\}$. Next we have to extend f^+ to the 3-ball bounding $S^2 \times \{1\}$. Since f^+ maps $S^2 \times \{1\}$ to a circle, this is the Lipshitz homotopy problem for maps from S^2 to S^1. By Proposition 1.9, we see that f^+ extends over the 3-ball bounding $S^2 \times \{1\}$ with the same Lipschitz constant.

To summarize, we have lifted $f : S^3 \times S^2$ to $f^+ : S^3 \to S^3$ with Lipschitz constant $\lesssim L^2$. By the long exact sequence of the Hopf fibration, f^+ is contractible. Proposition 1.6 tells us that there is a contraction $H^+ : B^4 \to S^3$ with Lipschitz constant $\lesssim L^2$. Finally, we define $H = h \circ H^+$. Since the Hopf fibration h has Lipschitz constant $\lesssim 1$, we see that H has Lipschitz constant $\lesssim L^2$. This finishes our sketch of Proposition 1.10. □

It is not known whether the bound $\operatorname{Lip} H \lesssim L^2$ given here can be improved. The Hopf fibration is a natural way to understand the homotopy theory of maps $S^3 \to S^2$, and the argument here seems a natural way to make it quantitative. But is there any principle that taking natural constructions in homotopy theory and rendering them quantitative gives the right quantitative bounds? Little is known about this. One striking example is the contrast between Gromov's tower-of-exponentials bound for the quantitative null-homotopy problem and the linear or quadratic bounds in Theorems 0.2 and 0.3. The tower of exponentials arose from trying to generalize the argument here to other fibrations, as we explain in the next subsection.

1.4. **On other fibrations.** Having seen the argument for the Hopf invariant in the last subsection, it's natural to apply similar ideas to other fibrations. Hopf fibrations exist only in a few special dimensions, but there are many other fibrations in the topology literature, such as the path space fibration or Postnikov towers, which are useful for studying homotopy groups in many situations. However, from the quantitative point of view, more general fibrations can have much worse quantitative estimates for the homotopy lifting property. This occurs already for fiber bundles where the holonomy is given by a diffeomorphism. For instance, consider the fiber bundle over the circle with fiber T^2 and with gluing map given by an Anosov diffeormorphism. If we think of T^2 as $\mathbb{R}^2/\mathbb{Z}^2$, then the gluing map is a linear map coming from a matrix M in $SL_2(\mathbb{Z})$, and the key feature of an Anosov diffeomorphism is that the entries of M^L grow exponentially in L. This occurs for most matrices $M \in SL_2(\mathbb{Z})$, for instance the matrix

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

Let us denote this fiber bundle as $E \to S^1$. Now consider the following homotopy lifting problem. We let $g : S^1 \times [0,1] \to S^1$ be given by $g(\theta,t) = 2\pi L t$, where the right-hand side is interpreted modulo 2π to give an element of S^1. We let $g^+ : S^1 \times \{0\} \to E$ be a map that sends S^1 to a homologically non-trivial curve in the fiber T^2 in homology class $a \in H_1(T^2) = \mathbb{Z}^2$. Notice that the Lipschitz constant of g is $2\pi L$ and the Lipshitz constant of $g^+ : S^1 \times \{0\} \to E$ is $\lesssim 1$. If we lift the homotopy g to extend g^+ to a map $g^+ : S^1 \times [0,1] \to E$, then g^+ maps $S^1 \times \{1\}$ to a 1-cycle in the fiber T^2 in the homology class $M^L a$. For most choices of a, the vector $M^L a \in \mathbb{Z}^2$ has size $\geq \exp(cL)$, and so the length of the image curve must be at least $\exp(cL)$. This shows that the Lipschitz constant of the homotopy lift g^+ is at least $\exp(cL)$.

The Hopf fibration was much better behaved because the holonomy of the connection around any loop is an isometry of the fiber, which avoids this issue. But once the holonomy of the connection around some loop is not an isometry, then there is a real danger that the holonomy given by going around the loop L times will have Lipschitz constant growing exponentially in L. In this situation, when we apply homotopy lifting to a map with Lipschitz constant L, we get a map with Lipschitz constant exponential in L.

Many arguments in homotopy theory involve a tower of fibrations. Suppose that $Y_j \to Y_{j-1} \to ... \to Y_1 \to Y_0 = Y$ is a tower of fibrations, and we begin with a map $f : S^m \to Y$ with Lipschitz constant L. If we use the homotopy lifting approach to lift f to a map $f_1 : S^m \to Y_1$, then unless we are lucky or clever, the Lipschitz constant of f_1 will be on the order of $\exp L$. Now if we lift f_1 to a map $f_2 : S^m \to Y_2$, then unless we are lucky or clever, the Lipschitz constant of f_2 will be on the order of $\exp(\exp L)$. This is the source of the tower of exponentials that Gromov discusses in [G].

Remark. The sketch in [G] does not include many details. So far I have not been able to fill in the details to give a proof of the bound with a tower of exponentials. Because of the much stronger bounds proven in [CDMW] and [CMW], the tower of exponentials bound is now obsolete. Nevertheless, I think it would be interesting to write down a full proof. The argument raises some natural questions, which I think are still interesting even after [CDMW] and [CMW]. For instance, for various interesting fibrations, what is the best quantitative version of the homotopy lifting property which can be proven?

1.5. **The state of affairs before the recent breakthrough.** We have now summarized the work on this Lipschitz homotopy problem that was done before the recent breakthrough by Chambers, Dotterrer, Ferry, Manin, and Weinberger. I was aware of this work and thought the problem was natural and interesting, but it seemed very difficult to me.

In the methods described so far, we begin with a proof from topology describing $\pi_m(S^n)$ for some m, n, which gives a criterion for a map $S^m \to S^n$ to be null-homotopic. Implicitly, many of these proofs actually construct a null-homotopy of any map which meets the criterion. To get a quantitatve null-homotopy, we start by making the construction from the topology literature explicit. If the construction involves some choices, we try to make those choices in a quantitatively efficient way. Finally, we measure the Lipschitz constant of the resulting map. There are many arguments in the literature to compute homotopy groups of different spaces, and for most of these arguments it is unknown what kind of quantitative bounds they can deliver. It would be interesting to carry out this approach for some torsion homotopy groups, like $\pi_4(S^3)$ or $\pi_6(S^3)$. Carrying out such an argument for maps $S^6 \to S^3$ sounded interesting and difficult to me. If m and n are bigger, the situation seemed even more complicated. For, say, $m = 1000$, $n = 100$, the group $\pi_{1000}(S^{100})$ has not been computed. In this scenario, it sounded even more difficult to me to understand the Lipschitz homotopy problem.

The recent papers [CDMW] and [CMW] deal with all m and n, including ones for which $\pi_m(S^n)$ has not been computed. They avoid this issue by taking a different approach to the problem which we describe in the next section.

2. The method of Chambers-Dotterrer-Manin-Weinberger

The proof of Theorem 0.2 uses two main tools, one from topology and one from geometry. The first tool is the classification of the infinite homotopy groups of spheres.

Theorem 2.1. *(Serre) The homotopy groups $\pi_m(S^n)$ are finite except for two cases:*

- *The case $m = n$, when $\pi_n(S^n) = \mathbb{Z}$ is given by the degree.*
- *The case when n is even and $m = 2n - 1$, when $\pi_{2n-1}(S^n) = \mathbb{Z} \oplus F$ where F is a finite group. The non-torsion part of $\pi_{2n-1}(S^n)$ is given by the Hopf invariant.*

The infinite homotopy groups of spheres are rare, and the non-torsion parts of the homotopy groups are much simpler than the torsion parts. The non-torsion part of the homotopy groups only involves the degree and the Hopf invariant, which we have already grappled with above. In the proof below, we will see that when we look at the problem from the right perspective, the effect of the torsion part of the homotopy groups can be controlled essentially by compactness.

The difference between even and odd n comes from this theorem. If n is odd, then $\pi_m(S^n)$ is finite except for $m = n$. If n is even, then there are two infinite homotopy groups of S^n. We will work in the easier case when n is odd.

The second tool comes from geometry and it is related to the isoperimetric inequality. The Lipschitz homotopy problem can be thought of as an inventive variation on the isoperimetric inequality. We are given a map f defined on the boundary of S^m with some quantitative bound, and we want to extend f to the ball B^{m+1} with a related quantitative bound. A key observation in the proof is that this variation of the isoperimetric inequality is closely related to classical isoperimetric inequalities. For instance, [CDMW] brings into play the following isoperimetric inequality due to Federer and Fleming.

Theorem 2.2. *(Federer-Fleming) Suppose that z is a k-cycle in S^m with $0 \le k < m$. Then z is the boundary of a $(k+1)$-chain y obeying the volume bound*

$$\mathrm{Vol}_{k+1}(y) \lesssim \mathrm{Vol}_k(z).$$

(This linear bound may look unfamiliar to some readers. Federer and Fleming also proved the bound $\mathrm{Vol}_{k+1}(y) \lesssim \mathrm{Vol}_k(z)^{\frac{k+1}{k}}$. If $\mathrm{Vol}_k(z)$ is smaller than 1, then the latter bound is better, but if $\mathrm{Vol}_k(z)$ is bigger than 1, then the linear bound is better. In the application here, the range where $\mathrm{Vol}_k(z)$ is bigger than 1 is the important range.)

This inequality is closely related to an estimate for primitives of differential forms which is what we will use in the proof below.

Proposition 2.3. *Suppose that β is a closed $(k+1)$-form with support in $(-1/2, 1/2)^m$. If $k+1 = m$, then assume also that $\int \beta = 0$. Then there is a k-form α with support in $(-1,1)^m$ solving the equation $d\alpha = \beta$ and obeying the bound*

$$\|\alpha\|_{L^\infty} \lesssim \|\beta\|_{L^\infty}.$$

We will work out the proof of this Proposition below, after we see how it is used in the proof of the Theorem.

2.1. The framework of induction on skeleta.

The proof is carried out using the framework of induction on skeleta which comes from obstruction theory. We begin with a map $f : S^m \to S^n$ with Lipschitz constant L. Using simplicial approximation (Lemma 1.5), we can assume that f is a simplicial map from (S^m, Tri_L) to $(S^n, \mathrm{Tri}_{S^n})$. (Recall that Tri_L is a triangulation of S^m into standard simplices of side length $\sim \frac{1}{L}$and Tri_{S^n} is the triangulation of S^n as the boundary of the $(n+1)$-simplex.)

We also know that f is null-homotopic. So we can extend f to a map $h : B^{m+1} \to S^n$, although we have no geometric information about this map h. We extend the triangulation Tri_L to a triangulation of B^{m+1}. It has $\sim L^{m+1}$ standard simplices of diameter $\sim 1/L$.

The general idea of the argument is to improve the map h until it obeys a Lipschitz bound. Roughly speaking, we want to "straighten out" the null-homotopy h. We will do this one skeleton at a time. First we homotope h (relative to the boundary) to a map H_1 so that the Lipschitz constant of H_1 on the 1-skeleton of B^{m+1} is $\lesssim L$. Next we homotope H_1 (relative to the boundary) to a map H_2 so that the Lipschitz constant of H_2 on the 2-skeleton of B^{m+1} is $\lesssim L$. We try to proceed in this way until we reach H_{m+1}. The resulting map H_{m+1} will be our desired homotopy: it will be equal to f on the boundary and have Lipschitz constant $\lesssim L$ on all of B^{m+1}.

Let's make a couple observations about this strategy. First of all, the reader may raise an objection. Suppose that $m = n-1$. The map h is a map from B^n to S^n. Suppose that the map $f : S^{n-1} \to S^n$ is a constant map. Of course in this case, we can make a constant null-homotopy H, so the theorem we want to prove is trivial. But suppose that $h : B^n \to S^n$ is a map of very high degree. (The degree is well-defined because h maps the boundary of B^n to a point.) No matter how we homotope h rel the boundary, the degree will still be very high, and so the Lipschitz constant of any such homotopy will have to be high also. This is true. On the other hand, if $m \le n-1$, then the Lipschitz null-homotopy theorem was proven above and it is fairly easy. From now on we will assume that we are in the interesting case $m \ge n$. In this case, it turns out that no such issue arises.

The next observation is that it is easy to construct H_1, ..., H_{n-1}. In fact, we can construct these maps so that H_j is simplicial on the j-skeleton of $(B^{m+1}, \mathrm{Tri}_L)$. First of all, we can homotope h to a map H_0 which sends vertices of $(B^{m+1}, \mathrm{Tri}_L)$ to vertices of $(S^n, \mathrm{Tri}_{S^n})$. Now we construct H_1. Pick an edge e of Tri_L. The map H_0 sends the endpoints of e to vertices of S^n. Any two (different) vertices of $(S^n, \mathrm{Tri}_{S^n})$ bound an edge. Therefore, there is a unique simplicial map $e \mapsto (S^n, \mathrm{Tri}_{S^n})$ which extends $H_0|_{\partial e}$. As long as $n > 1$, any two maps from e to S^n are relatively homotopic. So we can homotope $H_0|_e$ to a simplicial map. We do this for every edge e in Tri_L. Then by the homotopy extension theorem, we can homotope H_0 to our desired map H_1. The same procedure works to define H_2, ..., H_{n-1}. Our map H_{n-1} is now simplicial on the $(n-1)$-skeleton of the domain $(B^{m+1}, \mathrm{Tri}_L)$.

We run into a problem at the n-skeleton. Here is the issue. Let Δ be an n-simplex in Tri_L. We know that $H_{n-1}|_{\partial\Delta}$ is simplicial. This simplicial map extends to a unique simplicial map $S_\Delta : \Delta \to (S^n, \mathrm{Tri}_{S^n})$. However, $H_{n-1}|_\Delta$ and $S_\Delta|_\Delta$ may not be homotopic to each other (relative to the boundary). The possible homotopy classes of maps $\Delta \mapsto S^n$ with fixed boundary data are prescribed by a degree according to the Brouwer degree theory. We can define this degree as follows. Let $dvol_{S^n}$ be an n-form on S^n with integral 1. Then for a map $g : \Delta \to S^n$ which agrees with H_{n-1} on the boundary, we define

$$\text{Relative degree } (g) := \int_\Delta g^* dvol_{S^n}.$$

The relative degree is not an integer, but the difference between the relative degrees of any two different maps is an integer. Also, any two maps g_1 and g_2 are homotopic rel the boundary if and only if they have the same relative degree.

Let us abbreviate the relative degree of g on Δ by $\omega_g(\Delta)$. If $\omega_{H_{n-1}}(\Delta) = \omega_{S_\Delta}(\Delta)$ for every Δ, then we can homotope H_{n-1} rel the $(n-1)$-skeleton to a map H_n which is simplicial on the n-skeleton. However, there is no reason to believe that the relative degrees obey this property. We have no information at all about the relative degrees of H_{n-1}. They could be gigantic.

The size of the relative degree is the real issue. Suppose we knew that the relative degrees had absolute value at most 1000. In this special case, let us show that we could homotope H_{n-1} rel the $(n-1)$-skeleton to a map H_n with Lipschitz constant $\lesssim L$ on the n-skeleton. It's probably not hard to see this by hand, but here is a rather general finiteness argument to show it. The finiteness argument will be useful for us later on as well. Recall that H_{n-1} is simplicial on the $(n-1)$-skeleton. So there are only finitely many possible choices for $H_{n-1}|_{\partial\Delta}$. Since the relative degree has absolute value at most 1000, there are only finitely many choices for the relative degree. So there are only finitely many relative homotopy classes for $H_{n-1}|_{\Delta}$. For each of these relative homotopy classes a, we pick a map g_a. We homotope H_{n-1} to H_n so that the restriction of H_n to each n-simplex is one of the maps g_a. To finish, we just have to check that each map g_a has Lipschitz constant $\lesssim L$. This follows just because the number of maps g_a is finite! To see it, let us for the moment equip Δ with the unit simplex metric, g_{unit}, and let A be the largest Lipschitz constant of any of the maps $g_a : (\Delta, g_{unit}) \to (S^n, \mathrm{Tri}_{S^n})$. Now since Δ is really a standard simplex of side length $\sim 1/L$, the Lipschitz constant of H_n on each n-simplex is $\lesssim AL$. Since A is a constant depending only on n, $AL \lesssim L$ as desired.

The first hard step of the proof is to homotope H_{n-1} to a map H_n which agrees with H_{n-1} on the $(n-1)$-skeleton and so that the relative degrees of H_n are all $\lesssim 1$. By the argument above, we can then arrange that H_n has Lipschitz constant $\lesssim L$ on the n-skeleton of $(B^{m+1}, \mathrm{Tri}_L)$.

We first show that $\omega_{H_{n-1}}$ is a cocycle. By definition, $\omega_{H_{n-1}}$ assigns a real number to each n-simplex of Tri_L, so $\omega_{H_{n-1}}$ is a simplicial n-cochain: $\omega_{H_{n-1}} \in C^n(\mathrm{Tri}_L, \mathbb{R})$. But if Δ^{n+1} is an $(n+1)$-simplex of Tri_L, then $\delta\omega_{H_{n-1}}(\Delta^{n+1}) = \int_{\partial\Delta^{n+1}} H^*_{n-1} dvol_{S^n}$, which vanishes by Stokes theorem. (In other words, this integral is the degree of $H_{n-1} : \partial\Delta^{n+1} \to S^n$, which vanishes because H_{n-1} extends to Δ^{n+1}.)

Now we start to discuss how to modify H_{n-1}. Our modified map H_n will agree with H_{n-1} on the $(n-1)$-skeleton. Moreover, the whole homotopy from H_{n-1} to H_n will be constant on the $(n-2)$-skeleton. Let $\bar{H} : B^{m+1} \times [n-1, n] \to S^n$ denote the homotopy from H_{n-1} to H_n that we want to build. If we restrict $\bar{H}$ to $\Delta^{n-1} \times [n-1, n]$, then it has a relative degree, given by $\int_{\Delta^{n-1}\times[n-1,n]} \bar{H}^* dvol_{S^n}$. This relative degree is always an integer, because one option for $\bar{H}$ is the trivial homotopy. Since the homotopy $\bar{H}$ is constant along the boundary of B^{m+1}, the relative degree vanishes for each simplex Δ^{n-1} in the boundary of B^{m+1}. Any choice of integer degrees which vanishes on the boundary is possible. These relative degrees can be encapsulated as an integral $(n-1)$-cochain $\alpha \in C^{n-1}(\mathrm{Tri}_L, \mathbb{Z})$, vanishing on the boundary of B^{m+1}. Now α determines the relative degrees of the map H_n according to the formula

$$\omega_{H_n} = \omega_{H_{n-1}} - \delta\alpha.$$

This key formula again follows just by Stokes theorem. We apply Stokes theorem to $\bar{H}^* dvol_{S^n}$ on $\Delta^n \times [n-1, n]$. Since $\bar{H}^* dvol_{S^n}$ is closed, we get

$$0 = \int_{\partial(\Delta^n \times [n-1,n])} \bar{H}^* dvol_{S^n}.$$

Now $\partial(\Delta^n \times [n-1,n])$ has three parts: the top $\Delta^n \times \{n\}$, the bottom $\Delta^n \times \{n-1\}$, and the sides, $\partial\Delta^n \times [n-1,n]$. The top contributes $\omega_{H_n}(\Delta^n)$, the bottom contributes $-\omega_{H_{n-1}}(\Delta^n)$, and the sides contribute $\sum_{\Delta^{n-1} \subset \partial\Delta^n} \alpha(\Delta^{n-1}) = \delta\alpha(\Delta^n)$.

The construction of H_n now reduces to an estimate about cochains, which we state as a lemma. Recall that for a simplicial k-cochain β, $\|\beta\|_{L^\infty}$ is defined to be the supremum of $|\beta(\Delta^k)|$ over all k-simplices Δ^k in the simplicial complex.

Lemma 2.4. *Suppose ω is a real cocycle in $C^n(\mathrm{Tri}_L, \mathbb{R})$, and $\|\omega\|_{L^\infty(\partial B^{m+1})} \le 1$, then there is an integral $(n-1)$-cochain α which vanishes on ∂B^{m+1} so that*

$$\|\omega - \delta\alpha\|_{L^\infty} \lesssim 1.$$

The restriction that α is integral is not that important here. Suppose that α' is a real $(n-1)$-cochain that vanishes on ∂B^{m+1}. We can always round it off to get an integral $(n-1)$-cochain that vanishes on ∂B^{m+1} so that $\|\alpha - \alpha'\|_{L^\infty} \le 1$. Since each n-simplex of Tri_L borders $\lesssim 1$ $(n-1)$-simplices of Tri_L, we also get $\|\delta\alpha - \delta\alpha'\|_{L^\infty} \lesssim 1$. So it suffices to prove the lemma with a real cochain α'. As α' varies over all the real cochains that vanish on the boundary, $\omega - \delta\alpha'$ varies over all the real cocycles in $C^n(\mathrm{Tri}_L, \mathbb{R})$ that agree with ω on the boundary. Therefore, Lemma 2.4 reduces to the following extension lemma for cocycles:

Lemma 2.5. *Suppose ω_∂ is a real n-cocycle on (S^m, Tri_L) which can be extended to an n-cocycle on B^{m+1}. Then ω_∂ extends to a real n-cocycle ω_{ext} on $(B^{m+1}, \mathrm{Tri}_L)$ obeying the bound*

$$\|\omega_{ext}\|_{L^\infty} \lesssim \|\omega_\partial\|_{L^\infty}.$$

We will prove this lemma in the next subsection. This lemma is the more geometric part of the argument. The proof involves ideas in the spirit of the isoperimetric inequality.

We have now sketched the construction of H_n, except for the proof of Lemma 2.5. Before going on to higher skeleta, let us quickly review the properties of the map H_n that we have constructed.

(1) H_n is homotopic rel the boundary to h. Therefore, $H_n|_{\partial B^{m+1}} = f$.
(2) The restriction of H_n to any n-simplex of Tri_L comes from a finite list of maps $g_a : \Delta^n \to S^n$, independent of L. In particular, the restriction of H_n to the n-skeleton has Lipschitz constant $\lesssim L$.

We will homotope H_n to maps H_{n+1}, H_{n+2}, ..., H_{m+1} so that H_j enjoys these properties on the j-skeleton of Tri_L. In other words, we will inductively construct H_j so that

(1) H_j is homotopic rel the boundary to h. Therefore, $H_j|_{\partial B^{m+1}} = f$.
(2) The restriction of H_j to any j-simplex of Tri_L comes from a finite list of maps $g_a : \Delta^j \to S^n$, independent of L. In particular, the restriction of H_j to the j-skeleton has Lipschitz constant $\lesssim L$.

This setup comes from the paper [FW] of Ferry and Weinberger, who gave the following pithy argument. At this point we use Serre's theorem that $\pi_j(S^n)$ is finite for all $j > n$. (Recall that we are proving Theorem 0.2 in which n is odd.)

We know that H_n restricted to an n-simplex comes from a finite list of possible maps. If Δ^{n+1} is an $(n+1)$-simplex of Tri_L, then $H_n|_{\partial\Delta^{n+1}}$ comes from a finite list of possible maps. Now $\pi_{n+1}(S^n)$ is also finite. So for each map $g_a : \partial\Delta^{n+1} \to S^n$, there are only finitely many relative homotopy classes of maps $\Delta^{n+1} \to S^n$ with the given boundary data. For each g_a, and each homotopy class γ_b, choose a representative map $g_{a,b}$. Now we can homotope H_n relative to the boundary and the

n-skeleton to a map H_{n+1} so that the restriction of H_{n+1} to any $(n+1)$ simplex is one of the maps $g_{a,b}$. This defines H_{n+1} with the two desired properties.

By exactly the same argument, we can homotope H_j to H_{j+1}. Finally, we get a map H_{m+1} obeying the desired properties. In particular, H_{m+1} is a null-homotopy of f with Lipschitz constant $\lesssim L$. This finsihes our sketch of the proof of Theorem 0.2.

2.2. **Estimates for extensions and primitives.** In this section, we prove estimates for primitives and extensions, building up to the extension lemma for cocycles, Lemma 2.5. We will begin by proving estimates about differential forms, and then transfer them to estimates about cocycles. It would actually be possible to work with cocycles throughout, but I think the arguments we give will look more familiar in the context of differential forms. We will prove quantitative versions of some familiar qualitative facts about differential forms.

We start with the following classical qualitative theorem.

Theorem 2.6. *Suppose that β is a closed n-form with compact support in $(-1,1)^m$ If $n=m$ assume also that $\int \beta = 0$. Then there is an $(n-1)$-form α with compact support in $(-1,1)^m$ which obeys $d\alpha = \beta$.*

We are interested in adding to this theorem a quantitative estimate for α. In order to do this, we have to adjust a little our discussion of the support.

Proposition 2.7. *Suppose that $0 < s_1 < s_2$. Suppose that β is a closed n-form with compact support in $(-s_1,s_1)^m$. If $n=m$ assume also that $\int \beta = 0$. Then there is an $(n-1)$-form α with compact support in $(-s_2,s_2)^m$ which obeys $d\alpha = \beta$ and*

$$\|\alpha\|_{L^\infty} \le C(m,n,s_1,s_2)\|\beta\|_{L^\infty}.$$

We will approach this problem by studying one of the classical constructions that proves Theorem 2.6. The form constructed actually obeys the bound in Proposition 2.7.

It's worth mentioning that there are several approaches to the problem. Problems about primitives of differential forms (or cochains) are dual to problems about the boundaries of chains. This angle is explained and used in [CDMW], and it reduces Proposition 2.7 to an isoperimetric inequality for relative cycles, which was proven by Federer-Fleming. For context we state the relative isoperimetric inequality.

Theorem 2.8. *Suppose that z is a relative $(n-1)$-cycle in $[0,1]^m$, for some $n \le m$. Then z bounds an n-chain y obeying the volume bound*

$$\mathrm{Vol}_n\, y \lesssim \mathrm{Vol}_{n-1}\, z.$$

We mention this just to point out that Proposition 2.7 belongs to the general family of isoperimetric inequalities. But we give a different proof, following Lemma 7.13 in Gromov's book [G2].

Proof. In this proof, we write $A \lesssim B$ for $A \le C(m,n,s_1,s_2)B$.

We prove Proposition 2.7 by induction on the dimension. To frame the induction, it's convenient to use coordinates $x_1,...,x_{m-1},t$. Now we can write our closed n-form β in the form

$$\beta = dt \wedge \beta_1(t) + \beta_2(t),$$

where $\beta_1(t) \in \Omega^{n-1}(\mathbb{R}^{m-1})$ and $\beta_2(t) \in \Omega^n(\mathbb{R}^{m-1})$. The condition that β is closed reads

$$0 = d_{\mathbb{R}^m}\beta = -dt \wedge d_{\mathbb{R}^{m-1}}\beta_1(t) + dt \wedge \partial_t \beta_2(t) + d\beta_2(t).$$

(Here we use $d_{\mathbb{R}^m}$ denote the exterior derivative on $\mathbb{R}^m$ and $d_{\mathbb{R}^{m-1}}$ to denote the exterior derivative in the variables $x_1, ..., x_{m-1}$ only.) So saying that β is closed is equivalent to two conditions: for each t,

(1) $\beta_2(t)$ is a closed form on $\mathbb{R}^{m-1}$.
(2) $\partial_t \beta_2(t) = d_{\mathbb{R}^{m-1}}\beta_1(t)$.

Our first attempt to build a good primitive is the form α_1 defined by

$$\alpha_1(t) := \int_{-\infty}^{t} \beta_1(s) ds.$$

The form α_1 is indeed a primitive of β:

$$d_{\mathbb{R}^m}\alpha_1 = d_{\mathbb{R}^{m-1}}\alpha_1 + dt \wedge \partial_t \alpha_1 = \int_{-\infty}^{t} d_{\mathbb{R}^{m-1}}\beta_1(s)ds + dt \wedge \beta_1(t) =$$

$$= \int_{-\infty}^{t} \partial_s \beta_2(s) ds + dt \wedge \beta_1(t) = \beta_2(t) + dt \wedge \beta_1(t) = \beta.$$

Also, $\|\alpha_1\|_\infty$ obeys good bounds. If β is supported in $[-s_1, s_1]^m$, then

$$\|\alpha_1\|_\infty \leq 2s_1 \|\beta\|_\infty.$$

However, the support of α_1 is usually not finite. We know that for each t, $\beta_1(t)$ is supported in $[-s_1, s_1]^{m-1}$, and we also know that $\beta_1(t) = 0$ unless $-s_1 \leq t \leq s_1$. This tells us that α_1 is supported in $[-s_1, s_1]^{m-1} \times [-s_1, \infty]$. There are some special cases where α_1 is actually supported in $[-s_1, s_1]^m$, but in general α_1 is not compactly supported and we need to fix it. The key observation is that the problem of repairing α_1 is closely related to the original problem in lower dimensions, which lets us use induction. Notice that $\alpha_1(t)$ is unchanging for all t in the range $[s_1, \infty]$. For all such t,

$$\alpha_1(t) = \gamma = \int_{-\infty}^{\infty} \beta_1(s) ds.$$

Now γ is a closed $(n-1)$-form on $\mathbb{R}^{m-1}$ which has compact support in $[-s_1, s_1]^{m-1}$. Also, if $n = m$, then $\int \gamma = \int \beta = 0$. So we can apply induction to γ.

The base of our induction is the case when γ is a 0-form. If $n - 1 = 0$, then γ is a constant function with compact support. So if $m > n$, $\gamma = 0$. Also, if $n = m$, and $n - 1 = 0$, then γ is just a function on a point, so the fact that $\int \gamma = 0$ tells us again that $\gamma = 0$. When $\gamma = 0$, α_1 is actually supported on $[-s_1, s_1]^m$, and we are done.

When γ is not zero, we use induction to find an $(n-2)$-form η with support in $[-s_2, s_2]^{m-1}$ obeying $d\eta = \gamma$ and $\|\eta\|_{L^\infty} \lesssim \|\gamma\|_{L^\infty} \lesssim \|\beta\|_{L^\infty}$. Now we let $\psi(t)$ be a smooth function with $\psi(t) = 0$ for $t \leq s_1$ and $\psi(t) = 1$ for $t \geq s_2$, and we define

$$\alpha = \alpha_1 - d\left(\psi(t)\eta\right).$$

The form $\psi(t)\eta$ is supported in $[-s_2, s_2]^{m-1} \times [s_1, \infty)$. Now

$$d\left(\psi(t)\eta\right) = \psi'(t)\eta \wedge dt + \psi(t)\gamma.$$

In the range $t > s_2$, we have $\psi'(t) = 0$ and $\psi(t) = 1$, and so $d(\psi(t)\eta) = \gamma = \alpha_1(t)$. Therefore, our form α is indeed supported in $[-s_2, s_2]^m$. Moreover,

$$d\alpha = d\alpha_1 + dd(...) = d\alpha_1 = \beta.$$

Finally, $\|\alpha\|_\infty$ obeys a good bound:

$$\|\alpha\|_\infty \leq \|\alpha_1\|_\infty + \|\psi'(t)\|_\infty \|\eta\|_\infty + \|\gamma\|_\infty \lesssim \|\beta\|_\infty.$$

This finishes the proof of Proposition 2.7. □

Next, we need an estimate for extending closed forms based on the following classical qualitative theorem.

Theorem 2.9. *Suppose that ω_∂ is a closed n-form on S^m. If $m = n$, assume also that $\int_{S^m} \omega_\partial = 0$. Then ω_∂ extends to a closed n-form ω_{ext} on B^{m+1}.*

Here is a quantitative version of this extension theorem.

Lemma 2.10. *Suppose ω_∂ is a closed n-form on S^m which extends to a closed n-form on B^{m+1}. Then ω_∂ extends to a closed n-form ω_{ext} on B^{m+1} obeying the bound*

$$\|\omega_{ext}\|_{L^\infty} \lesssim \|\omega_\partial\|_{L^\infty}.$$

We prove Lemma 2.10 by going through the standard proof of Theorem 2.9 and giving quantitative bounds using Proposition 2.7.

Proof. We use polar coordinates (r, θ) on B^{m+1}. We choose $0 < r_1 < r_2 = 1$, and we let $\psi(r)$ be a smooth function with $\psi(r) = 0$ for $r < r_1$ and $\psi(r) = 1$ for $r > r_2$. Now we consider the form $\psi(r)\omega_\partial$ in polar coordinates on B^{m+1}. This is a smooth n-form, and its restriction to the boundary S^m is indeed ω_∂. It also has good L^∞ bounds, but it is not closed. Indeed

$$d\left(\psi(r)\omega_\partial\right) = \psi'(r) dr \wedge \omega_\partial.$$

The right-hand side is an $(n+1)$-form supported in the annulus $r_1 \leq r \leq r_2$. If we choose r_1, r_2 correctly, this form is supported in a small cube contained in B^{m+1}. Since it is exact it is clearly closed. Also, if $m = n$, then we are know by hypothesis that $\int_{S^m} \omega_\partial = 0$ and so

$$\int_{B^{m+1}} \psi'(r) dr \wedge \omega_\partial = 0.$$

Now we can use Proposition 2.7 to find a good primitive: an n-form α supported inside B^{m+1} so that

$$d\alpha = d\left(\psi(r)\omega_\partial\right), \text{ and}$$

$$\|\alpha\|_\infty \lesssim \|\psi'(r)\|_\infty \|\omega_\partial\|_\infty.$$

Now we define

$$\omega_{ext} = \psi(r)\omega_\partial - \alpha.$$

Since α has compact support inside of B^{m+1}, ω_{ext} restricts to ω_∂ on the boundary. Our equation for $d\alpha$ guarantees that ω_{ext} is closed. And finally our bound for $\|\alpha\|_\infty$ shows that $\|\omega_{ext}\|_\infty \lesssim \|\omega_\partial\|_\infty$. □

Remark: With minor adjustments, this same proof applies if we replace the ball B^{m+1} with a simplex Δ^{m+1}.

Finally, we prove Lemma 2.5,, the extension lemma for cocycles, using the extension lemma for differential forms we have just proven. For convenience, we restate Lemma 2.5:

Lemma. *Suppose ω_∂ is a real n-cocycle on (S^m, Tri_L) which can be extended to an n-cocycle on B^{m+1}. Then ω_∂ extends to a real n-cocycle ω_{ext} on $(B^{m+1}, \mathrm{Tri}_L)$ obeying the bound*

$$\|\omega_{ext}\|_{L^\infty} \lesssim \|\omega_\partial\|_{L^\infty}.$$

Proof. (proof sketch) Suppose that ω_∂ is a real n-cocycle on (S^m, Tri_L), and if $n = m$, assume in addition that $\omega_\partial(S^m) = 0$. We want to construct a closed n-form $\bar{\omega}_\partial$ so that for any n-face Δ^n, $\int_{\Delta^n} \bar{\omega}_\partial = \omega_\partial(\Delta^n)$, and so that

$$\|\bar{\omega}_\partial\|_{L^\infty} \sim L^n \|\omega_\partial\|_{L^\infty}.$$

We construct $\bar{\omega}_\partial$ by induction on skeleta. First we construct $\bar{\omega}_\partial$ on the n-skeleton, so that its support avoids the $(n-1)$-skeleton. Then using Lemma 2.10, we extend it to the $(n+1)$-skeleton, then to the $(n+2)$-skeleton, etc. (Technically speaking, to do this extension, we need a version of Lemma 2.10 for a simplex instead of the ball.) Once we have $\bar{\omega}_\partial$ in hand, we apply Lemma 2.10 to define a closed n-form $\bar{\omega}_{ext}$ on B^{m+1} obeying

$$\|\bar{\omega}_{ext}\|_{L^\infty} \lesssim \|\bar{\omega}_\partial\|_{L^\infty}.$$

Finally we define

$$\omega_{ext}(\Delta^n) = \int_{\Delta^n} \bar{\omega}_{ext}.$$

By Stokes theorem ω_{ext} is a cocycle. Also, we have the bound

$$|\omega_{ext}(\Delta^n)| \le |\Delta^n| \|\bar{\omega}_{ext}\|_{L^\infty} \lesssim L^{-n} L^n \|\omega_\partial\|_{L^\infty}.$$

□

2.3. **Final discussion.** Simplicial approximation and the framework of induction on skeleta allowed us to break the original problem into many smaller problems. The high-dimensional part – the part involving complicated homotopy groups, which seemed the most intimidating at the outset – is tamed by the Serre finiteness theorem and by this approach.

In hindsight, the hardest part of the proof of Theorem 0.2 is the step where we homotope H_{n-1} in order to control the relative degrees. This step only involves degrees - not more complicated homotopy invariants - but it involves many degrees because the relative degree is defined on each n-simplex of our fine triangulation of B^{m+1}. A key observation is that this subproblem relates to more classical estimates in metric geometry - like bounding the primitive of a differential form or like the Federer-Fleming isoperimetric inequality.

3. Open problems

We mention here some open problems in the general spirit of the theorems we have described here.

The main theorem implies that the space of contractible maps $S^m \to S^n$ with Lipschitz constant at most L is roughly connected – in the sense that any two such maps can be connected by a family of maps with Lipschitz constant $\lesssim L$ (whether n is odd or even!). One can ask similar questions about higher homotopy groups of the space of maps. Let $\mathrm{Maps}_L(X, Y)$ denote the space of maps $X \to Y$ with Lipschitz constant at most L, and $\mathrm{Maps}(X, Y)$ the space of continuous maps $X \to Y$. Suppose that $\gamma : S^1 \to \mathrm{Maps}_L(S^m, S^n)$, and that γ is contractible in $\mathrm{Maps}(S^m, S^n)$. How large do we need to choose L' to guarantee that γ is contractible in $\mathrm{Maps}_{L'}(S^m, S^n)$?

What would happen if we take other shapes besides unit spheres? For instance suppose that E^m is an m-dimensional ellipse and F^n is an n-dimensional ellipse. Recall that the n-dimensional ellipse with principal axes $R_0, ...R_n$ is the set defined by

$$\sum_{j=0}^{m} (x_j/R_j)^2 = 1.$$

The ellipses are a simple but non-compact family of Riemannian metrics on S^n. Now suppose that $f : E \to F$ is contractible with Lipschitz constant L. Can we homotope f to a constant map through maps of Lipschitz constant at most $L' = L'(m, n, L)$, independent of the dimensions of E and F? Can we even take $L' = C(m, n)L$ or $C(m, n)L^2$? What happens for other more complicated metrics?

We are concerned here with estimating the best Lipschitz constant of a map in a given homotopy class and with given boundary conditions. It would be interesting to understand more about the algorithmic aspects of this problem. For simplicity, let us forget for the moment about the boundary conditions. Suppose that Tri_1 is a triangulation of S^m into $\leq N$ simplices and Tri_2 is a triangulation of S^n into $\leq N$ simplices. Assign a metric to each sphere so that each simplex is the unit equilateral simplex. Fix a homotopy class $a \in \pi_m(S^n)$. It would be interesting to estimate the minimum Lipschitz constant of any map $f : (S^m, \mathrm{Tri}_1) \to (S^n, \mathrm{Tri}_2)$ in the homotopy class a. This already looks quite hard for $m = n = 2$, when a is the class of degree 1 maps! It would be interesting to know how accurate an estimate is possible with a polynomial-time algorithm. The paper [CKMVW] discusses the algorithmic difficulty of various problems in homotopy theory, such as determining the homotopy class of a given map $f : S^m \to S^n$.

One key point in the proof from [CDMW] is that we can deal with the torsion parts of the homotopy groups of spheres just by knowing that they are finite. It would be interesting to find questions that force us to engage more with their geometry. For instance, in order to estimate the constants $C(m, n)$ in the main theorems – Theorem 0.2 and Theorem 0.3 – it would be necessary to know something more about the torsion part of the homotopy groups. Some of the following questions may also push us to look harder at the torsion.

The Lipschitz constant is not the only measure of the geometric complexity of a map. Another option is the L^p norm of the derivative,

$$\|df\|_{L^p} = \left(\int_{S^m} |df(x)|^p dvol(x) \right)^{1/p}.$$

If $f : S^m \to S^n$ is a contractible map with $\|df\|_{L^p} \leq 100$, can we extend f to a map $H : B^{m+1} \to S^n$ with $\|dH\|_{L^p} \leq L'(m, n, p)$? If $p = +\infty$, then $\|df\|_{L^\infty}$ is just the Lipschitz constant of f, and some of the techniques we have discussed may apply for other p as well. Note, however, that if $p < n$, then the space of maps with $\|df\|_{L^p} \leq 100$ is not precompact in the C^0 topology. In this situation, it does not seem possible to do simplicial approximation, or at least it is not clear to me how one could do it. Simplicial approximation played a basic role in the work from [CDMW] that we discussed, and problems may become very different without it. If the answer to our last question is yes, we can go further and ask what is the smallest $L' = L'(m, n, p, L)$ so that any contractible map $f : S^m \to S^n$ with $\|df\|_{L^p} = L$ extends to a map $H : B^{m+1} \to S^n$ with $\|dH\|_{L^p} \le L'$? We can also ask about $\|dH\|_{L^q}$ for some $q \neq p$, and some other q may turn out to be more interesting than $q = p$. White [W] studied the connection between $\|df\|_{L^p}$ and the homotopy class of f, which is a related question.

Besides the Lipschitz constant of f and $\|df\|_{L^p}$, a third option is the k-dilation studied in [Gu]. We say that the k-dilation of a map f is at most L if, for every k-dimensional submanifold Σ^k in the domain,

$$\mathrm{Vol}_k(f(\Sigma)) \le L\,\mathrm{Vol}_k(\Sigma).$$

If $f : S^m \to S^n$ is a contractible map with k-dilation at most 1, can we extend f to a map $H : B^{m+1} \to S^n$ with a bound on the k-dilation of H? Or maybe the l-dilation of H for some $l \neq k$? Like with the Sobolev norms, the space of maps with k-dilation at most 1 is not precompact in C^0 for any $k \ge 2$ and so it looks hard to use simplicial approximation.

To conclude, we mention some open problems about isotopies of "thick ropes". Alexander Nabutovsky proved some striking theorems about a notion of thick ropes in [N]. He considers smooth submanifolds $M^k \subset \mathbb{R}^n$ with normal injectivity radius at least 1. Suppose that M is isotopic to the standard $S^k \subset \mathbb{R}^n$ and suppose that M has volume V. Is it possible to build an isotopy M_t, starting at M and ending at the standard sphere, so that M_t has normal injectivity radius at least 1 for all t and volume at most $C(k,n)V$? For $n \ge 6$ and $k = n-1$, [N] proves that the answer is no. In fact, he proves something much stronger. In some examples, there is no such isotopy with volume at most $C(k,n)e^{e^V}$. Indeed for any computable function $F(V)$, for V sufficiently large, there are examples where no such isotopy exists with $\mathrm{Vol}(M_t) \le F(V)$. The proof is based on logic – on the inability to recognize the n-sphere algorithmically.

Similar questions for curves embedded in $\mathbb{R}^3$ have been around for some time, but little is known about them. The paper [CH] by Coward and Hass gives good background about the question and also makes an important contribution. Is it possible to find an unknotted closed curve $\gamma \subset \mathbb{R}^3$ with normal injectivity radius 1 so that every isotopy from γ to the standard circle maintaining normal injectivity radius 1 goes through a curve of length strictly longer than γ? There is numerical evidence that the answer is yes, but no rigorous theorem. The paper [CH] rigorously proves a related theorem for links. They construct a two-component link $\lambda \subset \mathbb{R}^3$ with injectivity radius 1 in which the two components are topologically unlinked, and they prove that any isotopy of λ to a final position that separates the two links while maintaining normal injectivity radius 1 must increase the length of one of the two components. The following question is open. Given an unknotted closed curve $\gamma \subset \mathbb{R}^3$ with normal injectivity radius 1 and length L, can it be isotoped to the standard position through curves of normal injectivity radius at least one and length at most $100L$? At some point, I played around with a long piece of rope trying to find examples, and in all the examples I was able to produce, it was easy to find the desired isotopy.

Here is another variation on the theme of thick ropes in higher dimensions, where the method of [N] does not apply, and the answers may turn out very differently. Let $S^k(R)$ denote the k-sphere of radius R, and consider a 2-bilipschitz embedding I from $S^k(R) \times B^{n-k}(1)$ into $\mathbb{R}^n$. Suppose that I is isotopic to a standard (unknotted) embedding I_0 taking $\{0\}$ times $S^k(R)$ to a standard sphere of radius R in $\mathbb{R}^{k+1} \subset \mathbb{R}^n$. Is it possible to isotope I_0 to I via C-bilipschitz embeddings for some constant $C = C(k,n)$? If not, can we do it via $\beta_{k,n}(R)$-bilipschitz embeddings, and what is the best dependence on R?

One can also try to bound the behavior of the isotopy in time. We can think of the isotopy as a map

$$I : S^k(R) \times B^{n-k}(1) \times [0,T] \to \mathbb{R}^n \times [0,T],$$

where the last component of I is just the time coordinate of the domain. Then we can ask to construct an isotopy I so that I is β-bilipschitz and also T is not too big. Is it possible to find such an isotopy with $T = R$ and $\beta = C(k,n)$?

References

[A] V. Arnold, On A. N. Kolmogorov, in *Kolmogorov in perspective*, 89-108, Hist. Math., 20, Amer. Math. Soc., Providence, RI, 2000

[CKMVW] M. Cadek, M. Krcal, J. Matousek, L. Vokrinek, U. Wagner, Uli, Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. SIAM J. Comput. 43 (2014), no. 5, 1728-1780.

[CDMW] G. Chambers, D. Dotterrer, F. Manin, and S. Weinberger, Quantitative null-cobordism, arXiv:1610.04888

[CMW] G. Chambers, F. Manin, and S. Weinberger, Quantitative null-homotopy and rational homotopy type, arXiv:1611.03513

[CH] A. Coward, and J. Hass, Topological and physical link theory are distinct. Pacific J. Math. 276 (2015), no. 2, 387-400.

[FW] S. Ferry and S. Weinberger, Quantitative algebraic topology and Lipschitz homotopy. Proc. Natl. Acad. Sci. USA 110 (2013), no. 48, 19246-19250.

[G] M. Gromov, Quantitative homotopy theory. Prospects in mathematics (Princeton, NJ, 1996), 4549, Amer. Math. Soc., Providence, RI, 1999.

[G2] M. Gromov, *Metric structures for Riemannian and non-Riemannian spaces.* Based on the 1981 French original [MR0682063]. With appendices by M. Katz, P. Pansu and S. Semmes. Translated from the French by Sean Michael Bates. Progress in Mathematics, 152. Birkhuser Boston, Inc., Boston, MA, 1999.

[Gu] L. Guth, Contraction of areas vs. topology of mappings. Geom. Funct. Anal. 23 (2013), no. 6, 1804-1902.

[H] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2001.

[KB] A. Kolmogorov and Y. Barzdin, On the realization of networks in three-dimensional space, in *Selected Works of Kolmogorov, Volume 3*, ed. A. Shiryaev, Kluwer Academic Publishers, Dordrecht, 1993.

[N] A. Nabutovsky, Non-recursive functions, knots "with thick ropes", and self-clenching "thick" hyperspheres. Comm. Pure Appl. Math. 48 (1995), no. 4, 381-428.

[W] B. White, Homotopy classes in Sobolev spaces and the existence of energy minimizing maps. Acta Math. 160 (1988), no. 1-2, 1-17.