

PLATO’S CAVE AND DIFFERENTIAL FORMS

FEDOR MANIN

ABSTRACT. In the 1970s and again in the 1990s, Gromov gave a number of theorems and conjectures motivated by the notion that the real homotopy theory of compact manifolds and simplicial complexes influences the geometry of maps between them. The main technical result of this paper supports this intuition: we show that maps of differential algebras are closely shadowed, in a technical sense, by maps between the corresponding spaces. As a concrete application, we prove the following conjecture of Gromov: if X and Y are finite complexes with Y simply connected, then there are constants $C(X, Y)$ and $p(X, Y)$ such that any two homotopic L -Lipschitz maps have a $C(L + 1)^p$ -Lipschitz homotopy (and if one of the maps is constant, p can be taken to be 2.) We hope that it will lead more generally to a better understanding of the space of maps from X to Y in this setting.

1. INTRODUCTION

In 1996, Princeton University invited several prominent mathematicians, including Misha Gromov, to give a series of lectures entitled “Prospects in Mathematics”; each speaker would discuss their views on future directions in their field. Gromov’s talk was entitled “Quantitative homotopy theory”¹ and advanced the idea that the central questions of algebraic topology—such as, are two maps homotopic?—should be refined by asking about the sizes of the objects produced.

Indeed, a central weakness in the extremely powerful results of algebraic and geometric topology is their indirectness: they are obtained by reducing geometric problems to those of homotopy theory and homotopy problems to algebra, leaving us very little understanding of the shapes of the solutions, or whether finding the correlates of these solutions back in the geometric world is easy or hard. A beautiful example of this phenomenon is the result of Nabutovsky [Nab] that, although every embedded codimension-one sphere in the n -disk ($n > 4$) is isotopic to the boundary due to Smale’s h -cobordism theorem, the complexity of such an isotopy (measured by the size of an embedded normal bundle) cannot be bounded by any computable function of the complexity of the original embedded sphere.

Nabutovsky’s result follows from the unsolvability of the triviality problem for groups: while the embedded sphere is simply connected, it has no way of knowing that it is, and in particular it cannot know that it is isotopic to the boundary. With objects that are aware of their simple connectivity, homotopy theory tends to be more computationally tractable², and the solutions are correspondingly less complex. This seems to stem from the fact that the algebra describing them is commutative.

Indeed, in the setting of this paper, geometric complexity is controlled quite closely by the algebraic structure of maps. One of the most important tools for studying simply connected spaces is rational homotopy theory, first developed by Quillen and Sullivan in the 1970s. Almost immediately, Gromov realized that Sullivan’s machine has geometric consequences, providing lower bounds on the complexity of maps in a given homotopy class, and conjectured that these lower bounds are

¹Actually, due to a mistake somewhere along the way, the title of the talk was given as “Qualitative homotopy theory”—exactly the opposite of what Gromov meant! The notes were published in the proceedings of the conference as [GrQHT].

²But this does not mean tractable in an absolute sense; see [ČKMVW] and [FiVo] for two contrasting perspectives.

sharp. In the intervening years these ideas have been developed further by Gromov, and more recently by Weinberger, Ferry, Chambers, Dotterrer, Guth, and the author.

This paper seeks to strengthen the link between rational homotopy and geometry by showing a kind of inverse result to the reduction to algebra: that, at least in the world of compact spaces, the algebraic maps produced by the theory always have reasonably close geometric doppelgangers. This is robust enough to provide an almost immediate proof of one of the conjectures in Gromov's talk. We also provide applications which nibble at the margins of some other problems; it is the author's hope that with more effort and new techniques, other broader applications will be found.

Our results do not directly resolve any metric problems beyond homotopy theory; in settings such as cobordism or embedding, one also needs to analyze the reduction to homotopy theory. However, again the hope is that these results will simplify that task to some degree.

1.1. Gromov's conjectures. Gromov's 1990s questions have roots stretching back to the 1978 paper [GrHED], where he showed the following result:

Theorem. *Let Y be a simply connected finite complex with a (reasonable) metric. Then $\pi_n(Y)$ has polynomial growth, i.e. the number of elements which have a representative with Lipschitz constant at most L is polynomial in L .*

There are several reasons why the Lipschitz constant is a natural notion of complexity here. First, if two such spaces are homotopy equivalent, then they are Lipschitz homotopy equivalent; this means that Lipschitz invariants such as the asymptotics of the growth function of $\pi_n(Y)$ are as natural as homotopy invariants in this context. Conversely, every L -Lipschitz map is homotopic to a simplicial map on a subdivision with $\sim L^n$ simplices; see Prop. 2.1. This means that the amount of (homotopical) information in an L -Lipschitz map is about L^n . In this framing, Gromov's result says that for simply connected targets, this is a significant overestimate.³ On a more geometric level, the Lipschitz constant bounds the sizes of pullbacks of forms; this is its main property used in the proof.

Gromov's theorem vastly generalizes an observation about Hopf invariants. Suppose that $f : S^3 \rightarrow S^2$ is a smooth L -Lipschitz map between round unit spheres. Denote the volume form on S^2 by $d \text{vol}$; $f^*d \text{vol}$ is a closed 2-form in S^3 and therefore $f^*d \text{vol} = d\alpha$ for some 1-form α . Then, following J.H.C. Whitehead, the Hopf invariant of f is given by

$$H(f) = \int_{S^3} \alpha \wedge f^*d \text{vol}.$$

Now we look at the L^∞ -norms of these forms (i.e. the supremum of their values taken over all frames of unit vectors.) We know $\|f^*d \text{vol}\|_\infty \leq L^2$ and we can choose α so that $\|\alpha\|_\infty \leq C\|f^*d \text{vol}\|_\infty$.⁴ Therefore we have the inequality

$$H(f) \leq C \text{vol}(S^3)L^4.$$

Indeed, up to a constant, this is sharp: a map with Lipschitz constant $O(L)$ and Hopf invariant L^4 can be built via the composition

$$S^3 \xrightarrow{\text{Hopf map}} S^2 \xrightarrow{\text{degree } L^2} S^2.$$

To Gromov, this and other examples suggested the following conjecture:

Conjecture A ([GrQHT]). *The estimate on the growth of $\pi_n(Y)$ provided by the algorithm of [GrHED] is sharp.*

³On the other hand, it is sharp for examples like $\pi_n(T^n \vee S^n)$.

⁴The latter inequality is not entirely obvious and Gromov did not prove it in this paper, although he did give a sketch in [GrMS]; we prove it in Section 2.

As this algorithm is quite complex and not described in detail in the paper, the conjecture is necessarily impressionistic. Nevertheless, the intuition behind it is worth understanding; we attempt to demonstrate it with examples in §3.3.

Gromov returned to this theme in the 1990s in [GrMS, Ch. 7] and the conference paper [GrQHT]. In these works he presented two other conjectures which are relevant to the present work. The first concerns a cousin of the growth of homotopy groups. Given an element $\alpha \in \pi_n(Y)$, define the *distortion function*

$$\delta_\alpha(k) = \inf\{\text{Lip } f \mid f : S^n \rightarrow Y, [f] = k\alpha\}.$$

Note that in all cases $\delta_\alpha(k) = O(k^{1/n})$; this is because one can always find a representative of $k\alpha$ by precomposing a representative of α with a degree k map $f_k : S^n \rightarrow S^n$ with $\text{Lip } f_k = O(k^{1/n})$. On the other hand, for some α one can do better; we say such α are *distorted*, whereas those for which $\delta_\alpha(k) = \Theta(k^{1/n})$ are *undistorted*. In this language, what Gromov showed in [GrHED] is that the generator of $\pi_{2n+1}(S^n)$ is distorted, and also that δ_α is always bounded below by a polynomial.

Conjecture B ([GrQHT]). *An element $\alpha \in \pi_n(Y)$ is undistorted if and only if it has nonzero image under the Hurewicz map to $H_n(Y; \mathbb{Q})$. If it is distorted, then $\delta_\alpha(k) = O(k^{1/n+1})$.*

There is a “strong” but again impressionistic version of the conjecture which states that the bound implied by [GrHED] is sharp, and which is equivalent to Conjecture A.

The last conjecture⁵ can be seen as a relative analogue.

Conjecture C ([GrQHT]). *Let $f \simeq g : X \rightarrow Y$ be L -Lipschitz maps from a finite complex to a finite simply connected complex. Then there is a polynomially bounded function $P_{X,Y}(L) = O(L^{p(X,Y)})$ such that there is a homotopy between f and g through $P_{X,Y}(L)$ -Lipschitz maps.*

Gromov remarked that he could not think of an example where this polynomial had to be nonlinear.⁶

In the past 20 years, there has been some incremental progress on these conjectures. An unpublished result of Shmuel Weinberger (which appears in the author’s PhD thesis [M]) shows the following weak version of the distortion conjecture: there are no rationally nontrivial distorted elements in $\pi_*(Y)$ if and only if the Hurewicz map

$$\pi_*(Y) \otimes \mathbb{Q} \rightarrow H_*(Y; \mathbb{Q})$$

is injective. The proof uses the fact that distortion is well-understood for generalized Whitehead products. The homotopy conjecture is proven for target spaces Y whose rational homotopy structure is relatively simple (including spheres, H-spaces, and homogeneous spaces of Lie groups) in the series of papers [FW], [CDMW], and [CMW].

In this paper, we prove results about Lipschitz homotopies which generalize those of [CDMW] and [CMW] and are actually somewhat stronger than Conjecture C. Define the *length* of a homotopy (sometimes also referred to as *width*) to be the maximal Lipschitz constant of its restrictions to $\{x\} \times [0, 1]$, and its *thickness* to be the maximal Lipschitz constant of its restrictions to $X \times \{t\}$. Gromov’s conjecture only asks about thickness; here is a summary of the results of §5.2:

⁵It must be remarked that Gromov also made claims and conjectures about the set of mapping classes $[X, Y]$ for general X . However, these claims are based only on rational homotopy theory, whereas integral classes can be thought of in general as the integer points of some algebraic variety over \mathbb{Q} and are therefore extremely difficult to get a handle on. Thus, while some vague conjecture can be stated (“every integral homotopy class in $[X, Y]$ has a representative of the size predicted by its rational homotopy theory”), this does not translate into sharp bounds on the growth of $[X, Y]$, even conjectural ones. In the companion paper [MW], the author and Weinberger give a polynomial bound on the growth of $[X, Y]$; however, we make no claims about sharp bounds except in specific examples.

Perhaps stronger results can be obtained in instances where there is more structure, for example for $\text{Aut}(Y)$, which Sullivan [Sul] demonstrated is an arithmetic subgroup of an algebraic group of rational automorphisms.

⁶On the other hand, if we allow $\pi_1(Y)$ to be nontrivial and take $X = S^1$, this corresponds to the so-called isodiametric function of $\pi_1(Y)$, which for certain groups grows faster than any computable function.

Theorem. *Let Y be a finite simply connected complex and X a finite complex of dimension n .*

- (i) *There are constants $C(X, Y)$ and $p(X, Y)$ such that any homotopic L -Lipschitz maps $f \simeq g : X \rightarrow Y$ are homotopic via a homotopy of length C and thickness $C(L + 1)^p$.*
- (ii) *Moreover, any nullhomotopic L -Lipschitz map is nullhomotopic via a homotopy of length C and thickness $C(L + 1)^2$.*
- (iii) *If in addition Y has positive weights (an algebraic condition on the rational homotopy structure), then any nullhomotopic L -Lipschitz map is nullhomotopic via a homotopy of linear thickness and length $C(L + 1)^{n-1}$.*

The latter two bounds are sharp: there are spaces for which one parameter cannot be decreased without increasing the other. On the other hand, it's not clear whether linear thickness is achievable for some classes of maps not satisfying (iii).

The growth and distortion conjectures are more resistant for reasons which are explained later in the introduction, but we do prove that both hold for symmetric spaces.

1.2. Minimal models and DGA maps. To state more precisely the technical ideas in this paper, we must delve into Sullivan's model of rational homotopy theory. This is discussed in greater detail in §3 and we also refer the reader to [Sul] and [GrMo] for detailed exposition. More accurately, what we give here is *real* homotopy theory; the results are less impressive than those of rational homotopy theory in some respects that are irrelevant to the ideas in this paper, but this theory has the advantage of working with off-the-shelf differential forms which behave nicely with respect to smooth maps.

For our purposes, the main points of Sullivan's theory are that the algebra of smooth differential forms Ω^*Y on a compact manifold Y with boundary is a fairly good homotopy-theoretic model for the space Y itself; and that it in turn is modeled by a much smaller, easily described algebra closely related to the Postnikov tower of Y .

More precisely, we think of these as differential graded algebras (DGAs), that is, chain complexes (in this case over \mathbb{R}) equipped with a multiplication which satisfies the graded Leibniz rule. If Y is simply connected, then there is a homotopy equivalence (under a well-known notion of homotopy of DGAs which we define in §3) $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*Y$ where \mathcal{M}_Y^* is a DGA of finite type (i.e. generated by a finite vector space in every degree.) This *minimal model* has a number of nice properties, but all that matters for us is that given a map $f : X \rightarrow Y$ from some manifold X , we can describe the homomorphism $f^*m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*X$ using a finite number of differential-form-valued invariants. Indeed, up to homotopy, this description can be made finitary in a much stronger sense.

Write $[\mathcal{M}_Y^*, \Omega^*X]$ for the set of homotopy classes of DGA homomorphisms. Then $f \mapsto f^*m_Y$ induces a well-defined map $[X, Y] \rightarrow [\mathcal{M}_Y^*, \Omega^*X]$ which is finite-to-one by [Sul, Thm. 10.2(i)]. Moreover, in various cases where these sets have a group structure, this map is actually the homomorphism $\otimes \mathbb{R}$.

1.3. Existence of shadows. In this paper, we study the map $f \mapsto f^*m_Y$ more closely, as a continuous map

$$\text{Map}(X, Y) \rightarrow \text{Hom}(\mathcal{M}_Y^*, \Omega^*X),$$

where the latter object is equipped with a metric induced by its homotopy theory. We can think of these homomorphisms as "platonic forms" of maps. These include, of course, the pullbacks of genuine maps, just as a committed Platonist would have to admit that the world of concepts includes the concept of any particular object in the real world, as well as abstractions at various levels. But most platonic forms are indeed abstract. Moreover, the map is far from being a homotopy equivalence, even on connected components, since many algebraic homotopies have non-integer and even irrational invariants.

Nevertheless, the main technical theorem of this paper is that we can, in a weak sense, produce inverses. Suppose Y is compact and X has bounded geometry. If $\varphi : \mathcal{M}_Y^* \rightarrow \Omega^*X$ is in the connected component of a genuine map, then it has a shadow $f : X \rightarrow Y$ in the Plato's cave of genuine maps such that f^*m_Y is reasonably close to φ , as measured by the size of an (algebraic) homotopy between them. Moreover, the Lipschitz constant of f is closely related to a natural geometric functional on φ which we call the *formal dilatation*. Most of our applications actually use the relative form of this statement:

Theorem (Shadowing principle, informal version). *Let $A \subset X$ be a subcomplex and $u : A \rightarrow Y$ an L -Lipschitz map. Then any extension $\varphi : \mathcal{M}_Y^* \rightarrow \Omega^*X$ over X of u^*m_Y which is in the relative homotopy class of a genuine extension $\tilde{u} : X \rightarrow Y$ of u has a nearby shadow $f : X \rightarrow Y$ which is in the same relative homotopy class and has Lipschitz constant at most $CM + C$, where M is the formal dilatation of φ .*

The precise statement is given in Theorem 4.1.

The significance of this is that platonic maps are sometimes easier to construct than genuine maps, since they have fewer moving parts; this makes it easier to construct new geometrically bounded objects. For example, it is much easier to produce a homotopy in the algebraic world than the geometric one, and this is what gives us our powerful results about homotopies. Other new results are obtained by harnessing scaling automorphisms of DGAs. On the other hand, in order to realize the full potential of the shadowing principle, we need additional techniques for constructing such homomorphisms.

1.4. The method of Guth. The proof of the shadowing principle is inspired by Larry Guth's recent streamlined proof [Guth] of the main homotopical result of [CDMW]. We give an outline of this proof here.

Suppose we have a nullhomotopic L -Lipschitz map $f : S^m \rightarrow S^n$, where either n is odd or $m < 2n - 2$; we would like to construct a $C(m, n)L$ -Lipschitz nullhomotopy $F : S^m \times [0, 1] \rightarrow S^n$. First of all, we may assume, by a quantitative simplicial approximation result, that f is the composition of a simplicial map from some triangulation of S^m with edge lengths $\sim 1/L$ to $\partial\Delta^{n+1}$ and a smooth map that contracts all but one of the faces of $\partial\Delta^{n+1}$. Next, we choose some uncontrolled nullhomotopy F of this map. We will homotope this to a controlled homotopy.

First, choose a triangulation X of $S^m \times [0, 1]$ also at scale $\sim 1/L$, restricting to our triangulation of S^m at $t = 0$. We will proceed by induction on the skeleta of this triangulation. The key point is that at the k th step we will make sure that the k -simplices of X are mapped to S^n in one of a fixed set of ways, depending only on m and n . Then the Lipschitz constant is bounded by

$$(\max \text{ Lipschitz constant of a restriction to a simplex}) \cdot (\min \text{ edge length of } X)^{-1}.$$

For $k < n$, we do this simply by sending the whole k -skeleton to the basepoint of S^n . This may make the homotopy even worse than it was on higher simplices, but we will fix this in future steps. This gives us a homotopy F_{n-1} which sends $X^{(n-1)}$ to a point.

The n th step is the trickiest, and it is here that we use some algebra. Note that since $F_{n-1}|_{X^{(n-1)}}$ is constant, F_{n-1} has a well-defined degree on n -simplices. Let $c \in C^n(X)$ be the cochain whose value on simplices is this degree. Since F_{n-1} is defined on $(n+1)$ -cells, this is a cocycle.

We compare this to another cocycle that describes the "ideal" behavior of such a nullhomotopy. The piecewise smooth form $f^*d \text{vol} \in \Omega^n(S^m)$ is exact since f is nullhomotopic. Moreover, $\|f^*d \text{vol}\|_\infty \leq L^n$; by an isoperimetric result for forms, reproven in this paper as Lemma 2.4, we can find an $\alpha \in \Omega^{n-1}(S^m)$ such that $d\alpha = f^*d \text{vol}$ and $\|\alpha\|_\infty \leq C(m, n)L^n$. Let $\pi : X \rightarrow S^m$ be the obvious projection; then we define a cocycle $w \in C^n(X; \mathbb{R})$ by sending each n -simplex p to

$$w(p) = \int_p ((1-t)\pi^*f^*d \text{vol} + (-1)^n \pi^*\alpha \wedge dt).$$

The L^∞ bound then implies that $|w(p)| \leq 1 + C(m, n)$.

Note that $w = c$ on the simplices of $S^m \times \{0, 1\}$. Thus $c - w \in C^n(X, S^m \times \{0, 1\}; \mathbb{R})$ is a relative cocycle and hence a relative coboundary, $c - w = \delta b$ for some $b \in C^{n-1}(X, S^m \times \{0, 1\}; \mathbb{R})$. Now we homotope F_{n-1} to a map F_n as follows. The homotopy will be constant on $X^{(n-2)}$. On each $(n-1)$ -simplex q , we make the homotopy trace out a map of degree $[b(q)]$, i.e. the nearest integer to $b(q)$, and return to the constant map to the basepoint. This means that the degree of F_n on each n -simplex p is forced to be some number which is at most $(n+1)/2$ away from $(c-w)(p)$. This is bounded by a constant depending only on m and n ; for each degree below this bound, we fix a specific map on Δ^n and homotope to that map.

Now let $k > n$; by induction, we have a map F_{k-1} which takes a finite set of values on $(k-1)$ -simplices. In particular, there is a finite set of values that it can take on the boundary of any k -simplex p . Moreover, given $F_{k-1}|_{\partial p}$, the possible relative homotopy classes of $F_{k-1}|_p$ form a torsor for $\pi_k(S^n)$, which is finite by assumption. Thus we can fix a map in each such relative homotopy class and homotope to an F_k whose restriction to p is that map. Once $k = m+1$, we have completed the proof.

Let us return now to the n th step. In this paper, we reinterpret this as follows. The form

$$(1-t)\pi^* f^* d \text{vol} + (-1)^n \pi^* \alpha \wedge dt$$

should be thought of as an algebraic nullhomotopy of the form $f^* d \text{vol}$ which describes f up to finite uncertainty; this is made precise in §3. We construct our controlled nullhomotopy by pulling the uncontrolled homotopy F as close as we can to the controlled, but purely algebraic one.

In more general situations, the map and its nullhomotopy cannot be fully described by a single form. Instead, the description of a map $X \rightarrow Y$ is an algebra homomorphism $\mathcal{M}_Y^* \rightarrow \Omega^* X$. However, we can still use roughly the same procedure: take an uncontrolled geometric homotopy F and a controlled algebraic one Φ ; as long as they are homotopic to each other in the algebraic sense, we can gradually pull F towards Φ , skeleton by skeleton, until we get a geometric homotopy which is close to Φ , and therefore controlled. This works not only for homotopies but for maps in a relative homotopy class in general.

1.5. Seeing outside the cave. The method outlined in the previous section has an important weakness: in order to get the bound we want, we need to first find a DGA homomorphism that satisfies it. In the case of homotopies, there is an algorithm described in §3 which constructs such a homomorphism. The bound obtained this way, however, while sharp in some instances, is not, for example, sharp in the case of maps $S^3 \rightarrow S^2$. Here the algebraic method yields a quadratic bound, whereas I strongly suspect that the true bound is linear. In fact, we produce an only slightly superlinear bound in Theorem 5.7 using a somewhat mysterious ad hoc method.

Similarly, for elements of $\pi_n(Y)$ we can always produce not-too-large representatives algorithmically, but if we use the most general construction such representatives will not say anything nontrivial about distortion.

To highlight some of the uncertainties, we come back to maps $f : S^3 \rightarrow S^2$. To construct an algebraic nullhomotopy of such a map, it is enough to find a 1-form $\alpha \in \Omega^1(S^3)$ with $d\alpha = f^* d \text{vol}$ and $\eta \in \Omega^2(S^3)$ with $d\eta = \alpha \wedge f^* d \text{vol}$. By the aforementioned isoperimetric result, we can find η with $\|\eta\|_\infty \lesssim (\text{Lip } f)^4$. But is this the best we can do, or can we always make the choices so that $\|\eta\|_\infty \lesssim (\text{Lip } f)^3$? This is unclear; and even if this were ruled out, we would still not know that a geometric homotopy can't be linear, since these can be constructed in other ways.

Of course, the problem does reduce to a question about whether there are homomorphisms $\mathcal{M}_{S^2}^* \rightarrow \Omega(S^3 \times [0, 1])$ with certain L^∞ bounds on the images of the generators. The point is that the existence of such homomorphisms seems potentially just as hard to decide as the original questions about maps and homotopies. The same sort of questions bedevil any attempts at resolving Conjectures A and B purely through DGA methods; all the proofs we have use some kind of

self-maps that allow us to use one representative to generate a whole class of maps, whether geometrically or algebraically.

1.6. Extensions and generalizations. The shadowing principle has the advantage of being completely local. Therefore a number of extensions which are not shown in this paper nevertheless seem achievable. The author would like to thank David Kazhdan, Shmuel Weinberger, and Tali Kaufman for raising some of these points.

- (1) The results should hold for nilpotent targets as well as simply connected ones. This requires more complicated induction procedures and perhaps some stipulations regarding basepoints.
- (2) The results should hold for various extensions of rational homotopy theory, once one has a good understanding of the relevant algebra. This includes equivariant rational homotopy theory (see [Scull]) and perhaps the rational homotopy theory of more general diagrams of spaces à la [DF] (although this has never been explicitly developed) as well as sections of a fibration, or more generally for rational homotopy theory of maps fibered over some fixed base space.
- (3) The theorem holds for the case where the domain is an infinite complex of bounded geometry (although we do not give any applications that use this.) In such complexes, one could have DGA homomorphisms which are not bounded, but are controlled within an r -ball around some basepoint by some function $f(r)$. Then by rescaling or varying the sizes of subdivisions, we can get an honest map with similar control on the Lipschitz constant.
- (4) One interpretation of the shadowing principle is that in some sense, the map $\text{Map}(X, Y) \rightarrow \text{Hom}(\mathcal{M}_Y^*, \Omega^* X)$ induced by pullback of the minimal model is “almost dense” and induces a near-equivalence between the Lipschitz constant on $\text{Map}(X, Y)$ and a similar geometric functional on the other space. One could ask whether there is a stronger notion of connectivity between the Morse landscapes of these functionals; this needs to be done with some care since the map is not a homotopy equivalence. The π_0 version of this question is this: given a path in $\text{Hom}(\mathcal{M}_Y^*, \Omega^* X)$ between two genuine maps which is in the relative homotopy class of a genuine homotopy, can we find a genuine homotopy with similar geometry to the path? The π_n questions can be formed similarly. It seems that the answer must be yes, but to confirm this one needs to understand paths in the space of homomorphisms, most of which are not algebraic homotopies in the sense we use.

We can find a closer topological equivalence by restricting to homomorphisms of polynomial forms with rational coefficients. These are closely related to maps from X to the rationalization of Y , as discussed in [BrSz]. However, these do not usually come from pullbacks of maps, so we would still have to use the space $\text{Hom}(\mathcal{M}_Y^*, \Omega^* X)$ as a common refinement. Since smooth forms seem closely approximable by polynomials, it is likely that the geometry of this space is likewise quite similar.

1.7. Applications to geometric problems. One of the main motivations for studying quantitative algebraic topology is to try to understand the solutions to problems in geometric topology. The long-time method of doing geometric topology is to reduce it to problems in homotopy theory, then solve those problems using algebraic methods. One could therefore attempt to understand the solutions by putting geometric bounds on both the reduction and the homotopy theory. Here are some examples where this has been achieved.

- (1) In [CDMW], we gave a bound on the size of a nullcobordism of a nullcobordant manifold. This is a quantitative version of Thom’s cobordism theorem. Here, the algebraic problem was a special case of Conjecture C; the geometric problem was to get a bound on the size of Thom’s construction.
- (2) Already in [GrHED], Gromov uses his estimate on the growth of homotopy classes to bound the growth of embedding spaces. By a theorem of Haefliger, when $3n \geq 2(m + 1)$, isotopy classes

of embeddings of an m -manifold M in \mathbb{R}^n are in bijection with $\mathbb{Z}/2\mathbb{Z}$ -equivariant homotopy classes of maps $(M \times M) \setminus \Delta \rightarrow S^{2n-1}$. One direction is easy: every embedding $f : M \rightarrow \mathbb{R}^n$ is sent to the map

$$(m, n) \mapsto \frac{f(m) - f(n)}{|f(m) - f(n)|}.$$

Using (the free $\mathbb{Z}/2\mathbb{Z}$ -equivariant version of) Gromov's polynomial estimate and this explicit procedure, one sees that the number of homotopy classes of L -bilipschitz embeddings of M in \mathbb{R}^n is at most polynomial in L .

It would be interesting to investigate the space of such embeddings in greater detail. The methods of this paper provide solutions to some of the requisite algebraic problems. However, translating this into embedding theory requires a deeper, more geometric understanding of the correspondence going from equivariant maps back to embeddings.

Open problem. (1) Find a sharp estimate of the number of embeddings of some M in \mathbb{R}^n with a given bilipschitz constant or other geometric bound. (2) Find a bound on the difficulty of isotoping two isotopic embeddings (again, in terms of the bilipschitz constant or some other geometric bound.)

We hope that our results will induce more work on the geometric side of these problems and many others.

1.8. Structure of the paper. Section 2 introduces some technical results about the geometry of simplicial complexes which underpin the various proofs. In Section 3, we discuss rational homotopy theory in detail, including the geometric estimates introduced by Gromov. In Section 4, we state and prove the main technical result. Applications, including the proof of Conjecture C, are discussed in the last section.

1.9. Acknowledgements. I would like to first of all thank Shmuel Weinberger, who introduced me to these ideas and was present for every stage of their development, and without whom none of this would have been possible. The paper is part of a project that has been going on for several years; our collaborators Steve Ferry, Greg Chambers, and Dominic Dotterer also contributed many of the ideas which pervade this work.

I would also like to thank the Israel Institute for Advanced Studies, where this paper was conceived and largely written. Without the IIAS, I could not have benefitted from the wisdom of David Kazhdan and Tali Kaufman, who suggested various generalizations of the main result and directions for further study. I am grateful to Sasha Berdnikov for pointing out an important error in an earlier version. Finally, this paper owes its existence to Larry Guth and his incredibly thoughtful and clever reinterpretation of our earlier ideas. I cannot thank him enough.

2. GEOMETRIC PRELIMINARIES

2.1. Simplicial approximation. A key principle in this paper, carried over from [CDMW] and [CMW], is local standardization of maps. The simplest kind of such standardization is simplicial approximation. It was shown in [CDMW] that on the right sort of subdivision, simplicial approximation can be performed without increasing Lipschitz constants too much.

Let X be a simplicial complex with the standard metric. We say a subdivision of X is L -regular if the simplices have all edge lengths between $r^{-1}L$ and rL , for some r which perhaps depends on dimension. The most common notion of subdivision used is barycentric subdivision, which is not regular—the simplices get progressively skinnier. However, several regular subdivision schemes are available, including the following:

- Add a central vertex to each k -simplex to subdivide it into $k + 1$ cubes. Cubulate each such cube at scale $1/L$, then break each small cube into simplices in a standard way. This method was described in [FW].
- Slice each simplex into approximately L slices of equal width along planes parallel to each face. This subdivides it into a finite number of types of polyhedra. Apply a standard simplicial subdivision to each. A specific such method is given in [EdGr]; its advantage is that r can be taken to be a constant $\sqrt{2}$, not depending even on dimension.

Such subdivisions are useful for simplicial approximation of maps.

Proposition 2.1 (Quantitative simplicial approximation theorem). *For finite simplicial complexes X and Y with piecewise linear metrics, there is a constant C such that any L -Lipschitz map $f : X \rightarrow Y$ has a $C(L+1)$ -Lipschitz simplicial approximation via a homotopy of thickness $C(L+1)$ and length C .*

The main purpose of simplicial approximation in this paper, as in [CDMW] and [CMW], is to bound the behaviors of maps on simplices. If there are only finitely many things a map can do on a simplex, we can bound its Lipschitz constant by the maximum Lipschitz constant of these restrictions. Therefore it is useful to extract this more general property and give it a name.

Let \mathcal{F}_k be a finite set of maps $\Delta^k \rightarrow Y$, for some space Y . If X is a simplicial complex, a map $f : X \rightarrow Y$ is \mathcal{F} -mosaic if all of its restrictions to k -simplices are in \mathcal{F}_k . Here, \mathcal{F} is a semi-simplicial set whose simplices in degree k are \mathcal{F}_k , which can be formed naturally via restriction maps. We can think of this *shard complex* as a finite subcomplex of the singular simplicial set of Y .

We refer to a collection of maps as *uniformly mosaic* if they are all \mathcal{F} -mosaic with respect to a fixed unspecified shard complex \mathcal{F} .

The main advantage of this definition is that the property of being mosaic is preserved under postcomposition. Thus for example if we have a homotopy equivalence $\varphi : Z \rightarrow Y$ from a simplicial complex Z to a cell complex Y which contracts the 1-skeleton, then we can simplicially approximate a map $X \rightarrow Z$, then compose with φ to get an \mathcal{F}_φ -mosaic map for some fixed \mathcal{F}_φ whose 1-skeleton is a point.

2.2. Quantitative antidifferentiation. De Rham algebras exist in several variations, including smooth and piecewise polynomial. In this paper, we also use the algebra of simplexwise smooth forms on a simplicial complex. This has several advantages: such forms can be built skeleton-by-skeleton and this is the natural context for pullbacks of smooth forms on a manifold by simplexwise smooth functions. Given a simplicial complex X , this is the algebra we will denote Ω^*X ; for a manifold with boundary the same notation will denote the smooth forms.

A key step in both Gromov’s earliest arguments in [GrHED] and in this paper is quantitative antidifferentiation of forms: given an exact k -form with L^∞ -norm B , produce a $(k - 1)$ -form which it bounds with L^∞ -norm CB , with C depending on the space and perhaps some other requirements we impose. There are two methods for doing this in the literature. The original, due to Gromov, uses quantitative Poincaré lemmas to build antidifferentials skeleton by skeleton. Another method, which as far as we know was first used in [CDMW], uses the fact that forms are dual to currents, which obey the Federer–Fleming isoperimetric inequality. We introduce both methods as both produce results that will be useful in this paper.

Isoperimetric duality. We start by proving a result in infinite-dimensional linear algebra, a simple consequence of the Hahn–Banach theorem.

Theorem 2.2 (Isoperimetric duality). *Let $(V, \|\cdot\|_V)$ and $(W, \|\cdot\|_W)$ be normed vector spaces and $\varphi : V \rightarrow W$ a (not necessarily continuous) linear operator. There is an adjoint operator $\varphi^* : \Omega \rightarrow V^*$ where $\Omega \subseteq W^*$ is the space of operators $\omega : W \rightarrow \mathbb{R}$ such that ω and $\varphi^*\omega$ are both bounded. Then for any $C > 0$, the following are equivalent:*

- (1) For every $\varepsilon > 0$, every $w \in \text{im}(\varphi)$ has a preimage v with $\|v\|_V \leq C\|w\|_W + \varepsilon$.
(2) For every $\varepsilon > 0$, every $\nu \in \text{im}(\varphi^*)$ has a preimage ω with $\|\omega\|_{W^*} \leq C\|\nu\|_{V^*} + \varepsilon$.

Proof. For any constant $\sigma > C$, consider the operators

$$\varphi^{-1} : (\varphi(V), \|\cdot\|_W) \rightarrow (V/\ker \varphi, \|\cdot\|_\sigma),$$

where

$$\|\bar{v}\|_\sigma = \inf \{ \sigma \|\varphi(v_0)\|_W + \|v_1\|_V : v_0 + v_1 \in \bar{v} \},$$

and

$$(\varphi^*)^{-1} : (\varphi^*\Omega, \|\cdot\|_{\sigma^*}) \rightarrow (W^*/\ker \varphi^*, \|\cdot\|_{\text{inf}}),$$

where

$$\|\nu\|_{\sigma^*} = \max \left\{ \inf \left\{ \frac{1}{\sigma} \|\omega\|_{W^*} : \varphi^*\omega = \nu \right\}, \|\nu\|_{V^*} \right\},$$

$$\|\bar{\omega}\|_{\text{inf}} = \inf_{\omega \in \bar{\omega}} \|\omega\|_{W^*}.$$

Here, φ^{-1} is an isomorphism of vector spaces and a bounded operator, but not necessarily a bilipschitz equivalence; $(\varphi^*)^{-1}$ is injective but its image $\Omega/\ker \varphi^* \subseteq W^*/\ker \varphi^*$ is not necessarily the whole space. The conditions (1) and (2) are equivalent to the operator norms of φ^{-1} and $(\varphi^*)^{-1}$, respectively, being at most C . It is therefore enough to prove that φ^{-1} and $(\varphi^*)^{-1}$ are adjoint operators on dual normed vector spaces and therefore have the same norm.

First, any $\bar{\omega} \in W^*/\ker \varphi^*$ gives a well-defined operator on $\varphi(V)$: if $\bar{\omega} = \bar{\omega}'$, then $\varphi^*(\omega - \omega') = 0$ and so $\langle \omega, \varphi(v) \rangle = \langle \omega', \varphi(v) \rangle$. Conversely, by the Hahn–Banach theorem, any functional $\omega_0 : \varphi(V) \rightarrow \mathbb{R}$ which is continuous with respect to $\|\cdot\|_W$ has an extension to W^* . Thus $(W^*/\ker \varphi^*, \|\cdot\|_{\text{inf}}$ is the dual normed space to $(\varphi(V), \|\cdot\|_W)$. A similar argument shows that $(V/\ker \varphi, \|\cdot\|_\sigma)$ and $(\varphi^*\Omega, \|\cdot\|_{\sigma^*})$ are dual. Finally, it is clear that

$$\langle \nu, \varphi^{-1}(w) \rangle = \langle (\varphi^*)^{-1}\nu, w \rangle.$$

This completes the proof. □

We now apply this to forms and currents. In geometric measure theory, a *k-dimensional current* is simply a functional on the space of smooth differential k -forms, with a boundary operator ∂ defined to be dual to the differential. The *mass* of a current T , which may of course be infinite, is defined by $\text{mass}(T) = \sup_{\|\omega\|_\infty=1} T(\omega)$. Thus the space of currents of finite mass is dual to $(\Omega^n(X), \|\cdot\|_\infty)$. A *normal current* is a current T such that T and ∂T both have finite mass; in particular, any current of finite mass which is a cycle is normal. The space of normal k -currents in X is denoted $\mathbf{N}_k(X)$. We also define $\mathbf{N}_k(X, A) = \mathbf{N}_k(X)/\mathbf{N}_k(A)$, equipped with the quotient norm. Thus the following follows immediately from the duality theorem:

Lemma 2.3 (L^∞ isoperimetric duality). *The following are equivalent for any pair of simplicial complexes (X, A) and natural number k :*

- (1) Every exact form $\omega \in \Omega^k(X, A)$ has an antidifferential α with $\|\alpha\|_\infty \leq C\|\omega\|_\infty$.
(2) Every normal current $T \in \mathbf{N}_{k-1}(X, A)$ has a filling S with $\text{mass } S \leq C \text{ mass } T$.

Thus to show that the first inequality can be attained in the case when (X, A) is finite, we merely need to demonstrate the second. This is a version of the Federer–Fleming isoperimetric inequality, [FF60, Thm. 5.5]. In their original theorem, Federer and Fleming show that a k -current of mass T in \mathbb{R}^n whose boundary is in the k -skeleton of the unit cubical lattice can be pushed to a linear combination of k -cubes of this lattice through a $(k+1)$ -current of mass at most $C_{n,k} \text{ mass } T$; moreover, the resulting cubical k -chain has mass at most $C_{n,k} \text{ mass } T$ as well. Except for the precise constants, their proof can be used to push a current in a simplicial complex to its k -skeleton. Since it works by inductively pushing the current onto lower skeleta, it also works for a relative current

(when you reach A , stop pushing.) Finally, once we have deformed our current to a simplicial boundary in (X, A) , it is nullhomologous in a bounded way simply because the space of simplicial boundaries $B_k(X, A)$ is finite-dimensional. This shows the following:

Lemma 2.4. *Let (X, A) be a pair of finite simplicial complexes. There is an isoperimetric constant for boundaries in $\mathbf{N}_{k-1}(X, A)$. Equivalently, there is an isoperimetric constant for exact forms in $\Omega^k(X, A)$.*

Quantitative Poincaré lemmas. One important application of this type of isoperimetric inequality is to extending forms skeleton by skeleton. The following is a special case of Lemma 2.4 which is nevertheless worth stating on its own.

Lemma 2.5. *For every $0 < k \leq n$, there is a constant $C_{n,k}$ such that the following holds. Let $\omega \in \Omega^k(\Delta^n, \partial\Delta^n)$ be a closed k -form supported on the interior of the standard simplex. (If $k = n$, we require in addition that $\int_{\Delta^n} \omega = 0$.) Then there is a form $\alpha \in \Omega^{k-1}(\Delta^n, \partial\Delta^n)$ such that $d\alpha = \omega$ and $\|\alpha\|_\infty \leq C\|\omega\|_\infty$.*

From here, we show how to extend nonzero forms.

Lemma 2.6. *For every $0 < k \leq n$, there is a constant $C_{n,k}$ such that the following holds. Let $\omega \in \Omega^k(\Delta^n)$ be a closed k -form, and let $\alpha_\partial \in \Omega^k(\partial\Delta^n)$ be a $(k-1)$ -form such that $d\alpha_\partial = \omega|_{\partial\Delta^n}$. (If $n = k$, we also require that the pair satisfies Stokes' theorem, that is, $\int_{\Delta^k} \omega = \int_{\partial\Delta^k} \alpha_\partial$.) Then there is a $(k-1)$ -form $\alpha \in \Omega^{k-1}(\Delta^n)$ extending α_∂ such that $d\alpha = \omega$ and $\|\alpha\|_\infty \leq C(\|\omega\|_\infty + \|\alpha_\partial\|_\infty)$.*

Proof. Let U be the $1/2n$ -neighborhood of $\partial\Delta^n$ in Δ^n , and let $\varphi : U \rightarrow \partial\Delta^n$ be a smooth projection with Lipschitz constant L_π . Let $\epsilon : \Delta^n \rightarrow [0, 1]$ be a smooth bump function with Lipschitz constant L_ϵ which is 1 on $\partial\Delta^n$ and 0 outside U . Then $\epsilon\pi^*\alpha_\partial$ is an extension of α_∂ to Δ^n with

$$\begin{aligned} \|\epsilon\pi^*\alpha_\partial\|_\infty &\leq L_\pi^{k-1}\|\alpha_\partial\|_\infty \\ \|d(\epsilon\pi^*\alpha_\partial)\|_\infty &= \|d\epsilon \wedge \pi^*\alpha_\partial + \epsilon\pi^*d\alpha_\partial\|_\infty \leq L_\epsilon L_\pi^{k-1}\|\alpha_\partial\|_\infty + L_\pi^k\|\omega\|_\infty. \end{aligned}$$

Now we apply the previous lemma to $\omega - d(\epsilon\pi^*\alpha_\partial)$ to get an $\alpha' \in \Omega^k(\Delta^n, \partial\Delta^n)$ with

$$\|\alpha'\|_\infty \leq C_{n,k}(L_\epsilon L_\pi^{k-1}\|\alpha_\partial\|_\infty + (L_\pi^k + 1)\|\omega\|_\infty).$$

The form we are looking for is $\alpha = \alpha' + \epsilon\pi^*\alpha_\partial$. □

3. HOMOTOPY THEORY OF DGAs

In this section we sketch out the homotopy theory of differential graded algebras, following the treatment of [GrMo, Ch. IX and X]. The relatively explicit formulation helps us obtain quantitative bounds on the sizes of DGA homotopies, which we will later harness to obtain various geometric bounds. We also review Gromov's arguments bounding the homotopy classes of maps with a given Lipschitz constant.

A (commutative) differential graded algebra (DGA) will always denote a cochain complex of \mathbb{Q} - or \mathbb{R} -vector spaces equipped with a graded commutative multiplication which satisfies the (graded) Leibniz rule. The prototypical example of an \mathbb{R} -DGA is the algebra of smooth forms on a manifold or piecewise smooth forms on a simplicial complex.

The cohomology of a DGA is the cohomology of the underlying cochain complex. The relative cohomology of a DGA homomorphism $\varphi : \mathcal{A} \rightarrow \mathcal{B}$ is the cohomology of the cochain complex

$$C^n(\varphi) = \mathcal{A}^n \oplus \mathcal{B}^{n-1}$$

with the differential given by $d(a, b) = (da, \varphi(a) - db)$. This cohomology fits, as expected, into an exact sequence involving $H^*(\mathcal{A})$ and $H^*(\mathcal{B})$.

Given a coefficient vector space V , $H^*(\mathcal{A}, V)$ is the cohomology of the cochain complex $\text{Hom}(V, \mathcal{A}^n)$. By the universal coefficient theorem, this is naturally isomorphic to $\text{Hom}(V, H^*(\mathcal{A}))$, but we will frequently be using the cochain complex itself.

A *weak equivalence* between DGAs \mathcal{A} and \mathcal{B} is a homomorphism $\mathcal{A} \rightarrow \mathcal{B}$ which induces an isomorphism on cohomology.

An algebra \mathcal{A} is *simply connected* if $\tilde{H}^0(\mathcal{A}) = H^1(\mathcal{A}) = 0$. If \mathcal{A} is simply connected and of *finite type* (i.e. it has finite-dimensional cohomology in every degree) then it has a *minimal model*: a weak equivalence $m_{\mathcal{A}} : \mathcal{M}_{\mathcal{A}} \rightarrow \mathcal{A}$ where $\mathcal{M}_{\mathcal{A}}$ is freely generated as an algebra by finite-dimensional vector space V_n in degree n (we write

$$\mathcal{M}_{\mathcal{A}} = \bigwedge_{n=2}^{\infty} V_n)$$

and the differential satisfies

$$dV_n \subseteq \bigwedge_{k=2}^{n-1} V_k.$$

In other words, $\mathcal{M}_{\mathcal{A}}$ can be built up via a sequence of extensions

$$\mathcal{M}_{\mathcal{A}}(n+1) = \mathcal{M}_{\mathcal{A}}(n)\langle V_{n+1} \rangle,$$

with the differential on $\mathcal{M}_{\mathcal{A}}(n+1)$ extending that on $\mathcal{M}_{\mathcal{A}}(n)$, starting with $\mathcal{M}_{\mathcal{A}}(1) = \mathbb{Q}$ or \mathbb{R} . We refer to elements of the V_n as *indecomposables*. We will often describe finitely generated free DGAs by indicating the degree of generators as superscripts in parentheses: $a^{(3)}$ means that a is a generator in degree 3.

In particular, if Y is a simply connected manifold or simplicial complex, the algebra of forms Ω^*Y has a minimal model which we will call $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*Y$. This models the Postnikov tower of Y : each $V_n \cong \text{Hom}(\pi_n(Y), \mathbb{R})$ and the differential on V_n is dual to the k -invariant of the fibration $Y_{(n)} \rightarrow Y_{(n-1)}$. This can be shown inductively via obstruction theory.

3.1. Obstruction theory. Given a principal fibration $K(\pi, n) \rightarrow E \rightarrow B$ and a map $f : X \rightarrow B$ for some space X , obstruction theory gives an exact sequence of sets

$$H^n(X; \pi) \rightarrow [X, E] \rightarrow [X, B] \rightarrow H^{n+1}(X; \pi).$$

Moreover, over a given map $f : X \rightarrow B$, there is an exact sequence of groups

$$\dots \rightarrow \pi_1(E^X, \tilde{f}) \rightarrow \pi_1(B^X, f) \rightarrow H^n(X; \pi) \rightarrow \{\text{lifts of } f\} \rightarrow 0,$$

where the set of lifts is a torsor acted on by $H^n(X; \pi)$.

We now give DGA versions of these statements. First define homotopy of DGA homomorphisms as follows: $f, g : \mathcal{A} \rightarrow \mathcal{B}$ are homotopic if there is a homomorphism

$$H : \mathcal{A} \rightarrow \mathcal{B} \otimes \mathbb{R}\langle t^{(0)}, dt^{(1)} \rangle$$

such that $H|_{\frac{t=0}{dt=0}} = f$ and $H|_{\frac{t=1}{dt=0}} = g$. We think of $\mathbb{R}\langle t, dt \rangle$ as an algebraic model for the unit interval and this notion as an abstraction of the map induced by an ordinary smooth homotopy. In particular, it defines an equivalence relation [GrMo, Cor. 10.7]. Moreover, for any piecewise smooth space X there is a map

$$\rho : \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle \rightarrow \Omega^*(X \times [0, 1])$$

given by “realizing” this interval, that is, interpreting the t and dt the way one would as forms on the interval. We will use this realization map further in the paper.

We also introduce some notation which is useful for constructing homotopies between DGA homomorphisms. For any DGA \mathcal{A} , define an operator $\int_0^t : \mathcal{A} \otimes \mathbb{R}\langle t, dt \rangle \rightarrow \mathcal{A} \otimes \mathbb{R}\langle t, dt \rangle$ by

$$\int_0^t a \otimes t^i = 0, \int_0^t a \otimes t^i dt = (-1)^{\deg a} a \otimes \frac{t^{i+1}}{i+1}$$

and an operator $\int_0^1 : \mathcal{A} \otimes \mathbb{R}\langle t, dt \rangle \rightarrow \mathcal{A}$ by

$$\int_0^1 a \otimes t^i = 0, \int_0^1 a \otimes t^i dt = (-1)^{\deg a} \frac{a}{i+1}.$$

These provide a formal analogue of fiberwise integration; in particular, they satisfy the identities

$$(3.1) \quad d(\int_0^t u) + \int_0^t du = u - u|_{\substack{t=0 \\ dt=0}} \otimes 1$$

$$(3.2) \quad d(\int_0^1 u) + \int_0^1 du = u|_{\substack{t=1 \\ dt=0}} - u|_{\substack{t=0 \\ dt=0}}.$$

Now we state the main lemma of obstruction theory, which states the conditions under which a map can be extended over an elementary extension.

Proposition 3.3 (10.4 in [GrMo]). *Let $\mathcal{A}\langle V \rangle$ be an n -dimensional elementary extension of a DGA \mathcal{A} . Suppose we have a diagram of DGAs*

$$\begin{array}{ccc} \mathcal{A} & \xrightarrow{f} & \mathcal{B} \\ \downarrow & & \downarrow h \\ \mathcal{A}\langle V \rangle & \xrightarrow{g} & \mathcal{C} \end{array}$$

with $g|_{\mathcal{A}} \simeq hf$ by a homotopy $H : \mathcal{A} \rightarrow \mathcal{B} \otimes \mathbb{R}\langle t, dt \rangle$. Then the map $O : V \rightarrow \mathcal{B}^{n+1} \oplus \mathcal{C}^n$ given by

$$O(v) = \left(f(dv), g(v) + \int_0^1 H(dv) \right)$$

defines an obstruction class $[O] \in H^{n+1}(h : \mathcal{B} \rightarrow \mathcal{C}; V)$ to producing an extension $\tilde{f} : \mathcal{A}\langle V \rangle \rightarrow \mathcal{B}$ of f with $g \simeq h \circ \tilde{f}$ via a homotopy \tilde{H} .

When the obstruction vanishes, there are maps $(b, c) : V \rightarrow \mathcal{B}^n$ such that $d(b, c) = O$, i.e.

$$db(v) = f(dv)$$

$$dc(v) = h \circ b(v) - g(v) - \int_0^1 H(dv).$$

Then for $v \in V$ we can set $\tilde{f}(v) = b(v)$ and

$$(3.4) \quad \tilde{H}(v) = g(v) + \int_0^1 H(dv) + d(c(v) \otimes t).$$

This gives a specific formula for the extension.

This lemma has a relative analogue which is also quite useful.

Proposition 3.5 (10.5 in [GrMo]). *Let $\mathcal{A}\langle V \rangle$ be an n -dimensional elementary extension of a DGA \mathcal{A} . Suppose we have a diagram of DGAs*

$$\begin{array}{ccccc} \mathcal{A} & \xrightarrow{f} & \mathcal{B} & \xrightarrow{\mu} & \mathcal{D} \\ \downarrow & & \downarrow h & \nearrow \nu & \\ \mathcal{A}\langle V \rangle & \xrightarrow{g} & \mathcal{C} & & \end{array}$$

where

- (1) $g|_{\mathcal{A}} \simeq hf$ by a homotopy $H : \mathcal{A} \rightarrow \mathcal{B} \otimes \mathbb{R}\langle t, dt \rangle$,
- (2) ν is surjective,

- (3) $\nu \circ h = \mu$ on the nose, and
(4) $\mu \circ f = \nu \circ g|_{\mathcal{A}}$ on the nose.

Then the map $O : V \rightarrow \mathcal{B}^{n+1} \oplus \mathcal{C}^n$ given by

$$O(v) = \left(f(dv), g(v) + \int_0^1 H(dv) \right)$$

defines an obstruction class $[O] \in H^{n+1}(h : \mathcal{B} \rightarrow \mathcal{C}; V)$ to producing an extension $\tilde{f} : \mathcal{A}\langle V \rangle \rightarrow \mathcal{B}$ of f with $g \simeq h \circ \tilde{f}$ via a homotopy \tilde{H} where $\mu \circ \tilde{H}$ is constant (i.e. $\mu \circ \tilde{H} = (\mu \circ f) \otimes 1$.)

In effect, this says that in this case we can define b and c as above so that $\mu \circ b = \mu \circ g|_V$ and $\mu \circ c = 0$.

Finally, we give the DGA version of the exact sequence of groups.

Proposition 3.6 (14.4 in [GrMo]). *Let $\mathcal{A}\langle V \rangle$ be an n -dimensional elementary extension of a DGA \mathcal{A} . Then for any DGA \mathcal{B} and map $\varphi : \mathcal{A} \rightarrow \mathcal{B}$, there is an exact sequence of groups (and a torsor)*

$$[\mathcal{A}\langle V \rangle, \mathcal{B} \otimes \mathbb{R}\langle e^{(1)} \rangle]_{\tilde{\varphi}} \rightarrow [\mathcal{A}, \mathcal{B} \otimes \mathbb{R}\langle e^{(1)} \rangle]_{\varphi} \xrightarrow{\mathcal{O}} H^n(\mathcal{B}; V) \rightarrow \{\text{extensions of } \varphi\} \rightarrow 0,$$

where $\tilde{\varphi}$ is any extension of φ , the group structure on the first two sets is given by

$$(\varphi + \eta_1 \otimes e) \boxplus (\varphi + \eta_2 \otimes e) = \varphi + (\eta_1 + \eta_2) \otimes e,$$

and the arrow \mathcal{O} is given by

$$\varphi + \eta \otimes e \mapsto \eta|_{dV_n}.$$

3.2. Quantitative aspects. In this subsection, X will be a finite piecewise Riemannian simplicial complex and Y a compact simply-connected Riemannian manifold with boundary with minimal model $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*X$. The technical results of this paper largely concern homomorphisms $\varphi : \mathcal{M}_Y^* \rightarrow \Omega^*X$ for such X and Y . We would like to define a notion of size on such homomorphisms. Given a simplexwise Riemannian metric on X , we equip each $\Omega^k X$ with the L^∞ -norm; we also fix a norm on each of the vector spaces V_k of degree k indecomposables. Since the V_k are finite-dimensional, the choice of this norm affects anything that depends on finitely many of them only up to a constant. Given this data, we define the (formal) dilatation of φ by

$$\text{Dil}(\varphi) = \max_{k \in \{2, \dots, \dim X\}} \|\varphi|_{V_k}\|_{\text{op}}^{1/k}.$$

Note that if $f : X \rightarrow Y$ is an L -Lipschitz map, then f^* multiplies the L^∞ norm of k -forms by at most L^k . Therefore when $\varphi = f^*m_Y$ for some map $f : X \rightarrow Y$,

$$\text{Dil}(\varphi) \leq C \text{Lip } f,$$

where C depends only on m_Y and the norms on the V_k .

We define the dilatation of a homotopy in a similar way. Note that for any $\Phi : \mathcal{M}_Y^* \rightarrow \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle$, we can write

$$\Phi(v) = \sum_{i=0}^r \Phi_i^0(v) \otimes t^i + \sum_{j=0}^s \Phi_j^1(v) \otimes t^j dt.$$

Given a parameter τ , we set

$$\text{Dil}_\tau(\Phi) = \max_{k \in \{2, \dots, \dim X\}} \left(\sum_{i=0}^r (\|\Phi_i^0|_{V_k}\|_{\text{op}})^{1/k} + \sum_{j=0}^s \left(\frac{1}{\tau} \|\Phi_j^1|_{V_k}\|_{\text{op}} \right)^{1/k} \right).$$

One should think of τ as the “length” of the formal interval; then, since t varies between 0 and 1 over this whole length, we get $\|dt\|_\infty = 1/\tau$.

We will frequently want to apply the obstruction lemmas in such a way that we can say something quantitative about the extension. We give here a couple of specialized instances in which we can do this.

Proposition 3.7. *Suppose that $\Phi_k : \mathcal{M}_Y^*(k) \rightarrow \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle$ is a partially defined homotopy between $\varphi, \psi : \mathcal{M}_Y^* \rightarrow \Omega^*X$.*

(i) *The obstruction to extending Φ_k to a homotopy*

$$\Phi_{k+1} : \mathcal{M}_Y^*(k+1) \rightarrow \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle$$

is a class in $H^{k+1}(X; V_{k+1})$ represented by a cochain σ with (for any $\tau > 0$)

$$\|\sigma\|_{\text{op}} \leq \tau \text{Dil}_\tau(\Phi_k)^{k+2} + \text{Dil}(\varphi)^{k+1} + \text{Dil}(\psi)^{k+1}.$$

(ii) *If this obstruction class vanishes, then we can choose Φ_{k+1} so that*

$$\|(\Phi_{k+1})_i^j|_{V_{k+1}}\|_{\text{op}} \leq (C_{\text{IP}} + 2)(\tau \text{Dil}_\tau(\Phi_k)^{k+2} + \text{Dil}(\varphi)^{k+1} + \text{Dil}(\psi)^{k+1}),$$

where C_{IP} is the isoperimetric constant for $(k+2)$ -forms in X and $\tau > 0$ is arbitrary.

Proof. We let $\sigma = -\psi + \varphi + \int_0^1 \Phi_k(dv)$. This clearly satisfies the given bound. To see that it is indeed a representation of the obstruction to extending Φ_k , note that by Lemma 3.5 the obstruction can be directly formulated as

$$v \mapsto (\Phi_k(dv), \varphi(v) \oplus \psi(v)) \in \text{Hom}(V_{k+1}, C^{k+2}(\Omega^*X \otimes \mathbb{R}\langle t, dt \rangle, \Omega^*X \oplus \Omega^*X)).$$

We can zero out the second term by subtracting

$$d(\varphi(v) \otimes (1-t) + \psi(v) \otimes t, 0)$$

to get a homologous cochain $(p, 0)$. Finally, $(p, 0)$ is cohomologous to $((\int_0^1 p) \otimes dt, 0)$ via

$$d(\int_0^t p + (-1)^{k+1}(\int_0^1 p) \otimes t, 0).$$

Therefore $\int_0^1 p : V_{k+1} \rightarrow \Omega^{k+1}X$ is a cochain representing the obstruction class in $H^{k+1}(X; V_{k+1})$. But this is exactly σ .

Now, suppose the obstruction class vanishes. Then we can choose $c : V_{k+1} \rightarrow \Omega^*X$ such that $dc = -\sigma$ and $\|c\|_\infty \leq C_{\text{IP}}\|\sigma\|_{\text{op}}$ and set

$$\Phi_{k+1}(v) = \varphi(v) + \int_0^t \Phi_k(dv) + d(c(v) \otimes t),$$

similar to (3.4). This satisfies the bound in (ii). \square

In the specific instances we consider, we can often obtain better bounds. Suppose that φ and ψ both have dilatation $\leq L$, and that we can construct a homotopy Φ between them formally up to degree n without encountering obstructions. We claim that this can be done so that for $k \leq n$,

$$(3.8) \quad \|\Phi_i^j|_{V_k}\|_{\text{op}} \leq C(X, Y)L^{2k-2}, \text{ for } j = 0, 1.$$

Clearly this is true for $k = 2$, since degree 2 indecomposables have zero differential. Now suppose it's true up to $k - 1$. Then $a \in dV_k$ is sent to a form whose ∞ -norm is bounded by

$$C(a) \prod_{r_1 + \dots + r_\ell = k+1} L^{2r_i-2} \leq C(a)L^{2k-2}.$$

In particular, we end up with $\text{Dil}_1(\Phi) \lesssim L^{\frac{2n-2}{n}}$ and $\text{Dil}_{L^{-2}}(\Phi) \lesssim L^2$.

A second quantitative lemma concerns formal (that is, algebraic) concatenation of homotopies. The proof of [GrMo, Cor. 10.7] shows in a formal way that DGA homotopy is a transitive relation. We reproduce this proof with quantitative bounds on the size of the concatenation.

Proposition 3.9. *Suppose $\varphi, \psi, \xi : \mathcal{M}_Y^* \rightarrow \Omega^* X$ are homomorphisms and we are given homotopies Φ between φ and ψ and Ψ between ψ and ξ . Then we can find a homotopy*

$$\Xi : \mathcal{M}_Y^* \rightarrow \Omega^* X \otimes \mathbb{R}\langle t, dt \rangle$$

between φ and ξ such that

$$\text{Dil}(\Xi) \leq (\text{Dil}(\Phi) + \text{Dil}(\Psi)) \frac{\dim X + 1}{3}.$$

Proof. Roughly speaking, we will arrange Φ and Ψ along two sides of a formal square, extend the map to the rest of the square, and then restrict to the diagonal to get Ξ . Here is how this is done in detail.

Write Φ_i^0 and Ψ_i^0 for the coefficients of t^i of Φ and Ψ respectively, and Φ_i^1, Ψ_i^1 for the coefficients of $t^i dt$. We first note that the formula

$$“\Phi + \Psi” = \sum_{i=0}^{p_1} \Phi_i^0 \otimes t^i + \sum_{j=0}^{q_1} \Phi_j^1 \otimes t^j dt + \sum_{k=1}^{p_2} \Psi_k^0 \otimes s^k + \sum_{\ell=0}^{q_2} \Psi_\ell^1 \otimes s^\ell ds$$

(where k starts at 1 because the t parts already restrict to ψ when $t = 1$) defines a DGA map

$$\mathcal{M}_Y^* \rightarrow \Omega^* X \otimes \mathbb{R}\langle t, dt, s, ds \rangle / \langle s(t-1), (t-1)ds, (s-1)dt \rangle.$$

This should be thought of as the DGA of two sides of a square, and we want to lift to a map

$$\bar{\Xi} : \mathcal{M}_Y^* \rightarrow \Omega^* X \otimes \mathbb{R}\langle t, dt, s, ds \rangle$$

to the DGA of the whole square. We do this by induction on degree. The map is trivial on $\mathcal{M}_Y^*(1)$; then we extend from $\mathcal{M}_Y^*(n)$ to $\mathcal{M}_Y^*(n+1)$ by defining

$$\bar{\Xi}(v) = “\Phi + \Psi”(v) + \int_0^s (\bar{\Xi}(dv) - (\bar{\Xi}(dv))|_{t=1}).$$

It is easy to check that this has the right differential and the right restrictions to $t = 1$ and to $s = 0$. Finally, we take the “diagonal” $\Xi = \bar{\Xi}|_{s=t}$. Unfortunately, every time we extend from n to $n+1$ and $dv \neq 0$, we get $\|\Xi(v)\|_{\text{op}} \sim \|\Xi(dv)\|_{\text{op}}$. This means that this construction cannot give a better bound than that in the proposition. \square

3.3. Homotopy periods and Gromov’s results. Now let $f : S^n \rightarrow Y$ be a smooth map. If we attempt to nullhomotope $f^*m_Y : \mathcal{M}_Y^* \rightarrow \Omega^* S^n$ by the method of Prop. 3.7, the procedure does not fail until the very last step, where the obstruction

$$\alpha \in H^{n+1}(\Omega^* S^n \otimes \mathbb{R}\langle t, dt \rangle; V_n) \cong H^n(S^n; V_n)$$

given by the formula $\alpha = [f^*m_Y|_{V_n} + \int_0^1 \Phi_n|_{V_n}]$ may be nontrivial. This obstruction determines an element of $\pi_n(Y) \otimes \mathbb{R} \cong \text{Hom}(V_n, \mathbb{R})$ and can be computed algorithmically by repeated antidifferentiation. In other words, it generalizes Whitehead’s construction of the Hopf invariant and is another way of constructing the “homotopy periods” outlined by Sullivan in [Sul, §11].

Example 3.10. (i) Let $f : S^3 \rightarrow S^2$ be a map. The minimal DGA of S^2 is given by

$$\langle x^{(2)}, y^{(3)} \mid dx = 0, dy = x^2 \rangle;$$

clearly, $m_{S^2}y = 0$ in any minimal model. Thus the first stage of a nullhomotopy of $f^*m_{S^2}$ is given by

$$\Phi_2(x) = f^*m_{S^2}x \otimes (1-t) - c(x) \otimes dt$$

where $dc(x) = f^*m_{S^2}x$. The obstruction to extending this to b is given by $[f^*m_{S^2}x \wedge c(x)] \in H^3(S^3; \mathbb{R})$. This is exactly the Hopf invariant.

(ii) For a slightly more complicated example, we consider $Y = (S^3 \times S^3) \setminus D^6$. This is a 6-manifold homotopy equivalent to $S^3 \vee S^3$; the relevant part of the minimal model is

$$\langle x_1^{(3)}, x_2^{(3)}, y^{(5)}, z_1^{(7)}, z_2^{(7)}, \dots \mid dx_i = 0, dy = x_1 x_2, dz_i = x_i y, \dots \rangle.$$

Consider a map $f : S^7 \rightarrow Y$; again we try to nullhomotope $f^* m_Y$, and the first few stages are

$$\Phi(x_i) = f^* m_Y x_i \otimes (1 - t) + c(x_i) \otimes dt$$

$$\Phi(y) = f^* m_Y y \otimes (1 - t) + \frac{1}{2}(f^* m_Y x_1 \wedge c(x_2) + c(x_1) \wedge f^* m_Y x_2) \otimes (t^2 - t) + c(y) \otimes dt,$$

where $dc(x_i) = f^* m_Y x_i$ and $dc(y) = f^* m_Y y - \frac{1}{2}(f^* m_Y x_1 \wedge c(x_2) + c(x_1) \wedge f^* m_Y x_2)$. Since $m_Y z_i = 0$ for dimension reasons, the obstruction to extending to z_i is given by

$$c(x_i) \wedge \left(\frac{1}{2} f^* m_Y y + \frac{1}{12} (f^* m_Y x_1 \wedge c(x_2) + c(x_1) \wedge f^* m_Y x_2) \right) + \frac{1}{2} f^* m_Y x_i \wedge c(y).$$

Clearly, homotopy periods quickly become impractical to compute by hand for more complicated DGAs. Similar examples were computed by Richard Hain in his PhD thesis [Hain].

Suppose now that f is L -Lipschitz. Then by (3.8), we can make sure that the obstruction class α satisfies $\|\alpha\|_{\text{op}} \lesssim L^{2n-2}$. In particular, this proves the following results of Gromov:

Theorem 3.11. *Let Y be simply connected and Lipschitz homotopy equivalent to a finite complex.*

(i) *The distortion function of an element of $\pi_n(Y)$ is $\Omega(k^{1/(2n-2)})$.*

(ii) *The growth function of $\pi_n(Y)$ is polynomial and in fact $O(L^{(2n-2)\text{rk}(\pi_n(Y) \otimes \mathbb{Q})})$.*

While these were stated in various combinations in [GrHED], [GrQHT], and [GrMS, Ch. 7], the proofs are essentially omitted in the first two and incorrect in the last. This section is meant to close this gap.

The bounds above, however, are not sharp in most cases. The bounds whose sharpness is presumed by Gromov's conjectures are difficult, perhaps impossible, to summarize completely in all instances. The best we can say is that they are obtained by assuming that all pullbacks of genuine k -forms on Y have L^∞ norm $\lesssim L^k$ and inducting to obtain bounds on other forms. We illustrate this heuristic by way of yet another example.

Example 3.12. Let \mathbf{NF} be an 8-complex with the minimal model

$$\mathcal{M}_{\mathbf{NF}}^* = \langle x^{(3)}, y^{(3)}, z^{(5)}, T^{(10)}, \dots \mid dx = dy = 0, dz = xy, dT = xyz, \dots \rangle.$$

The geometry of this complex is discussed further in §5.1. Here we focus on the algebra. Note that $\pi_{10}(\mathbf{NF})$ has a single rational generator. Suppose $f : S^{10} \rightarrow \mathbf{NF}$ is L -Lipschitz; then we get the following partial nullhomotopy of $f^* m_{\mathbf{NF}}$:

$$\Phi(x) = f^* m_{\mathbf{NF}} x \otimes (1 - t) + c(x) \otimes dt$$

$$\Phi(y) = f^* m_{\mathbf{NF}} y \otimes (1 - t) + c(y) \otimes dt$$

$$\Phi(z) = f^* m_{\mathbf{NF}} z \otimes (1 - t) + \frac{1}{2}(f^* m_{\mathbf{NF}} x \wedge c(y) + c(x) \wedge f^* m_{\mathbf{NF}} y) \otimes (t^2 - t) + c(z) \otimes dt,$$

where

$$dc(x) = f^* m_{\mathbf{NF}} x$$

$$dc(y) = f^* m_{\mathbf{NF}} y$$

$$dc(z) = f^* m_{\mathbf{NF}} z - \frac{1}{2}(f^* m_{\mathbf{NF}} x \wedge c(y) + c(x) \wedge f^* m_{\mathbf{NF}} y).$$

Then our algorithm computes the obstruction to extending the nullhomotopy to T as

$$\frac{1}{3}(c(x) \wedge f^*m_{\mathbf{NF}}(y \wedge z) + c(y) \wedge f^*m_{\mathbf{NF}}(z \wedge x) + c(z) \wedge f^*m_{\mathbf{NF}}(x \wedge y)).$$

Now, $\|c(x)\|_\infty$ and $\|c(y)\|_\infty \lesssim L^3$ by Lemma 2.4, but the same argument only yields $\|c(z)\| \lesssim L^6$. This gives a bound of $O(L^{11})$ for the first two terms but $O(L^{12})$ for the last. On the other hand, the last term can be eliminated by subtracting the exact form $d(c(z) \wedge f^*m_{\mathbf{NF}}z)$. Thus we get an overall bound $\langle T, f \rangle = O(L^{11})$, although this is not immediately yielded by the algorithm above. As we will see below, this bound is sharp.

4. THE SHADOWING PRINCIPLE

Theorem 4.1 (The shadowing principle). *Let (X, A) be an n -dimensional simplicial pair with the standard metric on simplices and Y a simply connected compact Riemannian manifold with boundary which has a minimal model $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*Y$. Fix norms on the spaces V_k of k -dimensional indecomposables of \mathcal{M}_Y^* . Let $f : X \rightarrow Y$ be a map and $\varphi : \mathcal{M}_Y^* \rightarrow \Omega^*X$ a homomorphism such that*

- (1) $f^*m_Y|_A = \varphi|_A$ (i.e. the homomorphisms restrict to the same homomorphism $\mathcal{M}_Y^* \rightarrow \Omega^*A$.)
- (2) f^*m_Y and φ are homotopic rel A (i.e. via a homotopy whose restriction to A is constant.)
- (3) $f|_A$ is L -Lipschitz.
- (4) $\text{Dil}(\varphi) \leq L$.

*Then f is homotopic rel A to a $C(L+1)$ -Lipschitz map $g : X \rightarrow Y$ such that $g^*m_Y \simeq \varphi$ via a homotopy*

$$\Phi : \mathcal{M}_Y^* \rightarrow \Omega^*(X) \otimes \langle t, dt \rangle$$

whose restriction to A is constant, such that $\text{Dil}_{1/L}(\Phi) \leq C(L+1)$. The constant C depends on Y , m_Y , and the norms on indecomposables, as well as n (but not anything else about X .)

As stated in the introduction, we want to say that this means that pullbacks of genuine maps have reasonably high density in $\text{Hom}(\mathcal{M}_Y^*, \Omega^*X)$ when it is endowed with a metric of the form

$$d(\varphi, \psi) = \inf \{ \text{size}(\Phi) : \varphi \stackrel{\Phi}{\simeq} \psi \},$$

for some notion of size. However, there is some difficulty in defining an appropriate such notion— that is, we would like the size of a constant homotopy to be zero and the notion of distance to be nondegenerate, and this is already nontrivial. One notion that does the job (its nondegeneracy follows from (3.2)) is the *formal length*, given by

$$\text{length}(\Phi) = \text{Dil}(\int_0^1 \Phi).$$

Under this metric, the theorem states that there is a pullback of a genuine map within distance $O(L)$ of any homomorphism with dilatation L which lies in the homotopy class of the pullback of a genuine map. Put this way, this is a nontrivial statement since the set of all homomorphisms with dilatation $\leq L$ has diameter which is in general some polynomial in L ; this polynomial is linear only when Y has finite homotopy groups up to dimension n .

Unfortunately, the formal length does not correspond well to the length of a genuine homotopy, or, as far as I can tell, any geometric invariant of genuine maps. In particular, as shown in [CaSi] and again in this paper in Theorem 5.4, when the space $\text{Map}(X, Y)$ is equipped with the metric given by the optimal length of a homotopy, it actually has finite diameter.

Proof. By subdividing (X, A) at scale $1/L$ and rescaling so that simplices are unit size, we may assume $L = 1$; here we implicitly use the uniformity of the result with respect to the large-scale geometry of X .

At the cost of increasing the Lipschitz constant again to some $C_0 = C_0(n, Y)$, we may also assume that $f|_A$ is mosaic with respect to a fixed shard complex $Z \subset \Delta Y$ with $Z^{(1)} = *$. To reduce




FIGURE 1. A summary of the various maps, homotopies, and second order homotopies produced in the proof of Theorem 4.1. The bottom row consists of genuine maps $X \rightarrow Y$ and homotopies between them; the rest is on the level of DGAs. Maps become better controlled from left to right.

to this case, we modify both f and φ on the star $\text{st}(A)$, which we equip with a facewise linear deformation retraction to A ,

$$\pi : \text{st}(A) \times [0, 1] \rightarrow \text{st}(A),$$

and a simplexwise linear map $\tau : \text{st}(A) \rightarrow [0, 1]$ sending $A \mapsto 0$ and $\text{lk}(A) \mapsto 1$. Let $H : A \times [0, 1] \rightarrow Y$ be a $C(n, Y)$ -Lipschitz homotopy to a $C(n, Y)$ -Lipschitz mosaic map on some chosen shard complex. (Such a homotopy can be constructed by simplicially approximating on a complex which is homotopy equivalent to Y .) We use this homotopy on a collar of width $1/2$ around A , pushing $f|_{\text{st}A}$ to the outer part of the collar:

$$\hat{f}(x) = \begin{cases} H(\pi(x, 1), 1 - 2\tau(x)) & \tau(x) \leq 1/2 \\ f(\pi(x, 1 - 2\tau(x))) & \tau(x) \geq 1/2. \end{cases}$$

We push φ to the outer half of $\text{st}(A)$ by a similar formula, adding H^*m_Y on the inner half; this gives us an algebraic map $\hat{\varphi}$. Applying the rest of the proof to \hat{f} and $\hat{\varphi}$, we produce a map \hat{g} with the required properties such that $\hat{g}|_A = \hat{f}|_A$. To get the desired g we again push $\hat{g}|_{\text{st}(A)}$ to the outer $2/3$ of the star and add H , going in the opposite direction, to the collar. To show that the resulting g satisfies condition (2), note that it clearly has a short homotopy to the algebra map $\hat{\varphi}$ which is given by H^*m_Y on the inner third of $\text{st}(A)$, pushing $\hat{\varphi}$ out. But this map in turn has a short homotopy to φ .

We now give an overview of the induction on skeleta that characterizes the rest of the proof. At the $(k + 1)$ st step, we will produce an increasingly controlled intermediate map g_{k+1} which is homotopic to g_k via a homotopy H_{k+1} (and therefore homotopic to f). In particular, g_k will be equal to the final g on the k -skeleton of X and H_{k+1} will be a constant homotopy on the $(k - 1)$ -skeleton; its behavior on k -cells is crucial for establishing control over the behavior of g_{k+1} on $(k + 1)$ -simplices. Essentially, the behavior of g_k on $(k + 1)$ -simplices allows us to define an “almost coboundary” in $C^{k+1}(X; \pi_{k+1}(Y) \otimes \mathbb{R})$ and the homotopy H_{k+1} changes this cochain by the coboundary that it almost is, leaving a uniformly bounded remainder.

In order to figure out a recipe for doing this which can be continued further, we consult a homotopy Φ_k between g_k and φ over which we also have increasing control depending on k . We then construct Φ_{k+1} from Φ_k and H_{k+1} via a second-order homotopy Ψ_{k+1} . The objects we produce are summarized in Figure 1.

As a first step, we homotope f rel A to a map g_1 which sends $X^{(1)}$ to the basepoint of Y . We also fix a homotopy

$$\Phi_1 : \mathcal{M}_Y^* \rightarrow \Omega^*(X, A) \otimes \mathbb{R}\langle t, dt \rangle$$

between $g_1^*m_Y$ and φ . We will assume that on $v \in V_k$, Φ_1 takes the form

$$\Phi_1(v) = g_1^*m_Y(v) + \int_0^t \Phi_1(dv) + d(\beta_1(v) \otimes t),$$

where $\beta_1(v)$ is a $(k-1)$ -form satisfying $d\beta_1(v) = \varphi(v) - g_1^*m_Y - \int_0^1 \Phi_1(dv)$ and $\beta_1(v)|_A = 0$.

After the k th step, we assume that we have constructed the following:

- a map $g_k : (X, A) \rightarrow Y$, homotopic rel A to f , such that $g_k|_{X^{(k)}}$ is mosaic with respect to a shard complex $Z_k \subset \Delta Y$ which depends only on Y , m_Y , and the norms on the V_i ;
- a homotopy

$$\Phi_k : \mathcal{M}_Y^* \rightarrow A^*(X) \otimes \mathbb{R}\langle t, dt \rangle$$

from $g_k^*m_Y$ to φ such that $\text{Dil}_1((\Phi_k|_{\mathcal{M}_Y^*(k)})|_{X^{(k)}}) \leq C_k = C_k(n, Y)$, and which again, for $v \in V_i$, takes the form

$$\Phi_k(v) = g_k^*m_Y(v) + \int_0^t \Phi_k(dv) + d(\beta_k(v) \otimes t),$$

where $\beta_k(v)$ is an $(i-1)$ -form satisfying $d\beta_k(v) = \varphi(v) - g_k^*m_Y - \int_0^1 \Phi_k(dv)$ and $\beta_k(v)|_A = 0$.

We then construct the analogues one dimension higher. Let $b \in C^k(X, A; \pi_{k+1}(Y))$ be the simplicial cochain obtained by integrating $\beta_k|_{V_{k+1}}$ over k -simplices and choosing an element of $\pi_{k+1}(Y)$ whose image in V_{k+1} is as close as possible in norm (but which is otherwise arbitrary.) Note that the values of b are potentially unbounded. We use b to specify a homotopy $H_{k+1} : X \times [0, 1] \rightarrow Y$ from g_k to our new desired map g_{k+1} .

We start by setting H_{k+1} to be constant on $X^{(k-1)}$. On each k -simplex q , we set $H_{k+1}|_q$ to be a map such that

$$g_{k+1}|_q = H_{k+1}|_{q \times \{1\}} = H_{k+1}|_{q \times \{0\}} = g_k|_q,$$

but such that on the cell $q \times [0, 1]$, the map traces out the element $\langle b, q \rangle \in \pi_{k+1}(Y)$. This is well-defined since $H_{k+1}|_{\partial(q \times [0, 1])}$ is canonically nullhomotopic by precomposition with a linear contraction of the simplex.

Now, given that $g_{k+1} = g_k$ on the k -skeleton, the possible relative homotopy classes of the restriction of g_{k+1} to a $(k+1)$ -simplex p form a torsor for $\pi_{k+1}(Y)$. No matter how we fill in $g_{k+1}|_p$, we will get $g_{k+1}|_p - g_k|_p = \langle \delta b, p \rangle$ in this torsor. For each possible restriction $g_k|_{\partial p}$ (of which there are finitely many since they correspond to simplicial maps $\partial \Delta^{k+1} \rightarrow Z_k$) we fix representatives for each element of this torsor. We then set $g_{k+1}|_p$ to be the appropriate representative.

We then extend the homotopy in an arbitrary way to higher skeleta.

We now argue that, for a given restriction $g_k|_{\partial p} : \partial \Delta^{k+1} \rightarrow Y$, the subset of these maps which are actually $g_{k+1}|_p$ for some $(k+1)$ -simplex p are drawn from a finite set depending only on C_k , Y , and the norm on V_{k+1} . At various stages we will write “ $\lesssim 1$ ” for numbers that are bounded by a constant depending on these items and perhaps also the norms on V_i for $i \leq k$. In this language, it is enough to show the following lemma:

Lemma 4.2. *Fix a nullhomotopy $u : \Delta^{k+1} \rightarrow Y$ of $g_k|_{\partial p}$. The homotopy class of the map $\tilde{u} : S^{k+1} \rightarrow Y$ given by $g_k|_p$ on the northern hemisphere and u on the southern hemisphere is contained in a $\lesssim 1$ -ball around $\langle \delta b, p \rangle$ in V_{k+1}^* .*

Therefore, the homotopy class of the map obtained by gluing together $g_{k+1}|_p$ and u in a similar fashion is contained in a $\lesssim 1$ -ball around $0 \in V_{k+1}^*$. But since $\pi_{k+1}(Y) \rightarrow V_{k+1}^*$ is a homomorphism from a finitely generated group, there are finitely many elements in this ball.

Proof of the lemma. In §3.3, we described the real homotopy class of $\tilde{u} : S^{k+1} \rightarrow Y$ as the obstruction in V_{k+1}^* to homotoping \tilde{u}^*m_Y to zero. But equivalently, it is the obstruction to homotoping it to any other algebraically nullhomotopic map, for example the map which restricts to $\varphi|_p$ on each hemisphere.

So we build such a homotopy Ψ through degree k , then evaluate the obstruction to extending it to V_{k+1} . On the northern hemisphere, we simply use $\Psi = \Phi_k|_p$. On the southern hemisphere, we make sure $\text{Dil}_1(\Psi) \lesssim 1$ via an inductive construction: since $\text{Dil}_1((\Phi_k|_{\mathcal{M}_Y^*(k)})|_{\partial p}) \leq C_k$, $\text{Dil}(\varphi) \leq C_0$, and $\text{Dil}(u^*m_Y) \lesssim 1$, we can use quantitative Poincaré lemmas to find γ with $\|\gamma\|_{\text{op}} \lesssim 1$ such that

$$d\gamma(v) = \varphi|_p - u^*m_Y - \int_0^1 \Psi(dv)$$

and set

$$\Psi(v) = u^*m_Y(v) + \int_0^t \Psi(dv) + d(\gamma(v) \otimes t).$$

Now the obstruction to extending to V_{k+1} is given by

$$[\tilde{u}^*m_Y + \int_0^1 \Psi(dv)] \in V_{k+1}^*.$$

But this class is the sum of $\int_p d\beta_k|_{V_{k+1}}$ (coming from the northern hemisphere) and a $\lesssim 1$ error coming from the southern hemisphere. By Stokes' theorem, this is within $\lesssim 1$ of $\langle \delta b, p \rangle$. \square

This allows us to fix a new shard complex $Z_{k+1} = Z_k \cup Z^{(k+1)} \cup \mathcal{F}_{k+1}$, where \mathcal{F}_{k+1} is the finite set of restrictions to $(k+1)$ -cells we have produced.

It remains to define the homotopy Φ_{k+1} . We do this by applying a restriction to a second-order homotopy. Let $\pi : X \times [0, 1] \rightarrow X$ be the obvious projection. Then we will construct a homotopy

$$\Psi_{k+1} : \mathcal{M}_Y^* \rightarrow \Omega^*(X \times [0, 1]) \otimes \mathbb{R}\langle s, ds \rangle$$

between H_{k+1} and $\pi^*\varphi$ such that $\Psi_{k+1}|_{t=0} = \Phi_k$ and such that $\Phi_{k+1} := \Psi_{k+1}|_{t=1}$ has the properties we desire. Accordingly, we will use the notations Φ_{k+1} and $\Psi_{k+1}|_{t=1}$ interchangeably.

We build this homotopy by skeleta. First we set $\Psi_{k+1}|_{t=0} = \Phi_k$, $\Psi_{k+1}|_{X^{(k-1)} \times [0, 1]} = \pi^*\Phi_k|_{X^{(k-1)}}$. We also set $\Psi_{k+1}|_{t=1} = \Phi_k$ on k -simplices.

The crucial step is extending to $(k+1)$ -cells. Once we have done that in a compatible way, we can extend further any way we like, since we have only promised to control Φ_{k+1} on the $(k+1)$ -skeleton. We do this by induction on degree. For cells of the form $q \times [0, 1]$ where q is a k -simplex of X , we can extend arbitrarily through degree k . For cells of the form $p \times \{1\}$ where p is a $(k+1)$ -simplex, we have promised to control the size of the homotopy, which is part of Φ_{k+1} . Explicitly, for $v \in V_i$, $i \leq k$, we let

$$\Phi_{k+1}(v) = g_{k+1}^*m_Y(v) + \int_0^s \Phi_{k+1}(dv) + d(\beta_{k+1}(v) \otimes s);$$

we would like to define such a β_{k+1} on V_i which is bounded on $X^{(k+1)}$ and such that

$$d\beta_{k+1}(v) = \varphi(v) - g_{k+1}^*m_Y(v) - \int_0^1 \Phi_{k+1}(dv).$$

By induction, $d\beta_{k+1} = d\beta_k$ on $X^{(k)}$ and

$$\|d\beta_{k+1}|_{X^{(k+1)}}\|_{\text{op}} \lesssim 1.$$

By the second quantitative Poincaré lemma, we can therefore extend $\beta_k|_{X^{(k)}}$ to a $\beta_{k+1}|_{X^{(k+1)}}$ with

$$\|\beta_{k+1}|_{X^{(k+1)}}\|_{\text{op}} \lesssim 1.$$

Since $i \leq k$ all such choices differ by coboundaries.

Now we must extend to V_{k+1} in such a way that we can keep going. Given a k -simplex q , we will define

$$\Psi_{k+1}(v) = H_{k+1}^*m_Y(v) + \int_0^s \Psi_{k+1}(dv) + d(B(v) \otimes s)$$

where $dB(v) = \pi^*\varphi(v) - H_{k+1}^*m_Y(v) - \int_0^1 \Psi_{k+1}(dv)$. By a formal application of Stokes' theorem, we must have

$$\begin{aligned} 0 &= \int_{q \times [0,1]} dB(v) = \int_{q \times [0,1]} H_{k+1}^*m_Y(v) - \int_{\partial(q \times [0,1])} \Psi_{k+1}(v) \\ &= \langle b, q \rangle(v) - \int_q \beta_k(v) + \int_q \beta_{k+1}(v). \end{aligned}$$

Thus in order to define $B|_{q \times [0,1]}$, we need to set $\int_q \beta_{k+1}(v) = \int_q \beta_k(v) - \langle b, q \rangle(v)$; we can take $\beta_{k+1}|_q$ to be the volume form times a bump function scaled by this number. Since these values are $\lesssim 1$, we can use the second quantitative Poincaré lemma to extend $\beta_{k+1}(v)$ to $p \times \{1\}$, for every $(k+1)$ -simplex p , such that $d\beta_{k+1}(v)$ is as desired and

$$\|\beta_{k+1}|_{X^{(k+1)}}\|_{\text{op}} \lesssim 1.$$

This gives us $\Phi_{k+1}(v)|_p$ with the proper bound as well as $\Psi_{k+1}|_{q \times [0,1]}$.

Now we have a map

$$\Psi_{k+1} : \mathcal{M}_Y^*(k+1) \rightarrow \Omega^*(X \times \{0\} \cup (X \times [0,1])^{(k+1)}) \otimes \mathbb{R}\langle s, ds \rangle.$$

We can extend it to a map $\mathcal{M}_Y^*(k+1) \rightarrow \Omega^*(X \times [0,1]) \otimes \mathbb{R}\langle s, ds \rangle$ by the usual Poincaré lemma. In turn, we can extend this to \mathcal{M}_Y^* by applying the relative obstruction lemma to the diagram

$$\begin{array}{ccccc} \mathcal{M}_Y^*(k+1) & \xrightarrow{\Psi_{k+1}} & \Omega^*(X \times [0,1]) \otimes \mathbb{R}\langle s, ds \rangle & & \\ \downarrow & \dashrightarrow & \downarrow t=0 & \searrow & \\ \mathcal{M}_Y^* & \xrightarrow{\Phi_k} & \Omega^*X \otimes \mathbb{R}\langle s, ds \rangle & \xlongequal{\quad} & \Omega^*X \otimes \mathbb{R}\langle s, ds \rangle. \end{array}$$

This completes the construction of Φ_{k+1} and the inductive step. When $k = \dim X$, the result is the statement of the theorem. \square

5. APPLICATIONS

5.1. Distortion and growth. In theory, our results reduce Conjectures B and A to purely algebraic questions about the homotopy theory of maps between algebras of forms. In reality, however, it is not clear whether these questions are any easier to answer than the geometric questions they come from. In this section, we give some examples of geometric constructions that confirm Conjecture B for certain types of spaces, as well as a first attempt at a general theorem using our machinery. First, however, there is the following result, which is almost a triviality given the shadowing principle:

Theorem 5.1. *Let X be an n -dimensional simplicial complex with the standard simplexwise metric and Y a simply connected finite complex. Then there is $C(n, Y)$ such that if $\alpha, \beta \in [X, Y]$ are homotopy classes which are the same rationally, then $\|\alpha\|_{\text{Lip}} \leq C\|\beta\|_{\text{Lip}}$.*

The remarkable aspect is that this constant does not depend on the particular rational homotopy class or even on the topology of X , but only on its bounded geometry. This is although the *number* of distinct homotopy classes within a rational homotopy class may be unbounded, even for a fixed X ; see [MW] for examples of this phenomenon.

Proof. Let f be a (near-)optimal representative of β and g some representative of α . Choose a minimal model $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*Y$. The shadowing principle allows us to deform g to a map \tilde{g} such that \tilde{g}^*m_Y is near f^*m_Y ; in particular $\text{Lip}(\tilde{g}) \leq C(n, Y, m_Y) \text{Lip}(f)$. \square

Universal constructions. There are a number of cases in which the distortion of elements of homotopy groups can be determined geometrically, without using the machinery introduced in this paper. Gromov originally noted in [GrHED] that given a map $f : S^{2n-1} \rightarrow S^n$ with nonzero Hopf invariant h and Lipschitz constant L , a map with Hopf invariant $k^{2n}L$ and Lipschitz constant $\lesssim kL$ can be produced by composing with a self-map

$$S^{2n-1} \xrightarrow{f} S^n \xrightarrow{\deg=k^n} S^n.$$

More generally, a large number of homotopy group elements can be represented by the following universal construction. Given spheres S^{n_1}, \dots, S^{n_r} , their product can be given a cell structure with one cell for each subset of $\{1, \dots, r\}$. Define their *fat wedge* $\mathbb{V}_{i=1}^r S^{n_r}$ to be this cell structure without the top face. Let $N = -1 + \sum_{i=1}^r n_i$; the attaching map $\tau : S^N \rightarrow \mathbb{V}_{i=1}^r S^{n_r}$ of the missing face generates $\pi_N(\mathbb{V}_{i=1}^r S^{n_r})$. By definition, $\alpha \in \pi_N(Y)$ is contained in the *rth-order Whitehead product* $[\alpha_1, \dots, \alpha_r]$, where $\alpha_i \in \pi_{n_i}(Y)$, if it has a representative which factors through a map

$$S^{N-1} \xrightarrow{\tau} \mathbb{V}_{i=1}^r S^{n_i} \xrightarrow{f_\alpha} Y$$

such that $[f_\alpha|_{S^{n_i}}] = \alpha_i$. Note that there are many potential indeterminacies in how higher-dimensional cells are mapped, so $[\alpha_1, \dots, \alpha_r]$ is a set of homotopy classes rather than a unique class.⁷

Nevertheless, any class in this set has distortion $O(k^{1/(N+1)})$, for the following reason. Let $\sigma_i : S^{n_i} \rightarrow S^{n_i}$ be an $O(L)$ -Lipschitz map of degree L^{n_i} . Then the product of the σ_i 's induces a self-map of the fat wedge which has degree L^{N+1} on the missing cell. Thus the composition

$$S^{N-1} \xrightarrow{\tau} \mathbb{V}_{i=1}^r S^{n_r} \xrightarrow{\prod \sigma_i} \mathbb{V}_{i=1}^r S^{n_r} \xrightarrow{f_\alpha} Y$$

gives us an $O(L)$ -Lipschitz representative of $L^{N+1}\alpha$.

This class of examples has not been described in detail before, but it is not original to this paper. It was mentioned by Gromov in [GrQHT] and it also provides the tools to prove the following observation of Shmuel Weinberger:

Theorem 5.2. *The following are equivalent for a finite, simply connected complex Y :*

- (i) *All rationally nontrivial elements of $\pi_*(Y)$ are undistorted.*
- (ii) *The rational Hurewicz map $\pi_*(Y) \otimes \mathbb{Q} \rightarrow H_*(Y; \mathbb{Q})$ is injective.*
- (iii) *Y is rationally equivalent to a product of odd-dimensional spheres.*

This follows from the fact that the lowest-dimensional nonzero element in the kernel of the rational Hurewicz map is always a generalized Whitehead product. This is shown in [M, Lemma 5.2].

There are also more subtle examples of similar constructions. One of the simplest examples of a rational homotopy class which is not a generalized Whitehead product is contained in the π_{10} of an 8-dimensional, four-cell CW complex constructed as follows. Ordinary Whitehead products satisfy the relations for a Lie bracket: in particular, they are bilinear and satisfy the Jacobi identity. This can be demonstrated via universal topological constructions, as in [NaTo]. In particular, the rational homotopy groups of $S^3 \vee S^3$ are a free Lie algebra whose Lie bracket is the Whitehead product, generated by the identity maps of the two spheres, which we call f and g . So we attach two 8-cells killing

$$\pi_7(S^3 \vee S^3) \otimes \mathbb{Q} \cong \langle [[f, g], f], [[f, g], g] \rangle,$$

to get a space **NF**. (This stands for “non-formal”, as **NF** is also one of the simplest examples of a space which is not formal in the sense of Sullivan.) Then $\pi_{10}(\mathbf{NF}) \otimes \mathbb{Q}$ is generated by a single element. This can be seen by constructing the first few levels of the minimal model

$$\mathcal{M}_{\mathbf{NF}}^* = \langle x^{(3)}, y^{(3)}, z^{(5)}, T^{(10)}, \dots \mid dx = dy = 0, dz = xy, dT = xyz, \dots \rangle,$$

⁷See [AA78] for the relationship between generalized Whitehead products and Sullivan minimal models.

but a generator $\tau : S^{10} \rightarrow \mathbf{NF}$ can also be constructed explicitly as a sequence of homotopies between maps $S^9 \rightarrow \mathbf{NF}$, coned off at both ends:

$$* = [*_{S^7}, f] \xrightarrow{[\text{nullh.}, f]} [[f, g], g], f \xrightarrow{\text{Jacobi}} [[[f, g], f], g] \xrightarrow{[\text{nullh.}, g]} [*_{S^7}, g] = *.$$

To demonstrate that this element is distorted, we exhibit a representative of $L^{11}\tau$,

$$* = [*_{S^7}, L^3 f] \xrightarrow{[\text{nullh.}, L^3 f]} [[L^5[f, g], L^3 g], L^3 f] \xrightarrow{\text{Jacobi}} [[L^5[f, g], L^3 f], L^3 g] \xrightarrow{[\text{nullh.}, L^3 g]} [*_{S^7}, L^3 g] = *,$$

keeping track of the sizes of the intermediate maps and their homotopies. Clearly all the Whitehead products have Lipschitz constant at most L . Since the Jacobi identity is given by a universal construction, it can be done in linear space and time in terms of the Lipschitz constants of the entries. The nullhomotopy of $[L^5[f, g], L^3 g]$ can also be done in linear space and time using the composition

$$D^8 \xrightarrow{\tau} S^5 \times S^3 \xrightarrow{\sigma_5 \times \sigma_3} S^5 \times S^3 \rightarrow \mathbf{NF},$$

where τ is the attaching map of the top cell and σ_5 and σ_3 are maps of degree L^5 and L^3 , respectively. Likewise with the rest of the homotopies, which are also nullhomotopies of Whitehead products. This shows that τ has distortion $O(k^{1/11})$; the analysis of the minimal model in §3.3 shows that this is the best one can do.

Indeed, when one looks for homotopy group elements which are not generalized Whitehead products, such “nullhomotopies of Whitehead products in two different ways” come up naturally. It seems possible that one can build universal models for all rational homotopy classes (that is, all “higher rational homotopy operations”, as in [Blanc]) by an inductive application of this method.

Open problem. Can one prove Conjecture B for all spaces by applying self-maps and similar geometric methods to inductively built models?

Symmetric spaces. Our universal constructions generalize Gromov’s Hopf invariant example in one direction; we also generalize it in another, to a more general class of spaces that have self-maps with the right properties.

Theorem 5.3. *Suppose that the finite complex Y has the rational homotopy type of a Riemannian symmetric space. Then for any $\alpha \in \pi_n(Y)$, the distortion function is $\Theta(k^{1/(n+1)})$ if α is in the kernel of the Hurewicz map and $\Theta(k^{1/n})$ otherwise.*

Note that it is not clear whether the theorem contains any new results beyond the previous ones. Symmetric spaces are formal, meaning that their rational homotopy type is determined by their cohomology. From looking at presentations of the rational cohomology of nearly all symmetric spaces, it appears that their homotopy classes can always be represented as generalized Whitehead products. Nevertheless, there is also no obvious reason why this should be the case.

Open problem. Are all homotopy classes of symmetric spaces contained in generalized Whitehead product sets? Can this be shown other than by exhaustion?

The proof of the theorem heavily uses the fact that symmetric spaces admit a splitting homomorphism of algebras $H^*(Y; \mathbb{R}) \rightarrow \Omega^*(Y)$, induced by the harmonic forms. There has been some study of when the harmonic forms specifically induce such a splitting for Y a manifold [Kot], but besides formality it is not clear what the requirements are for such a splitting to exist.

Open problem. Give a topological characterization of all simplicial complexes Y for which the quotient map $\Omega^*(Y) \rightarrow H^*(Y; \mathbb{R})$ admits a splitting as a homomorphism of algebras.

It would not surprise me if the homotopy groups of such spaces were always generated by generalized Whitehead products.

Proof of Theorem 5.3. We use two topological properties of symmetric spaces. First, the indecomposables of the minimal model of a symmetric space split as $W_0 \oplus W_1$ where $W_0 = \ker d$ and $dW_1 \subset \mathbb{R}\langle W_0 \rangle$. This is true for all homogeneous spaces; see [FHT, §15(f)]. Second, symmetric spaces are *geometrically formal* [Kot]; that is, products of harmonic forms are harmonic, so in particular there is an algebra homomorphism $H^*(Y; \mathbb{R}) \rightarrow \Omega^*(Y)$. This induces a minimal model $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*(Y)$ such that $m_Y(y)$ is nonzero if and only if y is cohomologically nontrivial. This is the property we actually use.

Let us take a splitting $W_0 \oplus W_1$ as above. Note that there is an automorphism $\rho_L : \mathcal{M}_Y^* \rightarrow \mathcal{M}_Y^*$ which sends an indecomposable

$$w \mapsto \begin{cases} L^{\deg w} w & w \in W_0 \\ L^{\deg w + 1} w & w \in W_1. \end{cases}$$

Now, suppose $\alpha \in \pi_n(Y)$ is in the kernel of the Hurewicz map, and let $f : S^n \rightarrow Y$ be a representative of α . Then f^*m_Y is zero in degree n . By the method of §3.3, we build a homotopy $\Phi : \mathcal{M}_Y^* \rightarrow \Omega^*S^n$ from f^*m_Y to a map which sends $y \mapsto 0$ when $\deg y \leq n$ and $W_1 \rightarrow \mathbb{R}d \text{ vol}$ via the double dual α'' . Then $\Phi \circ \rho_L$ is a homotopy between $(L^{n+1}\alpha)''$ and a map φ_L whose image in degree k has operator norm $O(L^k)$, since it sends $W_1 \rightarrow 0$. Applying the shadowing principle, we get a CL -Lipschitz map $f_L : S^n \rightarrow Y$ whose homotopy class is $L^{n+1}\alpha$. This proves the theorem. \square

5.2. Lipschitz homotopies. The major application of the shadowing principle is in turning algebraic homotopies into geometric ones. Here previous geometric bounds were poor or nonexistent, and just producing some new ones is a big result. We produce several new results using the same schema:

- (1) Construct an algebraic homotopy between two maps $f, g : X \rightarrow Y$, with a bound on dilatation determined by their Lipschitz constants.
- (2) Concatenate this homotopy with an algebraic self-homotopy of g so that the result is homotopic rel ends to a genuine homotopy between f and g . This may increase the dilatation by an amount depending on the homotopy class of g .
- (3) Finally, apply the shadowing principle to obtain a genuine homotopy.

The latter two steps are encapsulated in the technical Theorem 5.6.

This schema can prove a number of different results, depending on the bound achieved in the first step. The most general and easily stated such result is as follows.

Theorem 5.4. *Let Y be a finite simply connected complex and X a finite complex of dimension n .*

- (i) *There are constants $C(X, Y)$ and $p(X, Y)$ such that any homotopic L -Lipschitz maps $f \sim g : X \rightarrow Y$ are homotopic via a $C(L + 1)^p$ -Lipschitz homotopy, which can in addition be taken to have length C .*
- (ii) *Moreover, any nullhomotopic L -Lipschitz map is nullhomotopic via a homotopy of length C and thickness CL^2 . (In particular, this is true for general homotopies if X has the rational homotopy type of a suspension.)*

Remarks. (a) Calder and Siegel [CaSi] and again Ferry and Weinberger [FW] gave proofs that constant-length homotopies can always be obtained in this context, but without any geometric bounds in the space direction. In this sense only the simultaneous bound on thickness is new.

- (b) Theorem 5.4 is a stronger statement than Conjecture C as given by Gromov, since Gromov did not ask for a bound on lengths of homotopies. On the other hand, since we give nonlinear thickness, it is weaker than Conjectures 2 and 3 in [CMW].

- (c) Part (ii) gives an almost sharp estimate of $O(L^{2n})$ on the volume of the nullhomotopy: for any completely general bound on length and thickness, we must have

$$(\text{length}) \cdot (\text{thickness})^n = \Omega(L^{2n-2}).$$

This is demonstrated by a sequence of examples first given in [CMW, §7.1]. Let X_n be a space constructed by attaching $(n+1)$ -cells to $S^2 \vee S^2$ to kill $\pi_n(S^2 \vee S^2) \otimes \mathbb{Q}$. Then $\pi_n(X_n)$ is finite, but since the generators of $\pi_n(S^2 \vee S^2) \otimes \mathbb{Q}$ have distortion $\sim k^{1/(2n-2)}$, we can find L -Lipschitz maps $S^n \rightarrow X_n$ for which every nullhomotopy has degree $\Omega(L^{2n-2})$ over some $(n+1)$ -cell.

In a more refined sense, the estimate is sometimes sharp: one cannot decrease the degree of the thickness bound while retaining constant length. This can be seen for maps $S^3 \rightarrow S^2$. Consider a nullhomotopic, $\Theta(L)$ -Lipschitz such map f which sends a solid torus inside S^3 to S^2 via a map whose cross-section has degree L^2 and which is constant on the circular fibers, and sends the complementary solid torus to the south pole of S^2 . Let C be a circle on the bounding torus which links nontrivially (hence with linking number L^2) with the preimage of the north pole. Then any nullhomotopy of f must have relative degree L^2 on $C \times [0, 1]$; therefore, if its length is constant, its thickness must be $\Omega(L^2)$. I would like to thank Sasha Berdnikov for pointing out this argument.

- (d) Moreover, in the case of maps $S^4 \times S^3 \rightarrow S^4$, the results of [CMW, §7.2] show that the exponent p from (i) cannot be less than $8/3$. Thus the bound of (ii) does not in general hold for non-nullhomotopies.
- (e) On the other hand, the estimate (ii) can be improved in various ways if we know more about the rational homotopy type of Y . For example, if Y is rationally k -connected, we can take the degree of the thickness bound to be $1 + 1/k$.
- (f) While the estimates (i) and (ii) look similar, they are actually different in certain crucial respects. The nullhomotopy estimate is very soft, using only facts about DGAs. On the other hand, the estimate for homotopies is actually false on the level of DGAs; indeed homotopies may be unbounded in the size of the original map. This can already be seen for homomorphisms modeling maps $S^4 \times S^3 \rightarrow S^4$, in terms of minimal models

$$\langle a^{(4)}, b^{(7)} \mid da = 0, db = a^2 \rangle \rightarrow \langle x^{(3)}, y^{(4)}, z^{(7)} \mid dx = dy = 0, dz = y^2 \rangle.$$

For small $\varepsilon > 0$, the pairs of homomorphisms $a \mapsto \varepsilon y$, $b \mapsto \varepsilon^2 z$ and $a \mapsto \varepsilon y$, $b \mapsto \varepsilon^2 z + xy$, which have norm bounded independent of ε , are homotopic via the homotopy

$$\begin{aligned} a &\mapsto \varepsilon y + (2\varepsilon)^{-1}x \otimes dt \\ b &\mapsto \varepsilon^2 z + xy \otimes t, \end{aligned}$$

whose size increases without bound as $\varepsilon \rightarrow 0$. Indeed, any homotopy must correctly resolve the obstruction class $(2\varepsilon)^{-1}x$ in degree 3. Thus to prove the polynomial bound we need to use the integral structure of the set of homotopy classes between X and Y . Unfortunately, the explicit estimate on the degree of the polynomial goes out the window in the process.

Our next theorem replicates and generalizes the results of [CDMW] and [CMW]. To do this requires a definition of spaces with positive weights, discussed in [BMSS]. A simply connected space Y has (\mathbb{Q} -)positive weights if the indecomposables of its minimal DGA split as $U_1 \oplus U_2 \oplus \cdots \oplus U_r$ so that for every $t \in \mathbb{Q}$ there is an automorphism φ_t sending $v \mapsto t^i v$, $v \in U_i$. Examples of spaces with \mathbb{Q} -positive weights include formal spaces [Shiga], coformal spaces [Dou], as well as homogeneous spaces and other “depth two” spaces whose indecomposables split as $W_1 \oplus W_2$, where $dW_1 = 0$ and $dW_2 \subset \bigwedge W_1$. A nonexample is a complex given in [MiT] which is constructed by attaching a 12-cell to $S^3 \vee \mathbb{C}P^2$; other, much higher-dimensional non-examples are given in [ArLu] and [Am].

For the below theorem, it is enough that Y has an $(n+1)$ -connected map to a space with positive weights. For example, the examples in Remark (c) above fit the bill since they have $(n+1)$ -connected maps to coformal spaces.

Theorem 5.5. *Suppose Y is a finite simply connected complex with positive weights equipped with automorphisms φ_t and X is a finite complex of dimension n . Then there are constants $C_1(n, Y)$ and $C_2(X)$ such that any nullhomotopic L -Lipschitz map $f : X \rightarrow Y$ has a nullhomotopy of length $C_1 C_2 (L+1)^d$ and thickness $C_1(L+1)$. Here d is the number of levels in a filtration of the indecomposables of $\mathcal{M}_Y^*(n)$,*

$$0 = W_0 \subset W_1 \subset \cdots \subset W_d \text{ with } \bigwedge W_d = \mathcal{M}_Y^*(n),$$

such that $dW_j \subseteq \bigwedge W_{j-1}$ and such that there is a basis for the indecomposables V_k in each degree k such that the subspaces $W_j \cap V_k$ for each j and

$$U_i \cap V_k = \{v \in V_k : \varphi_t(v) = t^i v\}$$

for each i are generated by subbases.

- Remarks.** (a) The number d is bounded above by $n-1$ since we can always take W_j to consist of all the indecomposables of $\mathcal{M}_Y^*(j+1)$. On the other hand, sometimes it can be much smaller: for example, for homogeneous spaces or any of the other depth two spaces discussed in [CMW], we can always choose $d=2$.
- (b) It remains unclear whether such linearly thick nullhomotopies are achievable when Y does not have positive weights. The example of [MiT] has an extremely complicated DGA; it may be worth looking for an example whose Sullivan model has a simpler presentation to test whether there is an obstruction to linear thickness. Nevertheless, the requirement that the space be neither formal nor coformal already forces a certain amount of complexity in the algebra.
- (c) Unlike the previous theorem, this one gives a sharp asymptotic estimate on the volume of nullhomotopies for the examples of [CMW, §7.1]. On the other hand, as we see in Theorem 5.7, it is not sharp in the case of maps $S^m \rightarrow S^{2n}$.

We now state the technical result which we use to convert estimates on algebraic homotopies to geometric ones.

Theorem 5.6. *Let Y be a simply connected compact Riemannian manifold with a minimal model $m_Y : \mathcal{M}_Y^* \rightarrow \Omega^*(Y)$ and X an n -dimensional finite simplicial complex with the standard metric. Let $f, g : X \rightarrow Y$ be homotopic Lipschitz maps and let*

$$\Phi : \mathcal{M}_Y^* \rightarrow X \otimes \mathbb{R}\langle t, dt \rangle$$

*be an algebraic homotopy between f^*m_Y and g^*m_Y with $\text{Dil}_\tau(\Phi) \leq \sigma$ for some τ and σ . Then for every $\alpha \in [0, 1]$ there is a homotopy between f and g of length $C(L_\lambda + 1)$, where*

$$L_\lambda = \sigma\tau + P(\|[f]\|_{\text{Lip}})^\alpha$$

and thickness $C(L_\theta + 1)$, where

$$L_\theta = \max\{\sigma, P(\|[f]\|_{\text{Lip}})^{1-\alpha}\},$$

and where C depends on Y , n , and the minimal model and P is a polynomial depending on X and Y .

Proof. We would like to construct a controlled homotopy by applying the shadowing principle to the pair

$$(L_\theta X \times [0, L_\lambda], L_\theta X \times \{0, L_\lambda\}),$$

homotoping an uncontrolled homotopy between f and g to be close to the realization

$$\rho\Phi : \mathcal{M}_Y^* \rightarrow \Omega^*(L_\theta X \times [0, L_\lambda]).$$

If we can do this, then we're done; but for this, the uncontrolled homotopy has to be actually DGA homotopic to $\rho\Phi$ rel the ends of the interval. Such a homotopy may not exist.

To resolve this, we concatenate $\rho\Phi$ with an algebraic self-homotopy of f^*m_Y ; this creates a new DGA map which is DGA homotopic to an honest homotopy. The numbers L_λ and L_θ are obtained by combining the measurements of Φ (algebraic ‘‘thickness’’ σ and ‘‘length’’ $\sigma\tau$) with those of this self-homotopy.

Let $H : X \times [0, 1] \rightarrow Y$ be an uncontrolled homotopy between f and g , and let $\Psi : \mathcal{M}_Y^* \rightarrow \Omega^*(X \times S^1)$ be a homomorphism which restricts to $\rho\Phi$ on one half of the circle and H^*m_Y on the other. Now let $\tilde{\Psi} : \mathcal{M}_Y^* \rightarrow \Omega^*X \otimes \mathbb{R}\langle e^{(1)} \rangle$ be obtained using Prop. 3.5 to lift through the diagram

$$\begin{array}{ccc} & \Omega^*X \otimes \mathbb{R}\langle e \rangle & \xrightarrow{e \mapsto 0} \Omega^*X. \\ & \tilde{\Psi} \nearrow & \downarrow e \mapsto d\theta \\ \mathcal{M}_Y^* & \xrightarrow{\Psi} \Omega^*(X \times S^1) & \searrow \Big|_{\theta=0} \end{array}$$

Next, let

$$\Xi = f^*m_Y + \eta \otimes e : \mathcal{M}_Y^* \rightarrow \Omega^*X \otimes \mathbb{R}\langle e \rangle$$

be a homomorphism such that

$$[\Xi \boxplus \tilde{\Psi}] \in [\mathcal{M}_Y^*, \Omega^*(X) \otimes \mathbb{R}\langle e \rangle]_{f^*m_Y}$$

is the rationalization of a class in $\pi_1(Y^X; f)$; let $F : X \times [0, 1] \rightarrow Y$ be a (uncontrolled) self-homotopy of f which is a representative of that class. By [MW, Lemma 5.2(i)] and the surrounding discussion, we can pick Ξ to be of polynomial length, i.e. so that for each k , $\|\eta|_{V_k}\|_{\text{op}} \leq P(\|[f]\|_{\text{Lip}})$, where P is a polynomial depending on X and Y . Note also that $\rho(f^*m_Y + \eta \otimes t)$ has dilatation ≤ 1 when scaled to be a map

$$\mathcal{M}_Y^* \rightarrow \Omega^*(\max\{\text{Lip } f, P(\|[f]\|_{\text{Lip}})^{1-\alpha}\}X \times [0, P(\|[f]\|_{\text{Lip}}^\alpha)]),$$

for any $\alpha \in [0, 1]$. Now we define

- $\tilde{\Phi} : \mathcal{M}_Y^* \rightarrow \Omega^*(L_\theta X \times [0, L_\lambda])$ to be the homomorphism obtained by concatenating appropriately scaled versions of $\rho(f^*m_Y + \eta \otimes t)$ and $\rho\Phi$;
- $\tilde{H} : L_\theta X \times [0, L_\lambda] \rightarrow Y$ to be the map obtained by concatenating scaled versions of F and H .

Now $\tilde{\Phi}$ and \tilde{H}^*m_Y are homotopic rel ends because concatenating either of them on both sides with H^*m_Y creates a map $\mathcal{M}_Y^* \rightarrow \Omega^*(X \times S^1)$ which is homotopic rel $X \times \{\theta = 0\}$ to $\rho(\Xi \boxplus \tilde{\Psi})$. Moreover, $\text{Dil}(\tilde{\Phi}) \leq 1$. Applying the shadowing principle to $\tilde{\Phi}$ and \tilde{H} therefore gives the controlled homotopy we desire. \square

We now prove Theorems 5.4 and 5.5.

Proof of Theorem 5.4. We first handle the case of nullhomotopies. First, notice that the map

$$[\mathcal{M}_Y^*(k+1), \Omega^*X \otimes \langle e \rangle]_0 \rightarrow [\mathcal{M}_Y^*(k), \Omega^*X \otimes \langle e \rangle]_0$$

is a surjection. This can be seen, for example, from the fact that maps in the latter group are only nonzero on indecomposables, and therefore the map

$$0 + \eta \otimes e \mapsto \eta|_{dV_{k+1}} \in H^{k+1}(X; V_{k+1})$$

is zero. So fix some uncontrolled algebraic nullhomotopy $\Psi : \mathcal{M}_Y^* \rightarrow \Omega^*X \otimes \langle t, dt \rangle$ of f^*m_Y , and let $\Phi : \mathcal{M}_Y^*(k) \rightarrow \Omega^*X \otimes \langle t, dt \rangle$ be a partial controlled nullhomotopy. The concatenation of $\rho(\Psi|_{\mathcal{M}_Y^*(k)})$

and $\rho\Phi$ is a map $\mathcal{M}_Y^*(k) \rightarrow \Omega^*(X \times S^1)$ whose restriction to $\theta = 0$ is zero. Therefore this map lifts to a map from $\mathcal{M}_Y^*(k+1)$ whose restriction to $\theta = 0$ is also zero; since Ψ is defined in $\mathcal{M}_Y^*(k+1)$, this implies that Φ also lifts.

So let $f : X \rightarrow Y$ be a nullhomotopic L -Lipschitz map. To choose a nullhomotopy Φ of f^*m_Y , we can use the following simple procedure. Set

$$\Phi : \mathbb{R} = \mathcal{M}_Y^*(1) \rightarrow \Omega^*(X) \otimes \mathbb{R}\langle t, dt \rangle$$

to be the trivial map. Then at the $(k+1)$ st stage, since there is no obstruction, we choose an extension as in Proposition 3.7(ii) to extend to $\mathcal{M}_Y^*(k+1)$. By (3.8), we get $\text{Dil}_{L^{-2}}(\Phi) \leq L^2$. By plugging this into Theorem 5.6, using $\alpha = 0$, we get an L^2 -Lipschitz nullhomotopy of f of constant length.

Now let $f \simeq g : X \rightarrow Y$ be non-nullhomotopic maps. In this general case, we still construct a homotopy

$$\Phi : \mathcal{M}_Y^* \rightarrow \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle$$

by lifting inductively from k to $k+1$. However, now the lift may be obstructed, so we will need to fix our partially-constructed homotopy

$$\Phi_k : \mathcal{M}_Y^*(k) \rightarrow \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle$$

using a self-homotopy of f^*m_Y . To get a bound on Φ_{k+1} in terms of Φ_k , we will use the following roadmap:

- (1) Using Proposition 3.7(i), estimate the size of the obstruction class in $H^{k+1}(\Omega^*X; V_{k+1})$ to extending the homotopy.
- (2) Find an element of $[\mathcal{M}_Y^*(k), \Omega^*X \otimes \mathbb{R}\langle e \rangle]_{f^*m_Y}$ which maps to this obstruction class, with an estimate on the size of a representative Ψ_k .
- (3) Algebraically concatenate the two homotopies; an estimate on the size of the new homotopy Φ'_k is provided by Proposition 3.9.
- (4) Finally, by Proposition 3.7(ii), Φ'_k lifts in a quantitative way to

$$\Phi_{k+1} : \mathcal{M}_Y^*(k+1) \rightarrow \Omega^*X \otimes \mathbb{R}\langle t, dt \rangle.$$

This provides an estimate on the dilatation of Φ_{k+1} in terms of that of Φ'_k .

Together, all these polynomial estimates will give a p such that $\text{Dil}_{L^{-p}} \Phi_k \leq L^p$.

We have explicit bounds on the degree to which steps (1), (3), and (4) distort the size of the homotopy. Step (2) is where we must use the fact that f and g are honest maps between spaces, and where we lose this explicit bound on the degree of the polynomial. The map

$$\mathcal{O} : [\mathcal{M}_Y^*(k), \Omega^*X \otimes \mathbb{R}\langle e \rangle]_{f^*m_Y} \rightarrow H^{k+1}(X, V_{k+1})$$

is given by $[f^*m_Y + \eta \otimes e] \mapsto [\eta|_{dV_{k+1}}]$. For a given homotopy class of f , this is a linear map in terms of the values of η on indecomposables. Moreover, if we fix a basis on the indecomposables and on $H^{k+1}(X, V_{k+1})$, the resulting matrix has entries given by polynomials in the values of f^*m_Y on indecomposables. Given a class c in the image of this matrix, we would like to find a bound on the minimal size of a preimage in terms of $\|c\|$.

Now, as observed by Sullivan, these groups are the result of tensoring

$$\pi_1((Y_{(k)})^X, f) \rightarrow H^{k+1}(X; \pi_{k+1}(Y))$$

with \mathbb{R} . Indeed, following [MW, §5], we can say somewhat more. If we equip $[\mathcal{M}_Y^*(k), \Omega^*X \otimes \mathbb{R}\langle e \rangle]_{f^*m_Y}$ with the norm which assigns to $\varphi + \eta \otimes e$ the operator norm of η restricted to the indecomposables of $\mathcal{M}_Y^*(k)$, then the map

$$\pi_1((Y_{(k)})^X, f) \rightarrow [\mathcal{M}_Y^*(k), \Omega^*X \otimes \mathbb{R}\langle e \rangle]_{f^*m_Y}$$

is $P_k(\text{Lip } f)$ -surjective for a polynomial P_k depending on X and Y , i.e. every point in the codomain is at most distance $P_k(\text{Lip } f)$ away from the image. In particular, this gives a basis \mathbf{b} for the image lattice whose vectors are polynomially bounded in terms of $\text{Lip } f$; we also have a polynomial bound on the vectors of $\mathcal{O}(\mathbf{b})$. Since the lattice is mapped to the image of $H^{k+1}(X; \pi_{k+1}(Y))$ in $H^{k+1}(X; V_{k+1})$, we also have a fixed lower bound (independent of f) on the $\dim(\text{im } \mathcal{O})$ -dimensional volume of the parallelotope spanned by $\mathcal{O}(\mathbf{b})$. This gives a polynomial lower bound on the shortest axis of this parallelotope. This completes step (2) and the proof. \square

Proof of Theorem 5.5. For spaces with positive weights, we can take advantage of automorphisms to give an alternate construction of an algebraic nullhomotopy. Fix a map $f : X \rightarrow Y$ and a family of automorphisms $\varphi_t : \mathcal{M}_Y^* \rightarrow \mathcal{M}_Y^*$ with the desired properties. Now, let $v \in V_k \cap U_i$ be an element of the basis mentioned in the statement of the theorem. Then we inductively define

$$\Phi(v) = f^* m_Y v \otimes t^i + c(v) \otimes it^{i-1} dt$$

where $c(v)$ is chosen so that $dc(v) = (-1)^{k+1} f^* m_Y v + c(dv)$. Here $c(dv)$ is defined by induction so that

$$\Phi(dv) = f^* m_Y(dv) \otimes t^i + c(dv) \otimes it^{i-1} dt;$$

we know $\Phi(dv)$ takes this form by positive weights and the definition of Φ on lower-degree indecomposables. Moreover, by the same argument as in the proof of Theorem 5.4(ii), there is no obstruction to finding such a $c(v)$.

The t^i -coefficients of Φ always have dilatation L ; moreover, by Lemma 2.4 and induction on j , the $t^{i-1} dt$ -coefficients $c(v)$ can be chosen so that for $v \in V_k \cap \mathcal{M}_j$, the operator norm is bounded by $C(X, Y) L^{\frac{k+j}{k}}$. This gives us $\text{Dil}_{C(X, Y) L^{d-1}} \Phi \leq L$; plugging this into Theorem 5.6, with $\alpha = 1$, gives the result. \square

5.3. Maps between spheres. The previous applications are to problems of great generality. But the shadowing principle can also be applied to yield new results in the much more specific situation of maps between spheres, beyond the results of [CDMW] and [CMW].

Sharper bounds on nullhomotopies. As noted before, the bounds of Theorem 5.5 are sharp for certain classes of examples. On the other hand, they turn out not to be sharp for example for maps $S^m \rightarrow S^n$. In that case, in the dimension range where Hopf invariants play a role—that is, when n is even and $m \geq 2n - 1$ —Theorem 5.5, for example, only yields a quadratic bound on length; but this seems to be an artifact of the way DGA homotopies are constructed, and one can do better by constructing and manipulating genuine homotopies. At the same time, the inductive approach given here forces the thickness of the homotopy to grow. I suspect that linear homotopies always exist in this situation but that finding them will require new tools. In fact, Sasha Berdnikov is now constructing such linear homotopies in the case $S^3 \rightarrow S^2$ using a purely geometric method.

Theorem 5.7. *Let n be even and $m \geq 2n - 1$. Then every L -Lipschitz nullhomotopic map $S^m \rightarrow S^n$ has a $\gamma(L)$ -Lipschitz nullhomotopy $S^m \times [0, 1] \rightarrow Y$, where $\gamma(L) = o(L^{1+\varepsilon})$ for every $\varepsilon > 0$.*

Proof. Let $\gamma(L)$ be the best possible such function. Fix some $\varepsilon > 0$; we will show using a recurrence relation that γ grows more slowly than $L^{1+\varepsilon}$.

Let $f : S^m \rightarrow S^n$ be a nullhomotopic L -Lipschitz map. The idea is to first nullhomotope a “slightly shrunken copy” of f . We then expand this nullhomotopy again to get a nullhomotopy of f . In a way, this is similar to Theorem 5.5; the key point is that in this case working with a genuine nullhomotopy lets us make this map not too much bigger.

Define

$$m_{S^n} : \mathcal{M}_{S^n}^* = \langle a^{(n)}, b^{(2n-1)} \mid da = 0, db = a^2 \rangle \rightarrow \Omega^* S^n$$

via $m_{S^n}(a) = d \text{ vol}$ and $m_{S^n}(b) = 0$, and let $C' = C'(m, S^n)$ be the constant given in the shadowing principle. We apply the principle to get a DGA homotopy

$$\Phi : \mathcal{M}_{S^n}^* \rightarrow \Omega^*(S^m) \otimes \mathbb{R}\langle t, dt \rangle$$

between the DGA map $\frac{1}{C'L^{\varepsilon n}} f^* m_{S^n}$ and a nullhomotopic $L^{1-\varepsilon}$ -Lipschitz map $g : S^m \rightarrow S^n$, such that $\text{Dil}_{C'L^{\varepsilon-1}} \Phi \leq L^{1-\varepsilon}$. This is possible since $f^* m_{S^n}$ is (algebraically) nullhomotopic, and hence so is its scaled version. Finally, we choose a $\gamma(L^{1-\varepsilon})$ -Lipschitz nullhomotopy $G : S^m \times [0, 1] \rightarrow Y$ of this g .

Now we produce a new homomorphism

$$\Psi : \mathcal{M}_{S^n}^* \rightarrow \Omega^*(S^m \times [0, 2])$$

satisfying $\Psi|_{t=0} = f^* m_{S^n}$ and $\Psi|_{t=1} = 0$, as well as $\text{Dil } \Psi \leq \max\{C'L^{1+\varepsilon}, C'L^\varepsilon \gamma(L^{1-\varepsilon})\}$, by setting

$$\Psi(a)|_{(x,s)} = \begin{cases} C'L^{n\varepsilon} \Phi(a)|_{t=s} & 0 \leq s \leq 1 \\ C'L^{n\varepsilon} G^* d \text{ vol}|_{(x,s-1)} & 1 \leq s \leq 2 \end{cases}$$

$$\Psi(b)|_{(x,s)} = \begin{cases} (C')^2 L^{2n\varepsilon} \Phi(a)|_{t=CL^\varepsilon s} & 0 \leq s \leq 1 \\ 0 & 1 \leq s \leq 2. \end{cases}$$

The key is the observation that since $G^* m_{S^n}(b) = 0$, the dilatation of $G^* m_{S^n}$ scales linearly as we expand. Meanwhile, Φ was small to begin with and so the fact that it scales superlinearly doesn't matter very much.

Finally, we apply the shadowing principle again to get a $C \text{ Dil } \Psi$ -Lipschitz nullhomotopy of f on this interval. Thus we obtain that

$$\gamma(L) \leq 2C \max\{C'L^{1+\varepsilon}, C'L^\varepsilon \gamma(L^{1-\varepsilon})\}.$$

Now choose L_0 large enough that $2CC' \leq (L_0)^{\varepsilon^2}$, and pick $K \geq 2CC'$ such that

$$\gamma((L_0)^{1-\varepsilon}) \leq K((L_0)^{1-\varepsilon})^{1+\varepsilon} = K(L_0)^{1-\varepsilon^2}.$$

Then $\gamma(L_0) \leq K(L_0)^{1+\varepsilon}$. By induction, we get $\gamma(L) = O(L^{1+\varepsilon})$. \square

Indeed, the only fact we used here about S^m is that $H^{n+1}(S^m)$ and $H^{2n}(S^m)$ are both zero; this allowed us to apply the shadowing principle to our algebraic nullhomotopy. Perhaps this requirement can be eliminated to obtain a similar bound for general $X \rightarrow S^n$. It seems plausible also that this method can be applied to various target spaces Y with positive weights to get bounds better than Theorem 5.5. This may be possible when Y has dimension considerably lower than X or even if it simply lacks homology in some relevant degrees.

Uniformity over the metric. In [Guth], Larry Guth asks the following question:

Question 5.8. *The n -dimensional ellipse with principal axes R_0, \dots, R_n is the set defined by*

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

Let E^m and F^n be m - and n -dimensional ellipses, respectively. If $f : E \rightarrow F$ is nullhomotopic and L -Lipschitz, 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 this be taken to be $C(m, n)L$ or $C(m, n)L^2$ as dictated by the rational homotopy? What about more complicated metrics on the sphere?

Such ellipses, and any metric on the sphere, can be closely approximated by simplicial complexes after sufficient scaling. Therefore Theorem 5.5 allows us to give a half-answer to this which is, however, less than half satisfying. As long as we fix the target metric on the sphere and L is larger

than some constant depending on the domain metric⁸, the nullhomotopy can go through maps of Lipschitz constant at most $C(F, m, n)L$ (if n is odd or $m < 2n - 1$) or $C(F, m, n)L^2$ (otherwise.) However, dependence on the target metric is a complete mystery, as it is in this whole paper.

Recent results in [FFWZ] suggest that the constants depending on the target space can be quite large even for relatively small target spaces in fixed dimension. On the other hand, it may be that things are less dire when we restrict to spaces of the same homeomorphism type.

REFERENCES

- [AA78] P. Andrews and M. Arkowitz, *Sullivan’s minimal models and higher order Whitehead products*, Canadian Journal of Mathematics **30** (1978), 961–982.
- [Am] Manuel Amann, *Degrees of self-maps of simply connected manifolds*, International Mathematics Research Notices (2015), no. 18, 8545–8589.
- [ArLu] Martin Arkowitz and Gregory Lupton, *Rational obstruction theory and rational homotopy sets*, Mathematische Zeitschrift **235** (2000), no. 3, 525–539.
- [Blanc] David Blanc, *Homotopy operations and rational homotopy type*, Algebraic Topology: Categorical decomposition techniques, Progress in Mathematics, no. 215, Birkhäuser, 2003, pp. 47–75.
- [BMSS] Richard Body, Mamoru Mimura, Hiroo Shiga, and Dennis Sullivan, *p -universal spaces and rational homotopy types*, Commentarii Mathematici Helvetici **73** (1998), no. 3, 427–442.
- [BrSz] Edgar Brown, Jr. and Robert Szczarba, *On the rational homotopy type of function spaces*, Transactions of the American Mathematical Society **349** (1997), no. 12, 4931–4951.
- [ČKMVW] Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner, *Extendability of continuous maps is undecidable*, Discrete & Computational Geometry **51** (2014), no. 1, 24–66.
- [CaSi] Allan Calder and Jerrold Siegel, *On the width of homotopies*, Topology **19** (1980), no. 3, 209–220.
- [CDMW] Gregory R Chambers, Dominic Dotterer, Fedor Manin, and Shmuel Weinberger, *Quantitative nullcobordism*, arXiv preprint arXiv:1610.04888 (2016).
- [CMW] Gregory R Chambers, Fedor Manin, and Shmuel Weinberger, *Quantitative nullhomotopy and rational homotopy type*, arXiv preprint arXiv:1611.03513 (2016).
- [Dou] Roy Douglas, *Positive weight homotopy types*, Illinois Journal of Mathematics **27** (1983), no. 4, 597–606.
- [DF] Emmanuel Dror Farjoun, *Homotopy and homology of diagrams of spaces*, Algebraic Topology: Proceedings of a Workshop held at the University of Washington, Seattle, 1985, Lecture Notes in Mathematics, vol. 1286, Springer, 1987, pp. 93–134.
- [EdGr] Herbert Edelsbrunner and Daniel R Grayson, *Edgewise subdivision of a simplex*, Discrete & Computational Geometry **24** (2000), no. 4, 707–719.
- [FF60] Herbert Federer and Wendell H Fleming, *Normal and integral currents*, Annals of Mathematics (1960), 458–520.
- [FHT] Yves Félix, Steve Halperin, and Jean-Claude Thomas, *Rational homotopy theory*, Graduate Texts in Mathematics, vol. 205, Springer, 2012.
- [FW] Steve Ferry and Shmuel Weinberger, *Quantitative algebraic topology and lipschitz homotopy*, Proceedings of the National Academy of Sciences **110** (2013), no. 48, 19246–19250.
- [FFWZ] Marek Filakovský, Peter Franek, Uli Wagner, and Stephan Zhechev, *Computing simplicial representatives of homotopy group elements*, arXiv preprint arXiv:1706.00380 (2017).
- [FiVo] Marek Filakovský and Lukáš Vokřínek, *Are two given maps homotopic? An algorithmic viewpoint*, arXiv preprint arXiv:1312.2337 (2013).
- [GrMo] Phillip A. Griffiths and John W. Morgan, *Rational homotopy theory and differential forms*, Birkhäuser, 1981.
- [GrHED] Mikhail Gromov, *Homotopical effects of dilatation*, Journal of Differential Geometry **13** (1978), no. 3, 303–310.
- [GrMS] ———, *Metric structures for Riemannian and non-Riemannian spaces*, vol. 152, Birkhäuser Boston, 1998.
- [GrQHT] ———, *Quantitative homotopy theory*, Invited Talks on the Occasion of the 250th Anniversary of Princeton University (H. Rossi, ed.), Prospects in Mathematics, 1999, pp. 45–49.
- [Guth] Larry Guth, *Recent progress in quantitative topology*, <http://math.mit.edu/~lguth/Exposition/quanthomotexpos2.pdf>, 2017.

⁸Roughly the inverse of the mesh size.

- [Hain] Richard M. Hain, *Iterated integrals and homotopy periods*, Memoirs of the American Mathematical Society, no. 291, January 1984.
- [Kot] Dieter Kotschick, *On products of harmonic forms*, Duke Mathematical Journal **107** (2001), no. 3, 521–531.
- [M] Fedor Manin, *Asymptotic invariants of homotopy groups*, Ph.D. thesis, University of Chicago, June 2015.
- [MW] Fedor Manin and Shmuel Weinberger, *Integral and rational mapping classes*, arXiv preprint arXiv:1802.05784 (2018).
- [MiT] Mamoru Mimura and Hiroshi Toda, *On p -equivalences and p -universal spaces*, Commentarii Mathematici Helvetici **46** (1971), no. 1, 87–97.
- [Nab] Alexander Nabutovsky, *Non-recursive functions, knots “with thick ropes,” and self-clenching “thick” hyperspheres*, Communications on Pure and Applied Mathematics **48** (1995), no. 4, 381–428.
- [NaTo] Minoru Nakaoka and Hiroshi Toda, *On Jacobi identity for Whitehead products*, Journal of the Institute of Polytechnics, Osaka City University. Series A: Mathematics **5** (1954), no. 1, 1–13.
- [Scull] Laura Scull, *Rational S^1 -equivariant homotopy theory*, Transactions of the American Mathematical Society **354** (2002), no. 1, 1–45.
- [Shiga] Hiroo Shiga, *Rational homotopy type and self maps*, Journal of the Mathematical Society of Japan **31** (1979), no. 3, 427–434.
- [Sul] Dennis Sullivan, *Infinitesimal computations in topology*, Publications Mathématiques de l’IHES **47** (1977), no. 1, 269–331.

DEPARTMENT OF MATHEMATICS, OHIO STATE UNIVERSITY, COLUMBUS, OHIO, USA
E-mail address: `manin@math.toronto.edu`