Area-Contracting Maps Between Rectangles

by

Lawrence Guth

Bachelor of Science, Yale University, June 2000

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2005

Author ....
Department of Mathematics
April 25, 2005

Certified by ....
Tomasz S. Mrowka
Professor of Mathematics
Thesis Supervisor

Accepted by ....
Pavel Etingof
Chairman, Departmental Committee on Graduate Students

Area-Contracting Maps Between Rectangles

by

Lawrence Guth

Submitted to the Department of Mathematics
on April 25, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this thesis, I worked on estimating the smallest k-dilation of all diffeomorphisms between two n-dimensional rectangles R and S. I proved that for many rectangles there are highly non-linear diffeomorphisms with much smaller k-dilation than any linear diffeomorphism. When k is equal to n-1, I determined the smallest k-dilation up to a constant factor.

For all values of k and n, I solved the following related problem up to a constant factor. Given n-dimensional rectangles R and S, decide if there is an embedding of S into R which maps each k-dimensional submanifold of S to an image with larger k-volume.

I also applied the k-dilation techniques to two purely topological problems: estimating the Hopf invariant of a map from a 3-manifold to a high-genus surface, and determining whether there is a map of non-zero degree from a 3-manifold to a hyperbolic 3-manifold.

Thesis Supervisor: Tomasz S. Mrowka
Title: Professor of Mathematics

Acknowledgments

I would like to thank all of my math teachers, including Michael Narosny, Deanna Wong, Joanne Lund, Priscilla Burbank-Schmitt, Mike Frantz, Terry Leverich, Serge Lang, Peter Jones, and my thesis advisor Tom Mrowka. They went beyond the call of duty, sharing their ideas and experiences with me.

I think that I put Tom in a somewhat awkward position. Instead of studying his field of expertise, I insisted pretty stubbornly on studying a different field in which few people work, and after three years I had no results. Tom gently made sure that I would graduate. (For the first few months of my fourth year, we met together privately for almost two hours a week.)

Tom was happy to talk about any kind of problem in geometry. Sometime in my second year, my friend Max revealed to him in the common room that I had spent several months thinking about crumpling paper, and so we once talked about that. When I would pass him in the hall, he would say "So?", and his eyebrows would rise in anticipation. Sometimes he added a few more words, which made the question clearer: "anything new and exciting?" For several years, I had nothing to say, but Tom continued to ask.

I thank Andre Henriques and Bobby Kleinberg for interesting conversations about math. Daniel Biss helped me figure out a topology problem that appears in section 7.8. Haynes Miller gave me some references to results in homotopy theory that I needed in section 7.9.

I want to thank my friend Ben Stephens for many kinds of support. We had conversations about math and about our lives. He edited the introduction of this thesis. In the difficult month of writing before the thesis deadline, he often cooked me dinner.

My most important role model as a scientist is my father. I want to thank him for teaching and encouraging me for twenty seven years. One of the little pleasures of my time at MIT was using the same office that my father had when he was a graduate student before I was born.

Finally, I want to thank my whole family for their love and support. Even with all this help, I was often frustrated in graduate school, and occasionally I thought I might have made the wrong choice. I also wondered what it would have been like to work in a well-established field, frequently talking with other students and more senior people in the same field as me. When I felt down, I would talk with my family at the kitchen table about the whole experience, with its good and bad parts, and I generally felt much better afterwards.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis is about estimating the k-dilation of mappings. The k-dilation of a smooth mapping measures how much the mapping stretches the k-dimensional volume of k-dimensional submanifolds in the domain. We say that a mapping has k-dilation at most λ if it maps every k-dimensional manifold of volume V to an image of volume at most λV. The k-dilation of f can be defined algebraically as the supremal value of the norm of $\Lambda^k df$. A smooth map is called k-contracting if its k-dilation is less than or equal to 1.

The 1-dilation of a smooth map is equal to its Lipshitz constant. We can think of the k-dilation as a variation on the Lipshitz constant, which measures how much k-dimensional volumes stretch instead of how much distances and lengths stretch.

We can now state the main problem of the thesis.

Problem 1. *Given two n-dimensional rectangles R and S, estimate the smallest k-dilation of all diffeomorphisms from R to S.*

When I first thought of this problem, I expected the linear diffeomorphism from R to S to have the smallest k-dilation, or at least close to the smallest k-dilation. Clearly, the linear diffeomorphism from R to S has the smallest n-dilation of any diffeomorphism. With a little more work, one can show that, up to a constant factor, the linear diffeomorphism has the smallest 1-dilation of any diffeomorphism from R to S. Also, if T_1 and T_2 are two flat tori, then each homotopy class of maps from T_1

to T_2 contains a linear map, and this linear map has the smallest k-dilation of all the maps in the homotopy class. It turns out, though, that linear diffeomorphisms can have far from the smallest k-dilation.

Theorem 1.1. *If k lies in the range $1 < k < n$, then there are n-dimensional rectangles R and S so that there is a k-contracting diffeomorphism from R to S and yet every linear diffeomorphism from R to S has arbitrarily high k-dilation.*

The 1-dilation and the n-dilation are probably more familiar to most people than the k-dilation for other values of k. This theorem shows that the k-dilation for the other values of k can be more complicated than the 1-dilation or the n-dilation.

To prove this proposition, we will construct some non-linear diffeomorphisms with small k-dilation, which we will call snake maps. For example, in three dimensions, we will construct a highly non-linear diffeomorphism from the rectangle $[0, \epsilon] \times [0, 1] \times [0, 1]$ to the rectangle $[0, \epsilon] \times [0, \epsilon] \times [0, \epsilon^{-1}]$ with 2-dilation less than 10,000. The 2-dilation of a linear map between these rectangles is ϵ^{-1}. Therefore, when ϵ is very small, our non-linear diffeomorphism has much smaller 2-dilation than any linear diffeomorphism.

It turns out that when k is equal to n-1, there is always a combination of linear diffeomorphisms and snake maps which gives the smallest possible k-dilation, up to a constant factor. As a result, we can solve Problem 1 when k is equal to n-1, up to constant factor.

To state the result, we make some conventions that we will use all through the thesis. If R is an n-dimensional rectangle, then we refer to the dimensions of R by R_i, for i from 1 to n. We always order the dimensions so that $R_1 \leq ... \leq R_n$.

Throughout the thesis, C and c will denote positive constants that depend only on the dimension n. Their value may change from line to line. When we use the letter c, the reader should imagine a very small number, and when we use the letter C, the reader should imagine a very large number. All the constants can be made explicit, but they are pretty bad.

Theorem 1.2. *Let R and S be n-dimensional rectangles. Define Q_i to be the quotient S_i/R_i. For each integer l in the range $1 \leq l \leq n-1$, we define a number $D(l)$ as* *follows.*

$$D_l = Q_1 ... Q_l (Q_{l+1} ... Q_n)^{\frac{n-l-1}{n-l}}.$$

Also, we define D_n to be $Q_2 ... Q_n$. Finally, we define D to be the largest of all the numbers D_l, for l from 1 to n.

Any diffeomorphism from R to S has (n-1)-dilation at least cD. On the other hand, we will construct a diffeomorphism from R to S with (n-1)-dilation less than CD.

The idea of k-width is very useful for estimating k-dilations. Roughly speaking, if we slice a Riemannian manifold M into k-dimensional slices, then the k-width of M is less than W if we can arrange that each of these slices has k-volume less than W. A little more formally, the k-width of a Riemannian n-manifold M is defined to be the infimal W so that there is a nice map from M to $\mathbb{R}^{n-k}$ each of whose fibers has k-volume less than W. (A nice map is defined to be a piecewise linear map whose restriction to each face has maximal rank. The reason we need to restrict to nice maps is that I don't even know how to prove that every smooth map from the unit n-ball to $\mathbb{R}^{n-k}$ has a fiber with k-volume at least ϵ.)

Since the k-width is defined using an infimum over a very large space of maps, it is not obvious how to estimate the k-width even for simple sets. For instance, the k-width of the unit n-cube is clearly at most 1, but is it actually 1? I don't know how to solve this problem. There are theorems of Almgren about families of cycles which show that the k-width of the unit n-sphere is exactly the volume of the unit k-sphere. As a corollary, one can give a lower bound for the k-width of the unit cube. There is also a more elementary argument due to Gromov which gives a lower bound for the k-width of the unit cube. (Both arguments appear in appendix 1 of [12].) Using one of these techniques, it is not hard to show that the k-width of a rectangle R is roughly $R_1 ... R_k$.

What can we say about the k-width of a very complicated set? Of course, if the set fits inside of the unit cube, its k-width is less than 1. Suppose, however, that we are dealing with a very diffuse set, consisting of thin tubes and membranes, but with volume less than 1. What can we say about the k-width of such a set? My geometric intuition suggests that it should be possible to fold such a set into the cube of side length 10 without distorting its internal geometry very much - just as one squeezes the water out of a sponge. If you could perform such a folding operation, you would get an upper bound for the k-width of the set. In fact, I don't know whether or not it is possible to perform this folding, but it still turns out that the k-width of a set is bounded by its volume.

Theorem 1.3. *If U is a bounded open set in $\mathbb{R}^n$, then the k-width of U is less than* $C\mathrm{volume}(\mathrm{U})^{\mathrm{k/n}}$.

The lower bounds for k-dilation in Theorem 1.2 come from understanding how many disjoint k-wide sets fit in a rectangle. For example, how many disjoint sets of 2-width (1/100) can fit into the unit cube in $\mathbb{R}^3$? More generally, if U_i are disjoint subsets contained in a rectangle R, what is the largest possible value of the sum $(\sum_i \text{k-width}(U_i)^q)^{1/q}$, where q is a real number between n/k and ∞? Using the width-volume inequality, we will answer this question. For each value of q, we will calculate a supremal value up to a constant factor C. The answer is a little painful to write in closed form, but it is given by taking the sets U_i to be disjoint sub-rectangles with dimensions $R_1 \times ... \times R_l \times R_l \times ... \times R_l$, for an integer l which depends on q. As q goes from n/k to ∞, this supremum interpolates between the volume of R and the k-width of R. The dependence on q is somewhat analogous to the L^q norm, which interpolates between the integral of a positive function and its supremum.

Our next theorem is a variation of the isoperimetric inequality which involves a sum of k-widths raised to a power.

Theorem 1.4. *If U is a bounded open set in $\mathbb{R}^n$ with smooth boundary, then there are disjoint subsets S_i in the boundary of U so that the following inequality holds.*

$$\mathrm{Volume(U)} < \mathrm{C} \sum_{\mathrm{i}} \textit{(n-2)-width}(\mathrm{S_i})^{\frac{n}{n-2}}.$$

To get an idea of what this theorem means, let us consider the case that the open set U is a rectangle R. The volume of U is equal to $R_1 ... R_n$. Following the analysis in the last full paragraph, we can estimate the right-hand side of this inequality up to a constant factor. If n is even, the right-hand side is less than $CR_2^{\frac{n}{n-2}} ... R_{n/2}^{\frac{n}{n-2}} R_{n/2+1} ... R_n$, and if n is odd, the right-hand side is less than $R_2^{\frac{n}{n-2}} ... R_{(n-1)/2}^{\frac{n}{n-2}} R_{(n+1)/2}^{\frac{n-1}{n-2}} R_{(n+3)/2} ... R_n$. For comparison, the standard isoperimetric inequality gives the estimate that the volume of R is less than $C(R_2 ... R_n)^{\frac{n}{n-1}}$. If the dimensions of R are very different from one another, then our isoperimetric inequality is sharper than the standard one.

The fact remains that our isoperimetric inequality gives an estimate which is much larger than the actual volume of R. To understand why we cannot expect our isoperimetric inequality to give a sharp estimate, let us first consider the standard isoperimetric inequality. The standard isoperimetric inequality gives a terrible estimate for the volume of long, thin rectangles because it is only allowed to consider the volume of the boundary. The standard isoperimetric inequality is really estimating the largest volume that can be enclosed by any hypersurface with the same volume as the boundary of R - which gives a much bigger number than the volume of R. Our isoperimetric inequality gives a large answer for the same reason. It only involves the (n-2)-widths of subsets of the boundary of R. Therefore, it is really estimating the largest volume that can be enclosed by an (n-2)-contracting embedding of the boundary of R into $\mathbb{R}^n$.

We will construct an (n-2)-contracting embedding from the boundary of R into $\mathbb{R}^n$ that shows that our isoperimetric inequality, considered from this point of view, is sharp up to a constant factor. If n is even, we will construct an (n-2)-contracting embedding which encloses a region of volume $cR_2^{\frac{n}{n-2}} ... R_{n/2}^{\frac{n}{n-2}} R_{n/2+1} ... R_n$. An analogous statement holds when n is odd.

This discussion suggests our next problem.

Problem 2. *Estimate the largest volume that can be enclosed by a k-contracting embedding of an (n-1)-dimensional ellipsoid E into* $\mathbb{R}^n$.

We have indicated the solution of this problem when $k = n - 2$. With a somewhat different technique, we will solve the problem up to a constant factor for all values of k.

Theorem 1.5. *Let E be an (n-1)-dimensional ellipsoid with principal axes $E_0 \leq ... \leq E_{n-1}$. We define a monomial $V_k(E) = E_1^{\frac{n}{k}}...E_{l-1}^{\frac{n}{k}}E_l^bE_{l+1}...E_{n-1}$, where the numbers b and l are determined by the condition that the total degree of $V_k(E)$ is n and the condition that $1 < b \leq \frac{n}{k}$.*

Any k-contracting embedding from E into $\mathbb{R}^n$ encloses a volume less than $CV_k(E)$. On the other hand, we will construct a k-contracting embedding that encloses a volume greater than $cV_k(E)$.

These results on isoperimetric inequalities and the volume enclosed by k-contracting maps can also be applied to estimate the k-dilation of diffeomorphisms. Suppose we have a k-contracting embedding I from an (n-1)-dimensional rectangle R' into an n-dimensional rectangle S, which takes the boundary of R' to the boundary of S, and which divides the volume of S roughly in half; and suppose we write S as $S' \times [0, S_n]$. Applying the techniques from the last theorem, it follows that $V_k(R') > cV_k(S')$. Now if $R = R' \times [0, R_n]$, and if there is a k-contracting diffeomorphism from R to S, then by restricting this diffeomorphism to an appropriate slice $R' \times \{x_n\}$, we get a k-contracting embedding I that cuts S in half. Therefore, it follows that $V_k(R') > cV_k(S')$.

In addition to studying k-contracting diffeomorphisms, we will also study k-expanding embeddings. An embedding from one domain into another is called k-expanding if it maps every k-dimensional manifold in the domain to a k-dimensional manifold of larger volume in the range. Here is an equivalent way to say it. There is a k-expanding embedding of S into R if there is some open subset U in R and a k-contracting diffeomorphism from U to S. Our first problem asked when there was a k-contracting diffeomorphism from R to S. We can also ask when there is a k-expanding embedding from S into R.

Problem 3. *Given two n-dimensional rectangles R and S, estimate the smallest k-dilation of any diffeomorphism from any subset U of R onto S.*

The main result of my thesis solves this problem up to a constant factor.

Theorem 1.6. *Let R and S be n-dimensional rectangles. Define Q_i to be the quotient S_i/R_i. For each number l in the range $0 \leq l \leq k$ and each number p in the range $k+1 \leq p \leq n$, define $D_k(l,p)$ by the following formula.*

$$D_k(l,p) = Q_1 \ldots Q_l (Q_{l+1} \ldots Q_p)^{\frac{k-l}{p-l}}.$$

We abbreviate the maximum of these numbers by D_k. Any diffeomorphism from any subset U of R onto S has k-dilation at least cD_k. On the other hand, we will construct an open subset U in R and a diffeomorphism from U to S with k-dilation less than CD_k.

The mappings involved in this theorem are not very complicated. The open set U is always quasi-isometric to some rectangle, and the diffeomorphism is just a linear map from this rectangle to S. The hard part is to prove that the k-dilation is greater than $cD_k(l,p)$ for every choice of l and p. The inequality $V_k(R') > cV_k(S')$ is just a special case of these inequalities. In fact, these inequalities include all of the inequalities we have mentioned so far as well as many others.

To prove Theorem 1.2, we estimated how many disjoint wide subsets fit into a rectangle R. We can think of a bunch of disjoint wide subsets as a thick neighborhood of a set of points. In order to prove Theorem 1.6, we need to make analogous estimates for thick neighborhoods of simplicial complexes of all dimensions. The first and most important step is to find a sensible definition for the k-width of a thick complex, which reduces to the original definition of k-width when the complex happens to be a point.

We will soon make some more estimates of the k-dilation of diffeomorphisms between rectangles, but in order to explain them we need to consider the k-dilations of some other maps. In [10], Gromov gave an estimate for the rational homotopy invariants of maps with bounded k-dilation. For example, if f is a map from (S^{4n-1}, g) to the unit 2n-sphere with 2n-dilation Λ, then Gromov proved that the Hopf invariant of f is bounded by $C(g)\Lambda^2$. For large values of Λ, this expression is sharp up to a constant factor, for an appropriate choice of $C(g)$. We consider the problem of estimating the constant $C(g)$, as g varies among the ellipsoids. We will prove the following theorem, which shows that for an ellipsoid E with principal axes $E_0 \leq ... \leq E_{4n-1}$, the constant C(g) is roughly $E_{2n}E_1...E_{4n-1}$.

Theorem 1.7. *If E is an ellipsoid of dimension 4n-1, and f is a 2n-contracting map from E to the unit 2n-sphere, then the Hopf invariant of f is bounded by* $CE_{2n}E_1...E_{4n-1}$. *On the other hand, if* $E_1 > C$, *then we will construct a 1-contracting map from E to* S^{2n} *with Hopf invariant greater than* $cE_{2n}E_1...E_{4n-1}$

Gromov's result can be generalized to all rational homotopy invariants. We will generalize our theorem to a number of more complicated rational homotopy invariants, but not all of them.

For our first generalization, let X be the bouquet $S^{k_1} \vee S^{k_2}$, where $2 \leq k_1 \leq k_2$. If E is an ellipsoid of dimension $n = k_1 + k_2 - 1$, then we can define a homotopy invariant as follows. Let α_1 be a top-form on S^{k_1} with $\int \alpha_1 = 1$, and the same for α_2. If f is a map from E to X, then we define $H(f) = \int_E f^*(\alpha_1) \wedge Pf^*(\alpha_2)$, where $Pf^*(\alpha_2)$ indicates any primitive of the exact form $f^*(\alpha_2)$. This number does not depend on the choice of primitive and it is a homotopy invariant of the map f. We can estimate the invariant H in terms of the k-dilation.

Theorem 1.8. *Give X a metric by putting the unit sphere metric on each sphere in the wedge. Assume that* $k_1 \leq k_2$. *If f is a* k_1*-contracting map from E to X, then H(f) is bounded by* $CE_{n-k_1+1}E_1...E_n$. *On the other hand, if* $E_1 > C$, *we will construct a 1-contracting map f with* $H(f)$ *greater than* $cE_{n-k_1+1}E_1...E_n$.

We can create more complicated homotopy invariants of the same kind by considering maps to the wedge of larger numbers of spheres. For example, let X be the wedge of three spheres $S^{k_1} \vee S^{k_2} \vee S^{k_3}$. We order the spheres so that $k_1 \leq k_2 \leq k_3$, and again we assume that $k_1 \geq 2$. If E is an ellipsoid of dimension $n = k_1 + k_2 + k_3 - 2$, then we can define two linearly independent homotopy invariants as follows.

$$H_1(f) = \int_E Pf^*(\alpha_1) \wedge f^*(\alpha_2) \wedge Pf^*(\alpha_3).$$

$$H_2(f) = \int_E f^*(\alpha_1) \wedge Pf^*(\alpha_2) \wedge Pf^*(\alpha_3).$$

Theorem 1.9. *Let f be a k_1-contracting map from E to X. Then the homotopy invariant $H_1(f)$ is less than $CE_{n-k_1+1}E_{n-k_1-k_2+2}E_1...E_n$. The homotopy invariant $H_2(f)$ is less than $CE_{n-k_2+1}E_{n-k_1-k_2+2}E_1...E_n$.*

On the other hand, if $E_1 > C$, we will construct a 1-contracting map f_1 with $H_1(f_1)$ greater than $cE_{n-k_1+1}E_{n-k_1-k_2+2}E_1...E_n$ and $H_2(f_1) = 0$; and we will also construct a 1-contracting map f_2 with $H_2(f_2)$ greater than $cE_{n-k_2+1}E_{n-k_1-k_2+2}E_1...E_n$ and $H_1(f_2) = 0$.

An interesting feature of this theorem is that the homotopy invariants H_1 and H_2 have different behaviors. As the number of spheres in the bouquet X increases, we can define more complicated homotopy invariants along these lines. For example, for maps to a wedge of four spheres, we will define six different homotopy invariants with six different behaviors. In many cases, but not in all cases, we will prove estimates analogous to the theorem above. We will give more details in the body of the paper.

Each of these estimates for homotopy invariants of maps from ellipsoids implies a lower bound on the k-dilation of diffeomorphisms between rectangles, for certain values of k. Some of these inequalities follow from the estimates in earlier theorems, but most of them are new. For example, we will prove that the smallest 3-dilation of a diffeomorphism between 5-dimensional rectangles is at least $(Q_1Q_2Q_3^2Q_4Q_5)^{1/2}$, and that the smallest 2-dilation of a diffeomorphism between 8-dimensional rectangles is at least $(Q_2Q_3Q_4^2Q_5^2Q_6Q_7^2Q_8)^{1/5}$.

Gromov also asked whether the k-dilation controls the torsion homotopy invariants of maps. In [10], he proved that any map from S^m to S^n with sufficiently small 2-dilation is null-homotopic (provided n is at least 2). We will show that the 3-dilation gives much less control of homotopy invariants.

Theorem 1.10. *For each n, there are infinitely many choices of m so that we can find homotopically non-trivial maps from S^m to S^n with arbitrarily small 3-dilation.*

The next simplest case of Problem 1 is to estimate the 2-dilation of diffeomorphisms between 4-dimensional rectangles. We will construct a variety of non-linear diffeomorphisms with small 2-dilation, generalizing the snake map in various ways. The results of Theorems 1.6 through 1.9 imply a variety of lower bounds for the 2-dilation of any diffeomorphism. Nevertheless, these results are far from solving Problem 1 up to a constant factor. I do not even have a guess of what the right answer might be. As a very partial result, we will solve the problem in the special case that R is the unit cube.

Theorem 1.11. *If there is a 2-contracting diffeomorphism from the unit 4-cube to S, then $S_2 S_3^3 S_4^2 < C$. On the other hand, if $S_2 S_3^3 S_4^2 < c$, then there is a 2-contracting diffeomorphism from the unit 4-cube to S.*

In the last chapter of the thesis, we pursue a connection between area-contracting maps and the topology of 3-manifolds. The bridge connecting these topics is the following result.

Theorem 1.12. *Let X be a 3-dimensional simplicial complex. We give X a metric by putting the standard metric on each simplex. Let M be a complete hyperbolic manifold (of any dimension). Then any continuous map f from X to M can be homotoped to a map $\bar{f}$ with 2-dilation bounded by C.*

In the 1970's, Thurston invented a simplex-straightening argument which shows that f can be homotoped to a map that has bounded 2-dilation on each 2-simplex and bounded 3-dilation on each 3-simplex. We prove our theorem by adding an additional argument to Thurston's, which allows us to bound the 2-dilation on each 3-simplex. Although we have made only a small, technical improvement I found the proof surprisingly difficult.

Given this theorem, we can apply geometrical estimates for 2-dilation to bound the homotopy invariants of arbitrary smooth maps to hyperbolic manifolds. By this method, we will prove two topological theorems.

The first theorem concerns estimates for a generalization of the Hopf invariant. Recall that if α is a 2-form on S^2 with integral 1, then the Hopf invariant of a map f from S^3 to S^2 can be defined by the formula $H(f) = \int_{S^3} f^*(\alpha) \wedge Pf^*(\alpha)$, where $Pf^*(\alpha)$ denotes any primitive of the exact form $f^*(\alpha)$. Now if X is any closed oriented 3-manifold and Σ is any closed oriented surface, then we can define an analogous invariant for maps from X to Σ. We pick a 2-form α on Σ with integral 1, and define the Hopf invariant of a map f to be $\int_X f^*(\alpha) \wedge Pf^*(\alpha)$. The form $f^*(\alpha)$ will always be closed but it may or may not be exact. If the form is not exact, the Hopf invariant of f is not defined. If the form is exact, however, then the integral above does not depend on the choice of primitive and defines a homotopy invariant of f, which we will call the Hopf invariant.

Thurston's straightening lemma was used to bound the degrees of maps to a hyperbolic manifold. Using our refined version, we can bound the Hopf invariant.

Theorem 1.13. *Let X be a closed oriented 3-manifold that can be triangulated by N simplices. Let f be a map from X to a surface of genus 2. If the Hopf invariant of f is defined, then it is bounded by C^N. On the other hand, we will give examples where the Hopf invariant is greater than c^N, for a constant $c > 1$.*

Using related techniques, we are able to bound the degrees of maps to a hyperbolic manifold with small injectivity radius.

Theorem 1.14. *Let X be a closed oriented 3-manifold that can be triangulated by N simplices. Let M be a closed oriented hyperbolic manifold with injectivity radius ϵ. If there is a map of non-zero degree from X to M, then ϵ is greater than C^{-N}. On the other hand, we will give examples of closed oriented hyperbolic manifolds with injectivity radius ϵ that can be triangulated by N simplices, where $\epsilon < c^{-N}$ for a constant $c > 1$.*

(Before I had proven any of the theorems in this thesis, my adviser suggested that I try to use area-contracting maps to bound the degrees of maps between 3-manifolds. Somewhat baffled, I looked at the literature on the degrees of maps between 3-manifolds. One of the interesting results that I found is the following theorem of Teruhiko Soma, from [19].

Theorem. *Let X be a closed oriented 3-manifold. Then there are only finitely many closed oriented hyperbolic 3-manifolds M which admit a map of non-zero degree from X.*

As a corollary of Theorem 1.14, we will give a new proof of this theorem of Soma.)

Chapter 2

The (n-1)-Dilation of Diffeomorphisms Between Rectangles

In this section, we solve the problem of (n-1)-contracting diffeomorphisms between rectangles up to a constant factor.

Theorem 2.1. *Let R and S be n-dimensional rectangles. If there is an (n-1)-contracting diffeomorphism from R to S, then the following inequalities hold. For every integer l in the range* $1 \leq l \leq n-1$,

$$A \qquad R_1 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} > c S_1 ... S_l (S_{l+1} ... S_n)^{\frac{n-l-1}{n-l}}.$$

$$B \qquad R_2 ... R_n > c S_2 ... S_n.$$

On the other hand, if these inequalities hold with a larger constant C in place of c, then we will construct an (n-1)-contracting diffeomorphism from R to S.

By scaling considerations, this theorem allows us to estimate the minimal (n-1)-dilation of any diffeomorphism from R to S.

Corollary. *Let R and S be n-dimensional rectangles. Define Q_i to be the quotient S_i/R_i. For each integer l in the range $1 \leq l \leq n-1$, we define a number $D(l)$ as follows.*

$$D_l = Q_1 ... Q_l (Q_{l+1} ... Q_n)^{\frac{n-l-1}{n-l}}.$$

Also, we define D_n to be $Q_2 ... Q_n$. Finally, we define D to be the largest of all the numbers D_l, for l from 1 to n.

Any diffeomorphism from R to S has (n-1)-dilation at least cD. On the other hand, we will construct a diffeomorphism from R to S with (n-1)-dilation less than CD.

Let us compare the best (n-1)-dilation of a diffeomorphism from R to S with the best (n-1)-dilation of a linear diffeomorphism from R to S. An easy calculation shows that the linear diffeomorphism from R to S with the best (n-1)-dilation is just a diagonal matrix which maps the interval $[0, R_l]$ to the interval $[0, S_l]$ for each l. Its (n-1)-dilation is the largest of the n quotients $(S_1 ... S_{l-1} S_{l+1} ... S_n)/(R_1 ... R_{l-1} R_{l+1} ... R_n)$ as l varies from 1 to n. If the largest of these numbers occurs for l=1 or l=n, then the linear map is at least roughly as good as any non-linear map. On the other hand, if the lth quotient in the list is much larger than the first quotient and the nth quotient, then there is a non-linear diffeomorphism from R to S with (n-1)-dilation much smaller than the (n-1)-dilation of any linear diffeomorphism.

The proof of the theorem has two parts. In the first part we construct some non-linear diffeomorphisms with small (n-1)-dilation. In the second part, we give some lower bounds for the (n-1)-dilation of an arbitrary diffeomorphism between rectangles.

2.1 The snake map

We will now construct our fundamental example of a non-linear diffeomorphism with smaller (n-1)-dilation than any linear diffeomorphism. This map is a diffeomorphism between 3-dimensional rectangles. We will call it the snake map.

Proposition 2.1.1. *If $R_1 = S_1$, $R_2 > S_2$, and $R_2R_3 > S_2S_3$, then there is a diffeomorphism from R to S with 2-dilation less than C.*

Proof. We let I be a smooth quasi-isometric embedding of the rectangle $[0, 3S_2] \times [0, 3S_3]$ into $[0, R_2] \times [0, R_3]$. Because R_2 is greater than S_2 and R_2R_3 is greater than S_2S_3, it is easy to construct an embedding with quasi-isometric constant 10. Let H be a smooth function on $[0, 3S_2] \times [0, 3S_3]$ which is equal to R_1 on the central rectangle $[S_2, 2S_2] \times [S_3, 2S_3]$, and which is equal to a tiny number δ on a neighborhood of the boundary of $[0, 3S_2] \times [0, 3S_3]$. Since $R_1 = S_1 \leq S_2$, we can choose H with Lipshitz constant less than 1.

The function $H \circ I^{-1}$ is defined on the image of I in $[0, R_2] \times [0, R_3]$, and it is equal to δ on the boundary of this image. We extend this function to all of $[0, R_2] \times [0, R_3]$ by setting it equal to δ on the complement of the image of I. We call the resulting function $\bar{H}$. The graph of the function $\bar{H}$ defines a surface in the rectangle R. The first step of our construction is to push R below this hypersurface by a Lipshitz map. Namely, we define $\Phi_1(x, y, z) = (x\bar{H}(y, z)R_1^{-1}, y, z)$. The map Φ_1 is a diffeomorphism from R to the region $0 \leq x \leq \bar{H}(y, z)$, $0 \leq y \leq R_2$, $0 \leq z \leq R_3$. Because the function H has Lipshitz constant less than 1, the function $\bar{H}$ has Lipshitz constant less than C, and the diffeomorphism Φ_1 has Lipshitz constant less than C as well.

The next step of our construction is to push this region into $[0, R_1] \times \text{image}(I)$. To do this, we first pick a diffeomorphism ϕ_2 from $[0, R_2] \times [0, R_3]$ to the image of I. This diffeomorphism will have the following properties. Its restriction to the image of the central rectangle $[S_2, 2S_2] \times [S_3, 2S_3]$ is the identity. Its restriction to a small open neighborhood of the image of I has Lipshitz constant less than 2. It maps the complement of the image of I to a δ-neighborhood of the boundary of the image of I, for some tiny number δ. It has Lipshitz constant less than $1000R_2/S_2$. At each point where the Lipshitz constant of ϕ_2 is more than 10, the smaller singular value of $d\phi_2$ is less than δ. (The diffeomorphism ϕ_2 is just a small perturbation of a retraction from $[0, R_2] \times [0, R_3]$ to the image of I.) Finally, we define $\Phi_2(x, y, z) = (x, \phi_2(y, z))$.

The map Φ_2 itself has very large 2-dilation, but the composition $\Phi_2 \circ \Phi_1$ has 2-dilation less than C. To prove this bound, we consider two cases. If the coordinates

(y, z) lie in the image of I, then we prove the bound as follows. The map Φ_1 has 1-dilation less than C at (x, y, z), and it maps (x, y, z) to $(x\bar{H}(y, z)R_1^{-1}, y, z)$. The map Φ_2 has 1-dilation less than C at this point. If the coordinates (y, z) do not lie in the image of I, then the derivative of $\Phi_2 \circ \Phi_1$ is given by the following matrix.

$$d(\Phi_2 \circ \Phi_1) = \begin{pmatrix} \delta & 0 & 0 \\ 0 & \frac{\partial \phi_2^y}{\partial y} & \frac{\partial \phi_2^y}{\partial z} \\ 0 & \frac{\partial \phi_2^z}{\partial y} & \frac{\partial \phi_2^z}{\partial z} \end{pmatrix}.$$

(In this equation ϕ_2^y denotes the y-coordinate of the mapping ϕ_2.) If the Lipshitz constant of ϕ_2 at (y, z) is less than 10, then the 2-dilation of this matrix is certainly less than C. On the other hand, if the Lipshitz constant of ϕ_2 at (y, z) is more than 10, then the smaller singular value of $d\phi_2$ is less than δ, and the larger singular value of $d\phi_2$ is less than $1000R_2/S_2$. Therefore, the matrix $d(\Phi_2 \circ \Phi_1)$ has two singular values less than δ and the largest singular value less than $1000R_2/S_2$. If δ is sufficiently small, the matrix has 2-dilation less than C.

Next, we define a map Φ_3 from the region $[0, R_1] \times \text{image}(I)$ to $[0, S_1] \times [0, 3S_2] \times [0, 3S_3]$. This map is defined by $\Phi_3(x, y, z) = (x, I^{-1}(y, z))$. It is quasi-isometric. The composition $\Phi_3 \circ \Phi_2 \circ \Phi_1$ is an embedding of R into $[0, S_1] \times [0, 3S_2] \times [0, 3S_3]$, with 2-dilation less than C. The image of this embedding contains the sub-rectangle $[0, S_1] \times [S_2, 2S_2] \times [S_3, 2S_3]$, which is isometric to S. Since this sub-rectangle is convex, there is a 1-contracting retraction from the image of R to it. This retraction is not a diffeomorphism, but it can easily be approximated by a diffeomorphism Φ_4 with Lipshitz constant as close as we like to 1.

The composition $\Phi_4 \circ \Phi_3 \circ \Phi_2 \circ \Phi_1$ is a diffeomorphism from R to S with 2-dilation bounded by C. □

Using the snake map, we can quickly construct a 2-contracting diffeomorphism between three-dimensional rectangles R and S which obey the conditions in Theorem 2.1. Suppose $R_1 R_2 > S_1 S_2$, $R_1^2 R_2 R_3 > S_1^2 S_2 S_3$, and $R_2 R_3 > S_2 S_3$. It suffices to construct a diffeomorphism from R to S with 2-dilation less than C.

If $R_1 < S_1$, then we define a 2-contracting linear diffeomorphism from R to T, with $T_1 = S_1$, $T_2 = R_2R_1/S_1$, and $T_3 = R_3R_1/S_1$. (The length T_2 is indeed bigger than T_1 because $R_1R_2 > S_1S_2$.) Using the first two equations in the list above, we see that $T_2 > S_2$ and $T_2T_3 > S_2S_3$. Therefore, there is a snake map from T to S with 2-dilation less than C.

If $R_1 \geq S_1$ but $R_2 < S_2$, then we define a 2-contracting linear diffeomorphism from R to T, with $T_1 = R_1R_2/S_2$, $T_2 = S_2$, and $T_3 = R_3R_2/S_2$. (The length T_3 is indeed bigger than T_2 because $R_2R_3 > S_2S_3$.) Since $R_1R_2 > S_1S_2$, $T_1 > S_1$. Since $R_2R_3 > S_2S_3$, $T_3 > S_3$. Therefore, there is a 1-contracting linear diffeomorphism from T to S.

If $R_1 \geq S_1$ and $R_2 \geq S_2$, there is a snake map from R to S with 2-dilation less than C.

To generalize the snake map to higher dimensions, we take the Cartesian product of a snake map from a three-dimensional rectangle R to a three-dimensional rectangle S with an identity map from an (n-3)-dimensional rectangle R' to itself. The resulting map is an (n-1)-contracting diffeomorphism from $R \times R'$ to $S \times R'$, which we will also call a snake map. All of the (n-1)-contracting diffeomorphisms needed for Theorem 2.1 can be constructed by composing a sequence of snake maps with a linear map. To prove this fact requires some tedious algebra, which we divide into a sequence of propositions.

Proposition 2.1.2. *If the rectangles R and S obey the conditions $R_1 \geq S_1$, and $R_2 ... R_b \geq S_2 ... S_b$ for each number b in the range $2 \leq b \leq n$, then there is a diffeomorphism from R to S with (n-1)-dilation less than C. This diffeomorphism is a composition of snake maps and contracting linear maps.*

Proof. If $R_i \geq S_i$ for every value of i, then there is a contracting linear diffeomorphism from R to S. Let a denote the smallest number so that $R_a < S_a$. We perform a snake map from T to S, where the three active directions are $[0, S_1] \times [0, S_2] \times [0, S_a]$. More precisely, write $S = [0, S_1] \times [0, S_2] \times [0, S_a] \times S'$, and $T = [0, S_1] \times [0, \lambda S_2] \times [0, \lambda^{-1} S_a] \times S'$. We choose λ to be the smaller of the two ratios R_2/S_2 and S_a/R_a. Then there is a snake map from T to S with (n-1)-dilation less than C.

If the ratio R_2/S_2 is smaller than S_a/R_a, then T is equal to $[0, R_2]$ times a lower dimensional rectangle T'. We let R' be equal to the (n-1)-dimensional rectangle $R_1 \times R_3 \times ... \times R_n$. Now it is easy to check that the rectangles R' and T' satisfy the hypotheses of this proposition. By induction on the dimension, we can assume that there is a diffeomorphism from R' to T' with (n-2)-dilation less than C. Taking the direct product of this diffeomorphism with the identity map from $[0, R_2]$ to itself gives a diffeomorphism from R to T with (n-1)-dilation less than C. Composing with the snake map from T to S gives a diffeomorphism from T to S with (n-1)-dilation less than C.

If the ratio S_a/R_a is smaller than R_2/S_2, then T is equal to $[0, S_1] \times [0, S_2 S_a/R_a] \times [0, R_a] \times S'$. In this case, the rectangle T obeys $R_1 \geq T_1$, $R_2 ... R_b \geq T_2 ... T_b$ for all b, and $R_i \geq T_i$ for all i less than or equal to a. By induction on the number a, we can assume that there is a diffeomorphism from R to T with (n-1)-dilation less than C. Composing with the snake map from T to S, we get a diffeomorphism from T to S with (n-1)-dilation less than C. □

Proposition 2.1.3. *If $R_1 \geq S_1$ and $R_2 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} \geq S_2 ... S_l (S_{l+1} ... S_n)^{\frac{n-l-1}{n-l}}$, for each l in the range $1 \leq l \leq n-1$, then there is a rectangle T with $T_1 = S_1$, $R_2 ... R_b \geq T_2 ... T_b$ for each b in the range $2 \leq b \leq n$, and an (n-1)-contracting linear diffeomorphism from T to S.*

Proof. We prove this proposition by induction. More generally, we will prove that for every p in the range $1 \leq p \leq n-1$, the conclusion of the proposition follows from the following conditions, which we call C(p).

1. $R_1 \geq S_1$.
2. $R_2 ... R_a \geq S_2 ... S_a$ for a in the range $2 \leq a \leq p$,
3. $R_2 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} \geq S_2 ... S_l (S_{l+1} ... S_n)^{\frac{n-l-1}{n-l}}$ for l in the range $p \leq l \leq n-1$,

and

4. $R_2 ... R_n \geq S_2 ... S_n$.

When p=1, these hypotheses are just the hypotheses of the proposition. On the other hand, the conclusion of the proposition follows immediately from $C(n-1)$.

Therefore, it suffices to show that if the conclusion follows from $C(q)$ for each q greater than p, then it also follows from $C(p)$. We now carry out that induction.

Suppose that R and S satisfy $C(p)$. Let b be the smallest number for which $R_2 ... R_b < S_2 ... S_b$. If there is no such b, then the conclusion of the proposition clearly holds for R and S. Because of condition 4 we know that b is not equal to n, and because of the condition 3 with $l = n - 1$, we know that b is not equal to n-1. Therefore, b lies in the range $2 \leq b \leq n - 2$.

The case $b = 2$ is special. If $b = 2$, then we must have had $p = 1$. There is an (n-1)-expanding linear diffeomorphism from S to S', where $S_1' = S_1$, $S_2' = R_2$, and $S_i' = S_i (S_2/R_2)^{1/(n-3)}$ for all $i \geq 3$. (The length S_2' is at least S_1' because $S_1' = S_1 \leq R_1 \leq R_2 = S_2'$.) We check that S' obeys $C(2)$. Condition 1 follows because $R_1 \geq S_1 = S_1'$. Condition 2 follows because $R_2 = S_2'$. Condition 3 follows for l in the range $2 \leq l \leq n - 1$, by a straightforward calculation. Finally, condition 4 follows from Condition 3 for l=2 along with the equality $S_2' = R_2$. Since $p = 1$, our inductive hypothesis is that $C(2)$ implies the conclusion of the proposition. Therefore the conclusion of the proposition holds for R and S in this case.

Now we deal with the more general case that $b > 2$. In this case, we apply an (n-1)-expanding linear transformation to S that leaves S_1 through S_{b-1} invariant, decreases S_b, and increases all the other directions equally, until either $R_2 ... R_b = S_2 ... S_b$, or $S_{b-1} = S_b$. In the latter case, we then apply an (n-1)-expanding linear transformation to S that leaves S_1 through S_{b-2} invariant, decreases S_{b-1} and S_b equally, and increases all the other directions equally until either $R_2 ... R_b = S_2 ... S_b$ or $S_{b-2} = S_{b-1} = S_b$. In the latter case, we then apply a linear transformation that decreases S_{b-2}, S_{b-1}, and S_b, and so on. Because $R_2 > S_2$, this process terminates without decreasing S_2. At the end of the process, we have an equality $R_2 ... R_b = S_2' ... S_b'$. We check that the rectangles R and S' obey the condition $C(b)$. Condition 1 follows because $S_1' = S_1 \leq R_1$. Condition 2 follows because $R_2 ... R_b = S_2' ... S_b'$, and because for each a less than b, $S_2' ... S_a' \leq S_2 ... S_a \leq R_2 ... R_a$. A calculation shows that $S_2' ... S_l' (S_{l+1}' ... S_n')^{\frac{n-l-1}{n-l}} = S_2 ... S_l (S_{l+1} ... S_n)^{\frac{n-l-1}{n-l}}$ for $l \geq b$. Therefore, condition 3 holds for R and S'. Finally condition 4 follows from the case l=b of condition 3 along with the equality $R_2...R_b = S'_2...S'_b$. Since b is greater than p, our inductive hypothesis tells us that the conclusion of the proposition holds for R and S'. Since there is an (n-1)-contracting linear diffeomorphism from S' to S, the conclusion of the proposition holds for R and S.

This argument proves the inductive step and hence the proposition. $\square$

Proposition 2.1.4. *Suppose the rectangles R and S obey the following inequalities for each l in the range* $1 \leq l \leq n-1$.

$$R_1...R_l(R_{l+1}...R_n)^{\frac{n-l-1}{n-l}} \geq S_1...S_l(S_{l+1}...S_n)^{\frac{n-l-1}{n-l}}.$$

Also suppose that $R_2...R_n \geq S_2...S_n$. *Then there is a rectangle* S' *with an (n-1)-contracting linear diffeomorphism to S, so that R and* S' *obey the hypotheses of Proposition 2. These are* $R_1 \geq S'_1$ *and* $R_2...R_l(R_{l+1}...R_n)^{\frac{n-l-1}{n-l}} \geq S'_2...S'_l(S'_{l+1}...S'_n)^{\frac{n-l-1}{n-l}}$*, for each l in the range* $1 \leq l \leq n-1$.

Proof. If $R_1 < S_1$, then there is an (n-1)-expanding linear diffeomorphism from S to S', where $S'_1 = R_1$ and $S'_i = S_i(S_1/R_1)^{1/(n-2)}$ for i at least 2. A short calculation shows that $S'_1...S'_l(S'_{l+1}...S_n)^{\frac{n-l-1}{n-l}} = S_1...S_l(S_{l+1}...S_n)^{\frac{n-l-1}{n-l}}$ for each l between 1 and n-1. Therefore, $R_1...R_l(R_{l+1}...R_n)^{\frac{n-l-1}{n-l}} \geq S'_1...S'_l(S'_{l+1}...S'_n)^{\frac{n-l-1}{n-l}}$. Since $R_1 = S'_1$, $R_2...R_l(R_{l+1}...R_n)^{\frac{n-l-1}{n-l}} \geq S'_2...S'_l(S'_{l+1}...S'_n)^{\frac{n-l-1}{n-l}}$, and the proposition follows in this case.

If $R_1 \geq S_1$, then we apply an (n-1)-expanding linear transformation to S which decreases S_2 and increases all other directions of S equally until either $S_1 = R_1$ or $S_1 = S_2$. In the second case, we apply an (n-1)-expanding linear transformation to S which decreases S_3 and increases all other directions of S equally until either $S_1 = R_1$ or $S_1 = S_2 = S_3$, and so on. We continue this process until either $S_1 = R_1$ or $S_1 = S_n < R_1$. In the latter case, the proposition follows trivially.

In the former case, we call the final rectangle in this chain of (n-1)-expanding diffeomorphisms S'. Suppose that the last diffeomorphism was decreasing S_{b+1} and increasing S_i for all other i. We have $R_1 = S'_1 = S'_2 = ... = S'_b$. A short calculation shows that $S'_1...S'_l(S'_{l+1}...S'_n)^{\frac{n-l-1}{n-l}} = S_1...S_l(S_{l+1}...S_n)^{\frac{n-l-1}{n-l}}$ as long as l is at least b+1.

Since $R_1 = S'_1$, it follows that $R_2 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} \geq S'_2 ... S'_l (S'_{l+1} ... S'_n)^{\frac{n-l-1}{n-l}}$ as long as l is at least b+1.

Since $S'_2 = ... = S'_b = R_1$, we know that $R_2 ... R_l \geq S'_2 ... S'_l$ for l less than or equal to b. Also, a short calculation shows that $S'_2 ... S'_n = S_2 ... S_n$. By hypothesis, $S_2 ... S_n \leq R_2 ... R_n$. Therefore $R_2 ... R_n \geq S'_2 ... S'_n$. Combining this inequality with the first inequality of this paragraph, it follows that $R_2 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} \geq S'_2 ... S'_l (S'_{l+1} ... S'_n)^{\frac{n-l-1}{n-l}}$ as long as l is at most b. Therefore, the conclusion of the proposition holds in this case also. □

Given these three propositions, we can construct all the (n-1)-contracting diffeomorphisms which we need to prove Theorem 2.1. Suppose that the side lengths of R and S obey the inequality $R_1 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} \geq S_1 ... S_l (S_{l+1} ... S_n)^{\frac{n-l-1}{n-l}}$ for each l in the range $1 \leq l \leq n-1$ and that $R_2 ... R_n \geq S_2 ... S_n$. It suffices to construct a diffeomorphism from R to S with (n-1)-dilation less than C. According to Proposition 2.1.4, there is a linear (n-1)-contracting diffeomorphism from T to S, where T obeys the inequalities $R_1 \geq T_1$ and $R_2 ... R_l (R_{l+1} ... R_n)^{\frac{n-l-1}{n-l}} \geq T_2 ... T_l (T_{l+1} ... T_n)^{\frac{n-l-1}{n-l}}$ for each l in the range $1 \leq l \leq n-1$. According to Proposition 2.1.3, there is a linear (n-1)-contracting diffeomorphism from U to T, where U obeys the inequalities $R_1 \geq U_1$, and $R_2 ... R_b \geq C U_2 ... U_b$ for each number b in the range $2 \leq b \leq n$. According to Proposition 2.1.2, there is a non-linear diffeomorphism from R to U with (n-1)-dilation less than C.

2.2 Packings by wide sets

We define the width of an open set U in $\mathbb{R}^n$ to be the infimal number W so that there is a piecewise-linear function π from U to $\mathbb{R}$ so that each level set of π has volume less than W. This definition has a good relationship with (n-1)-contracting diffeomorphisms. Namely, if there is an (n-1)-contracting diffeomorphism from U to V, then the width of U is greater than or equal to the width of V. (To see this, approximate the diffeomorphism by a piecewise linear map with a negligible increase in the (n-1)-dilation.)

In this section we will prove a sequence of estimates for the widths of subsets of $\mathbb{R}^n$. First we estimate the width of the unit ball. This proof is due to Gromov. We let $\omega(n)$ denote the volume of the unit ball in $\mathbb{R}^n$.

Proposition 2.2.1. *The width of the unit ball in $\mathbb{R}^n$ is at least $[(1/2)^{\frac{n-1}{n}} - 1/2]n\omega(n)$.*

Proof. The isoperimetric inequality in $\mathbb{R}^n$ can be written as follows.

$$[\text{Volume}(U)\omega(n)^{-1}]^{\frac{n-1}{n}} \leq \text{Volume}(\partial U)n^{-1}\omega(n)^{-1}.$$

Let π be a piece-wise linear map from the unit ball in $\mathbb{R}^n$ to $\mathbb{R}$. If there is some value of x so that $\pi^{-1}(x)$ has positive n-dimensional volume, then the width of f is infinite. If not, we can find a value of x so that $\pi^{-1}(x)$ is a polyhedral chain in the unit ball which cuts it into two pieces of equal volume. One of these two pieces meets the boundary in less than half of its volume. Therefore, this piece has volume $(1/2)\omega(n)$, and its boundary has volume at most $(n/2)\omega(n) + \text{Volume}(\pi^{-1}(x))$. According to the isoperimetric inequality, the volume of $\pi^{-1}(x)$ is at least $[(1/2)^{\frac{n-1}{n}} - 1/2]n\omega(n)$. □

Remark: As far as I know, this proof was first written down by Gromov in appendix 1 of Filling Riemannian Manifolds, [12]. It may well be older. Probably, the exact value of the width of the unit ball is known to be $\omega(n-1)$, but I am not sure how to prove it. In Appendix B, we give a short description of Almgren's work on families of cycles, which gives a sharp estimate for the width of the unit sphere.

Corollary. *The width of the unit cube is at most 1 and is greater than c.*

Corollary. *The width of a rectangle R is at least $cR_1...R_{n-1}$ and at most $R_1...R_{n-1}$.*

Proof. If the rectangle R is a cube, then this estimate follows from the previous corollary by scaling. For general R, there is an (n-1)-contracting linear map from R to a cube with side-length $(R_1...R_{n-1})^{1/(n-1)}$. Therefore, the (n-1)-width of R is at least the (n-1)-width of this cube, which is greater than $cR_1...R_{n-1}$. On the other hand, the function x_n has level sets which are parallel to the face of R with dimensions $R_1 \times ... \times R_{n-1}$, and so the width of R is at most $R_1...R_{n-1}$. □

Using this corollary, we can prove one of the inequalities in Theorem 2.1.

Proposition 2.2.2. *If there is an (n-1)-contracting diffeomorphism from R to S, then* $R_1 \ldots R_{n-1} > cS_1 \ldots S_{n-1}$.

Proof. By the last proposition, we know that $R_1 \ldots R_{n-1}$ is greater than or equal to the width of R. Since there is an (n-1)-contracting diffeomorphism from R to S, the width of R is greater than or equal to the width of S. But according to the last proposition, the width of S is greater than $cS_1 \ldots S_{n-1}$. □

This proposition proves inequality A of Theorem 2.1 in the case that $l = n - 1$. In order to prove the other cases of inequality A, we will need upper bounds on the widths of subsets of rectangles. The first case of an upper bound which we will prove is the following.

Proposition 2.2.3. *(Width-Volume Inequality) If U is a bounded open set in* $\mathbb{R}^n$*, then the width of U is less than* $C\, Volume(U)^{\frac{n-1}{n}}$.

This upper bound follows from a width-volume inequality for functions. Let f be a non-negative function on $\mathbb{R}^n$. Then the width of f is defined to be the smallest number W so that there is a piecewise linear map π from $\mathbb{R}^n$ to $\mathbb{R}$ with $\int_{\pi^{-1}(y)} f < W$ for each y in $\mathbb{R}$.

Theorem 2.2. *Suppose that f is a continuous, compactly supported function in* $\mathbb{R}^n$*, and that* $0 \leq f \leq 1$ *pointwise. Then the width of f is bounded by* $C(\int f)^{\frac{n-1}{n}}$.

Proof. By rescaling the coordinates, we can reduce this theorem to the case that $\int f = 1$.

Let $S(x)$ be the (n-1)-skeleton of the unit lattice in $\mathbb{R}^n$, centered at the point x. For a random point x in the unit cube, the average value of $\int_{S(x)} f$ is n. Therefore, we can pick a point x so that $\int_{S(x)} f \leq n$. From now on, we call this $S(x)$ simply S.

We are going to construct a piecewise linear map to $\mathbb{R}$ so that each fiber closely hugs the skeleton S except for a set of bounded volume. Since $\int_S f$ is bounded, we can bound the integral of f over the part of the fiber near to S. On the other hand, since $f \leq 1$, we can bound the integral of f over the remainder of the fiber. We now turn to the details.

We will first construct a piecewise linear isomorphism Ψ of $\mathbb{R}^n$ which fixes S and which pushes most of space into a small neighborhood of S. The complement of S consists of disjoint unit cubes. Let K be one of these cubes. Since the boundary of K lies in S, Ψ fixes the boundary of K. Let k be a cube with side length ϵ and the same center as K, for some very small number ϵ which we will choose later. The restriction of Ψ to k is the dilation whose image is the cube with side $(1-\epsilon)$ and with the same center as K. While Ψ magnifies the tiny central cube k, it squeezes the region between k and K into the $\epsilon/2$-neighborhood of the boundary of K. Once we have defined the restriction of Ψ to K, we extend Ψ to a periodic function on $\mathbb{R}^n$.

It is a little bit tedious to write down the extension of Ψ to the region between K and k, so we will postpone it to the end of the proof. We will divide the region between K and k into convex polyhedra K_F, where F is a hyperface of K. We will construct a triangulation T of K so that each simplex lies in one of the K_F. The extension of Ψ is linear on each simplex of this triangulation. The map Ψ takes each simplex of T in K_F into a very small neighborhood of F. Finally, the restriction of Ψ to each simplex in K_F takes almost every (n-1)-plane to be almost parallel to F.

Next let π_0 be a linear projection from $\mathbb{R}^n$ to $\mathbb{R}$, which we will choose later. Our map π will be equal to the composition $\pi_0 \circ \Psi^{-1}$. The fibers of π_0 are parallel hyperplanes. Let P be one of these hyperplanes. The corresponding fiber of π is $\Psi(P)$. We only need to worry about the part of the fiber which lies in the support of f. For generic choice of π_0, if ϵ is sufficiently small, each plane P will hit the central ϵ cube k for at most one cube K in the support of f. The part of P through this central region is mapped to a piece of plane inside of that cube K, with volume less than C. Since f is no more than 1, the integral of f over $\Psi(P \cap k)$ is less than C.

The remainder of P lies in the union of the convex polyhedra K_F. For each simplex Δ of our triangulation, the map Ψ takes $P \cap \Delta$ to a piece of plane which is very close to the face F. For each Δ and for most choices of π_0, this piece of plane is also very close to being parallel to F. Since Ψ is periodic, we only have to consider finitely many linear maps. Therefore, we can choose π_0 so that the restriction of Ψ to every Δ in every K_F maps the fibers of π_0 to (n-1)-planes that are almost parallel to F.

Because f is continuous, the integral of f over $\Psi(P \cap \Delta)$ is bounded by the integral of f over the face F. Therefore, the integral of f over $\Psi(P \cap K_F)$ is bounded by $C \int_F f$. Finally, applying this to each region K_F, we can bound the integral of f over the complement of the central ϵ-cubes k by $C \int_S f$, which is less than C.

Assembling these two bounds, we see that the integral of f over $\Psi(P)$ is less than C, for each plane P which is a fiber of π_0. In particular, the integral of f over any fiber of π is less than C. The map π is clearly piecewise linear. For a generic map π_0, the map π is generic also. Therefore, the width of f is bounded by C.

To finish the proof, we now have to construct the extension of Ψ to the region between K and k. Any reasonable extension of Ψ to all of K will probably work. Over the course of the thesis we will have to construct a number of mappings in similar situations, so we are going to write down a standard approach now. In a first reading of the proof, the reader should not take these details too seriously.

The first step is to divide the region between K and k into convex polyhedra. For each face F of K, we can define a pyramid with base equal to F and apex equal to the center of K. This pyramid meets the boundary of k in exactly one face. The intersection of the pyramid with the region between K and k is a convex polyhedron. (This intersection is the convex hull of the union of F and the corresponding face of k.) As F varies among the faces of K, these convex polyhedra tile the region between K and k. The region between K and $(1-\epsilon)K$ has a combinatorially equivalent tiling, consisting of the intersection of the above pyramids with this smaller region. The map which we will construct will map each convex polyhedron in the first tiling to the corresponding convex polyhedron in the second tiling.

The second step is to define the map on each convex polyhedron. We will often have to define piecewise linear maps between combinatorially equivalent convex polyhedra. It turns out be convenient to define such maps using barycentric subdivisions.

Let P be a convex polyhedron. For each face F of P, let $c(F)$ be a point in the interior of F. Using $c(F)$ we can define the barycentric subdivision of P inductively on the faces of P. The barycentric subdivision of any polyhedron is a triangulation. Suppose that we have defined the barycentric triangulation on the boundary of a face F. Then the barycentric subdivision of F is the cone from $c(F)$ to the barycentric subdivision of the boundary. (If Δ is a p-simplex in the barycentric subdivision of the boundary of F, then the barycentric subdivision of F contains a (p+1)-simplex with base Δ and apex $c(F)$.)

Now suppose that P_1 and P_2 are combinatorially equivalent convex polyhedra. More precisely, suppose that we are given a particular combinatorial equivalence from P_1 to P_2. Also, suppose that we have chosen a "center" $c(F)$ for each face of each polyhedron. Given this data, there is a unique map from P_1 to P_2, which maps the center of each face of P_1 to the center of the corresponding face of P_2, and which is linear on each simplex of the barycentric subdivision of P_1. This map is a PL isomorphism from P_1 to P_2, taking each simplex of the barycentric subdivision of P_1 onto the corresponding simplex of the barycentric subdivision of P_2. We will call this map the barycentric map from P_1 to P_2.

The restriction of Ψ to K_F will be a barycentric map. We need to set up some coordinates. Put the center of the cube K at 0. The face F is given by the inequalities $-1/2 \leq x_i \leq 1/2$ for i from 1 to n-1, and the equation $x_n = 1/2$. The cone from F to the center of K is given by the inequalities $0 \leq x_n \leq 1/2$ and $-x_n \leq x_i \leq x_n$ for i from 1 to n-1. The convex polyhedron K_F is equal to this cone, minus a small cube centered at the origin with side length ϵ. The polyhedron K_F is given by the equations $\epsilon/2 \leq x_n \leq 1/2$ and $-x_n \leq x_i \leq x_n$. The polyhedron K_F is combinatorially equivalent to a rectangle. We are going to map K_F to an analogous polyhedron that lies between K and $(1-\epsilon)K$. This image polyhedron is defined by the equations $1/2 - \epsilon/2 \leq x_n \leq 1/2$ and $-x_n \leq x_i \leq x_n$. It is also combinatorially equivalent to a rectangle. The polyhedra have the face F in common. There is exactly one combinatorial equivalence from K_F to the other polyhedron which restricts to the identity on F.

To define the barycentric map, we need to pick a "center" for each face of K_F and each face of the other polyhedron. We define the center of each face to be the average value of its vertices. Given this choice, there is a unique barycentric map from K_F onto the other polyhedron defined in the last paragraph. Since the definition of the center of a face only depends on the face itself, the map that we defined on two neighboring polyhedra agrees on the boundary. It is easy to check that this barycentric map restricts to the identity on the face F and that it restricts to Ψ on the boundary between K_F and the small central cube k. Therefore, this barycentric map gives a globally well-defined PL isomorphism Ψ. The barycentric map takes K_F into the $\epsilon/2$ neighborhood of the face F.

The only thing that we still have to do is to check that the barycentric map on each barycentric simplex of K_F maps almost every (n-1)-plane to be almost parallel to F. An explicit calculation shows that each linear map is given by a matrix $M_{i,j}$ with the following form. The diagonal elements $M_{i,i}$ are at least on the order of 1 when i is between 1 and n-1. (Depending on which simplex we are looking at, some of these entries will be on the order of ϵ^{-1} and others will be on the order of 1.) The entry $M_{n,n}$ is on the order of ϵ. The off-diagonal entries are all 0, except for the entries $M_{i,n}$. The absolute value of $M_{i,n}$ is less than 2. For ϵ very small, this matrix maps almost every (n-1)-plane to be almost parallel to F. $\square$

To prove the width-volume inequality for a bounded open set U, we apply the theorem to a continuous function f approximating the characteristic function of U. (All we need to assume about f is the following: $\int f$ is roughly the volume of U, f is greater than or equal to the characteristic function of U, f is between 0 and 1, and f is compactly supported.)

In order to bound the (n-1)-dilation of diffeomorphisms between rectangles, we need to understand the following problem. If U is an open set in a rectangle R with a given volume, what is the largest possible width of U? According to the width-volume inequality, the width of U is less than $CVolume(U)^{\frac{n-1}{n}}$. Also, the width of U is at most the width of R, which is at most $R_1 ... R_{n-1}$. These two inequalities are not sufficient to give a good estimate for the width of U. We will prove a sequence of inequalities which interpolate between them.

Proposition 2.2.4. *Let U be a subset of the rectangle R. For each l in the range $0 \leq l \leq n-1$, we have the following inequality.*

$$\mathit{Width}(U) < C(R_1...R_l)^{\frac{1}{n-l}}\,\mathit{Volume}(U)^{\frac{n-l-1}{n-l}}.$$

Proof. When $l = 0$, this inequality reduces to the width-volume inequality. When $l = n-1$, this inequality says that the width of U is less than $CR_1...R_{n-1}$. Since U is a subset of R, the width of U is at most the width of R, and this inequality follows. Now we turn to the intermediate values of l.

Let f be a continuous approximation to the characteristic function of U. More precisely, f is a non-negative continuous function supported in R in a small neighborhood of U, with $0 \leq f \leq 1$, $f = 1$ on U, and $\int f$ approximately equal to the volume of U.

Let g be the function on the (n-l)-dimensional rectangle $R_{l+1} \times ... \times R_n$ defined by $g(y) = (R_1...R_l)^{-1} \int_{R_1 \times ... \times R_l} f(x,y)dx$. The function g is continuous, compactly supported, and $0 \leq g \leq 1$. Applying our theorem to the function g, it follows that the width of g is bounded by $C(\int g)^{\frac{n-l-1}{n-l}}$. This expression is roughly equal to $C[(R_1...R_l)^{-1}\text{Volume}(U)]^{\frac{n-l-1}{n-l}}$. According to the definition of width, there is some smooth function π from $\mathbb{R}^{n-l}$ to $\mathbb{R}$ so that the integral of g over each level set is bounded by this expression. We define a smooth map $\bar{\pi}$ on R by $\bar{\pi}(x,y) = \pi(y)$. The volume of U intersected with a fiber $\bar{\pi}^{-1}(z)$ is bounded by $(R_1...R_l)$ times the integral of g over the the fiber $\pi^{-1}(z)$. Therefore, the width of U is bounded by $C(R_1...R_l)^{\frac{1}{n-l}}\text{Volume}(U)^{\frac{n-l-1}{n-l}}$. □

We are now ready to prove all the inequalities in Theorem 2.1. We restate all the inequalities as a proposition.

Proposition 2.2.5. *If R and S are n-dimensional rectangles, and there is an (n-1)-contracting diffeomorphism from R to S, then the following inequalities hold.*

A. $R_1...R_l(R_{l+1}...R_n)^{\frac{n-l-1}{n-l}} > cS_1...S_l(S_{l+1}...S_n)^{\frac{n-l-1}{n-l}}$, *for each l in the range* $1 \leq l \leq n-1$.

B. $R_2...R_n > cS_2...S_n$.

Proof. Let l be any integer in the range $1 \leq l \leq n-2$. Inside of S, we choose disjoint subrectangles with dimensions $S_1 \times ... \times S_{l-1} \times S_l \times S_l \times ... \times S_l$. We can choose roughly N of these disjoint rectangles, where N $= S_{l+1}...S_n/S_l^{n-l}$. Each rectangle has width at least $cS_1...S_lS_l^{n-l-1}$.

According to the last proposition, any subset of R with width W has volume at least $[W(R_1...R_l)^{\frac{-1}{n-l}}]^{\frac{n-l}{n-l-1}}$. Therefore, the inverse image of each of the above rectangles has volume at least $c(S_1...S_l)^{\frac{n-l}{n-l-1}}S_l^{n-l}(R_1...R_l)^{\frac{-1}{n-l-1}}$. These inverse images are disjoint, and so their total volume is less than the volume of R. Since there are N inverse images, the following inequality holds.

$$R_1...R_n > cN(S_1...S_l)^{\frac{n-l}{n-l-1}}S_l^{n-l}(R_1...R_l)^{\frac{-1}{n-l-1}}.$$

Plugging in the value of N, which is roughly $S_{l+1}...S_n/S_l^{n-l}$, we get the following inequality.

$$R_1...R_n > c(S_1...S_l)^{\frac{n-l}{n-l-1}}S_{l+1}...S_n(R_1...R_l)^{\frac{-1}{n-l-1}}.$$

Finally, bringing the factor $(R_1...R_l)^{\frac{-1}{n-l-1}}$ to the other side leaves the following inequality.

$$(R_1...R_l)^{\frac{n-l}{n-l-1}}R_{l+1}...R_n > c(S_1...S_l)^{\frac{n-l}{n-l-1}}S_{l+1}...S_n.$$

Taking the $(n-l-1)/(n-l)$ power of each side leaves inequality A for any integer l in the range $1 \leq l \leq n-2$. We already proved inequality A in case l is equal to n-1. This inequality followed because the width of R is at least the width of S. Therefore, we have proven inequality A in all cases.

Inequality B is completely elementary. The number $2nR_2...R_n$ is at least as great as the volume of the boundary of R. Since there is an (n-1)-contracting diffeomorphism from R to S, the volume of the boundary of R is at least as great as the volume of the boundary of S. But the volume of the boundary of S is greater than $2S_2...S_n$. Therefore, $R_2...R_n > (1/n)S_2...S_n$. □

This finishes the proof of Theorem 2.1.

Chapter 3

The Width-Volume Inequality

In the previous chapter, we gave several lower bounds for (n-1)-dilation. All of these lower bounds generalize in a straightforward way to give lower bounds for the k-dilation. In particular, we will define a k-width which generalizes the width that we defined in the last chapter. The k-width of a set U is small if the set can be swept out by a family of k-dimensional surfaces with small k-volume. The main result of the chapter is that the k-width of an open set in $\mathbb{R}^n$ is controlled by its volume. Given this inequality we will prove lower bounds for the k-dilation of diffeomorphisms between rectangles. When k lies in the range $1 < k < n-1$, the resulting estimates are far from good enough to calculate the minimal k-dilation up to a constant factor.

3.1 The definition of k-width

In this section, we define a k-width, which measures the k-dimensional volume needed to sweep out a set by k-dimensional slices. The k-width generalizes the width defined in the last section, which measured the (n-1)-dimensional volume needed to sweep out a set by (n-1)-dimensional slices. Gromov defined a k-volume width in appendix 1 to the paper Filling Riemannian Manifolds [12], and similar definitions have probably been around a long time at least implicitly. For technical reasons, we need a definition which differs slightly from previous definitions.

It turns out to be convenient, for technical reasons, to work with piecewise linear maps. A piecewise linear map from a domain to $\mathbb{R}^m$ is called generic if its restriction to each simplex (of any dimension) has maximal rank. The generic PL maps have several convenient features. The inverse image of each point is a polyhedral chain of the expected dimension. Moreover, the inverse image of a polyhedral chain is again a polyhedral chain of the expected dimension. Also, if π is a generic PL map, then the integral of a continuous function over the fiber $\pi^{-1}(y)$ is a continuous function of y.

We define the k-width of U to be the infimal W so that there is a generic PL map π from U to $\mathbb{R}^{n-k}$ with fibers of k-volume less than W. We will sometimes abbreviate the k-width of U by $W_k(U)$.

The most basic estimate of k-width is the following.

Proposition 3.1.1. *(Gromov, Almgren) The k-width of the unit n-cube is at least some positive constant $c(n)$.*

Proof. Here is Gromov's proof. Suppose that π is a generic PL map from the unit n-cube to $\mathbb{R}^{n-k}$, and that each fiber has k-volume less than δ. Take a very fine triangulation T of $\mathbb{R}^{n-k}$. Since π is generic PL, the inverse image of each vertex of T is a (relative) cycle in the unit n-cube. Using the isoperimetric inequality, the cycle can be filled by a (k+1)-chain of volume less than $C\delta$. For each vertex v of T, define F(v) to be a filling of $\pi^{-1}(v)$ with (k+1)-volume bounded by $C\delta$.

Since π is generic PL and the triangulation T is very fine, we can assume that the inverse image of each edge of T has (k+1)-volume less than ϵ, which is much less than δ. For each edge E of T, we define a (k+1)-cycle C(E) in the unit n-cube, which is equal to the union of $\pi^{-1}(E)$, and $F(v_1)$ and $F(v_2)$, where v_1 and v_2 are the endpoints of E. The (k+1)-volume of C(E) is bounded by $C\delta$. Next, choose a (k+2)-chain F(E) with boundary C(E), and with (k+2)-volume bounded by $C\delta$.

Iterating this construction we finally define a cycle $C(\Delta)$ for each (n-k)-simplex Δ of T. Each cycle $C(\Delta)$ has n-volume less than $C\delta$. On the other hand, the sum of the cycles $C(\Delta)$ is equal to the sum of the chains $\pi^{-1}(\Delta)$, which is homologous to the fundamental cycle of the unit cube. Therefore one of the cycles $C(\Delta)$ has volume at least 1, and $\delta \geq 1/C > 0$.

This finishes the proof of the theorem. It is essentially identical to Gromov's proof in appendix 1 of Filling Riemannian Manifolds [12]. □

In fact, the proof actually gave a stronger result.

Corollary. *If M admits a k-contracting map of non-zero degree to the unit n-sphere, then the k-width of M is at least c.*

Given Gromov's theorem, it is easy to estimate the k-width of a rectangle up to a constant factor.

Corollary. *The k-width of a rectangle R is greater than $cR_1 ... R_k$ but less than $R_1 ... R_k$.*

Proof. Since there is a linear diffeomorphism from the rectangle R to the unit cube with k-dilation equal to $(R_1 ... R_k)^{-1}$, it follows that the k-width of R is at least $cR_1 ... R_k$. On the other hand, since the projection to the last $(n-k)$ coordinates is an admissible map, the k-width of R is at most $R_1 ... R_k$. □

3.2 The width-volume inequality

The width-volume inequality which we proved in the first section generalizes to the k-width for all k. For any bounded open set U in $\mathbb{R}^n$, we will prove the bound k-width$(U) < C\text{volume}(U)^{k/n}$.

As before, our inequality will follow from a slightly more general inequality involving the k-widths of functions. Let f be a function on $\mathbb{R}^n$ which is greater than or equal to zero. We say that the k-width of f is less than W if there is a generic PL map Ψ from $\mathbb{R}^n$ to $\mathbb{R}^{n-k}$ so that the integral of f over each fiber of Ψ is less than W.

Theorem 3.1. *If f is a continuous function with compact support on $\mathbb{R}^n$, and $0 \leq f \leq 1$, then k-width$(f) < C(\int f)^{k/n}$.*

The original inequality for sets U follows from this theorem by taking f to be a continuous approximation to the characteristic function of U. We now prove the theorem.

Proof. After rescaling the coordinates, it suffices to prove that k-width$(f) < C$ when $\int f = 1$. From now on, we will assume that $\int f = 1$.

We begin by finding a k-dimensional skeleton S with $\int_S f$ bounded. The reason for finding S is that each fiber of the map π which we are going to construct will hug the skeleton S except along regions of controlled volume.

Let S(x) be the k-skeleton of the unit lattice with center x. Since the integral of f is 1, the average value of $\int_{S(x)} f$ as x varies over the unit cube is equal to $\binom{n}{k}$. We can choose a point x so that $\int_{S(x)} f$ is no more than the average value $\binom{n}{k}$. From now on, we refer to S(x) simply as S. The set S is the k-skeleton of a translated unit lattice, and we have proven that $\int_S f \leq \binom{n}{k}$.

We will tile the space $\mathbb{R}^n$ with polyhedra in a way that fits nicely with S. First we need to make some definitions. Let T be the (n-k-1)-skeleton dual to S. If A is a k-dimensional face in S, then we define the link of A in the following way. The set A is defined by equations $x_i = a_i$ for (n-k) coordinates i, and equations $a_j \leq x_j \leq a_j + 1$ for the other k coordinates. There is an (n-k) cube transverse to A given by the equations $a_i - 1/2 \leq x_i \leq a_i + 1/2$ for the (n-k) coordinates i above, and $x_j = a_j + 1/2$ for the other k coordinates. This cube is simply the (n-k) cube centered at the center of A, perpendicular to A, and parallel to the coordinate axes. The link of A is defined to be the boundary of this (n-k) cube. It consists of $2(n-k)$ (n-k-1) cubes, each of which is an (n-k-1) dimensional face of T. If B is an (n-k-1) dimensional face of T, we define the link of B in an analogous way. It is a topological k-sphere consisting of 2(k+1) k-dimensional faces of S. We let A denote a k-dimensional face of S and B an (n-k-1)-dimensional face of T. A quick calculation shows that A is in the link of B if and only if B is in the link of A. For each pair (A, B) of faces with A in the link of B, we define $K(A, B)$ to be the convex hull of the union of A and B. These sets $K(A, B)$ are the tiles of our tiling.

We check that the polyhedra $K(A, B)$ tile $\mathbb{R}^n$. The hyperfaces of the tile $K(A, B)$ correspond to pairs (A, b) where b is an (n-k-2)-face in the boundary of B, or pairs (a, B), where a is a (k-1)-face in the boundary of A. (The corresponding face is just the convex hull of A and b, or of a and B.) Each face borders exactly two tiles in our tiling. Given a face (A, b), let B' be the (n-k-1)-face in the link of A which lies on the other side of b from B. Then $K(A, B')$ is the only other tile with (A, b) as a face. Therefore, the tiles form a pseudo-manifold, and the embedding of the tiles is an orientation preserving proper map from the tile space to $\mathbb{R}^n$. The intersection of $K(A, B)$ with the skeleton S is equal to A. In particular, the only tiles that come near to the center of A are tiles $K(A, B)$ for some B in the link of A. It is easy to check that a typical point very close to the center of A lies in exactly one of the tiles $K(A, B)$. Therefore, the tiles have disjoint interiors and cover all of space.

Next we define a PL isomorphism Ψ which leaves each tile invariant, and which maps the ϵ-neighborhood of T to the complement of the ϵ-neighborhood of S, for a very small number ϵ which we will choose later.

After renumbering the coordinates, translating, and reflecting, we can assume that A and B have the following simple form. The face A is given by the inequalities $0 \leq x_i \leq 1$ for i from 1 to k, $x_i = 0$ for i from $k+1$ to n. The face B is given by inequalities $-1/2 \leq x_i \leq 1/2$ for i from $k+1$ to $n-1$, the equations $x_i = 1/2$ for i from 1 to k, and $x_n = 1/2$. The convex set $K(A, B)$ is given by the inequalities $0 \leq x_n \leq 1/2$, $-x_n \leq x_i \leq x_n$ for i from $k+1$ to $n-1$, and $|1/2 - x_i| \leq |1/2 - x_n|$ for i from 1 to k.

We slice K into two polyhedra with the plane $x_n = t$, where $0 < t < 1/2$. Let $K(t)$ be the intersection of K with the set $x_n \leq t$, and let $K'(t)$ be the intersection of K with the set $x_n \geq t$. Up to combinatorial equivalence, the resulting two polyhedra do not depend on the choice of t.

The restriction of Ψ to $K(1/2-\epsilon)$ is the barycentric map from $K(1/2-\epsilon)$ to $K(\epsilon)$. The restriction of Ψ to $K'(1/2-\epsilon)$ is the barycentric map from $K'(1/2-\epsilon)$ to $K'(\epsilon)$. To define these barycentric maps, we need to pick a "center point" for each face of each polyhedron. It turns out that a rather ad hoc choice of center point simplifies the calculations that we have to do later. Here is the definition for the center point of faces of $K(t)$. When i is in the range $1 \leq i \leq n-1$, the x_i coordinate of the center point for a face F is just the average value of the x_i coordinate of the vertices of F. The x_n coordinate of the center point for F is equal to t if all of the vertices of F lie in the plane $x_n = t$; it is equal to 0 if all of the vertices of F lie in A; and it is equal to $t/2$ if some vertices of F lie on the plane $x_n = t$ and some vertices lie on A. For faces that lie in $K'(t)$, we define the center to be the average of the vertices. A short calculation shows that if a face lies in both $K(t)$ and $K'(t)$, or if it lies in several different tiles $K(A, B)$, then all of the definitions of its center agree. Therefore, the map Ψ is consistently defined, and it gives a periodic PL isomorphism of $\mathbb{R}^n$.

We will call the simplices in the barycentric subdivision of $K(1/2 - \epsilon)$ good simplices, and the simplices in the barycentric subdivision of $K'(1/2 - \epsilon)$ bad simplices. (Recalling what we have done so far, we see that good simplices are rather big and bad simplices are small. Bad simplices lie near to T, and good simplices lie fairly near to S. The map Ψ has Lipshitz constant 1 on the good simplices, but very large Lipshitz constant on the bad simplices.)

The restriction of Ψ to each good n-simplex Δ in $K(A, B)$ is a linear map which takes Δ into the ϵ-neighborhood of the face A. Moreover, this linear map takes almost every k-plane to a k-plane which is almost parallel to A. To check this last property, we let $M_{i,j}$ be the matrix of this linear map. An explicit calculation shows that the matrix elements have the following form. The diagonal elements $M_{i,i}$ are at least on the order of 1, when i goes from 1 to k. The diagonal elements $M_{i,i}$ have absolute value less than 2ϵ when i goes from $k+1$ to n. All of the off-diagonal elements vanish except for the entries $M_{i,n}$. Finally, the absolute value of $M_{i,n}$ is bounded by $2M_{i,i}$. This matrix maps almost every k-plane to a k-plane almost parallel to A.

Next, we will choose a linear projection π_0 from $\mathbb{R}^n$ to $\mathbb{R}^{n-k}$ which we will choose later. The fibers of π_0 are parallel k-planes. Let Δ be a good n-simplex simplex in $K(A, B)$. We have just seen that the linearization of Ψ on Δ maps almost every k-plane to be almost parallel to A. Since Ψ is periodic, there are less than C different linear maps among all the good simplices in the barycentric triangulation of $\mathbb{R}^n$. If ϵ is sufficiently small, we can choose π_0 so that the linearization of Ψ on each good simplex in each $K(A, B)$ maps the fibers of π_0 to k-planes that are almost parallel to A.

We define π to be the composition $\pi_0 \circ \Psi^{-1}$. Since Ψ^{-1} is a periodic PL isomorphism, for almost every choice of π_0, the composition π is generic PL. The fibers of π_0 are parallel k-planes. Let P be any one of these k-planes. Then the corresponding fiber of π_0 is $\Psi(P)$. Therefore, we need to prove that the integral of f over $\Psi(P)$ is less than C.

First we bound the contribution to our integral from the good simplices. We consider the intersection of P with a good simplex Δ, which lies in $K'(A, B)$. The image of this intersection under Ψ is a piece of k-plane which lies very close to A. It is also very close to being parallel to A. Since f is continuous, if ϵ is sufficiently small, the integral of f over $\Psi(P \cap \Delta)$ is bounded by $2 \int_A f$. Since there are less than C good simplices lying in tiles K which border A, the sum of the contribution from all good simplices is bounded by $C \int_S f$, which is bounded by C.

Next we bound the contribution to our integral from the bad simplices. For each bad simplex Δ, $\Psi(P \cap \Delta)$ is a piece of k-plane lying in a tile K(A,B). Therefore, this piece of k-plane has volume less than C. Since the function f is less than or equal to 1, the integral of f over $\Psi(P \cap \Delta)$ is less than C. Now fix a very large ball which easily contains the support of f. We claim that for ϵ sufficiently small, the number of bad simplices which lie in this ball and which intersect P is less than C. Let T_0 be the intersection of T with this large ball. Since π_0 is generic, the images of the faces of T_0 intersect transversely in $\mathbb{R}^{n-k}$. Therefore, the number of faces of T_0 which meet any fiber $\pi_0^{-1}(y)$ is less than C. For ϵ sufficiently small, the number of faces of T_0 which come within ϵ of $\pi_0^{-1}(y)$ is less than C. Since there are less than C bad simplices lying in tiles K which border a given face B of T_0, the number of bad simplices which meet the support of f and which intersect P is less than C. Therefore, the total contribution of all the bad simplices is less than C.

Combining these two estimates, we see that the integral of f over $\Psi(P)$ is less than C. Since the sets $\Psi(P)$ are the fibers of π, we see that the k-width of f is less than C. □

In order to bound the k-dilation of diffeomorphisms between rectangles, we will use a variation of this inequality, which holds for sets U lying in a rectangle R. We might expect to find a stronger inequality for a set U confined to a rectangle for the following reason. For a given volume V, the set U with the largest k-width is approximately a round ball of radius $V^{1/n}$. If a round ball of radius $V^{1/n}$ does not fit in the rectangle R, then we must substitute a different shape. It seems reasonable to guess that a rectangle of dimensions $R_1 \times ... \times R_l \times S \times ... \times S$, where $S > R_l$, is the "widest" shape of its volume that fits in R. This intuition turns out to be correct, at least up to a constant factor.

Proposition 3.2.1. *If U is an open set contained in R, then for each integer l between 0 and k, the following inequality holds.*

$$k\text{-}width(U) < C_n(R_1...R_l)^{(n-k)/(n-l)}\, Volume(U)^{(k-l)/(n-l)}.$$

Proof. When $l = 0$, this inequality reduces to the width-volume inequality. When $l = k$, this inequality says that the width of U is less than $CR_1...R_k$. Since U is a subset of R, the width of U is at most the width of R, and this inequality follows. Now we turn to the intermediate values of l.

Let f be a continuous approximation to the characteristic function of U. More precisely, f is a non-negative continuous function supported in R in a small neighborhood of U, with $0 \leq f \leq 1$, $f = 1$ on U, and $\int f$ approximately equal to the volume of U.

Let g be the function on the (n-l)-dimensional rectangle $R_{l+1} \times ... \times R_n$ defined by $g(y) = (R_1...R_l)^{-1} \int_{R_1 \times ... \times R_l} f(x, y)dx$. The function g is continuous, compactly supported, and $0 \leq g \leq 1$. Applying our theorem to the function g, it follows that the (k-l)-width of g is bounded by $C(\int g)^{\frac{k-l}{n-l}}$. This expression is roughly equal to $C[(R_1...R_l)^{-1}\text{Volume}(U)]^{\frac{k-l}{n-l}}$. According to the definition of (k-l)-width, there is some generic PL function π from $\mathbb{R}^{n-l}$ to $\mathbb{R}^{n-k}$ so that the integral of g over each level set is bounded by this expression. We define a smooth map $\bar{\pi}$ on R by $\bar{\pi}(x, y) = \pi(y)$. The volume of U intersected with a fiber $\bar{\pi}^{-1}(z)$ is bounded by $(R_1...R_l)$ times the integral of g over the the fiber $\pi^{-1}(z)$. Therefore, the k-width of U is bounded by $C(R_1...R_l)^{\frac{n-k}{n-l}}\text{Volume}(U)^{\frac{k-l}{n-l}}$. $\square$

3.3 Application to the k-dilation of diffeomorphisms

We will apply our estimates for k-width to lower bound the k-dilation of diffeomorphisms, or more generally to control k-expanding embeddings. An embedding is called k-expanding if it maps each k-dimensional submanifold of the domain to an image with larger k-dimensional volume. If there is a k-contracting diffeomorphism from R to S, then its inverse is a k-expanding embedding of S into R.

In section 1, we showed that if the k-width of a rectangle R is roughly $R_1...R_k$. If there is a k-expanding embedding of S into R, then the k-width of R is at least the k-width of S, and so $R_1...R_k > cS_1...S_k$. A k-expanding embedding is also p-expanding for each p greater than k. Therefore, if there is a k-expanding embedding of S into R then $R_1...R_p > cS_1...S_p$ for each p in the range $k \leq p \leq n$. Next we turn to some harder inequalities based on counting disjoint wide subsets of R and S.

Theorem 3.2. *Suppose that there is a k-expanding embedding of S into R. Then the following inequality holds for each l in the range* $1 \leq l \leq k-1$.

$$(R_1...R_l)^{\frac{n-l}{k-l}} R_{l+1}...R_n > c(S_1...S_l)^{\frac{n-l}{k-l}} S_{l+1}...S_n.$$

Proof. For each l in the range $1 \leq l \leq k-1$, the rectangle S contains N disjoint sub-rectangles with dimensions

$S_1 \times ... \times S_l \times ... \times S_l$,

$N = (S_{l+1}...S_n)/S_l^{n-l}$.

Each of these rectangles has k-width roughly $S_1...S_l S_l^{k-l}$.

Suppose that there is k-expanding embedding of S into R and consider the images of these rectangles in R. Since our embedding is k-expanding, the image of each rectangle has k-width at least $S_1...S_l S_l^{k-l}$. Let V be the smallest volume of any of these images. Next we apply the version of the width-volume inequality in Proposition 3.2.1 to the image with the smallest volume, which gives us the following inequality.

$$S_1...S_l S_l^{k-l} < C(R_1...R_l)^{(n-k)/(n-l)} V^{(k-l)/(n-l)}.$$

Rearranging this inequality gives the following lower bound for V.

$$V > c(S_1...S_l)^{(n-l)/(k-l)} S_l^{n-l} (R_1...R_l)^{-(n-k)/(k-l)}.$$

We began with N disjoint sub-rectangle in S. Their images are N disjoint subsets of R, each with volume at least V. Their total volume is at most the volume of R, which equals $R_1...R_n$. Therefore, we have the following inequality.

$$R_1...R_n > cNV > cN(S_1...S_l)^{(n-l)/(k-l)} S_l^{n-l} (R_1...R_l)^{-(n-k)/(k-l)}.$$

After plugging in the value of N, we get the following inequality.

$$R_1...R_n > c(S_{l+1}...S_n)(S_1...S_l)^{(n-l)/(k-l)} (R_1...R_l)^{-(n-k)/(k-l)}.$$

Moving all of the terms involving R to the left hand side, we are left with the inequality we were trying to prove.

$$(R_1...R_l)^{(n-l)/(k-l)} R_{l+1}...R_n > c(S_1...S_l)^{(n-l)/(k-l)} S_{l+1}...S_n.$$

□

These inequalities are all necessary for a k-expanding embedding of S into R, but they are far from sufficient. We will give a complete answer to the embedding problem in Chapter 6.

When there is a k-contracting diffeomorphism from R to S, then we can apply the above analysis to the boundaries of R and S. With a small amount of extra work, the above proofs adapt to give the following inequalities.

For each p in the range $k+1 \leq p \leq n$, $R_2...R_p > cS_2...S_p$.

For each l in the range $2 \leq l \leq k$, we have the following inequality.

$$(R_2...R_l)^{\frac{n-l}{k-l+1}} R_{l+1}...R_n > c(S_2...S_l)^{\frac{n-l}{k-l+1}} S_{l+1}...S_n.$$

3.4 A generalization of the snake map

The construction of the snake map in the last chapter easily generalizes to produce some non-linear diffeomorphisms with small k-dilation for all k in the range $1 < k < n$. In particular, we will prove the following theorem.

Theorem 3.3. *For every value of k in the range $1 < k < n$, there are n-dimensional rectangles R and S, with a k-contracting diffeomorphism from R to S, such that any linear diffeomorphism from R to S has arbitrarily high k-dilation.*

We begin by constructing some non-linear maps with small 2-dilation, which are just a small generalization of the snake map.

Lemma 3.1. *There is a diffeomorphism from the rectangle R with dimensions $R_1 \times \ldots \times R_n$ to the rectangle S with dimensions $R_1 \times \ldots \times R_{n-2} \times \lambda^{-1}R_{n-1} \times \lambda R_n$ with 2-dilation less than C, for any λ in the range $1 \leq \lambda \leq (R_{n-1}/R_{n-2})$.*

Proof. We pick a quasi-isometric embedding I of $[0, 3\lambda^{-1}R_{n-1}] \times [0, 3\lambda R_{n-2}]$. Let A be the image of the central rectangle $[\lambda^{-1}R_{n-1}, 2\lambda^{-1}R_{n-1}] \times [\lambda R_n, 2\lambda R_n]$ under the embedding I. We define a subset U of R which is equal to the union of $[0, R_1] \times \ldots \times [0, R_{n-2}] \times A$ with $\{0\}^{n-2} \times [0, R_{n-1}] \times [0, R_{n-2}]$. There is a Lipshitz retraction of R to this set U, which takes the boundary of R to the boundary of U. Next we pick a retraction of $[0, R_{n-1}] \times [0, R_n]$ to A, which takes the complement of A to the boundary of A. Taking the direct product of this retraction with the identity gives a retraction from U to $[0, R_1] \times \ldots \times [0, R_{n-2}] \times A$. This latter retraction is the identity on the first part of U, and on the second part of U it maps a 2-dimensional set to a 1-dimensional set in the boundary . Therefore it is 2-contracting. Finally, using the inverse of I, we construct a quasi-isometry from $[0, R_1] \times \ldots \times [0, R_{n-2}] \times A$ to our target rectangle S. The composition of these maps has 2-dilation less than C and takes the boundary of R to the boundary of S with degree 1. Slightly perturbing this map gives a diffeomorphism from R to S. $\square$

Using these maps, along with a little algebra, we construct non-linear diffeomorphisms with small k-dilation.

Corollary. *There is a k-contracting diffeomorphism from R to S if the side lengths of R and S obey the following inequalities.*

1. $R_i > CS_i$ *for* $1 \leq i \leq n-k$.
2. $R_{n-k+1}...R_{n-k+a} > CS_{n-k+1}...S_{n-k+a}$ *for each a in the range* $1 \leq a \leq k$.

Proof. If $R_i \geq S_i$ for every value of i, then there is a contracting linear diffeomorphism from R to S. Let a denote the smallest number so that $R_a < S_a$. We perform a snake map from T to S, where the active directions are $[0, S_1] \times ... \times [0, S_{n-k+1}] \times [0, S_a]$. More precisely, write $S = [0, S_1] \times ... \times [0, S_{n-k+1}] \times [0, S_a] \times S'$, and $T = [0, S_1] \times ... \times [0, S_{n-k}] \times [0, \lambda S_{n-k+1}] \times [0, \lambda^{-1} S_a] \times S'$. We choose λ to be the smaller of the two ratios R_{n-k+1}/S_{n-k+1} and S_a/R_a. Then there is a snake map from T to S with k-dilation less than C.

If the ratio R_{n-k+1}/S_{n-k+1} is smaller than S_a/R_a, then T is equal to $[0, R_{n-k+1}]$ times a lower dimensional rectangle T'. We write R as $[0, R_{n-k+1}] \times R'$. Now it is easy to check that the rectangles R' and T' satisfy the hypotheses of this proposition with $k-1$ in place of k. By induction on the dimension, we can assume that there is a diffeomorphism from R' to T' with (k-1)-dilation less than C. Taking the direct product of this diffeomorphism with the identity map from $[0, R_{n-k+1}]$ to itself gives a diffeomorphism from R to T with k-dilation less than C. Composing with the snake map from T to S gives a diffeomorphism from T to S with k-dilation less than C.

If the ratio S_a/R_a is smaller than R_{n-k+1}/S_{n-k+1}, then T is equal to $[0, S_1] \times ... \times [0, S_{n-k}] \times [0, S_{n-k+1} S_a / R_a] \times [0, R_a] \times S'$. In this case, the rectangle T obeys $R_i \geq T_i$ for i from 1 to $n-k$, $R_{n-k+1}...R_b \geq T_{n-k+1}...T_b$ for all b, and $R_i \geq T_i$ for all i less than or equal to a. By induction on the number a, we can assume that there is a diffeomorphism from R to T with k-dilation less than C. Composing with the snake map from T to S, we get a diffeomorphism from T to S with k-dilation less than C. $\square$

In particular, for every k in the range $1 \leq k \leq n$, there are pairs of rectangles R and S with k-contracting diffeomorphisms from R to S, and so that the k-dilation of any linear diffeomorphism from R to S is arbitrarily high.

Combining the results of the last two sections gives certain upper and lower bounds for the smallest k-dilation of a diffeomorphism from R to S. When $k = n - 1$, the upper bound and the lower bound agree up to a constant factor. (For $k = n - 1$, we get the same upper and lower bounds as in Chapter 2.) On the other hand, when $1 < k < n - 1$, there is an enormous gap between these upper and lower bounds. In Chapters 6-8, we will reduce the gap somewhat, but we will not be able to close it.

Chapter 4

An Isoperimetric Inequality Involving Packings by Wide Sets

The main result of this chapter is the following variation of the isoperimetric inequality.

Theorem 4.1. *Let U be a bounded open set in $\mathbb{R}^n$ with smooth boundary. Then the boundary of U contains disjoint sets S_i so that the following inequality holds.*

$$\text{Volume}(U) < C \sum \text{(n-2)-width}(S_i)^{\frac{n}{n-2}}.$$

4.1 A ball covering argument

This section gives the first half of the proof of Theorem 4.1.

We need to find some disjoint sets S_i in the boundary of U with large (n-2)-width. The sets S_i will be the intersections of the boundary of U with some disjoint balls. We will call a ball $B(p, R)$ good if the (n-2)-width of $\partial U \cap B(p, R)$ is at least cR^{n-2}.

We will prove that for each point p in U, there is a positive radius R so that $B(p, R)$ is a good ball. Fix a point p in U. For each positive number R, we define a map ϕ_R from Euclidean space to the unit n-sphere, which maps $B(p, R)$ to the complement of the south pole using the exponential map and the rest of $\mathbb{R}^n$ to the south pole. The

Lipshitz constant of ϕ_R is π/R. Next we define a metric g on $\partial U \times (0, \infty)$, where we parameterize $(0, \infty)$ with the variable R. The metric g restricted to the slice $\partial U \times \{R\}$ is very small away from $\partial U \cap B(p, R)$, and is equal to π/R times the given metric on $\partial U \cap B(p, R)$. The metric g is very large transverse to the slices $\partial U \times \{R\}$. Finally, we define the map Φ from $\partial U \times (0, \infty)$ to the unit n-sphere by $\Phi(x, R) = \phi_R(x)$. Using the metric g on the domain and the standard metric on the unit sphere, the map Φ is a contracting map. For some very large number λ, the map Φ takes $\partial U \times \{\lambda^{-1}\}$ to the south pole and $\partial U \times \{\lambda\}$ into a very small neighborhood of the North pole. We compose Φ with a degree 1 map that pinches this small neighborhood of the North pole to the North pole and leaves the South pole invariant. This pinching map can have Lipshitz constant as close as we like to 1. The composition is a degree 1 map from $\partial U \times [\lambda^{-1}, \lambda]$ to the unit n-sphere, which maps each boundary component of the domain to a single point of the range.

In the next section, we will prove a slicing lemma, which guarantees that one of the slices $\partial U \times \{R\}$ has (n-2)-width at least c, using the restriction of the metric g. The metric g restricted to this slice is essentially equal to the given metric on $\partial U \cap B(p, R)$ rescaled by a factor π/R. Therefore, the (n-2)-width of $\partial U \cap B(p, R)$ is at least cR^{n-2}. This proves that the ball $B(p, R)$ is good for at least one value of R.

According to the Vitali covering lemma, there are disjoint good balls $B(p_i, R_i)$ with $\sum R_i^n > c$ Volume(U). We define S_i to be the intersection of the boundary of U with $B(p_i, R_i)$. Since the (n-2)-width of S_i is at least cR_i^{n-2}, we see that $\sum$ (n-2)-width$(S_i)^{n/(n-2)}$ is at least $c \sum R_i^n \geq c$ Volume(U).

This finishes the proof of our theorem, except for the slicing lemma, which we will prove in the next section.

4.2 Slicing inequalities

Let me first explain what I mean by a slicing inequality in a fairly general setting. Suppose that we have a Riemannian metric g on $\bar{M} = M \times [0, 1]$, for some manifold M. Next, suppose that we know that $\bar{M}$ is large in some sense. For example, we might know that there is a contracting map of non-zero degree from $\bar{M}$ to the unit sphere which takes the boundary of $\bar{M}$ to a point. For each t between 0 and 1, we consider the restriction of g to $M \times \{t\}$ to get a metric g_t on M. The Riemannian manifolds (M, g_t) are "slices" of $\bar{M}$. What can we conclude about these slices? Is it true that at least one slice is large in some sense?

Now we state our main inequality.

Lemma 4.1. *Let M be a manifold of dimension m. Let g be a Riemannian metric on $\bar{M} = M \times [0,1]$ and define g_t to be the restriction of g to $M \times \{t\}$. Then the following inequality holds.*

$$W_{m-1}(\bar{M}, g) \leq 2 \sup_t W_{m-1}(M, g_t).$$

In this equation, W_{m-1} denotes the (m-1)-width.

Warning: This inequality does NOT generalize to k-width for $k < m-1$.

Proof. By assumption, for each t we can find a generic PL map π_t from the slice $M \times \{t\}$ to $\mathbb{R}$ whose fibers have volume less than W. If the maps π_t varied continuously in t, then we could define a continuous map π from $M \times [0,1]$ to $\mathbb{R}^2$ by $\pi(x,t) = (\pi_t(x), t)$, and all the fibers would have volume less than W. The maps π_t may not vary continuously, and our proof consists of a trick to turn them into a continuous family. We will choose generic PL maps π_i from M to $\mathbb{R}$ for i from 1 to some large number N, so that the following holds. We define g_i to be the supremum of g_t for all t in the range $(i-1)/N \leq t \leq (i+1)/N$. For sufficiently big N, we can arrange that for each y in $\mathbb{R}$, $\pi_i^{-1}(y)$ has volume less than $W + \epsilon$ with respect to the metric g_i.

Therefore, it suffices to solve the following homotopy problem. Given two generic PL maps, π_0 and π_1, from an m-dimensional Riemannian manifold (M, g) to the real numbers, all of whose fibers have volume less than W, we will construct a homotopy F_t with $F_0 = \pi_0$ and $F_1 = \pi_1$ so that for each t, each fiber of F_t has volume less than 2W. We first reparameterize the ranges so that the images of π_0 and π_1 are both contained in the interval $(0,1)$. Next, we define $F_t(x) = min(\pi_0(x) + t, \pi_1(x) + 1 - t)$. The map F is continuous in both t and x. Moreover, F is a piecewise linear map from

$(M, g) \times [0,1]$ to $\mathbb{R}^2$. After we slightly perturb π_0 and π_1, the map F will be generic. Finally, for each t, each fiber of F_t is contained in the union of a fiber of π_0 and a fiber of π_1. Therefore, each fiber of F_t has volume less than 2W. $\square$

We would like to apply this lemma to the manifold $\partial U \times [\lambda^{-1}, \lambda]$ with its metric g, which were defined in section 1. We know from section 1 that there is a degree 1 map from $(\partial U \times [\lambda^{-1}, \lambda], g)$ to the unit sphere, with Lipshitz constant less than 2, and which collapses each boundary component to a point. According to the corollary after Proposition 3.1.1, the (n-2)-width of of $(\partial U \times [\lambda^{-1}, \lambda], g)$ is at least c. Our lemma now guarantees guarantees that one of the slices $\partial U \times \{R\}$ has (n-2)-width at least c/2 (where the width is measured using the restriction of the metric g). This estimate is exactly the one we used in section 1.

In the paper [14], Gromov proved a different slicing inequality involving the Uryson width. Recall that the Uryson p-width of a Riemannian manifold (M, g) is the smallest number W so that there is a continuous map f from M to a p-dimensional polyhedron whose fibers have diameter less than W. (Warning: The k-width in this paper refers to the sizes of k-dimensional fibers in a manifold, whereas the Uryson p-width refers to the sizes of fibers of maps to a p-dimension polyhedron.)

Lemma. *(Gromov) If f is a continuous function from (M, g) to $\mathbb{R}$ and each level set of f has Uryson k-width less than 1, then (M, g) has Uryson (2k+1)-width less than 1.*

If we use Gromov's slicing lemma instead of the slicing lemma at the beginning of this section, we can prove a different isoperimetric inequality.

Corollary. *Let U be a bounded open set in $\mathbb{R}^{2n}$ with smooth boundary. Then the boundary of U contains disjoint sets S_i so that the following inequality holds.*

$$\textit{Volume}(U) < C \sum UW_{n-1}(S_i)^{2n}.$$

In this formula, UW_{n-1} denotes the Uryson (n-1)-width. If U is a bounded open set in $\mathbb{R}^{2n+1}$ with smooth boundary, then the boundary of U contains disjoint open sets S_i so that the following inequality holds.

$$Volume(U) < C \sum UW_{n-1}(S_i)^{2n+1}.$$

Proof. The proof follows exactly the proof in section 1, except that each reference to k-width needs to be replaced by an appropriate reference to an Uryson width. We call a ball $B(p, R)$ good if the Uryson (n-1)-width of $\partial U \cap B(p, R)$ is at least cR. We need to show that for each point p inside U, there is a radius R so that $B(p, R)$ is good.

Following the arguments in section 1, we know that there is a degree 1 map from $(\partial U \times [\lambda^{-1}, \lambda], g)$ to the unit sphere with Lipshitz constant 2. According to Gromov's paper [14], the Uryson p-width of this manifold must be at least c, where p is one less than the dimension of $\partial U \times [\lambda^{-1}, \lambda]$. Applying Gromov's slicing inequality, it follows that one of the slices $(\partial U \times \{R\}, g)$ has Uryson (n-1)-width at least c. Since this slice is essentially the intersection $\partial U \cap B(p, R)$ rescaled by the factor π/R, the intersection itself must have Uryson (n-1)-width at least cR. Therefore, for each point p in U, there is a radius R so that $B(p, R)$ is good.

Applying the Vitali covering lemma as in section 1 proves the corollary. ☐

There may well be other interesting slicing inequalities.

Chapter 5

The Volume Enclosure Problem

In this section, we estimate how much volume can be enclosed by a k-contracting embedding of an ellipsoid.

Theorem 5.1. *Let E be an n-dimensional ellipsoid with principal axes $E_0 \le ... \le E_n$. We will define below a monomial $V_k(E)$ in the principal axes of E so that any k-contracting embedding of E into $\mathbb{R}^{n+1}$ bounds a region of volume less than $CV_k(E)$. On the other hand, we will construct a k-contracting embedding of E which bounds a volume greater than $cV_k(E)$.*

The monomial $V_k(E)$ has the following form.

$$V_k(E) = E_1^{\frac{n+1}{k}} ... E_{l-1}^{\frac{n+1}{k}} E_l^b E_{l+1} ... E_n.$$

To specify $V_k(E)$ exactly, we need to pick a value of l and a value of b. These values are uniquely determined by two conditions. First, the total degree of $V_k(E)$ is $n+1$. Second, the number b lies in the range $1 < b \le \frac{n+1}{k}$.

The proof of the theorem follows the same lines as the isoperimetric inequality in Chapter 4. In place of the slicing inequality that we used in Chapter 4, we will use a version of the width-volume inequality for families.

In section 1, we give the proof of the main theorem assuming a lemma. In section 2, we explain and prove a width-volume inequality for families of functions. The inequality in section 2 is the simplest and most natural version of the width-volume inequality for families, but the lemma that we need is a small modification of it. We prove the lemma in section 3.

5.1 The outline of the proof

We begin with the easy direction of the proof, constructing a k-contracting embedding of E into $\mathbb{R}^{n+1}$ which encloses a volume greater than $cV_k(E)$.

When $k = 1$, the monomial $V_k(E)$ is simply $E_1^2 E_2 ... E_n$. It is not difficult to construct a 1-contracting map bounding this much volume. The ellipsoid E is C-bilipshitz equivalent to the double of a rectangle with side lengths $E_1 \leq ... \leq E_n$. The dimensions of this rectangle do not depend on E_0 at all, and so E is in particular bilipshitz to the ellipsoid E' with principal axes $E'_0 = E_1$ and $E'_i = E_i$ for all other i. The ellipsoid E' encloses a volume proportional to $E_1^2 E_2 ... E_n$. After rescaling the bilipshitz map from E to E', we get a 1-contracting embedding of E which bounds a volume greater than $cE_1^2 E_2 ... E_n$.

For all k less than or equal to $(n+1)/2$, the monomial $V_k(E)$ is $E_1^2 E_2 ... E_n$. When k increases beyond $(n+1)/2$, the volume $V_k(E)$ finally increases, and we need to construct new maps to enclose volume $V_k(E)$. The ellipsoid E contains approximately $N = E_{l+1} ... E_n / E_l^{n-l}$ disjoint rectangles of dimension $E_1 \times ... \times E_{l-1} \times E_l \times E_l \times ... \times E_l$. (The number l here has the same value as in the statement of Theorem 5.1.) Each of these rectangles admits a k-contracting diffeomorphism to the cube with sidelength $S = (E_1 ... E_l E_l^{k-l})^{1/k}$. Therefore, each rectangle admits a k-contracting degree 1 map to the n-sphere of radius cS. There is a k-contracting embedding of E into $\mathbb{R}^{n+1}$ which takes each rectangle in our collection to an n-sphere of radius cS and which takes the rest of E to a bunch of very thin tubes connecting them. A little more precisely, there will be one thin tube leaving each n-sphere and running to a central small sphere. The total volume enclosed by the image of E is greater than cNS^{n+1}, which is equal to $cV_k(E)$.

It remains to prove that the volume enclosed by the image of E is less than $CV_k(E)$. The proof has two main ideas: a version of the width-volume inequality for families, and the ball packing argument from Chapter 4. The version of the width-volume inequality that we need is the following lemma.

Lemma 5.1. *Let R be an n-dimensional rectangle, and let f be a continuous compactly supported function on $R \times \mathbb{R}$. Suppose that $0 \leq f \leq 1$. Let x be a coordinate on R and y a coordinate on $\mathbb{R}$. Let δ be any number greater than zero. Then there exists a generic PL map from $R \times \mathbb{R}$ to $\mathbb{R}^{n-k} \times \mathbb{R}$ with the following properties.*

1. *For each l in the range $0 \leq l \leq k$,*

$$(\int_{\pi^{-1}(z,y)} f)^{\frac{n-l}{k-l}} < C(R_1 ... R_l)^{\frac{n-l}{k-l}-1}(\delta + \int_{R\times\{y\}} f).$$

2. *The fiber $\pi^{-1}(z,y)$ lies within δ of the plane $R \times \{y\}$.*

We defer the proof of this lemma to sections 2 and 3. In order to keep track of the complicated exponents in the inequality, it helps to notice the scaling of each side. Each length R_i has degree 1. The integral of f over a k-dimensional fiber $\pi^{-1}(z,y)$ has degree k. The integral of f over the n-dimensional rectangle $R \times \{y\}$ has degree n. Using these degrees, each side has total degree $k(\frac{n-l}{k-l})$.

The inequality that we will need in our proof is the geometric average of two of the inequalities in 1. First, we choose the integer l so that $\frac{n-l+1}{k-l+1} < \frac{n+1}{k} \leq \frac{n-l}{k-l}$. We can write $\frac{n+1}{k}$ as an average $(1-a)(\frac{n-l+1}{k-l+1}) + a(\frac{n-l}{k-l})$, where a lies in the range $0 < a \leq 1$. Now we take the geometric average of inequality 1 using l-1 and l, weighted by $1-a$ and a. We get the following inequality.

$$(\int_{\pi^{-1}(z,y)} f)^{\frac{n+1}{k}} < C(R_1 ... R_{l-1})^{\frac{n+1}{k}-1} R_l^{b-1}(\delta + \int_{R\times\{y\}} f).$$

In this equation, the number b is short for the ugly expression $(1-a)+a(\frac{n-l}{k-l})$. The numbers l and b in this formula are the same numbers that appear in the definition of $V_k(E)$. An easy calculation shows that b lies in the range $1 < b \leq \frac{n+1}{k}$. By scaling considerations, the total degree of the monomial $(R_1 ... R_{l-1})^{\frac{n+1}{k}-1} R_l^{b-1}$ is equal to 1. Therefore, we have the following formula.

$$(\int_{\pi^{-1}(z,y)} f)^{\frac{n+1}{k}} < C(R_1 ... R_n)^{-1} V_k(R)(\delta + \int_{R\times\{y\}} f). \quad (1)$$

The proof of our theorem now follows from the ball-packing argument of Chapter 4, after substituting this inequality for the slicing lemma used there.

Let f be a k-contracting embedding of E into $\mathbb{R}^{n+1}$, and let U be the open set with boundary the image of f. Pick a point p in U. We define $V(p,R)$ to be the volume of $f^{-1}(B(p,R))$ in E. We say that the ball $B(p,R)$ is good if the volume obeys the following inequality.

$$V(p,R) > c(E_1...E_n)V_k(E)^{-1}R^{n+1}.$$

For each point p in U, we will prove that for some radius R the ball $B(p,R)$ is good. For each R, we consider the (modified) exponential map from $\mathbb{R}^{n+1}$ to the unit (n+1)-sphere which takes $B(p,R)$ to the complement of the south pole and takes the rest of $\mathbb{R}^{n+1}$ to the south pole. Composing with f we get a map Φ_R from E to the unit (n+1)-sphere. We view this family of maps as a single map from $E \times (0,\infty)$ to the unit (n+1)-sphere. Next we put a metric g on $E \times (0,\infty)$ as follows. The restriction of g to each slice $E \times \{R\}$ is very small away from $f^{-1}(B(p,R))$ and is equal to the standard metric rescaled by π/R on $f^{-1}(B(p,R))$. Transverse to the slices, g is very large. We now have a map Φ from $(E \times (0,\infty), g)$ to the unit (n+1)-sphere. This map Φ is k-contracting. Because p lies inside of U, the map Φ has degree 1. For some large number λ, we can slightly alter Φ to get a degree 1 map from $(E \times [\lambda^{-1},\lambda], g)$ to the unit (n+1)-sphere taking each boundary component to a point.

Let U be the open set whose intersection with $E \times \{R\}$ is equal to $f^{-1}(B(p,R))$. We think of E as the double of a rectangle R with side lengths $E_1 \leq ... \leq E_n$, and we let f be a continuous approximation to the projection of U onto $R \times [\lambda^{-1},\lambda]$. We apply Lemma 5.1 to the function f. The lemma gives us a generic PL map π from $\mathbb{R}^n \times [\lambda^{-1},\lambda]$ to $\mathbb{R}^{n-k} \times \mathbb{R}$. The integral of f over $\mathbb{R}^n \times \{R\}$ is approximately $V(p,R)$. According to equation 1, the integral of f over the fiber $\pi^{-1}(y,R)$ is less than the following expression.

$$C[(E_1...E_n)^{-1}V_k(E)(\delta + V(p,R))]^{\frac{k}{n+1}}.$$

By abuse of notation, we will also use π to indicate the PL map from $E \times [\lambda^{-1},\lambda]$ which we get by first projecting from the E to R and then applying π. Since f is at least one on the image of U, the volume of $\pi^{-1}(y,R)$ is bounded by the same expression. Lemma 5.1 also tells us that the fiber $\pi^{-1}(y,R)$ lies within δ of the ellipsoid $E \times \{R\}$. For sufficiently small δ, we conclude that the volume of the fiber $\pi^{-1}(y,R)$ with respect to the metric g is bounded by the following expression.

$$CR^{-k}[(E_1...E_n)^{-1}V_k(E)(\delta + V(p,R))]^{\frac{k}{n+1}}. \quad (2)$$

As we noted above, the manifold $(E \times [\lambda^{-1}, \lambda], g)$ admits a k-contracting degree 1 map to the unit sphere. According to the corollary after Proposition 3.1.1, it has k-width at least c. Therefore, one of the fibers $\pi^{-1}(y,R)$ must have volume at least c. Given the upper bound in equation 2, it follows that for some R in the range $[\lambda^{-1}, \lambda]$, the volume of $f^{-1}(B(p,R))$ obeys the following inequality.

$$V(p,R) > c(E_1...E_n)V_k(E)^{-1}R^{n+1} - \delta.$$

Finally, taking δ very small compared to λ^{-1}, we see that the ball $B(p,R)$ is good for some radius R in the range $[\lambda^{-1}, \lambda]$.

Now by the Vitali covering lemma, we can find disjoint good balls $B(p_i, R_i)$ with $\sum R_i^{n+1} > c$Volume(U). By the definition of a good ball, the region $f^{-1}(B(p_i,R_i))$ in E has volume at least $cR_i^{n+1}(E_1...E_n)V_k(E)^{-1}$. Since the balls $B(p_i,R_i)$ are disjoint, their inverse images are disjoint subsets of E, and so the total volume is bounded by $CE_1...E_n$. This bound shows that $\sum R_i^{n+1} < CV_k(E)$. Therefore, Volume(U) $<$ $\mathrm{C}\sum \mathrm{R}_i^{n+1} < \mathrm{CV_k(E)}$.

This finishes the proof of our theorem, except for the proof of the lemma, which we will give in the next sections.

5.2 The width-volume inequality for families

In this section, we will prove an analog of the width-volume inequality that holds for families of functions.

Here is the setup. Let f be a compactly supported continuous function on the product $\mathbb{R}^n \times \mathbb{R}^q$, $0 \leq f \leq 1$. Let x be a coordinate on $\mathbb{R}^n$ and y be a coordinate on $\mathbb{R}^q$. We can think of f as a family of functions on $\mathbb{R}^n$ parameterized by $\mathbb{R}^q$. Similarly, we give coordinates (z, y) to $\mathbb{R}^{n-k} \times \mathbb{R}^q$ where z is a coordinate on $\mathbb{R}^{n-k}$ and y is a coordinate on $\mathbb{R}^q$.

Theorem 5.2. *For each $\delta > 0$, there is a generic PL map π from $\mathbb{R}^n \times \mathbb{R}^q$ to $\mathbb{R}^{n-k} \times \mathbb{R}^q$ obeying the following estimate for the integral of f over fibers.*

$$\int_{\pi^{-1}(z,y)} f < C(\delta + \int_{\mathbb{R}^n \times \{y\}} f)^{k/n}.$$

Moreover, the fiber $\pi^{-1}(z, y)$ lies within δ of the plane $\mathbb{R}^n \times \{y\}$.

(The constant C depends on n and q.)

Before we prove this theorem, let us mention a simpler corollary. If $\int_{\mathbb{R}^n \times \{y\}} f < 1$ for every y, then the k-width of f is less than C.

Proof. The first step of the proof is to set up a scaffold S_0 analogous to the k-skeleton of the unit lattice in the proof of the width-volume inequality. The scaffold S resembles a lattice on large balls, but the lattice spacing needs to vary from one region to another because the bound we want to prove varies from one region to another.

Instead of the usual Euclidean metric, we are going to work with the metric $dx^2 + Ady^2$, for some number A which is enormous compared to δ^{-1} or to the size of the support of f. We define a function $L(y) = (\delta + \int_{\mathbb{R}^n \times \{y\}} f)^{1/n}$. Because A is so large, the function $L(y)$ is almost constant on very large balls in this metric. Many parts of the proof will take place on bounded regions, and up to a small error, it is possible to treat L as being constant on these regions.

The metric $dx^2 + Ady^2$ is a Euclidean metric, and we consider the unit lattice in this Euclidean space, with axes parallel to the x and y coordinates. Next we consider the barycentric subdivision of this unit lattice. We define a mapping on the vertices of this triangulation, which maps a vertex with coordinates (x, y) to the point $(L(y)x, y)$. This mapping extends to a PL mapping which is linear on each simplex of the barycentric subdivision. Because L varies so slowly as y changes, this PL mapping is a PL isomorphism. We define S_0 to be the image of the k-skeleton of the unit lattice under this barycentric map.

Next, we define a Riemannian metric g on $\mathbb{R}^N = \mathbb{R}^n \times \mathbb{R}^q$ with $ds^2 = dx^2 + AL(y)^2dy^2$. Because $L(y)$ changes so slowly, this metric appears very close to Euclidean on very large balls. With respect to this metric, the k-skeleton S_0 appears on large balls to be very close to the k-skeleton of the cubical lattice of side length L. Moreover, because the number A is so large, this metric is strictly larger than the original Euclidean metric $dx^2 + dy^2$.

The next step of the proof of the width-volume inequality was to translate the k-skeleton S_0 to a k-skeleton S with $\int_S f$ bounded in terms of the integral of f. In the present situation, we need a localized version of this construction. Namely, we will "wiggle" the skeleton S_0 to get a k-skeleton S, so that the integral of f over each face F of S is bounded in terms of the integral of f on a ball near F. Next we give an exact description of the wiggling procedure, which I call barycentric wiggling.

The k-skeleton S_0 was defined as the image of the k-skeleton of the unit lattice by a PL map. We are going to perturb this PL map. The original PL map took a vertex of the barycentric triangulation located at (x, y) to the point $(L(x), y)$. A perturbation of this map is allowed to take the vertex at (x, y) to any point in a ball of radius $2^{-N}L$ around the original target $(L(x), y)$. (From here on, every ball mentioned in the proof is defined using the metric g.) This perturbed map on the vertices also extends to a PL map which is linear on each simplex of the barycentric subdivision. Because L is essentially constant and because $2^{-N}L$ is small compared to L, this map is still a PL isomorphism, and it is bilipshitz equivalent to the original map. We call this map Φ. The skeleton S is defined to be the image of the k-skeleton of the unit lattice under Φ. We claim that for a carefully chosen map Φ, the integral of f over each k-face of S obeys the following inequality.

$$\int_{\Phi(F)} f < CL^{-(N-k)} \int_{B(\Phi(c(F)),2^N L)} f.$$

(In this equation, F denotes a k-face of the unit lattice, $c(F)$ denotes its center, and B denotes a ball defined using the metric g.)

Before we get started, let me make a few remarks. If we take a given k-face F and simply translate its image $\Phi(F)$ by a random vector with length on the order of $2^{-N}L$, then on average the integral of f over the translation will be bounded by the expression on the right-hand side of the last inequality. If we could translate all the faces independently of one another then we could easily choose Φ to make the above inequality hold for every face F. We cannot, however, move the k-faces independently of each other, because they intersect in (k-1)-faces. Roughly speaking, the barycentric perturbation is the best perturbation that we can do without tearing the skeleton apart.

To define our perturbed PL isomorphism Φ, we have to define its value on each vertex of the barycentric triangulation. The vertices of the barycentric triangulation correspond to the faces of the unit lattice. We are going to define Φ first for the vertices corresponding to 0-faces of the unit lattice, then for the vertices corresponding to 1-faces, and so on. After we have defined Φ on the vertices corresponding to all faces of dimension up to p, the restriction of Φ to the p-skeleton is determined.

First, we define Φ on the vertices of the unit lattice. For each vertex v, we have to pick a value w for $\Phi(v)$ which lies in certain allowed ball of radius $2^{-N}L$. For an average choice of w, the following inequality holds.

$$\int_{B(w,2^{N-1}L)} dist(x,w)^{-(N-k)} f(x)dx < CL^{-(N-k)} \int_{B(w,2^{N}L)} f(x)dx.$$

The average value of the left hand side as w varies over a ball of radius $2^{-N}L$ is bounded by $CL^{-N}\int_{B(2^{-N}L)}(\int_{B(w,2^{N-1}L)} dist(x,w)^{-(N-k)}f(x)dx)dw$. But the integral $\int_{B(2^{-N}L)} dist(x,w)^{-(N-k)}dw$ is less than CL^k. Plugging this fact into the expression for the average gives the inequality above.

Next, we proceed inductively to define Φ on the center of each p-face of the unit lattice, for p from 1 to k-1. Suppose that we have to define Φ on the center of a q-face F. Let the boundary of F be E_i. Let w be the value of Φ at the center of F. The point w is constrained to lie in a certain ball of radius $2^{-N}L$. Since we have already defined Φ on the (p-1)-skeleton of the unit lattice, choosing w defines Φ on the face F. For an average choice of w, the following inequality holds.

$$\int_{B(w,2^{N-p-1}L)} dist(x,\Phi(F))^{-(N-k)} f(x)dx$$

$$< C\sum_i \int_{B(w,2^{N-p}L)} dist(x,\Phi(E_i))^{-(N-k)} f(x)dx.$$

To see this, we consider the barycentric subdivision of the face F. For each p-simplex Δ, we can consider $\int dist(x,\Phi(\Delta))^{-(N-k)} f(x)dx$. If we can bound this integral for each Δ we are done. The p-simplex Δ can be thought of as a pyramid whose base is a (p-1)-simplex Δ' lying in $\Phi(E_i)$ and whose apex is at w. The plane spanned by Δ depends only on the projection of w to the (N-p+1)-plane normal to Δ', and in particular it only depends on the direction of the projected vector on S^{N-p}, which we call ϕ. If we write the projection of x onto the plane normal to Δ' in polar coordinates (r,θ), then $dist(x,\Phi(\Delta))$ is at least $crdist(\theta,\phi)$. The average value that we are trying to calculate is bounded by $C\int d\phi(\int r^{-(N-k)} dist(\theta,\phi)^{-(N-k)} f(x)d\theta)$. Since $(N-p)$ is strictly bigger than $(N-k)$, the integral $\int dist(\theta,\phi)^{-(N-k)} d\phi$ converges, and so our average value is bounded by $C\int r^{-(N-k)} f(x)dx$. Unwinding the definitions, we see that this expression is bounded by $C\int dist(x,\Phi(E_i))^{-(N-k)} f(x)dx$.

Composing all the inequalities that we have proven so far, we see that for each (k-1)-face E of the unit lattice, we have the following inequality.

$$\int_{B(2^{N-k}L)} dist(x,\Phi(E))^{-(N-k)} f(x)dx < CL^{-(N-k)} \int_{B(2^N L)} f(x)dx.$$

Finally, we have to define Φ at the center point of a k-face F. The value w of Φ at this center point is constrained to lie in a ball of radius $2^{-N}L$. Let the boundary of F be equal to the union of some (k-1)-faces E_i. We have already defined Φ on the E_i. For an average value of w, the integral of f over $\Phi(F)$ is bounded by the following inequality.

$$\int_{\Phi(F)} f(x)dx < C\sum_i \int_{B(2^{N-k}L)} dist(x,\Phi(E_i))^{-(N-k)} f(x)dx.$$

To see this, we again consider the barycentric subdivision of the face F. It suffices to bound the average value of $\int_{\Phi(\Delta)} f(x)dx$ for each k-simplex Δ in this subdivision of

F. The image $\Phi(\Delta)$ is a k-simplex with base Δ' lying in one of the images $\Phi(E_i)$ and with apex w. By a routine calculation, the average value of $\int_\Delta f(x)dx$ is bounded by $C\int_{B(2^{N-k}L)} dist(x, \Phi(E_i))^{-(N-k)} f(x)dx$.

Composing the last two inequalities gives us the bound that we wanted to prove.

$$\int_{\Phi(F)} f < CL^{-(N-k)} \int_{B(\Phi(c(F)),2^N L)} f.$$

This estimate finishes the construction of the scaffold S. The next step in the proof of the width-volume inequality is to define a PL isomorphism Ψ of $\mathbb{R}^N$ that squeezes most of space into a small neighborhood of S. In Chapter 3, we constructed a periodic PL isomorphism Ψ that squeezed most of space into a small neighborhood of the unit lattice. The map Ψ squeezed the complement of an ϵ-neighborhood of the (N-k-1)-skeleton of the dual lattice into the ϵ-neighborhood of the k-skeleton of the unit lattice. The ϵ-neighborhood of the dual (N-k-1)-skeleton is called the bad region, and the complement is called the good region. The good region inside of a fundamental cube is divided into C simplices. On each good simplex Δ, the map Ψ is a linear map which takes that simplex into the ϵ-neighborhood of one of the k-faces of the k-skeleton. Finally, the map Ψ on Δ maps almost every k-plane to a k-plane almost parallel to the k-skeleton of the unit lattice.

We will use a squeezing map Ψ' defined by $\Psi' = \Phi \circ \Psi \circ \Phi^{-1}$. Since $\Phi^{-1}(S)$ is the k-skeleton of the unit lattice, Ψ' squeezes most of space into an $L\epsilon$ neighborhood of S.

The next step in the proof of the width-volume inequality was to choose a projection π_0 from $\mathbb{R}^n$ to $\mathbb{R}^{n-k}$ at a good angle with respect to Ψ. In our case, we would like to choose a projection at a good angle with respect to Ψ'. Globally, it is not clear that such a good angle exists, so we will have to define a PL map π_0.

In a given region where L is almost constant, we call a generic linear map to $\mathbb{R}^{n-k}$ good if its fibers are at an angle that gets mapped to be almost tangent to S by the restriction of Ψ to each good simplex. An orthogonal projection to a random (N-k)-plane, chosen with respect to the locally almost constant metric $dx^2 + AL(y)^2dy^2$ is good with high probability. We look at the barycentric subdivision of some triangulation of $\mathbb{R}^N$ with simplices that are fairly standard and quite large compared to $L(y)$ but quite small compared to A. Then we define the PL projection π_0 by taking a barycentric perturbation of the projection onto the last $(N-k)$ coordinates. (We perturb the values of the barycentric subdivision at a given vertex within a small region shaped like the image of the projection map from simplices bordering the given vertex.) To define this barycentric perturbation, we must first choose its value on the vertices corresponding to 0-faces, then on the vertices corresponding to 1-faces, and so on. We have defined a good map on the n-dimensional simplices. Now we inductively define a good map on lower dimensional simplices by saying that a map on a q-face is good if almost all of its extensions to the neighboring (q+1)-faces are good. When ϵ is very small, it follows inductively that most maps from a given simplex are good. Then we define π_0 starting at the 0-faces so that it is good at each step. In this way, we produce a good map π_0.

Let us pick a large cube $[-S, S]^n$ in $\mathbb{R}^n$ so that the support of f is easily contained in $[-S, S]^n \times \mathbb{R}^q$. The fiber $\pi^{-1}(z, y)$, intersected with this cylinder, lies within a ball of radius CS around the point $(0, y)$. The projection from the fiber intersect this region to the first k coordinates has image $[-S, S]^k$ and is bilipshitz, with bilipshitz constant less than C.

Now let P be the fiber $\pi_0^{-1}(z, y)$. According to our wiggling inequality, the sum of the integral of f over all the k-faces of S that come within CL of the fiber P is less than the following expression.

$$C \sum_{\text{faces near P}} L^{-(N-k)} \int_{B(\text{face,2L})} f.$$

Since f is essentially constant in the y direction, this expression is less than $CL^{-(N-k)} L^q \int_{\mathbb{R}^n \times \{y\}} f$, which is less than CL^k.

On the other hand, when ϵ is very small, any fiber of π_0 meets less than C bad simplices.

Finally, we define our projection π to be the composition $\pi_0 \circ \Phi \circ \Psi^{-1} \circ \Phi^{-1}$. As usual, we need only be concerned about the part of the fibers of π near to the support of f. The fiber $\pi^{-1}(z,y)$ lies within a distance $L(y)$ of the fiber $\pi_0^{-1}(z,y)$. It is pressed close to the k-skeleton S except on a region of volume bounded by CL^k. The integral of f along the part of $\pi^{-1}(z,y)$ which is pressed close to S is bounded by a constant times the sum of the integral of f on all the faces of S lying within CL of the fiber $\pi_0^{-1}(z,y)$. We have just seen that this sum is bounded by CL^k. On the other hand, since f is no more than 1, the integral of f on the remainder of the fiber $\pi^{-1}(z,y)$ is bounded by CL^k. Therefore, the total integral of f over the fiber $\pi^{-1}(z,y)$ is bounded by CL^k. Recalling the definition of L, we see that this integral is bounded by $C(\delta + \int_{\mathbb{R}^n \times \{y\}} f(x)dx)^{k/n}$. So far, we have bounded the integral of f over the fiber $\pi^{-1}(z,y)$ in the metric g. Since the metric g is bigger than the original Euclidean metric on $\mathbb{R}^n$, it follows that the same estimate holds using the Euclidean metric. This estimate is the bound that we wanted to prove.

The intersection of the fiber $\pi_0^{-1}(z,y)$ with the region $[-S,S]^n \times \mathbb{R}^q$ lies in the ball of radius CS around $(0,y)$. Therefore, the intersection of the fiber $\pi_0^{-1}(z,y)$ with $[-S,S]^n \times \mathbb{R}^q$ lies in the ball of radius $C(S+L)$. These balls are defined using the metric $g = dx^2 + AL(y)^2 dy^2$. Since L is at least δ and A is very very large, it follows that the fiber $\pi^{-1}(z,y)$ is actually as close as we like to the plane $\mathbb{R}^n \times \{y\}$ in the original Euclidean metric. □

5.3 The width-volume inequality for families inside a rectangle

In the last section, we proved a families version of the width-volume inequality in Theorem 3.1. In order to prove Theorem 5.1, we need a families version of the width-volume inequality for subsets of rectangles in Proposition 3.2.1. Let us recall the statement of Lemma 5.1.

Lemma. *Let R be an n-dimensional rectangle, and let f be a continuous compactly supported function on $R \times \mathbb{R}$. Suppose that $0 \leq f \leq 1$. Let x be a coordinate for R and y a coordinate for $\mathbb{R}$. Let δ be any number greater than zero. Then there exists a* *generic PL map from $R \times \mathbb{R}$ to $\mathbb{R}^{n-k} \times \mathbb{R}$ with the following properties.*

1. *For each l in the range $0 \le l \le k$,*

$$\int_{\pi^{-1}(z,y)} f < C(R_1 ... R_l)^{\frac{n-k}{n-l}} (\delta + \int_{R \times \{y\}} f)^{\frac{k-l}{n-l}}.$$

2. *The fiber $\pi^{-1}(z, y)$ lies within δ of the plane $R \times \{y\}$.*

Proof. This proof is a small modification of the proof of the width-volume inequality for families. We will go through the steps of that proof and describe the alterations that need to be made at each step.

The first step is to construct a rough scaffold S_0. We need to choose the size of this scaffold in a different way, and we need to do a bit of work to make sure that the scaffold is transverse to the rectangle R.

Define $L(y)$ to be the smallest value of $[(\delta + \int_{R \times \{y\}} f)/(R_1 ... R_l)]^{\frac{1}{n-l}}$, as l varies in the range $0 \le l \le k$. A short calculation shows that the choice of l that gives the smallest value to this expression obeys $R_l \le L \le R_{l+1}$. (There may be two different choices of l that give the smallest value to this expression. In this case, these two values are consecutive, say l_0 and $l_0 + 1$, and $L = R_{l_0+1}$.)

We begin with the unit lattice in $\mathbb{R}^n$, with respect to the metric $dx^2 + Ady^2$. We consider the barycentric subdivision of this lattice, and we move a vertex of this subdivision located at (x, y) to $(L(y)x, y)$ as before. We extend this map to a piecewise linear isomorphism. As before, we define S_0 to be the image of the k-skeleton of the unit lattice. From now on, we work in the metric $g = dx^2 + AL(y)^2dy^2$.

We translate R so that it is equal to the product $[R_1, 2R_1] \times ... \times [R_m, 2R_m]$. As a result, each k-face of S_0 intersects $R \times \mathbb{R}$ in a region with volume at most $2R_1 ... R_l L^{k-l}$. This inequality holds for any value of l.

Next, we apply barycentric wiggling to get a skeleton S. In the original proof, we wiggled each vertex in a ball of radius $2^{-n}L$. Now we need to use a ball of radius cL for some small constant c that we will choose later. As c gets smaller, the map Φ which takes the unit lattice to S rotates the angles of planes by less and less. The same barycentric wiggling proof now gives the following estimate for the integral of f over a k-face F of S.

$$\int_F f < CL^{-(n-k+1)} \int_{B(CL)} f.$$

Next, we construct a PL isomorphism Ψ which presses most of space against the unit k-skeleton. We will use a perturbation of the map Ψ from Chapter 3. We call the map from chapter 3 Ψ_0. Now we define Ψ by following Ψ_0 with with a barycentric perturbation that fixes the unit k-skeleton but moves around the rest of space a little bit. The map $\Phi \circ \Psi \circ \Phi^{-1}$, which pushes most of space to the k-skeleton S, expands certain k-planes that pass through the bad region. As a result of our perturbation, most k-planes which hit the bad region get mapped to k-planes that are slightly, but quantitatively, transverse to R.

Finally, we construct a PL projection π_0 from $\mathbb{R}^n$ to $\mathbb{R}^{n-k+1}$, which is a barycentric perturbation of the projection onto the last (n-k+1) coordinates, with respect to a triangulation of size much larger than $\sup_y L(y)$ but still very tiny with respect to A. Locally, a good projection is one for which the k-plane fibers on the good simplices get pushed almost parallel to S and the k-plane fibers on the bad simplices get pushed slightly transverse to $R \times \mathbb{R}$. With respect to the local metric, almost every orthogonal projection is good. By our standard barycentric trick, there is a globally good PL projection π_0.

As before, we define $\pi = \pi_0 \circ \Phi \circ \Psi^{-1} \circ \Phi^{-1}$. The intersection of each fiber of π with $R \times \mathbb{R}$ is bilipshitz to $R_1 \times ... \times R_k$, with the projection onto the first k coordinates being the bilipshitz map. Each fiber hugs the k-skeleton S except along a set of volume less than $CR_1...R_l L^k$ for each value of l. The integral of f over the part of S within L of the fiber is bounded by $CL^{-(n-k+1)} L \int_{R\times\{y\}} f$, which equals $CL^{k-l} L^{l-n} \int_{R\times\{y\}} f$. When l takes the value which minimizes the expression in the definition of L, then this quantity is less than $CR_1...R_l L^{k-l}$. We saw above that for this value of l, $R_l \leq L \leq R_{l+1}$. Therefore, the integral of f over the fiber $\pi^{-1}(z, y)$ is less than $CR_1...R_l L^{k-l}$ for every value of l in the range $0 \leq l \leq k$. Plugging in the definition of L, we conclude that the integral of f over the fiber $\pi^{-1}(z, y)$ is bounded by $C(R_1...R_l)^{\frac{n-k}{n-l}} (\int_{R\times\{y\}} f)^{\frac{k-l}{n-l}}$ for every value of l.

The fiber $\pi^{-1}(z,y)$ lies within CR_k of the k-plane defined by having last (n-k+1) coordinates equal to (z,y). This statement holds in the metric g. Therefore, in the original metric, $\pi^{-1}(z,y)$ lies within δ of the plane $R \times \{y\}$. □

Chapter 6

Area-Expanding Embeddings of Rectangles

In this chapter we consider the problem of deciding whether there is a k-expanding embedding of one rectangle into another. We are able to solve this problem, up to a constant factor, for all values of k and n. This theorem is the main result of my thesis.

Theorem 6.1. *Let R and S be n-dimensional rectangles. If there is a k-expanding embedding of S into R, then the following inequalities hold.*

For each integer l in the range $0 \leq l \leq k$, and each integer p in the range $k+1 \leq p \leq n$,

$$R_1 ... R_l (R_{l+1} ... R_p)^{(k-l)/(p-l)} > c S_1 ... S_l (S_{l+1} ... S_p)^{(k-l)/(p-l)}.$$

On the other hand, if the same inequalities hold with a large constant C in place of the small constant c, then there is a k-expanding embedding of S into R.

Over the course of the first three sections, we will prove the inequalities in the theorem. In fact, we will prove something a little bit stronger. Define $S(p,l)$ to be the complement in S of the (p-1)-skeleton of the cubical axes-parallel lattice with side length S_l. If there is a k-expanding embedding from $S(p,l)$ into R, then we will prove that $R_1 ... R_l (R_{l+1} ... R_p)^{(k-l)/(p-l)} > c S_1 ... S_l (S_{l+1} ... S_p)^{(k-l)/(p-l)}$.

The region $S(p,l)$ is a thick neighborhood of the (n-p)-skeleton of a lattice with side length S_l. The main idea of the proof is to find a good definition for the k-width of the open set $S(p,l)$ around this skeleton. We explain this idea in the first section, and in the next two sections we apply it to prove the inequalities above.

Once we have proven the inequalities, we have to construct the k-expanding embeddings described in the theorem. The embeddings involved are geometrically very simple - they are standard quasi-isometric foldings followed by linear maps. Checking that such an embedding exists whenever the side lengths of R and S obey the above inequalities requires some tedious algebra. We do the tedious algebra in section 4. Finally, in section 5, we apply our results to estimate the k-dilation of diffeomorphisms between rectangles.

6.1 The k-width of an open set around a complex

Let X be a finite complex embedded in $\mathbb{R}^n$, and let U be a neighborhood of X. We are going to give a definition for the k-width of U around X.

More generally, for each face F of X, suppose that we have a function $B(F)$, non-negative and compactly supported. We are going to define a k-width of the set of functions B around X.

When we defined k-width in Chapter 3, we used generic PL maps from $\mathbb{R}^n$ to $\mathbb{R}^{n-k}$. The fibers of such a map form a family of k-cycles. It would have been reasonable to define k-width using more general families of k-cycles. At the time, we did not need to use more general families, but we will need them in this chapter.

A nice family of k-cycles in $\mathbb{R}^n$ consists of the following data. We have a pseudomanifold W of dimension w, a generic PL map π from W to $\mathbb{R}^{w-k}$, and a generic proper PL map P from W to $\mathbb{R}^n$. If y is a point in $\mathbb{R}^{w-k}$, then $\pi^{-1}(y)$ is a k-cycle in W, and $P(\pi^{-1}(y))$ is a k-cycle in $\mathbb{R}^n$. For example, if π' is a generic PL map from $\mathbb{R}^n$ to $\mathbb{R}^{n-k}$, π' defines a nice family of cycles by taking $W = \mathbb{R}^n$, $\pi = \pi'$, and P equal to the identity. We will work with oriented families of cycles, which means that W is oriented, but non-oriented families of cycles also make sense. Families of cycles of all dimensions make sense, but we will assume throughout the chapter that the dimension of W is n, which means that the parameter space is $\mathbb{R}^{n-k}$.

We need some simple vocabulary to talk about nice families of k-cycles. Let C be a nice family of k-cycles. For a point y in $\mathbb{R}^{n-k}$, we define $C(y)$ to be the k-cycle $P(\pi^{-1}(y))$. More generally, if Δ is an m-simplex in $\mathbb{R}^{n-k}$, then we define $C(\Delta)$ to be the $(k+m)$-chain $P(\pi^{-1}(\Delta))$. If C and C' are two nice families of k-cycles, then we say that C and C' agree on an open set U if $C(y) \cap U$ is equal to $C'(y) \cap U$ for every point y in $\mathbb{R}^{n-k}$. The degree of a family C is defined to be the degree of the proper map P from W to $\mathbb{R}^n$. If two nice families agree outside of a compact set, then they have the same degree.

Now we can state the definition of the k-width of B around X. We say that the k-width of B around X is at least W, if for every nice family of cycles C in $\mathbb{R}^n$, there is a nice family C' which agrees with C outside of a compact set, and with the property that if $\int_{C(y)} B(F) < w$, then $C'(y)$ is disjoint from F. (This last property holds for every face F of X and every point y in $\mathbb{R}^{n-k}$.)

This definition may seem confusing at first. I'm going to try to make it a little clearer. Suppose we pick a nice family C of k-cycles in $\mathbb{R}^n$ with degree 1. For each function $B(F)$, we define the w-thick shadow of $B(F)$ to be the set of points y in $\mathbb{R}^{n-k}$ so that $\int_{C(y)} B(F) > w$. (One good feature of nice families of cycles is that this integral varies continuously in y.) We denote the w-thick shadow of $B(F)$ by $T_w(F)$. If the width of B around X is at least w, then first of all each of these sets $T_w(F)$ is not empty. Moreover, we can deduce something about the way these sets overlap.

Let C' be the nice family of k-cycles guaranteed by the definition above. We have a map P' from W' to $\mathbb{R}^n$ with degree 1. If y lies in the image of $\pi'(P'^{-1}(F))$, then the face F meets $C'(y)$, and so y lies in $T_w(F)$. In other words, the map π' takes $P'^{-1}(F)$ into $T_w(F)$. Since C has degree 1, the map P' also has degree 1. After putting the map in general position with respect to X, the inverse image of each 0-face of X is a sum of points in W with multiplicity, and with total weight 1. For each 0-face v in X, $\pi'(P^{-1}(v))$ is a 0-chain in $T_w(v)$ homologous to a point. For each 1-face E in X with boundary $v_1 - v_2$, $\pi'(P^{-1}(E))$ is a 1-chain in $T_w(E)$ with boundary

$\pi'(P^{-1}(v_1)) - \pi'(P^{-1}(v_2))$. For each 2-face F of X with boundary $E_1 + E_2 + E_3 + E_4$, $\pi'(P^{-1}(F))$ is a 2-chain in $T_w(F)$ with boundary equal to $\sum_i \pi'(P^{-1}(E_i))$. And the pattern continues for higher-dimensional faces.

These comments about overlapping thick shadows seem to me a little more intuitive than the definition of k-width of B around X, but the actual definition is a little stronger, and we need it.

Now, if U is an open set containing X, we define the k-width of U around X to be the supremal value of the k-width of B around X, where $B(F)$ are functions supported in U and $\sum_F B(F)$ is less than 1 pointwise.

Let us consider the simplest example: a rectangular lattice. Let X be the q-skeleton of an m-dimensional rectangular lattice of dimensions $N_1 \times ... \times N_m$, where $q < m$. Embed X in $\mathbb{R}^n$ as part of a cubic lattice of side length L, where $m \leq n$. Let U be the (L/3)-neighborhood of X. (Although we don't need this fact, X is a deformation retract of U.)

Theorem 6.2. *The k-width of U around X is greater than cL^k. (As usual, the constant c depends only on n, and does not depend on the dimensions of the complex X.)*

This theorem is the analog for k-width around complexes of the fact that the k-width of the cube of side length L is at least cL^k.

Proof. By scaling it suffices to prove this theorem when L is equal to 1.

Our first task is to define some functions B(F), which requires a few steps. For each face F of X, we define a region $U(F)$, which is equal to a thickened copy of the center of the face F. Suppose that F is an l-dimensional face of X. Then $U(F)$ is defined to be the product of $C_1 \times C_2$, where C_1 is an l-dimensional sub-cube of F, with the same center as F, and with side length $1 - 2^{-4l-1}$, and C_2 is a cube perpendicular to F, centered on F, and with side length 2^{-4l-3}. Inside of $U(F)$, we define a region $U'(F)$, which is the product $C_1' \times C_2'$, where C_1' is the sub-cube of F with the same center as F and with side length $1 - 2^{-4l-2}$ and C_2' is a cube perpendicular to F, centered on F, and with side length 2^{-4l-4}. A neighborhood of $U'(F)$ lies inside of

$U(F)$. The regions $U(F)$ and $U(G)$ overlap only if F is a sub-face of G or vice versa. The union of $U'(G)$ for all sub-faces G of F contains the face F. Next we define bump functions $b(F)$ which are continuous approximations to the characteristic function of $U(F)$, multiplied by some constant c. Finally, we define $B(F)$ to be the sum of $b(G)$ over all the sub-faces G of F. We choose the constant c so that $\sum B(F)$ is less than 1 pointwise.

Now we suppose that we have a nice family C of cycles in $\mathbb{R}^n$. We need a procedure to alter C on a compact set to get a family C'. First of all, we choose a very fine triangulation T of the range $\mathbb{R}^{n-k}$. We can choose this triangulation sufficiently fine that if $\int_{C(y)} B(F) > 2w$ for any y in Δ, then the same integral is greater than w for every y in Δ.

Recall that we defined a function b(F) which was a continuous approximation of the characteristic function of $U(F)$. We define the w-thick shadow of b(F) to be the set of y in $\mathbb{R}^{n-k}$ so that $\int_{C(y)} b(F) > w$. We denote the w-thick shadow of b(F) by $t_w(F)$. (To set the notation, the thick shadow of B(F) is T(F), and the thick shadow of b(F) is t(F).)

The number w will be a small constant, depending only on n, which we will choose later. We are going to alter C on a compact set to get a family C' with the following property.

1. If a simplex Δ of T does not meet $t_{2w}(F)$, then the chain $C'(\Delta)$ does not intersect $U'(F)$.

We check that this property suffices. Suppose that y does not lie in $T_w(F)$. Let Δ be a simplex of T containing y. It follows that Δ does not intersect $T_{2w}(F)$. Since $B(F) \geq b(G)$ for each sub-face G of F, it follows that Δ does not intersect $t_{2w}(G)$. According to Property 1, $C'(y)$ does not intersect $U'(G)$ for any subface G of F. But F lies in the union of the sets $U'(G)$. Therefore, $C'(y)$ does not intersect F.

We are going to construct a sequence of nice families C_l for l from -1 to q. The family C_{-1} is just the initial family C, and the family C_q will be C'. Each family is formed by altering the previous one on a compact set. The family C_l will obey property 1 for faces F of dimension less than or equal to l. Since X is a q-dimensional complex C_q will obey property 1 entirely.

In order to construct a nice family C_l, it suffices to define a (k+m)-chain $C_l(\Delta)$ in $\mathbb{R}^n$ for each m-simplex Δ of T, obeying two simple properties. The boundary of $C_l(\Delta)$ should be equal to $C_l(\partial\Delta)$. (After picking an orientation of Δ, this statement should hold as integral chains.) Also, $C_l(\Delta)$ should agree with $C_{l-1}(\Delta)$ outside of a compact set.

For inductive purposes, we will prove that the chains $C_l(\Delta)$ satisfy two more conditions.

2. If Δ is a 0-simplex which does not lie in $t_{4w}(F)$, then $C_l(\Delta)$ meets $U(F)$ in a region with k-volume less than Cw. On the other hand, if Δ is any m-simplex for $m \geq 1$, then $C_l(\Delta)$ meets $U(F)$ in a region with (k+m)-volume less than Cw.

3. If Δ is an m-simplex and F is an l-face of X, then $C_l(\Delta)$ agrees with $C_{l-1}(\Delta)$ in the region inside of $U(F)$ but outside the $\frac{m+1}{n}2^{-4l-8}$-neighborhood of $U'(F)$.

Since the triangulation T is very fine, we can assume that $C_{-1}(\Delta) = C(\Delta)$ obeys property 2. Properties 1 and 3 only assert something for the later families C_l, when $l \geq 0$. This gives the base for our induction.

The main point of the proof is an inductive procedure to modify the chain $C_{l-1}(\Delta)$ into the chain $C_l(\Delta)$. We will define $C_l(\Delta)$ first for all the 0-simplices Δ, then for all the 1-simplices, and so on, up to (n-k)-simplices.

The construction of $C_l(\Delta)$ for 0-simplices

Since we will be talking so much about $t_{2w}(F)$, we abbreviate t_{2w} by t.

For each l-face F, we repeat the following procedure.

If the 0-simplex Δ lies in $t(F)$, then we don't need to do anything. If Δ does not lie in $t(F)$ then we will modify $C_{l-1}(\Delta)$ as follows. First notice that by property 2, the intersection of $C_{l-1}(\Delta)$ with $U(F)$ has k-volume less than Cw.

Let $H(r)$ denote the boundary of the r-neighborhood of $U'(F)$. As long as r is less than 2^{-4l-8}, the surface $H(r)$ lies in $U(F)$. We can choose a value of r in the range $[0, (1/n)2^{-4l-8}]$ so that $H(r)$ meets $C_{l-1}(\Delta)$ transversely in a (k-1)-cycle C of volume less than Cw. The cycle C lies in $U(F)$, but outside of $U'(F)$. By induction, $C_{l-1}(\Delta)$ lies outside of $U'(G)$ for each sub-face G of F for which Δ^0 lies outside of $t(G)$. Let

V be the region in $U(F)$, outside of $U'(F)$, and outside of $U'(G)$ for each sub-face G of F for which Δ^0 lies outside of $t(G)$. Since C is a subset of $C_{l-1}(\Delta)$, C lies in V.

We claim that there is a filling F of C in V, with k-volume less than Cw. For any compact polyhedron such as V in $\mathbb{R}^n$, there is some small constant v so that each (k-1)-cycle C in V of volume less than v is homologically trivial and can be filled by a k-chain with k-volume less than $C_V$volume(C). Since there are only finitely many possible polyhedra V which occur in our construction, we can take a uniform values for v and C_V. If w is sufficiently small, we will have $Cw < v$, and we can apply this estimate. This proves the claim.

We construct $C_l(\Delta)$ from $C_{l-1}(\Delta)$ by deleting the portion of $C_{l-1}(\Delta)$ inside of $H(r)$ and replacing it with the filling F we have just constructed. (We repeat this procedure for each face F. Since the regions $U(F)$ are disjoint, these surgeries don't interfere with one another.)

Next we have to check that the family of cycles $C_l(\Delta)$ obeys all the properties which we laid out above. We carefully followed property 1 when we filled the cycle C inside the polyhedron V. By induction, we can assume that if Δ does not lie in $t_{2w}(G)$, then the k-volume of $C_{l-1}(\Delta) \cap U(G)$ is bounded by Cw. The surgery which we performed in $U(F)$ affects this intersection only if F is a subface of G or G is a sub-face of F. This occurs for less than C l-faces F. Each of these surgeries adds less than Cw to the volume of $C_{l-1}(\Delta)$. Therefore the k-volume of $C_l(\Delta) \cap U(G)$ is less than Cw. Finally, since we chose r in the range $[0, (1/n)2^{-4l-8}]$ and performed our surgery inside of the r-neighborhood of $U'(F)$, our surgery obeys Property 3.

The construction of $C_l(\Delta)$ for m-simplices

Since we already constructed $C_l(\Delta)$ for the 0-simplices, we can assume that m is at least 1, and that we have already constructed $C_l(\Delta)$ for all the (m-1)-simplices of T.

For each l-face F we repeat the following operation.

Let Δ be an m-simplex. We check how many vertices of Δ lie in $t(F)$.

If each vertex of Δ lies in $t(F)$, then we just put $C_l(\Delta) = C_{l-1}(\Delta)$. If some vertices of Δ do not lie in $t(F)$ then we will modify $C_{l-1}(\Delta)$ as follows. By Property 2 applied to $C_{l-1}(\Delta)$, the volume of $C_{l-1}(\Delta)$ intersected with $U(F)$ is bounded by Cw. As before, we define $H(r)$ to be the boundary of the r neighborhood of $U'(F)$. For r in the range $[0, 2^{-4l-8}]$, the surface $H(r)$ lies in $U(F)$. Because of the volume bound on $C_{l-1}(\Delta)$, we can choose r in the range $[\frac{m}{n}2^{-4l-8}, \frac{m+1}{n}2^{-4l-8}]$, so that the intersection of $C_{l-1}(\Delta)$ with $H(r)$ has (k+m-1)-volume less than Cw. By property 3, in the region between $H(\frac{m}{n}2^{-4l-8})$ and $H(\frac{m+1}{n}2^{-4l-8})$, $C_l(\partial\Delta)$ agrees with $C_{l-1}(\partial\Delta)$. Therefore, we can form a cycle C consisting of the union of the intersection of $C_{l-1}(\Delta)$ with $H(r)$ and the intersection of $C_l(\partial\Delta)$ with the r-neighborhood of U'(F). According to property 2, the (k+m-1)-volume of $C_l(\partial\Delta)$ intersected with $U(F)$ is also bounded by Cw. (If m=1, this is still true for the following reason. One of the vertices in the boundary of Δ does not lie in $t(F) = t_{2w}(F)$. Since the triangulation is very fine, the other vertex does not lie in $t_{4w}(F)$. Then by property 2, the volume of $C_l(\partial\Delta)$ intersected with $U(F)$ is bounded by Cw.) Therefore, the volume of C is bounded by Cw.

The cycle C lies in $U(F)$. If every vertex of Δ is disjoint from $t(F)$, according to property 1 applied to $C_l(\partial F)$, the cycle C avoids $U'(F)$. Similarly, according to property 1, the cycle C also avoids $U'(G)$ for each proper subface G of F so that all vertices of Δ are disjoint from $t(G)$. Therefore, just like for 0-simplices, we can fill C by a chain with (k+m)-volume less than Cw which lies in $U(F)$ and avoids all of the $U'(G)$ which it's supposed to avoid. We define $C_l(\Delta)$ to be $C_{l-1}(\Delta)$, minus the part of the chain inside the r-neighborhood of U, plus the filling of C which we just constructed.

Again we have to check that the nice family $C_l(\Delta)$ obeys the properties that we stated above. We carefully performed our filling to obey Property 1. We can assume by induction that the (k+m)-volume of $C_{l-1}(\Delta)$ intersected by $U(G)$ is less than Cw. The chain $C_l(\Delta)$ intersected with $U(G)$ differs by a bounded number of surgeries, each of which adds at most Cw volume. Therefore, the (k+m)-volume of $C_l(\Delta)$ intersected with $U(G)$ is bounded by Cw, and so our construction satisfies Property 2. Since we chose r in the range $[\frac{m}{n}2^{-4l-8}, \frac{m+1}{n}2^{-4l-8}]$ and performed our surgery inside of the r-neighborhood of $U'(F)$, our construction obeys Property 3.

Therefore, by induction, we get a nice family of cycles $C_q(\Delta)$ which satisfies property 1. Therefore, the k-width of U around X is at least w. The number w is greater than zero. It depended on n, but it did not depend on the choice of X. □

The set U relates to the rectangle S in the following way. Suppose that X in the (n-p)-skeleton of an (n-l)-dimensional rectangular lattice with dimensions roughly $(S_{l+1}/S_l) \times ... \times (S_n/S_l)$. Set $L = (S_1 ... S_l S_l^{k-l})^{1/k}$. Now there is a k-expanding diagonal linear embedding of U into S, taking the point x with coordinates $(x_1, ..., x_n)$ to $((S_1/L)x_1, ..., (S_{l-1}/L)x_{l-1}, (S_l/L)x_l, (S_l/L)x_{l+1}, ...(S_l/L)x_n)$. The image of U lies in the set $S(p, l)$ defined in the introduction to this chapter.

6.2 A width-volume inequality for multiple families of functions

Suppose that I is a k-expanding embedding of $S(p, l)$ into R. By a slight perturbation, we can assume that I is piecewise linear. Since we already had a k-expanding embedding of the set U into $S(p, l)$, we can compose them to get a k-expanding embedding of U into R. In the last section, we defined compactly supported bump functions $B(F)$ on U. By abuse of notation, we will also use $B(F)$ to refer to the pushforward of these functions in R. Since I is k-expanding, the functions $B(F)$ have k-width at least cL^k around the complex $I(X)$.

Let R' be the rectangle with dimensions $R_{l+1} \times ... \times R_n$. We then define bump functions $B'(F)$ on R' by the following formula.

$$B'(F)(y) = (R_1 ... R_l)^{-1} \int_{[0,R_1] \times ... \times [0,R_l]} B(F)(x, y) dx.$$

The function $B'(F)$ is just the average value of $B(F)$ over the l-plane perpendicular to R'. In the last section, we constructed the functions $B(F)$ so that the sum $\sum_F B(F)$ is at most 1 pointwise. By an easy averaging argument, the sum $\sum_F B'(F)$ is also at most 1 pointwise.

Consider the restriction of the functions $B'(F)$ to a subrectangle with dimensions $R_{l+1} \times ... \times R_p$. By the width-volume inequality, if one of these restrictions has (k-l)-width at least W, then it has integral at least $cW^{\frac{p-l}{k-l}}$. Since the integral of the sum of $B'(F)$ is at most $R_{l+1}...R_p$, it follows that the number of faces F so that the restriction of $B'(F)$ to our sub-rectangle has width at least W is bounded by $CW^{-\frac{p-l}{k-l}} R_{l+1}...R_p$. With a little bit more thought, it is possible to show that we can find a single generic PL map π from our sub-rectangle of R' to $\mathbb{R}^{p-k}$ so that the integral $\int_{\pi^{-1}(y)} B'(F)$ is less than W for every z for all but N faces F, where $N < CW^{-\frac{p-l}{k-l}} R_{l+1}...R_p$. We need a version of this result for families, which we now formulate.

Proposition 6.2.1. *Let $B'(F)$ be functions on $\mathbb{R}^{n-l}$. We will use coordinates (x, y) for $\mathbb{R}^{n-l}$, where x is in $\mathbb{R}^{p-l}$ and y is in $\mathbb{R}^{n-p}$. Each function $B'(F)$ is compactly supported, continuous, and obeys $0 \leq B'(F) \leq 1$. Suppose that $\int_{\mathbb{R}^{p-l} \times \{y\}} \sum_F B'(F)$ is less than M for every choice of y. We will use coordinates (z, y) for $\mathbb{R}^{n-k}$ where z lies in $\mathbb{R}^{p-k}$ and y lies in $\mathbb{R}^{n-p}$. If π is a generic PL map from $\mathbb{R}^{n-l}$ to $\mathbb{R}^{n-k}$, we define $N_W(y)$ be the number of faces F for which there exists some point z in $\mathbb{R}^{p-k}$ so that the integral of $B'(F)$ over $\pi^{-1}(z, y)$ is at least W. Then for each W, there exists a generic PL map π so that $N_W(y) < CW^{-\frac{p-l}{k-l}} M$.*

Proof. The proof of the lemma is only a small modification of the proof of the width-volume inequality for families, Theorem 5.2. After rescaling, this lemma reduces to the special case that W=1. After rescaling y, we can assume that each function $B'_F(x, y)$ is practically constant in y on the scale 1.

We let S_0 be a cubical lattice of side length s, for a number s on the order of 1 which we will choose later. We choose a linear projection π_0 from $\mathbb{R}^{n-l}$ to $\mathbb{R}^{n-k}$. We want to choose π_0 to have two properties. The first property is that the inverse image $\pi_0^{-1}(\mathbb{R}^{p-k} \times \{y\})$ is fairly close to being parallel to the plane $\mathbb{R}^{p-l} \times \{y\}$. More precisely, the orthogonal projection from the fiber to the plane should be bilipshitz with bilipshitz constant C. After changing coordinates on the range, this condition holds for most linear projections. The second property is that the fibers of π_0 get mapped to k-planes effectively tangent to the k-skeleton of S_0 by the map

Ψ constructed in Chapter 3, and that this remains true for any k-planes tilted by a small angle ϵ relative to the fibers of π_0.

Next, we cover $\mathbb{R}^{n-l}$ with overlapping regions $R(y_i)$, where each region is the $2^n s$ neighborhood of one of the plane $\pi_0^{-1}(\mathbb{R}^{p-k}, y_i)$, and so that any point of $\mathbb{R}^n$ lies in less than C regions.

Let δ be a positive number that we will choose later. We pick a barycentric wiggle, Φ, based on the barycentric subdivision of the cubical lattice with side length s. We wiggle each vertex of the barycentric subdivision inside of a very small ball, so that we distort the angles little enough that π_0 still enjoys the second property above. We set $S = \Phi(S_0)$. For a good choice of the wiggle Φ, we can guarantee that the skeleton S obeys the following inequalities. For each region $R(y_i)$, the number of faces F so that the integral $\int_{S \cap R(y_i)} B'(F) > \delta$ is bounded by $C\delta^{-1} \int_{\mathbb{R}^{p-l} \times \{y_i\}} B'(F)$. Put another way, for each region $R(y_i)$, we can pick out a set $X(y_i)$ of faces of F, with cardinality less than $C\delta^{-1}M$, so that $\int_{S \cap R_i} B'(F) < \delta$ unless F is a member of $X(y_i)$.

Now we define a map $\pi = \pi_0 \circ \Phi \circ \Psi \circ \Phi^{-1}$. The integral of $B'(F)$ over each fiber $\pi^{-1}(z, y)$ is controlled by the sum of two terms. To define the first term, we let $R(z, y)$ be the $2^n s$ neighborhood of $\pi^{-1}(z, y)$. The first term is $C \int_{S \cap R} B'(F)$. The second term is Cs^k. By choosing s sufficiently small, we can guarantee that the second term is less than $1/2$.

To control the first term, we need to use the barycentric wiggling estimate. The region R lies in the union of less than C regions $R(y_i)$. We define $X(y)$ to be the union of the sets $X(y_i)$ corresponding to these regions $R(y_i)$. The cardinality of $X(y)$ is less than $C\delta^{-1}M$. If a face F is not contained in $X(y)$, then the integral $\int_{S \cap R} B'(F)$ is less than $C\delta$. Therefore, $\int_{\pi^{-1}(z,y)} B'(F)$ is less than $C\delta + (1/2)$. We now choose δ to be $1/(4C)$ for this last constant C. Therefore, $\int_{\pi^{-1}(z,y)} B'(F) < 1$ unless F lies in $X(y)$. But the cardinality of $X(y)$ is less than $C\delta^{-1}M$, which is less than CM. $\square$

To apply this lemma, we need to pick a value of W. We are going to choose $W = c(R_1 ... R_l)^{-1} L^k$. Applying the last proposition gives us a map, which we will call π', from R' to $\mathbb{R}^{n-k}$. By precomposing with the orthogonal projection from R to R', we get a map π from R to $\mathbb{R}^{n-k}$. The integral of $B(F)$ over the fiber $\pi^{-1}(z, y)$ is equal to $(R_1...R_l)\int_{\pi'^{-1}(z,y)} B'(F)$.

Let us write down what the result of the proposition says about π. Recall that $N_W(y)$ is defined to be the number of faces F so that the integral of $B'(F)$ over the fiber $\pi'^{-1}(z,y)$ is at least W for some z. Equivalently, $N_W(y)$ is the number of faces F so that the integral of $B(F)$ over the fiber $\pi^{-1}(z,y)$ is at least cL^k. Equivalently, $N_W(y)$ is the number of faces F, so that the thick shadow $T_{cL^k}(F)$ meets the plane $\mathbb{R}^{p-k} \times \{y\}$. In any case, the last proposition tells us that this number $N_W(y)$ is less than $CW^{-\frac{p-l}{k-l}}M$.

The fibers of the map π form a nice family of k-cycles C, of degree 1. According to the theorem in the first section, the functions $B(F)$ have k-width at least cL^k around the complex I(X). Therefore, we can alter C on a compact set to get a nice family C', so that each cycle $C'(z,y)$ intersects a face $I(F)$ of $I(X)$ only if (z,y) lies in $T_{cL^k}(F)$. Now we define a family of p-cycles by taking $C''(y)$ to be the union of $C'(z,y)$ for all z. Since the projection from $\mathbb{R}^{n-k}$ to $\mathbb{R}^{n-p}$ is generic PL, C'' is a nice family of p-cycles of degree 1. Each cycle $C''(y)$ intersects at most $N_W(y)$ different faces of I(X).

Now we write down our estimate for this number $N_W(y)$. We showed that $N_W(y)$ is less than $CW^{-\frac{p-l}{k-l}}M$. But W is equal to $c(R_1...R_l)^{-1}L^k$, and M is equal to $R_{l+1}...R_p$. Therefore, $N_W(y)$ is less than $C(R_1...R_l)^{\frac{p-l}{k-l}}L^{-k\frac{p-l}{k-l}}R_{l+1}...R_p$. But L was defined to be $(S_1...S_lS_l^{k-l})^{1/k}$. Therefore, $N_W(y)$ is less than the following expression.

$$C(R_1...R_l)^{\frac{p-l}{k-l}}R_{l+1}...R_p(S_1...S_l)^{-\frac{p-l}{k-l}}S_l^{-(p-l)}.$$

We have proven that there is a nice family C" of p-cycles, of degree 1, such that each p-cycle meets the complex $I(X)$ in less than the above number of faces. We recall that X is the (n-p)-skeleton of an (n-l)-dimensional rectangular lattice with dimensions roughly $(S_{l+1}/S_l) \times ... \times (S_n/S_l)$.

6.3 Combinatorial width

We now have to deal with certain combinatorial questions which are closely analogous to the questions about k-width which we have studied earlier. Let X be a cubical (or simplicial) complex. We define the combinatorial width of X over $\mathbb{R}^{n-p}$ to be at most W if there is a generic PL map from X to $\mathbb{R}^{n-p}$ so that each fiber intersects at most W faces of X. For example, suppose that X is the n-skeleton of the unit lattice in $\mathbb{R}^n$, restricted to a rectangle of dimensions $N_1 \times ... \times N_n$. The combinatorial width of this complex over $\mathbb{R}^{n-p}$ gives a kind of coarse analogue of the geometric p-width of the rectangle. As we will will see, the combinatorial width of this complex over $\mathbb{R}^{n-p}$ is roughly $N_1...N_p$. The proof we will give is basically a coarse analogue of the standard geometric proof.

Now suppose that instead of the n-skeleton of the unit lattice, we take only the q-skeleton of the unit lattice, restricted to the same rectangle. If q is less than n-p, then each fiber of a generic PL map from X to $\mathbb{R}^{n-p}$ meets only C faces of X, regardless of the size of the N_i. We will prove that the width of the (n-p)-skeleton over $\mathbb{R}^{n-p}$ is roughly $N_1...N_p$.

Proposition 6.3.1. *Let X be the (n-p)-skeleton of a rectangular lattice in $\mathbb{R}^m$ of dimension $N_1 \times ... \times N_m$. Suppose that X is embedded into $\mathbb{R}^n$. Let C be a nice family of p-cycles, parameterized by $\mathbb{R}^{n-p}$, with non-zero degree. Then one of the p-cycles in C must intersect X in at least $cN_1...N_{m+p-n}$ faces.*

Proof. In the situation where we will use this proposition, the embedding of X in $\mathbb{R}^n$ is standard. This point is not important, however, and we begin by reducing the general case to the special case that $m = n$ and that X is embedded in $\mathbb{R}^n$ as a rectangular region in the standard unit lattice.

Let X be embedded into $\mathbb{R}^n$, and let F be a family of p-cycles, parameterized by $\mathbb{R}^{n-p}$, so that each cycle meets X in at most W faces. If n is not at least 4m, then we embed $\mathbb{R}^n$ into $\mathbb{R}^{4m} = \mathbb{R}^n \times \mathbb{R}^{4m-n}$. Then we replace each cycle C in our family by $C \times \mathbb{R}^{4m-n}$. After this argument, we may assume that n is at least 4m. In this case, any embedding of X into $\mathbb{R}^n$ is isotopic to a standard embedding. In particular, the embedding of X into $\mathbb{R}^n$ extends to an embedding of $\mathbb{R}^m$, where X is considered as a subset of the (n-p)-skeleton of the unit lattice in $\mathbb{R}^m$. Next, we move the embedding of $\mathbb{R}^m$ into $\mathbb{R}^n$ into general position, and restrict each cycle to $\mathbb{R}^m$. In this way, we obtain a family of (m+p-n)-cycles in $\mathbb{R}^m$, parameterized by $\mathbb{R}^{n-p}$, which sweep out $\mathbb{R}^m$, and so that each cycle meets at most W faces of X. The complex X is just the (m-p)-skeleton of the unit lattice in the rectangle with dimensions $N_1 \times ... \times N_m$.

From now on we assume that X is the (n-p)-skeleton of the unit lattice in the rectangle R with dimensions $N_1 \times ... \times N_n$ in $\mathbb{R}^n$, and that C is a nice family of p-cycles with non-zero degree. Suppose that each cycle in C meets X in less than W faces. It suffices to prove that W is greater than $cN_1...N_p$. We pick a very fine triangulation T of the parameter space $\mathbb{R}^{n-p}$, so that for each simplex Δ in T, $C(\Delta)$ meets X in less than CW faces.

The first main step of the proof is to surger the family of cycles so that each cycle is covered by CW cubes of the unit lattice.

To do this, we will construct a sequence of nice families of cycles $C(l, m)$, where l goes from (n-p) up to n, and m goes from 0 to n-p. Our original nice family is $C(n-p, n-p)$. The next family is $C(n-p+1, 0)$, then $C(n-p+1, 1)$, and so on up to $C(n-p+1, n-p)$. The next family after that is $C(n-p+2, 0)$, and so on up to $C(n, n-p)$, which is the last family in the sequence. Each family is obtained by doing surgery on the previous family, supported in a compact subset, and so each family has non-zero degree.

The main property of these families is the following. For any simplex Δ of the triangulation T, $C(l, m)(\Delta)$ meets the (l-1)-skeleton of the unit lattice in our rectangle in less than CW faces. Moreover, if the dimension of Δ is less than or equal to m, then $C(l, m)(\Delta)$ meets the l-skeleton of the unit lattice in our rectangle in less than CW faces.

We will construct each $C(l, m)$ inductively, by doing surgery on the previous one. To anchor the induction, we have defined $C(n-p, n-p)$ to be our original nice family C. We have assumed that for this family, $C(\Delta)$ meets X in less than CW faces, for any simplex Δ in T. Since X is the (n-p)-skeleton of the unit lattice in R, $C(n-p, n-p)$

does obey the condition above.

The construction of $C(l,0)$

We are going to construct $C(l,0)$ by doing surgery on $C(l-1,n-p)$. Fix a 0-simplex Δ in T. By induction, $C(l-1,n-p)(\Delta)$ meets the (l-1)-skeleton of the unit lattice in less than CW faces. We define an l-face of the unit lattice in R to be legal if its boundary intersects $C(l-1,n-p)(\Delta)$. The rest of the l-faces of the unit lattice in R we call illegal. There are less than CW legal l-faces. We are going to surger $C(l-1,n-p)(\Delta)$ so that it no longer intersects any illegal l-faces of the unit lattice in the rectangle R. The surgery will take place in a small neighborhood of these illegal l-faces.

Suppose F is an illegal l-face. By definition, $C(l-1,n-p)(\Delta)$ does not intersect the boundary of F. Therefore, after putting the family of cycles in general position, the intersection of $C(l-1,n-p)(\Delta)$ with F is a (p-n+l)-cycle K in F. If the cycle K is empty, we don't need to do anything, but if it's not empty we will perform a surgery. Pick a chain L in F with boundary K. The cycle $C(l-1,n-p)(\Delta)$ in a neighborhood of F is equal to $K \times B^{n-l}$. To define $C(l,0)(\Delta)$, we remove the region $K \times B^{n-l}$ and replace it with $L \times \partial B^{n-l}$. The resulting p-cycle avoids the illegal face F. We repeat this operation for every illegal face F, and the resulting cycle is $C(l,0)(\Delta)$. It meets the l-skeleton of the unit lattice in R in at most CW faces.

Next we have to define $C(l,0)(\Delta)$ for simplices Δ of dimension greater than 0. For simplices of dimension at least 2, we simply define $C(l,0)(\Delta) = C(l-1,n-p)(\Delta)$. For 1-simplices, we make the following surgery. Let Δ be a 1-simplex with boundary $v_1 - v_2$. For each of these vertices, $C(l,0)(v) - C(l-1,n-p)(v)$ is given by a sum of p-cycles, each of the form $L \times \partial B^{n-l} - \partial L \times B^{n-l}$. Each of these p-cycles is the boundary of a (p+1)-chain $L \times B^{n-l}$. Call the sum of these (p+1)-chains $M(v)$. We define $C(l,0)(\Delta)$ as $C(l-1,n-p)(\Delta)+M(v_1)-M(v_2)$. The boundary of $C(l,0)(\Delta)$ is indeed equal to $C(l,0)(v_1)-C(l,0)(v_2)$. Since each of these (p+1)-chains is supported away from the (l-1)-skeleton of the unit lattice in R, the chain $C(l,0)(\Delta)$ still meets the (l-1)-skeleton of the unit lattice in R in less than CW faces. Finally, we need to check that for a 2-simplex Δ, the boundary of $C(l,0)(\Delta)$ is equal to $C(l,0)(\partial\Delta)$. If

Δ is a 2-simplex of T with vertices v_1, v_2, and v_3, and with edges E_{12}, E_{23}, and E_{31}, then we perform this calculation as follows.

$C(l,0)(\partial\Delta) = C(l,0)(E_{12}) + C(l,0)(E_{23}) + C(l,0)(E_{31}) =$

$C(l-1,n-p)(E_{12}) + M(v_1) - M(v_2) + C(l-1,n-p)(E_{23}) +$

$M(v_2) - M(v_3) + C(l-1,n-p)(E_{31}) + M(v_3) - M(v_1) =$

$\partial C(l-1,n-p)(\Delta) = \partial C(l,0)(\Delta)$.

This finishes the construction of $C(l,0)$.

The construction of $C(l,m)$

Next we construct $C(l,m)$ when m is bigger than 0 by performing surgery on $C(l,m-1)$. The construction is similar in many ways to the last one.

For simplices Δ of dimension less than m, $C(l,m)(\Delta) = C(l,m-1)(\Delta)$. The next step is to define $C(l,m)(\Delta)$ for m-simplices (Δ).

Fix an m-simplex Δ in T. By induction, $C(l,m-1)(\Delta)$ meets the (l-1)-skeleton of the unit lattice in less than CW faces. Also by induction, $C(l,m-1)(\partial\Delta)$ meets the l-skeleton of the unit lattice in less than CW faces. We define an l-face of the unit lattice in R to be legal if its boundary intersects $C(l,m-1)(\Delta)$ or if it intersects $C(l,m-1)(\partial\Delta)$. The rest of the l-faces of the unit lattice in R we call illegal. There are less than CW legal l-faces. We are going to surger $C(l,m-1)(\Delta)$ so that it no longer intersects any illegal l-faces of the unit lattice in the rectangle R. The surgery will take place in a small neighborhood of these illegal l-faces.

Suppose F is an illegal l-face. By definition, $C(l,m-1)(\Delta)$ does not intersect the boundary of F, and $C(l,m-1)(\partial\Delta)$ does not intersect F. Therefore, after putting the family of cycles in general position, the intersection of $C(l,m-1)(\Delta)$ with F is a (m+p-n+l)-cycle K in F. If the cycle K is empty, we don't need to do anything, but if it's not empty we will perform a surgery. Pick a chain L in F with boundary K. The cycle $C(l,m-1)(\Delta)$ in a neighborhood of F is equal to $K \times B^{n-l}$. To define $C(l,m)(\Delta)$, we remove the region $K \times B^{n-l}$ and replace it with $L \times \partial B^{n-l}$. The resulting (p+m)-cycle avoids the illegal face F. We repeat this operation for every illegal face F, and the resulting cycle is $C(l,m)(\Delta)$. It meets the l-skeleton of the unit lattice in R in at most CW faces.

Next we have to define $C(l,m)(\Delta)$ for simplices Δ of dimension greater than m. For simplices of dimension at least m+2, we simply define $C(l,m)(\Delta) = C(l,m-1)(\Delta)$. For (m+1)-simplices, we make the following surgery. Let Δ be an (m+1)-simplex with boundary $\sum \Delta_i$. For each face of the boundary, $C(l,m)(\Delta_i) - C(l,m-1)(\Delta_i)$ is given by a sum of (p+m)-cycles, each of the form $L\times\partial B^{n-l} - \partial L\times B^{n-l}$. Each of these (p+m)-cycles is the boundary of a (p+m+1)-chain $L\times B^{n-l}$. Call the sum of these (p+1)-chains $M(\Delta_i)$. We define $C(l,m)(\Delta)$ as $C(l,m-1)(\Delta)+\sum M(\Delta_i)$. The boundary of $C(l,m)(\Delta)$ is indeed equal to $C(l,m)(\partial\Delta)$. Since each of these (p+m+1)-chains is supported away from the (l-1)-skeleton of the unit lattice in R, the chain $C(l,m)(\Delta)$ still meets the (l-1)-skeleton of the unit lattice in R in less than CW faces. Finally, we need to check that for an (m+2)-simplex Δ, the boundary of $C(l,m)(\Delta)$ is equal to $C(l,m)(\partial\Delta)$. Suppose that the boundary of Δ is equal to $\sum \Delta_i$. Then $C(l,m)(\partial\Delta)$ is equal to the sum $\sum C(l,m)(\Delta_i)$. Now, suppose that the boundary of Δ_i is equal to $\sum_{j\neq i}\Delta_{i,j}$. Then $\sum_i C(l,m)(\Delta_i) = \sum_i (C(l,m-1)(\Delta_i) + \sum_j M(\Delta_{i,j}))$. Now $\Delta_{i,j}$ is equal to $\Delta_{j,i}$ with opposite orientation, so the terms $M(\Delta_{i,j})$ and $M(\Delta_{j,i})$ cancel, leaving $C(l,m)(\partial\Delta) = \sum_i C(l,m-1)(\Delta_i)$. This sum is equal to the boundary of $C(l,m-1)(\Delta)$, which is equal to the boundary of $C(l,m)(\Delta)$.

This finishes the construction of $C(l,m)$. By induction, each chain $C(n,n-p)(\Delta)$ intersects less than CW unit cubes of the rectangle R.

The second main step of the proof is to prove an isoperimetric inequality which says that any k-cycle in our rectangle which can be covered by M cubes is the boundary of a (k+1)-cycle that can be covered by CMN_{k+1} cubes.

Let C be a relative k-cycle in our rectangle, lying in the union of M cubes. For each cube, we can push the cycle C out to the boundary of that cube. Proceeding inductively, we construct a cobordism from C to a union of k-faces of the unit lattice, the cobordism staying inside the union of M cubes that cover C. Therefore, without loss of generality we may assume that the cycle C consists of some union of k-faces of the unit lattice, counted with multiplicity.

Next we pick a point p with non-integer coefficients and use it to push our cycle out to the boundary by the following recipe. Let F be a k-face of the unit lattice included in C. The face F is defined by n equations, one for each coordinate. For k of the n coordinates, we have an equation of the form $a_i \leq x_i \leq a_i + 1$, and for the other n-k coordinates we have an equation of the form $x_i = b_i$. We define an integer j by saying that the coordinates from x_1 up to x_{j-1} have equations of the first form, but the coordinate x_j has an equation of the second form. (The integer j is at least 1 and at most k+1.) Now, we define a chain H(F) by the following recipe. If $a_i \leq p_i \leq a_i + 1$ for every i in the range $1 \leq i \leq j - 1$, then the cycle $H(F)$ is defined by the following equations. (Technical point: if j=1, then this hypothesis is vacuously fulfilled.) For the first j-1 coordinates, there are no equations. If $p_j < a_j$, then we use the equation $x_j > a_j$. If $p_j > a_j$, then we use the equation $x_j < a_j$. For i in the range $j + 1 \leq i \leq m$, we include the equation for the coordinate x_i in the definition of F. On the other hand, if it is not the case that $a_i \leq p_i \leq a_i + 1$ for every i in the range $1 \leq i \leq j - 1$, then H(F) is empty. We assign H(F) an orientation and multiplicity consistent with the orientation and multiplicity of F in C.

We now check that the chains H(F) counted with multiplicity, fit together to form a relative chain with boundary C. For any sum D of k-faces F of the unit lattice, we define $H(D)$ to be the sum of $H(F)$. By an easy calculation, $H(\partial H(F) - F) = 0$. It follows that $H(\partial H(C) - C) = 0$. We abbreviate $\partial H(C) - C$ by K. Since C is a relative cycle, K is also a relative cycle. Since $H(K) = 0$, every face of K is parallel to the x_1 axis. But since K is a relative cycle, it must have the form $[0, N_1] \times K'$, for some relative (k-2)-cycle K' in the rectangle R' with dimensions $N_2 \times ... \times N_n$. Let p' be the orthogonal projection of p to R'. Now $H(K)$ is equal to $[0, N_1] \times H_{p'}(K')$. Since $H(K) = 0$, we conclude that $H_{p'}(K') = 0$. By induction, we can conclude that K' is zero, and hence that K is zero. In other words, the boundary of $H(C)$ is indeed C.

It is easy to check that for a fixed face F and a random point p, the number of cubes necessary to cover H(F) is on average less than CN_{k+1}. Therefore, for some point p, the filling H(C) can be covered by less than CMN_{k+1} cubes.

Given the isoperimetric inequality, Gromov's method for bounding the k-width of the unit cube shows that W is at least $cN_1...N_p$. We give the details of this proof.

From now on, we denote $C(n, n-p)$ by C'. For each 0-simplex Δ of T, choose a (p+1)-chain $F(\Delta)$ with boundary $C'(\Delta)$, so that $F(\Delta)$ meets less than CWN_{p+1} unit cubes of the unit lattice in the rectangle R. Then, for each 1-simplex Δ of T, define a (p+1)-cycle $D(\Delta)$ to be equal to $C'(\Delta) + F(\partial\Delta)$. The cycle $D(\Delta)$ is contained in less than CWN_{p+1} unit cubes of the rectangle R. For any (p+1)-cycle $\sum \Delta_i$ in T, $\sum D(\Delta_i) = \sum C(\Delta_i)$, because the terms involving F cancel. Now we continue inductively. Once we have defined $D(\Delta)$ for m-simplices Δ, as a (p+m)-dimensional relative cycle in R which meets less than $CWN_{p+1}...N_{p+m}$ unit cubes, then we define a (p+m+1)-dimensional chain $F(\Delta)$ which fills this cycle, so that $F(\Delta)$ meets only $CWN_{p+1}...N_{p+m+1}$ unit cubes. Once we have defined $F(\Delta)$ for each m-dimensional simplex Δ, we define a (p+m+1)-dimensional cycle $D(\Delta)$ for each (m+1)-simplex Δ by the formula $D(\Delta) = C'(\Delta) + F(\partial\Delta)$. The cycle $D(\Delta)$ meets only $CWN_{p+1}...N_{p+m+1}$ unit cubes. As before, for any (m+1)-cycle $\sum \Delta_i$ in T, $\sum D(\Delta_i) = \sum C'(\Delta_i)$. At the end of the induction, we have defined an n-cycle $D(\Delta)$ for each (n-p)-simplex Δ of T. Each of these n-cycles is contained in less than $CWN_{p+1}...N_n$ unit cubes. If $W < (1/C)N_1...N_p$, then each cycle is contained in a proper subset of R and so has degree 0. But $\sum D(\Delta)$ over all the (n-p)-simplices Δ of T is equal to $\sum C'(\Delta)$ over all the (n-p)-simplices of Δ. This cycle has non-zero degree because C' has non-zero degree. Therefore, $W > (1/C)N_1...N_p$. This is the inequality which we were trying to prove. □

Now we return to the situation of a k-expanding embedding of S into R. We recall that X is the (n-p)-skeleton of an (n-l)-dimensional rectangular lattice with dimensions roughly $(S_{l+1}/S_l) \times ... \times (S_n/S_l)$. Applying the last proposition, we see that any family of p-cycles of non-zero degree must include a cycle which intersects $I(X)$ in at least $cS_{l+1}...S_pS_l^{-(p-l)}$ faces. At the end of the last section, we constructed a family of cycles of degree 1, so that each cycle met $I(X)$ in less than $C(R_1...R_l)^{\frac{p-l}{k-l}}R_{l+1}...R_p(S_1...S_l)^{-\frac{p-l}{k-l}}S_l^{-(p-l)}$ faces. Therefore, we get the following inequality.

$$C(R_1...R_l)^{\frac{p-l}{k-l}}R_{l+1}...R_p(S_1...S_l)^{-\frac{p-l}{k-l}}S_l^{-(p-l)} > cS_{l+1}...S_pS_l^{-(p-l)}.$$

Moving all terms involving S to the right-hand side and raising the equation to the power $(k-l)/(p-l)$, we get the inequality that we wanted to prove.

$$R_1...R_l(R_{l+1}...R_p)^{(k-l)/(p-l)} > cS_1...S_l(S_{l+1}...S_p)^{(k-l)/(p-l)}.$$

6.4 Constructing area-expanding embeddings

When all of the inequalities stated in Theorem 6.1 hold with the constant C, we will construct a k-expanding embedding of S into R. The embedding has a very simple form. It is just a k-expanding linear map followed by a simple quasi-isometric folding map.

Let us define the folding maps that we will use. If R and S are two dimensional rectangles, with $R_1 > S_1$ and $R_1R_2 > S_1S_2$, then there is a quasi-isometric embedding of S into R. This embedding simply winds S around inside of R like a snake.

Next, let $a < b$ be integers between 1 and n. Suppose that $R_i = S_i$ except when i is equal to a or b, and that $R_a > S_a$ and $R_aR_b > S_aS_b$. Again there is a quasi-isometric embedding of S into R. This embedding is simply the direct product of the embedding that we constructed above for the coordinates a and b with the identity in the other coordinates.

Composing these quasi-isometric embeddings proves the following baby lemma.

Lemma 6.1. *There is a quasi-isometric embedding of S into R if, for each p between 1 and n,* $R_1...R_p > CS_1...S_p$.

Next we have a little algebra lemma.

Lemma 6.2. *There is a rectangle T with* $S_1...S_p \leq T_1...T_p$ *for all p and a k-contracting linear map from R onto T if and only if the following inequalities hold. For each integer l in the range* $0 \leq l \leq k$, *and each integer p in the range* $k+1 \leq p \leq n$,

$$R_1...R_l(R_{l+1}...R_p)^{(k-l)/(p-l)} \geq S_1...S_l(S_{l+1}...S_p)^{(k-l)/(p-l)}.$$

Proof. First, let us suppose that the above inequalities hold. If it is also true that $S_1...S_p \leq R_1...R_p$ for every p, then we are done. Let b be the smallest integer so that $S_1...S_b > R_1...R_b$. Because of all the inequalities in the hypothesis of the lemma, we know that $b < k$.

We will define a sequence of linear diffeomorphisms $R = R(0) \to R(1) \to ... \to R(c)$, for some integer c between 1 and k-1. The diffeomorphism to $R(j)$ is called L_j. When j is less than c, the rectangle $R(j)$ has $R(j)_1 = ... = R(j)_{j+1}$. The linear map L_j increases each $R(j-1)_i$ for i between 1 and j by a factor of λ_j and decreases every other $R(j-1)_i$ by a factor of $\lambda_j^{-j/(k-j)}$, for some number $\lambda_j > 1$. From the last sentence, it follows that each L_j is k-contracting. Finally, if c is not bigger than b, then $R(c)_1...R(c)_b = S_1...S_b$. If c is bigger than b, then $R(c)_1...R(c)_c = S_1...S_c$.

Now we define the maps L_j. It suffices to define λ_j. There is a maximum value of λ_j which increases $R(j-1)_j$ and decreases $R(j-1)_{j+1}$ until they meet. If there is a lesser value of λ_j which makes $R(j)_1...R(j)_m = S_1...S_m$, where m is the maximum of b and j, then use that value and take $c = j$. If not, use the maximal value. As we increase j, $R(j)_1...R(j)_b$ increases. If $R(b)_1...R(b)_b < S_1...S_b$, then $R(b)_1...R(b)_{b+1} < S_1...S_{b+1}$, because $R(b)_1 = R(b)_{b+1}$. More generally, for j at least b, if $R(j)_1...R(j)_j < S_1...S_j$, then $R(j)_1...R(j)_{j+1} < S_1...S_{j+1}$ also.

From the formula for the map L_j, it follows that $R(j)_1...R(j)_k = R_1...R_k$ for every j, and by hypothesis $R_1...R_k \geq S_1...S_k$. Therefore, the above construction terminates with c less than or equal to k-1.

We write $R(c) = R(c)_1 \times ... \times R(c)_m \times R'$, where m is the maximum of b and c. We have proven above that $R(c)_1...R(c)_m = S_1...S_m$. Moreover, for every p less than m, $R(c)_1...R(c)_p \geq S_1...S_p$. If b is greater than or equal to c, this follows because $R_1...R_p \geq S_1...S_p$, and the definition of L_j shows that $R(j)_1...R(j)_p \geq R_1...R_p$ for every p less than k. If c is greater than b, this follows because $R(c)_1 = R(c)_m$ and $R(c)_1...R(c)_m = S_1...S_m$. In either case it is true.

The maps L_j preserve many of the inequalities that we have assumed. In particular, if $l \geq j$, then the following equality holds.

$$R(j)_1 ... R(j)_l (R(j)_{l+1} ... R(j)_p)^{(k-l)/(p-l)} = R_1 ... R_l (R_{l+1} ... R_p)^{(k-l)/(p-l)}.$$

Therefore, if $l \geq m$ then

$$R(c)_1 ... R(c)_l (R(c)_{l+1} ... R(c)_p)^{(k-l)/(p-l)} \geq S_1 ... S_l (S_{l+1} ... S_p)^{(k-l)/(p-l)}.$$

Since $R(c)_1 ... R(c)_m = S_1 ... S_m$, we can divide the above inequality on both sides, leaving the following inequality for all $l \geq m$.

$$R(c)_{m+1} ... R(c)_l (R(c)_{l+1} ... R(c)_p)^{(k-l)/(p-l)} \geq S_{m+1} ... S_l (S_{l+1} ... S_p)^{(k-l)/(p-l)}.$$

Now we define $S' = S_{m+1} \times ... \times S_n$. We can rewrite the above inequalities in terms of R' and S'. The result is that for each l in the range $0 \leq l \leq k - m$ and each p in the range $k - m + 1 \leq p \leq n - m$, the following holds.

$$R'_1 ... R'_l (R'_{l+1} ... R'_p)^{(k-m-l)/(p-l)} \geq S'_1 ... S'_l (S'_{l+1} ... S'_p)^{(k-m-l)/(p-l)}.$$

By induction on the dimension, we can assume that there is a (n-m)-dimensional rectangle T', so that $T'_1 ... T'_p \geq S'_1 ... S'_p$, and a (k-m)-contracting linear map from R' to T'. We finally define T to be $R(c)_1 \times ... \times R(c)_m \times T'_1 \times ... \times T'_{n-m}$. The direct product of the (k-m)-contracting linear map from R' to T' with the identity map is a k-contracting linear diffeomorphism from $R(c)$ to T. Since we already have a k-contracting linear map from R to R(c), we can compose the two maps to get a k-contracting linear map from R to T. Also, we already know that $T_1 ... T_a = R(c)_1 ... R(c)_a \geq S_1 ... S_a$ when a is less than or equal to m. But for larger a, $T_1 ... T_a = T_1 ... T_m T'_1 ... T'_{m-a} \geq S_1 ... S_m S'_1 ... S'_{m-a} = S_1 ... S_m$. Therefore, the rectangle T satisfies all the conditions which we wanted to prove.

The proof in the opposite direction is much easier. This time we suppose that there is a rectangle T with $T_1 ... T_p \geq S_1 ... S_p$ for each p from 1 to n, and a k-contracting linear diffeomorphism from R to T. Since there is a k-contracting linear diffeomorphism from R to T, the inequality $R_1 ... R_l R_{i_1} ... R_{i_{k-l}} \geq T_1 ... T_l T_{i_1} ... T_{i_{k-l}}$ holds for any indices

$l < i_1 < ... < i_{k-l} \leq p$. By taking geometric averages of these inequalities for various choices of indices $i_1 < ...i_{k-l}$, it follows that $R_1...R_l(R_{l+1}...R_p)^{(k-l)/(p-l)} \geq T_1...T_l(T_{l+1}...T_p)^{(k-l)/(p-l)}$. Since $T_1...T_l \geq S_1...S_l$, and $T_1...T_p \geq S_1...S_p$, the inequality $T_1...T_l(T_{l+1}...T_p)^{(k-l)/(p-l)} \geq S_1...S_l(S_{l+1}...S_p)^{(k-l)/(p-l)}$ holds as well. Combining these two inequalities shows that $R_1...R_l(R_{l+1}...R_p)^{(k-l)/(p-l)} \geq S_1...S_l(S_{l+1}...S_p)^{(k-l)/(p-l)}$, which is what we wanted to prove. □

Combining our two lemmas, we can construct a k-expanding embedding of S into R whenever the rectangles obey the following inequalities. For each l in the range $0 \leq l \leq k$ and each p in the range $k+1 \leq p \leq n$,

$$R_1...R_l(R_{l+1}...R_p)^{(k-l)/(p-l)} > CS_1...S_l(S_{l+1}...S_p)^{(k-l)/(p-l)}.$$

This finishes the proof of Theorem 6.1.

6.5 Application to the k-dilation of diffeomorphisms

We can apply our estimates to give lower bounds for the k-dilation of diffeomorphisms. First we define some terms.

For each l in the range $0 \leq l \leq k$ and each p in the range $k+1 \leq p \leq n$, define $D_{l,p}$ as follows.

$$D_{l,p} = Q_1...Q_l(Q_{l+1}...Q_p)^{\frac{k-l}{p-l}}.$$

For each l in the range $1 \leq l \leq k+1$ and each p in the range $k+2 \leq p \leq n$, define $D'_{l,p}$ as follows.

$$D'_{l,p} = Q_2...Q_l(Q_{l+1}...Q_p)^{\frac{k-l+1}{p-l}}.$$

Proposition 6.5.1. *Each diffeomorphism from R to S has k-dilation greater than $cD_{l,p}$ for each l in the range $0 \leq l \leq k$ and each p in the range $k+1 \leq p \leq n$. Also, each diffeomorphism from R to S has k-dilation greater than $cD'_{l,p}$ for each l in the range $1 \leq l \leq k+1$, and each p in the range $k+2 \leq p \leq n$.*

Proof. Suppose there is a diffeomorphism from R to S with k-dilation D. Then, there is a k-contracting diffeomorphism from $D^{1/k}R$ to S. Therefore, there is a k-expanding embedding of S into $D^{1/k}R$. By the inequalities in Theorem 6.1, $D_{l,p}/D$ is less than C. Therefore, D is greater than $cD_{l,p}$.

We get the other inequalities by considering the map from the boundary of R to the boundary of S. We take a set U in the largest hyperface of S, and define functions $B(F)$ as in the proof of Theorem 6.1. We pull the functions $B(F)$ back to the boundary of R. This boundary is bilipshitz to the double of the largest hyperface of R. We now define functions $B'(F)$ on the rectangle with dimensions $R_{l+1} \times ... \times R_n$ by averaging $B(F)$ over the double of a rectangle with dimensions $R_2 \times ... \times R_l$ in the double of the largest hyperface of R. The rest of the proof then proceeds just as before. As a result, one gets the same inequalities that would follow if there were a k-contracting diffeomorphism from the largest hyperface of R to the largest hyperface of S - namely, the k-dilation is greater than $cD'_{l,p}$. □

Chapter 7

Rational Homotopy Invariants of Area-Contracting Maps

In this chapter, we will estimate the rational homotopy invariants of k-contracting maps from ellipsoids to wedges of unit spheres. Here is a typical example of the estimates we will prove.

Let E be a 7-dimensional ellipsoid with principal axes $E_0 \leq ... \leq E_7$. Let f be a map from E to the bouquet $S^2 \vee S^3 \vee S^4$. Let α_2 be a 2-form supported on S^2 away from the basepoint and with $\int_{S^2} \alpha_2 = 1$, and define α_3 and α_4 similarly. For any exact k-form α on E, let $P\alpha$ denote any primitive of α. The integral $H(f) = \int_E Pf^*(\alpha_2) \wedge f^*(\alpha_3) \wedge Pf^*(\alpha_4)$ is a homotopy invariant of the map f. In order to define the k-dilation of f, give each sphere in the bouquet the standard unit sphere metric.

Theorem. *For any 2-contracting map f, H(f) is bounded by $CE_1E_2E_3E_4^2E_5E_6^2E_7$. On the other hand, if E_1 is sufficiently large, we will construct a 1-contracting map f with $H(f)$ at least $cE_1E_2E_3E_4^2E_5E_6^2E_7$.*

Similarly, we can define $H'(f) = \int_E f^*(\alpha_2) \wedge Pf^*(\alpha_3) \wedge Pf^*(\alpha_4)$. The integral $H'(f)$ is also a homotopy invariant of f, which is different from $H(f)$.

Theorem. *For any 2-contracting map f, H'(f) is bounded by $CE_1E_2E_3E_4^2E_5^2E_6E_7$. On the other hand, if E_1 is sufficiently large, we will construct a 1-contracting map f* *with $H'(f)$ at least $cE_1E_2E_3E_4^2E_5^2E_6E_7$.*

If X is any simply connected finite complex, then the dual of $\pi_n(X)\otimes\mathbb{Q}$ is a vector space of invariants defined analogously to H using pullbacks of differential forms and primitives. We will prove estimates for the primitives of an exact form on an ellipsoid E in terms of its principal axes. If H is an invariant in the dual of $\pi_n(X)\otimes\mathbb{Q}$ which is defined using forms of degree at least k, and if f is a k-contracting map from E to X, then our estimates for primitives imply an upper bound on $H(f)$. Next we will investigate how sharp this upper bound is in the special case that X is a bouquet of spheres. For many invariants H, but not all of them, we will show that this upper bound is sharp up to a constant factor when E_1 is sufficiently large.

These estimates can be applied to prove estimates about k-contracting maps between rectangles. For example, we can apply the theorems above to 2-contracting diffeomorphisms between 8-dimensional rectangles. If f is a 2-contracting diffeomorphism from R to S then it restricts to a 2-contracting diffeomorphism between their boundaries, which are bilipshitz to ellipsoids. The theorems above imply that the following inequalities hold in the side-lengths of R and S.

$$R_2R_3R_4R_5^2R_6R_7^2R_8 > cS_2S_3S_4S_5^2S_6S_7^2S_8.$$

$$R_2R_3R_4R_5^2R_6^2R_7R_8 > cS_2S_3S_4S_5^2S_6^2S_7S_8.$$

At the end of the chapter, we will briefly consider torsion homotopy invariants of k-contracting maps. We will construct some homotopically non-trivial maps with very small k-dilation. In particular, we will prove the following theorem.

Theorem For each n, there are infinitely many choices of m so that there are homotopically non-trivial maps from S^m to S^n with arbitrarily small 3-dilation.

7.1 Review of rational homotopy invariants and differential forms

We give a quick elementary review of the rational homotopy invariants of maps from S^n to a bouquet of spheres of dimensions between 2 and n. Rational homotopy invariants of maps from S^n to a space X are defined to be elements of the dual space to $\pi_n(X) \otimes \mathbb{Q}$. In this section, we will define a bunch of rational homotopy invariants for maps from S^n to a wedge of spheres using differential forms. We will review the proof that these expressions are rational homotopy invariants, and in particular we will check that the proof goes through using forms that are only assumed to be bounded and measurable.

Let X be a wedge of spheres $S^{k_1} \vee ... \vee S^{k_r}$. Let α_i be a top-dimensional form on S^{k_i} with $\int_{S^{k_i}} \alpha_i = 1$. We only assume that α_i are bounded and measurable. Let f be a map from S^n to X or from $S^n \times [0,1]$ to X, and let a_i be the pullback $f^*(\alpha_i)$.

We will construct inductively a list of "legal" expressions involving the forms a_i, wedge products, and primitives. First, each form a_i is a legal expression of degree k_i. Next, if A and B are legal expressions of degrees a and b, and if $a + b \leq n + 1$, then $PA \wedge B$ is a legal expression of degree $a + b - 1$. (If A is exact, the expression PA denotes any primitive of A. If A is not exact, the expression PA is not defined.)

Proposition 7.1.1. *If A is a legal expression of degree n, then $\int_{S^n} A$ is independent of the choices of primitive, and defines a rational homotopy invariant of the map f. This remains true even if we choose primitives which are only bounded and measurable.*

Proof. By induction, every legal expression has degree at least 2. Therefore, if A and B are legal expressions with $a + b \leq n + 1$, then the degree a of A lies in the range $2 \leq a \leq n - 1$.

The wedge product $a_i \wedge a_j = 0$ pointwise for any i, j, including $i = j$. By induction $A \wedge B = 0$ for any two well-defined legal expressions.

Next, we will show by induction that any well-defined legal expression is closed. We have to compute $d(PA \wedge B) = A \wedge B \pm PA \wedge dB$. By induction, we can assume

$dB = 0$. But we proved in the previous paragraph that $A \wedge B = 0$.

In particular, each well-defined legal expression A of degree between 2 and n-1 is exact. Therefore, if A and B are well-defined legal expressions with total degree at most $n+1$, then the expression $PA \wedge B$ is well-defined. It follows by induction that every legal expression is well-defined.

Next, we prove by induction that if we change the choice of primitives in a legal expression, then the legal expression changes by $d(\omega \wedge L)$, where ω is any form and L is a form given by a legal expression.

To start the induction, suppose that we change the choice of primitive for the last primitive in a legal expression $C = PA \wedge B$. Changing the choice of primitive means changing PA by a closed form. Since the degree of A is between 2 and n-1, the degree of this closed form is between 1 and n-2. Since our manifold is homotopic to S^n, this form must be exact, and we write it as $d\omega$. Then the form $C = PA \wedge B$ changes by $d(\omega \wedge B)$.

Now we prove inductively that if we change a primitive anywhere in a legal expression, the resulting form changes by $d(\omega \wedge L)$. Suppose we are looking at the legal expression $PA \wedge B$, and that we have changed a primitive in A. By induction, we can assume that A changed by $d(\omega \wedge L)$. Therefore, PA changes by $\omega \wedge L$ plus an exact form $d\omega_2$. The term $\omega \wedge L \wedge B$ vanishes, because the wedge product of any two legal expressions vanishes. Therefore $PA \wedge B$ changes by $d(\omega_2 \wedge B)$. Now suppose that we have changed a primitive in B. By induction, we can assume that B changed by $d(\omega \wedge L)$. Therefore $PA \wedge B$ changes by $PA \wedge d(\omega \wedge L)$, which is equal to $\pm d(\omega \wedge PA \wedge L)$. (Here we have used the fact that $A \wedge L$ vanishes.)

Let us define an invariant of f by taking an n-form A in our list and defining $A(f) = \int_{S^n} A$. The choice of primitives only changes A by an exact form, so A(f) does not depend on the choice of primitives. Thus A(f) is an invariant of the map f. Moreover, applying the whole discussion to the domain $S^n \times [0,1]$ and using the fact that A is closed, it follows that A(f) is a homotopy invariant of f.

Finally, if S is a self map of S^n of degree D then $A(f \circ S)$ is equal $\int_{S^n} S^*(A)$, which is equal to $\int_{S_*(S^n)} A$, which is equal to $DA(f)$. Therefore, A defines a rational homotopy invariant. □

According to Dennis Sullivan's theory of minimal models, the invariants that we have just constructed generate the rational homotopy invariants of X. The theory of minimal models is explained in [9]. (The reader should be aware that the invariants we have just constructed have many linear dependences.)

7.2 Review of primitives and isoperimetric inequalities

If α_i is a form on X of degree at least k, and if the map f is k-contracting, then $f^*(\alpha_i)$ is bounded pointwise by the C^0 norm of α_i. Let RH be a rational homotopy invariant constructed in terms of forms of degree at least k. In order to bound RH, we need to understand how large the primitive of a bounded form can be.

In a different work, Dennis Sullivan explained how the problem of estimating primitives is related to the isoperimetric inequality.

Lemma. *(Sullivan) Let β be an exact measurable (k+1)-form in a Riemannian manifold M with L^∞ norm bounded by 1. Suppose that every exact k-cycle C with volume V admits a filling by a (k+1)-chain with volume less than $IsoV$. Then β has a primitive α which is a measurable k-form with L^∞ norm bounded by Iso.*

Proof. Here is the proof of the lemma, which we borrow from Gromov's book [15]. We record the proof here because it is quite short and we will have to use a more complicated version of it later. The proof is based on Whitney's theory of flat chains and cochains. (This theory is explained in Whitney's book *Geometric Integration Theory* [20].) For a polyhedral k-chain C in M we define the flat norm of C to be the infimum of $Iso|C - \partial D| + |D|$, taken over all real polyhedral (k+1)-chains D. (If C is a real polyhedral p-chain, equal to $\sum a_i P_i$ where a_i is a real number and P_i is a polyhedron, then $|C|$ denotes the mass of C, which is defined to be $\sum |a_i| \text{volume}(P_i)$.) The flat k-chains are the completion of the space of polyhedral k-chains in the flat norm.

Consider the space of flat k-chains in M, and the subspace of exact flat k-cycles. We have a linear functional defined on the vector space of flat exact k-cycles which assigns to each exact k-cycle C the integral of β over any filling of C. (This linear functional is well-defined because β is exact.) Since $|\beta|$ is bounded by 1 pointwise, the value of the functional on an exact cycle with filling volume less than V is bounded by V. By our isoperimetric inequality, the value of the functional on an exact cycle with volume V is bounded by $IsoV$. Therefore, the value of the functional on an exact cycle is bounded by the flat norm of the cycle. By the Hahn-Banach theorem, this linear functional has an extension to the space of all flat k-chains in M, with norm bounded by 1. Such a linear functional is called a flat cochain with comass 1.

By a theorem of Whitney, each flat k-cochain corresponds to a measurable differential k-form α. This theorem appears as Theorem 5A of Chapter 9 of his book *Geometric Measure Theory* [20]. Since our flat k-cochain has norm 1, it returns a value at most 1 to any k-chain with volume $1/Iso$. Hence the measurable differential k-form α has norm at most Iso. But the integral of α over any exact k-cycle is equal to the integral of β over a filling of that k-cycle. Therefore $d\alpha = \beta$. □

7.3 Linear isoperimetric inequalities

Because of Sullivan's lemma, we would like to know the best linear isoperimetric constants in an ellipsoid.

Proposition 7.3.1. *Any (k-1)-cycle of volume V in the ellipsoid E has a filling with volume bounded by* $C max(E_k, E_{n-k+1})V$.

This estimate is sharp up to a constant factor. We will give the proof in the next section, as a special case of Lemma 7.1. It follows as a special case of our next proposition. Combining this proposition with Sullivan's Lemma we get the following corollary.

Corollary. *Any exact measurable k-form β on E with L^∞ norm bounded by B has a measurable primitive α with L^∞ norm bounded by* $C max(E_k, E_{n-k+1})B$.

We can use this estimate to prove upper bounds for rational homotopy invariants as follows. Suppose that f is a k-contracting map from an ellipsoid E to a wedge of spheres $X = S^{k_1} \vee ... \vee S^{k_r}$, where each k_i is at least k. For each sphere S^{k_i} in the wedge, let α_i be the volume form of S^{k_i}. Since f is k-contracting, each of the pullbacks $f^*(\alpha_i)$ is bounded by 1 pointwise. Let H(f) be a rational homotopy invariant given by the integral of a form, which is constructed from the pullbacks $f^*(\alpha_i)$ by taking primitives and wedge products. Whenever we take the primitive of a form we have already bounded, we can find a bound on the primitive using the corollary above. Whenever we take the wedge product of two bounded forms, the result is bounded by the product of the two bounds. As a result, we get a pointwise bound on the integrand, which gives a bound on H(f).

We will carry out this procedure in detail in section 6. It will turn out that for first order invariants, this method is sharp up to a constant factor. It will also turn out that for higher order invariants, this method is not sharp up to a constant factor.

We can hope to improve this estimate for the following reason. When we take the primitive of a bounded form, some of its coefficients may need to be as large as the bound above, but other coefficients may be much smaller. It may happen, for instance, that two forms are both large in the same direction, and so when we take their wedge product, the largest term cancels. It turns out that such cancellations occur, and that they lead to sharper upper bounds for rational homotopy invariants. In order to exploit these cancellations systematically, we need to give directionally dependent estimates for the primitives of differential forms.

7.4 Directionally dependent isoperimetric inequalities

The ellipsoid E is bilipshitz to the double of the rectangle with side lengths $E_1 \leq ... \leq E_n$, with a bilipshitz constant independent of E. Therefore, it suffices to prove estimates for the double of this rectangle. On each rectangle in the double we use coordinates $x_1, ..., x_n$. If β is a k-form, we expand it on each of the two rectangles as $\beta = \sum \beta_I dx^I$, where I is a set of k distinct numbers between 1 and n. We define b_I to the supremum of the function β_I on E.

Theorem 7.1. *Let β be an exact k-form on E, which may only be bounded measurable. For each set J of (k-1) distinct indices, we define a number a_J as follows. If 1 is in J, and d is the smallest index not in J, then $a_J = E_d b_{J \cup d}$. If 1 is not in J, and e is the smallest index in J, then $a_J = \sum_{d=1}^{e-1} E_d b_{J \cup d}$. Then β has a bounded measurable primitive $\alpha = \sum \alpha_J dx^J$, so that $|\alpha_J| < Ca_J$ for each J.*

The proof of this theorem will be very similar to the proof of Sullivan's lemma. In place of the isoperimetric inequality used there, we will use a directionally dependent isoperimetric inequality which we now formulate.

Put a very fine lattice on the rectangle $E_1 \times ... \times E_n$, so that each face of the rectangle lies in the lattice. Taking the double of the lattice gives a cubical complex in E. We call a k-cycle rectilinear if it consists of a union of faces of this cubical complex for some sufficiently fine lattice. The rectilinear cycles are clearly dense in the space of flat cycles, and they are convenient to work with in the proof of the following proposition. If I is a set of k distinct indices and C is a rectilinear k-chain, then the I-volume of C is defined to be the total volume of the faces in C parallel to I.

Lemma 7.1. *Let C be a rectilinear (k-1)-cycle in E with J-volume V_J. For each set I of k distinct indices define W_I as follows. If 1 is in I and e is the smallest index not in I, then $W_I = \sum_{d=1}^{e-1} E_d V_{I-d}$. If 1 is not in I and d is the smallest index in I, then $W_I = E_d V_{I-d}$. Then there is a filling of C with I-volume less than CW_I for each I.*

Proof. We construct our filling in two steps. In the first step, we push C to a cycle that lies in one rectangle of the two rectangle of the double. In the second step, we fill this cycle inside this rectangle.

Step 1. Let C be a (k-1) relative rectilinear cycle in the rectangle R, with dimensions $E_1 \times ... \times E_n$. We will construct a k-chain F whose boundary is the union of C and B, where B is a rectilinear chain in the boundary of R. Let W_I be the I-volume of F, and let V'_J be the J-volume of B. Our filling with obey the following estimates.

1. If e is the smallest index not in I, then $W_I < C \sum_{d=1}^{e-1} E_d V_{I-d}$.

2. If e is the smallest index not in J, and e_2 is the second smallest index not in J, then $V'_J < C \sum_{d=e}^{e_2-1} (E_d/E_e) V_{J+e-d}$.

Now we construct the filling. Let p be any point in the rectangle R. Let A^{k-1} be a (k-1)-face of our very fine lattice, parallel to the coordinates J. The set A is given by inequalities $a_j \leq x_j \leq a'_j$ for each j in J, and by equalities $x_i = a_i$ for each i not in J. Let d be the lowest index not in J.

Now we define a k-dimensional rectilinear chain in R called $F_p(A)$. If $a_j \leq p_j \leq a'_j$ for each j from 1 to d-1, then $F_p(A)$ is given by the following equations. For j from 1 to d-1, x_j can take any value in $[0, R_j]$. If $a_d > p_d$, then we have the inequality $x_d \geq a_d$; and if $a_d < p_d$, then we have the inequality $x_d \leq a_d$. For the other coordinates, we have the same inequalities and equations as in the definition of A. On the other hand, if p_j does not lie in $[a_j, a'_j]$ for some j between 1 and d-1, then $F_p(A)$ is empty.

We define $B_p(A)$ to be the intersection of the boundary of $F_p(A)$ with the boundary of R. This intersection is a rectilinear (k-1)-chain in the boundary of R.

We extend these definitions linearly to all of the rectilinear (k-1)-chains. The k-chain F and the (k-1)-chain B will be $F_p(C)$ and $B_p(C)$ for a point p that we will choose later.

We first claim that if C is a relative (k-1)-cycle, then the boundary of $F_p(C)$ is equal to $C + B_p(C)$. In other words, we want to show that $\partial F_p(C) - B_p(C) - C$ is equal to zero. We first consider $F_p(\partial F_p(C) - B_p(C) - C)$. A straight-forward calculation shows that for each face A, $F_p(\partial F_p(A) - B_p(A) - A)$ is equal to zero. By linearity, $F_p(\partial F_p(C) - B_p(C) - C)$ is also equal to zero. Now $\partial F_p(C) - B_p(C)$ is a relative (k-1)-cycle, and C is also a relative (k-1)-cycle, so their difference is a relative (k-1)-cycle. Therefore, it suffices to check that if K is a rectilinear relative (k-1)-cycle and $F_p(K) = 0$, then K vanishes. Since $F_p(K)$ is zero, it follows that every face of K contains x_1 as a tangent coordinate. Since K is a relative cycle, it follows that K is equal to $[0, E_1] \times K'$, where K' is a relative (k-2)-cycle in the quotient rectangle with dimensions $E_2 \times ... \times E_n$. Let p' be the orthogonal projection of p down to this quotient rectangle. Then $F_p(K)$ is equal to $[0, E_1] \times F_{p'}(K')$. Since $F_p(K)$ vanishes, it follows that $F_{p'}(K')$ vanishes. Now it is easy to check that if K' is zero-dimensional and $F_{p'}(K')$ is equal to zero, then K' vanishes. Therefore, by induction on the dimension of the rectangle, it follows that K is zero, proving our claim.

For a fixed A with axes parallel to J and a random p, we can bound the expectation value of the I-volume of $F_p(A)$. Let d be the smallest index not in J. Then the expectation value of the $(J+d)$-volume of $F_p(A)$ is at most $1/2E_d\text{volume}(A)$. For every other I, the I-volume of $F_p(A)$ is zero. Similarly, for each e in the range $1 \leq e \leq d$, the expectation value of the $(J+d-e)$-volume of $B_p(A)$ is at most $(E_d/E_e)\text{volume}(A)$. Therefore, we can choose p so that the I-volume of $F_p(A)$ and the J-volume of $B_p(A)$ obey estimates 1 and 2 above. This finishes the proof of step 1.

Step 2. Let C be a closed (k-1)-cycle in the rectangle R. Suppose that C has J-volume V_J. Then C has a filling F whose I-volume obeys the following inequality. If d is the smallest index in I, then $W_I < E_d V_{I-d}$.

First, compress C to the plane $x_1 = 0$, by filling vertically. The fillings have appropriately bounded volumes, (where 1 is the smallest member of I), and the compression kills all J-volume of C when J contains 1, and leaves the rest of the J-volume unaltered. Then we proceed inductively, compressing the result to the sub-plane $x_2 = 0$ again by filling vertically. The fillings have appropriately bounded volumes, (where 2 is the smallest member of I), and the compression kills all J-volume of C when J contains 2. We continue to proceed in this way until we have compressed C into a (k-1)-plane, which fills it. This finishes the proof of Step 2. $\square$

As a special case of this proposition, we see that any rectilinear (k-1)-cycle of volume V can be filled by a k-chain with volume less than $C\max(E_k, E_{n-k+1})V$. Since any flat (k-1)-chain of volume V can be approximated by a rectilinear one with volume less than CV, the same inequality holds for all (k-1)-cycles. This proves proposition 7.3.1.

Now we have the geometrical tools to prove Theorem 7.1. Again let β be an exact k-form with the given bounds. We define a linear functional L on rectilinear (k-1)-cycles, where $L(C)$ is defined to be the integral of β over any filling of C. Next we define a norm on the rectilinear (k-1)-chains, which is equivalent to the flat norm, but which has a particular directionally dependent form which reflects the inequality that we want to prove. Our norm on a rectilinear (k-1)-chain C is equal to the infimum of $|C - \partial D|_A + |D|_B$ over all rectilinear chains D, in terms of norms $||_A$ and $||_B$ which we have to define.

Write a rectilinear (k-1)-chain C in the form $\sum_J C_J$, where J varies over the (k-1)-tuples of coordinates, and C_J is parallel to the corresponding (k-1)-plane. The norm $|C|_A$ is defined to be $\sum a_j$ volume(C_J). (The numbers a_J were defined in the statement of this theorem.) Similarly, write a rectilinear k-chain D in the form $\sum_I D_I$, where I varies over the k-tuples of the coordinates, and D_I is parallel to the corresponding k-plane. Our norm $|D|_B$ is defined to be $\sum b_I$ volume(D_I). We have to check that the functional L is bounded in terms of our variant of the flat norm. Clearly if a cycle has a rectilinear filling D, then the integral of β over this filling is at most $\sum b_I$ volume(D_I). According to Lemma 7.1, if C is a rectilinear (k-1)-cycle, then it admits a rectilinear filling D with volumes bounded as follows. If 1 is in I and e is the smallest index not in I, then volume$(D_I) < C \sum_{d=1}^{e-1} E_d$ volume(C_{I-d}), and if 1 is not in I and d is the smallest index in I, then volume$(D_I) < CE_d$ volume(C_{I-d}). The evaluation of β on this filling is bounded by $\sum b_I$ volume(D_I), which is bounded by $C \sum a_J$ volume(C_J). Therefore, the linear functional L has norm less than C with respect to our norm. By the Hahn-Banach theorem, L extends to a linear functional on the space of all rectilinear chains with norm less than C.

Since our norm is equivalent to the standard flat norm, and since the rectilinear chains are dense in the space of flat chains, this extension of L is a flat cochain. By Whitney's theorem, it corresponds to a bounded measurable differential form α. Because the integral of α over any rectilinear (k-1)-cycle is equal to the integral of β over a filling of that cycle, $d\alpha = \beta$. As a member of the dual to the space of rectilinear cochains, α has norm less than C using our norm. It follows from this that the integral of α over a (k-1)-cube in the J-plane with side length s is less than

$Ca_J s^{k-1}$. This inequality shows that the coefficients of the differential form α obey the stated bounds.

7.5 Lipshitz maps with large rational homotopy invariants

In this section, we will construct Lipshitz maps from ellipsoids to wedges of spheres with large values of certain rational homotopy invariants.

Theorem 7.2. *Let X be a wedge of spheres $S^{k_1} \vee ... \vee S^{k_{d+1}}$, with each sphere of dimension at least 2. Let $n = -d + \sum k_i$. Let E be an n-dimensional ellipsoid with principal axes $E_1 \leq ... \leq E_n$, and suppose that E_1 is greater than 1. Then there is a map f from E to X with Lipshitz constant less than C so that the following holds. Let a_i be the pullback of the volume form of S^{k_i} by f.*

$$\int_E Pa_1 \wedge ... \wedge Pa_d \wedge a_{d+1} > (E_{n-k_1+1} E_{n-k_1-k_2+2} ... E_{n-k_1-...-k_d+d})(E_1 ... E_n).$$

Also, certain other rational homotopy invariants of the map f vanish. In particular, the integral $\int_E Pa_1 \wedge ... \wedge Pa_{l-1} \wedge a_l \wedge \Omega = 0$ when the form $\Omega = P\Omega_1 \wedge ... \wedge P\Omega_r$, is made from the forms $a_{l+1}, ..., a_{d+1}$ using $(d+1-l)$ primitives, and if r is at least 2.

Proof. By induction, it suffices to prove this proposition for n-dimensional ellipsoids, assuming that it holds for ellipsoids of dimension less than n.

Actually, we need to use a slightly different inductive hypothesis. If X is a wedge of spheres, let X' denote the complement of the base point in X. If f is a map from E to X, define the support of f to be the closure of $f^{-1}(X')$. If E is an ellipsoid given by the equation $\sum_{i=0}^n x_i^2/E_i^2 = 1$, then we call the regions of E where $x_n < -(9/10)E_n$ or where $x_n > (9/10)E_n$ the tips of E. We will prove inductively that we can find a map f supported away from the tips of E, and with the properties listed in the proposition. By induction, we can assume that this is the case for ellipsoids of dimension less than n.

We think of E as the double of a rectangle. We need to describe two open sets U_1 and U_2 in E. We can write each set U_i as a product of the double of the rectangle with dimensions $E_1 \times ... \times E_{n-k_1}$ with an open region R_i properly contained in the k_1-dimensional rectangle with dimensions $E_{n-k_1+1} \times ... \times E_n$. The region R_1 is equal to a central sub-rectangle given by the equations $1/4E_i < x_i < 3/4E_i$ for each i from (n-k+1) to n. The region R_2 is equal to an "annulus" around R_1. It is given by the conditions $1/10E_i < x_i < 9/10E_i$ for each i from (n-k+1) to n, and $x_i > 4/5E_i$ or $x_i < 1/5E_i$ for one i between (n-k+1) and n. Clearly R_1 and R_2 are disjoint.

Since $E_1 > 1$, we can easily construct a compactly supported map f_1 from R_1 to S^{n-k_1+1} of degree $E_{n-k_1+1}...E_n$, with Lipshitz constant less than C. We define the restriction of f to U_1 by projecting to R_1 and then applying f_1. This restriction has Lipshitz constant less than C.

The region U_1 is a thick neighborhood of an (n-k)-sphere S_1, given by the product of the double of $E_1 \times ... \times E_{n-k_1}$ with the central point of R_1. The region U_2 is bilipshitz to $E' \times S_2$, where E' is an $(n-k_1+1)$ dimensional ellipsoid of principal axes $E_0 \leq E_1 \leq ... \leq E_{n-k_1+1}$ minus its tips, and S_2 is a (k_1-1)-dimensional sphere lying in the k_1-plane given by the equations $x_1 = ... = x_{n-k_1} = 0$, and which wraps around R_1. In particular, the linking number of S_1 and S_2 is 1.

By induction, there is a compactly supported map f_2 from E' to the wedge $S^{k_2} \vee ... \vee S^{k_{d+1}}$ which satisfies the conclusions of the proposition. In particular, the map f_2 has Lipshitz constant less than C, and the integral $\int_{E'} Pa_2 \wedge ... \wedge Pa_d \wedge a_{d+1}$ is at least $E_{n-k_1-k_2+2}...E_{n-k_1-...-k_d+d}(E_1...E_{n-k_1+1})$. We define the restriction of f to U_2 by projecting from U_2 to E' and then applying f_2. This restriction is compactly supported and has Lipshitz constant less than C.

Finally, we extend f to the part of E away from U_1 and U_2 by mapping to the base point of X. The resulting map f has Lipshitz constant less than C and is supported away from the tips of E. Since the spheres S_1 and S_2 have linking number 1, the integral $\int_E Pa_1 \wedge (Pa_2 \wedge ... \wedge Pa_d \wedge a_{d+1})$ is equal to the following expression.

$(E_{n-k_1+1}...E_n)E_{n-k_1-k_2+2}...E_{n-k_1-...-k_d+d}(E_1...E_{n-k_1+1})$.

This expression reduces to $E_{n-k_1+1}...E_{n-k_1-...-k_d+d}(E_1...E_n)$. Therefore, f obeys the lower bound stated in the proposition.

It remains to prove that the other rational homotopy invariants stated in the proposition vanish. We will also prove this by induction on the dimension. Recall that a_i is defined to be the pullback of the volume form of S^{k_i} to E by the map f. Analogously, we can define a_i' to be the pull-back of the volume form to E' by the map f_2.

The rational homotopy invariant which we have to deal with is $\int_E Pa_1 \wedge ... \wedge Pa_{l-1} \wedge a_l \wedge \Omega$.

First we deal with the case that l is greater than 1. In this case, our integral is equal to $(\int_{R_1} a_1)(\int_E' Pa_2' \wedge ... \wedge Pa_{l-1}' \wedge a_l' \wedge \Omega')$, where Ω' has the same expression as Ω but with each form a_i replaced by a_i'. By induction, the second integral is zero, and we are done.

Next we deal with the case that $l = 1$. In this case, our integral is equal to $\int_E a_1 \wedge P\Omega_1 \wedge ... \wedge P\Omega_r$. Each of the forms Ω_i has degree at least 2. Since r is at least 2, each of the forms Ω_i has degree at most $n - k_1$. Again, let Ω_i' be the form on E' that has the same expression as Ω, but with each form a_i replaced by a_i'. Each of the primitives inside of Ω_i and also each primitive $P\Omega_i'$ can be defined on E'. Therefore, the form $P\Omega_1 \wedge ... \wedge P\Omega_r$ is the pullback of a compactly supported form on E' by the projection from U_2 to E'. In particular, this form is supported on U_2, for an appropriate choice of primitives. But the form a_1 is supported on U_1. Therefore, our integral is equal to zero. □

7.6 Estimates for rational homotopy invariants

In this section, we apply the results of the previous two sections to prove our main estimates for rational homotopy invariants of maps to bouquets of spheres. Using the directionally dependent primitive estimate Theorem 7.1, we will get upper bounds for rational homotopy invariants of k-contracting maps. On the other hand, we can use Theorem 7.2 to produce Lipshitz maps with fairly large rational homotopy invariants. For many invariants, but not all of them, we will prove that the invariant of a map guaranteed by Theorem 7.2 is within a constant factor of the maximum value of that invariant for any k-contracting map.

We will work mostly with non-repeating RH invariants - namely, those in which a given form $f^*(\alpha_i) = a_i$ appears no more than once.

Invariants of order 1

In this subsection, we consider invariants of order 1. We define X to be the bouquet $X = S^{k_1} \vee S^{k_2}$, where $2 \leq k_1 \leq k_2$ and $n = k_1 + k_2 - 1$. Let f be a map from S^n to X, and define a_i to be the pullback of the volume form of S^{k_i}. There is only one rational homotopy invariant of maps from S^n to X.

$H(f) = \int_E Pa_1 \wedge a_2$.

Proposition 7.6.1. *For any k_1-contracting map f from E^n to X, the invariant $H(f)$ is less than $CE_{n-k_1+1}E_1 \ldots E_n$. On the other hand, if $E_1 > C$, then there is a 1-contracting map from E to X with $H(f)$ at least $cE_{n-k_1+1}E_1 \ldots E_n$.*

Proof. Since f is k_1-contracting, the norms of a_1 and a_2 are bounded by 1 pointwise. According to Proposition 7.3.1, the norm of Pa_1 is bounded by $Cmax(E_{n-k_1+1}, E_{k_1})$. Since $n = k_1+k_2-1$, $n-k_1+1$ is equal to k_2 which is at least as large as k_1. Therefore, Pa_1 is bounded by CE_{n-k_1+1}. It follows that $H(f)$ is bounded by $CE_{n-k_1+1}E_1 \ldots E_n$.

On the other hand, if E_1 is sufficiently large, then Theorem 7.2 constructs a 1-contracting map f from E to $S^{k_1} \vee S^{k_2}$ with RH(f) greater than $cE_{n-k_1+1}E_1 \ldots E_n$. □

By the same method, we can get an analogous estimate for the Hopf invariant. Suppose that f is a 2k-contracting map from E^{4k-1} to S^{2k}. By the same argument we used above, Hopf(f) is at most $CE_{2k}E_1 \ldots E_{4k-1}$. Using Theorem 7.2, if E_1 is sufficiently big, we get a 1-contracting map from E to $S^{2k} \vee S^{2k}$ with $\int Pa_1 \vee a_2$ greater than $cE_{2k}E_1 \ldots E_{4k-1}$. Composing this map with maps from $S^{2k} \vee S^{2k}$ to S^{2k} of bidegree (1,0), (0,1), and (1,1), we get 1-contracting maps from E to S^k with Hopf invariants $\int Pa_1 \wedge a_1$, $\int Pa_2 \wedge a_2$, and $\int P(a_1+a_2) \wedge (a_1+a_2)$. But a quick calculation shows that $\int P(a_1 + a_2) \wedge (a_1 + a_2) - \int Pa_1 \wedge a_1 - \int Pa_2 \wedge a_2 = 2 \int Pa_1 \wedge a_2 > cE_{2k}E_1 \ldots E_{4k-1}$. Therefore, one of our three 1-contracting maps has Hopf invariant at least $cE_{2k}E_1 \ldots E_{4k-1}$.

Invariants of order 2

In this subsection we consider non-repeating invariants of order 2. We define X to be the bouquet $X = S^{k_1} \vee S^{k_2} \vee S^{k_3}$, with $2 \le k_1 \le k_2 \le k_3$, and we define n to be $k_1 + k_2 + k_3 - 2$.

The non-repeating RH invariants of order 2 for maps from S^n to X are a 2-dimensional vector space spanned by the following two invariants.

$$H_1(f) = \int Pa_1 \wedge a_2 \wedge Pa_3.$$

$$H_2(f) = \int a_1 \wedge Pa_2 \wedge Pa_3.$$

A general rational homotopy invariant in this vector space is given by a combination $c_1 H_1 + c_2 H_2$.

Proposition 7.6.2. *If f is a k_1-contracting map from E to X, then the rational homotopy invariants of f obey the following bounds.*

$H_1(f) < C E_{n-k_1+1} E_{n-k_1-k_2+2} E_1 ... E_n.$

$H_2(f) < C E_{n-k_2+1} E_{n-k_1-k_2+2} E_1 ... E_n.$

On the other hand, if $E_1 > C$, then we will construct 1-contracting maps f_1 and f_2 which show that these upper bounds are sharp.

$H_1(f_1) > c E_{n-k_1+1} E_{n-k_1-k_2+2} E_1 ... E_n.$

$H_2(f_2) > c E_{n-k_2+1} E_{n-k_1-k_2+2} E_1 ... E_n.$

Similar estimates hold for any combination $c_1 H_1 + c_2 H_2$. For any k_1-contracting map f, this rational homotopy invariant obeys the following upper bounds.

$(c_1 H_1 + c_2 H_2)(f) < C(|c_1| E_{n-k_1+1} + |c_2| E_{n-k_2+1}) E_{n-k_1-k_2+2} E_1 ... E_n.$

On the other hand, if $E_1 > C$, then we will construct a 1-contracting map f which shows that this upper bound is sharp.

$(c_1 H_1 + c_2 H_2)(f) > c(|c_1| E_{n-k_1+1} + |c_2| E_{n-k_2+1}) E_{n-k_1-k_2+2} E_1 ... E_n.$

Proof. If f is k_1-contracting, then each form a_i is bounded by 1 pointwise. If we apply Proposition 7.3.1, we can get an upper bound for each primitive Pa_i. The primitive Pa_1 will be bounded by $C max(E_{k_1}, E_{n-k_1+1})$, and Pa_3 will be bounded by

$Cmax(E_{k_3}, E_{n-k_3+1})$. In the first case, we can say that E_{n-k_1+1} is larger than E_{k_1}, but in the second case, either E_{k_3} or E_{n-k_3+1} could be larger depending on the choice of the dimensions k_i. In any case, $H_1(f)$ is bounded by $CE_{n-k_1+1}max(E_{k_3}, E_{n-k_3+1})$. This estimate turns out not to be sharp.

To improve the estimate, we use the directionally dependent estimate Theorem 7.1 in place of Proposition 7.3.1. We write Pa_1 as a sum $f_I dx^I$, where I varies over the $(k_1 - 1)$-tuples of coordinates. Similarly, we write Pa_3 as a sum $g_J dx^J$, where J varies over the $(k_3 - 1)$-tuples of coordinates. Theorem 7.1 gives us upper bounds on each function f_I and g_J. For example, if I does not contain 1, and if d is the smallest coordinate in I, then f_I is less than CE_{d-1}. If I does contain 1, and if e is the smallest coordinate not in I, then f_I is less than CE_e. Identical estimates hold for g_J.

Now we expand $H_1(f)$ as $\int_E a_2 \wedge (\sum_{I,J} f_I g_J dx^I \wedge dx^J)$. The main point of this exercise is that many of the terms $dx^I \wedge dx^J$ vanish. The maximum of the norm of $f_I g_J$ over all pairs (I, J) is equal to the upper bound for $Pa_1 \wedge Pa_3$ that we got using Proposition 7.3.1. Because of the vanishing terms, however, the norm of $Pa_1 \wedge Pa_3$ is bounded by the maximum of the norm of $f_I g_J$ over all the *disjoint* pairs (I, J). Given our bounds on f_I and g_J, it follows that this maximum is less than $CE_{n-k_1+1}E_{n-k_1-k_2+2}$. Therefore, $H_1(f)$ is bounded by $CE_{n-k_1+1}E_{n-k_1-k_2+2}E_1 ... E_n$.

Applying the same analysis to Pa_2 and Pa_3, we see that $H_2(f)$ is bounded by $CE_{n-k_2+1}E_{n-k_1-k_2+2}E_1 ... E_n$. Combining these bounds, we see immediately that $(c_1H_1 + c_2H_2)(f)$ is less than $C(|c_1|E_{n-k_1+1}|c_2|E_{n-k_2+1})E_{n-k_1-k_2+2}E_1 ... E_n$. We have now established all the upper bounds in the proposition.

Next we have to construct maps with large rational homotopy invariants. From now on, we can assume that $E_1 > C$. Then we can construct 1-contracting maps using Theorem 7.2, after a suitable reordering of the spheres in the target X.

The map f_1 is the map given by Theorem 7.2 applied to $X = S^{k_1} \vee S^{k_2} \vee S^{k_3}$. According to Theorem 7.2, we have the following lower bound.

$\int_E Pa_1 \wedge Pa_2 \wedge a_3 > cE_{n-k_1+1}E_{n-k_1-k_2+2}E_1 ... E_n$.

We also know that $\int_E a_1 \wedge Pa_2 \wedge Pa_3 = 0$. Combining these equations, we see that $|H_1(f_1)| > cE_{n-k_1+1}E_{n-k_1-k_2+2}E_1 ... E_n$. If necessary, we can precompose f_1 with a reflection so that $H_1(f_1)$ is positive.

The map f_2 is the map given by Theorem 7.2 applied to $X = S^{k_2} \wedge S^{k_1} \wedge S^{k_3}$. According to Proposition 7.5.1, we have the following lower bound.

$\int_E Pa_2 \wedge Pa_1 \wedge a_3 > cE_{n-k_2+1}E_{n-k_1-k_2+2}E_1...E_n$.

We also know that $\int_E a_2 \wedge Pa_1 \wedge Pa_3 = 0$. Combining these equations, we see that $|H_2(f_2)| > cE_{n-k_2+1}E_{n-k_1-k_2+2}E_1...E_n$. If necessary, we can precompose f_2 with a reflection so that $H_2(f_2)$ is positive.

So far, we have constructed the maps f_1 and f_2 described in the proposition. We still have to construct a map f which gives a large value to a more general rational homotopy invariant $(c_1H_1 + c_2H_2)$. When we were discussing f_1, we noted that according to Theorem 7.2, $\int_E a_1 \wedge Pa_2 \wedge Pa_3 = 0$. This equation means that $H_2(f_1) = 0$. Similarly, in discussing f_2, we noted that $\int_E a_2 \wedge Pa_1 \wedge Pa_3 = 0$, which means that $H_1(f_2)$ equals zero. Therefore, we can estimate $(c_1H_1 + c_2H_2)(f_i)$.

$|(c_1H_1 + c_2H_2)(f_1)| > c|c_1|E_{n-k_1+1}E_{n-k_1-k_2+2}E_1...E_n$.

$|(c_1H_1 + c_2H_2)(f_2)| > c|c_2|E_{n-k_2+1}E_{n-k_1-k_2+2}E_1...E_n$.

Therefore, either f_1 or f_2 satisfies the lower bound for the value of $(c_1H_1 + c_2H_2)$. $\square$

At this point, I would like to give an example to put these results into some perspective. Let's consider $\pi_{13}(S^4 \vee S^5 \vee S^6) \otimes \mathbb{Q}$, which is a 2-dimensional vector space. The invariants H_1 and H_2 define coordinates on this vector space. Let $D_k(E)$ denote the subset of $\pi_{13}(S^4 \vee S^5 \vee S^6)$ which can be achieved by k-contracting from an ellipsoid E. Let $R(E)$ denote the rectangle in $\pi_{13}(S^4 \vee S^5 \vee S^6) \otimes \mathbb{Q}$ given by the following equations.

$|H_1| < E_{10}E_6E_1...E_{13}$.

$|H_2| < E_9E_6E_1...E_{13}$.

We have proven that $D_4(E)$ is contained in CR. We have also proven that if $E_1 > C$, then the convex hull of $D_1(E)$ contains cR. Probably, with a little more work, we could show that if $E_1 > C$, then $D_1(E)$ actually contains cR. If this were true, then we would know that $cR \subset D_1(E) \subset D_2(E) \subset D_3(E) \subset D_4(E) \subset CR$. Also, if $E_1 > C$, then $D_5(E)$ is clearly the whole homotopy group. In short, these techniques give a fairly good picture of $D_k(E)$ under the assumption that $E_1 > C$. On the other hand, if $E_1 << 1$, then we still know that $D_4(E)$ is contained in CR, but I don't know whether $D_4(E)$ is much smaller.

Invariants of order 3

In the next few pages, we will try to indicate how far these techniques can be pushed. In the last subsection, we proved a pretty sharp estimate for each non-repeating invariant of order 2. In this subsection, we will prove an analogous estimate for each non-repeating invariant of order 3. The idea of the proof is the same, only the algebra is more tedious. In the next subsection, I will explain why these techniques do not give an analogous estimate for each non-repeating invariant of order 4.

We define X to be a bouquet of four spheres $S^{k_1} \vee ... \vee S^{k_4}$, with $2 \leq k_1 \leq k_2 \leq k_3 \leq k_4$, and we define n to be $k_1 + k_2 + k_3 + k_4 - 3$. The space of non-repeating rational homotopy invariants of maps from S^n to X is six-dimensional. If f is a map from S^n to X, we define a_i to be the pullback of the volume form of S^{k_i}. In terms of these forms, we can write down a basis for the rational homotopy invariants.

$H_1(f) = \int Pa_1 \wedge Pa_2 \wedge a_3 \wedge Pa_4$.

$H_2(f) = \int Pa_1 \wedge a_2 \wedge Pa_3 \wedge Pa_4$.

$H_3(f) = \int a_1 \wedge Pa_2 \wedge Pa_3 \wedge Pa_4$.

$H_4(f) = \int P(a_1 \wedge Pa_2) \wedge a_3 \wedge Pa_4$.

$H_5(f) = \int P(a_1 \wedge Pa_3) \wedge a_2 \wedge Pa_4$.

$H_6(f) = \int P(a_1 \wedge Pa_4) \wedge a_2 \wedge Pa_3$.

A general rational homotopy invariant is given by a combination of these invariants $\sum c_i H_i$. We are going to prove a proposition which gives fairly sharp estimates for any of these invariants provided that $E_1 > C$. In order to state the estimates, we need to make some definitions.

$Z_1 = E_{n-k_1+1} E_{n-k_1-k_2+2} E_{n-k_1-k_2-k_3+3}$.

$Z_2 = E_{n-k_1+1} E_{n-k_1-k_3+2} E_{n-k_1-k_2-k_3+3}$.

$Z_3 = E_{n-k_2+1} E_{n-k_2-k_3+2} E_{n-k_1-k_2-k_3+3}$.

$Z_4 = E_{n-k_2+1} E_{n-k_1-k_2+2} E_{n-k_1-k_2-k_3+3}$.

$Z_5 = E_{n-k_3+1} E_{n-k_3-k_1+2} E_{n-k_1-k_2-k_3+3}$.

$Z_6 = max(E_{n-k_3+1} E_{n-k_3-k_2+2} E_{n-k_1-k_2-k_3+3}, E_{n-k_4+1} E_{n-k_4-k_1+2} E_{n-k_4-k_1-k_2+3})$.

Proposition 7.6.3. *If f is a k_1-contracting map from E to X, then the homotopy invariant $(\sum c_i H_i)(f)$ has norm less than $C \sum |c_i| Z_i E_1 ... E_n$. On the other hand, if $E_1 > C$, then there is a 1-contracting map from E to X with $(\sum c_i H_i)(f)$ greater than $c \sum |c_i| Z_i E_1 ... E_n$.*

Proof. The first step is to give the upper bounds. We suppose that f is a k_1-contracting map from E to X, and we need to prove that $H_i(f) < C Z_i E_1 ... E_n$. This inequality follows by repeatedly applying Theorem 7.1 to estimate each coefficient of each primitive. Since the calculations are quite tedious, we will omit most of them. We include the calculations for H_6, which is the most difficult case.

Estimating $\int P(Pa_4 \wedge a_1) \wedge Pa_3 \wedge a_2$

Let $\alpha = \sum \alpha_I dx^I$ be a primitive of a_4. In this equation, I is a set of $(k_4 - 1)$ numbers in the range $1...n$. We define I^+ to be the smallest number which in not in I. We define I^- to be one less than the smallest number in I. (If 1 is in I, I^- then E_{I^-} is defined to be zero.) According to Theorem 7.1, we can choose α to obey the following inequalities.

$$|\alpha_I| < C(E_{I^+} + E_{I^-}). \tag{1}$$

We define $\beta = \sum \beta_J dx^J$ to be $\alpha \wedge a_1$. In this equation, J is a set of $(k_1 + k_4 - 1)$ numbers. We define $J(1)$ to be the largest number in J, $J(2)$ to be the second largest number, and so on. The coefficients of β obey the following inequalities.

$$|\beta_J| < C \sum_{I \subset J} |\alpha_I|$$

$$< C(E_{\min(J^+, k_4)} + E_{J(k_4-1)-1}). \tag{2}$$

Let $\gamma = \sum \gamma_K dx^K$ be a primitive of β. According to Theorem 7.1, we can choose γ to obey the following bounds.

$$|\gamma_K| < C(E_{K^+} |\beta_{K \cup K^+}| + \sum_{i=2}^{K^-} E_i |\beta_{K \cup i}|). \tag{3}$$

(Remark: We have left out the case $i = 1$ in the second term of the sum because it occurs in the first term.)

Next we plug our bounds for β into this equation. The first term is controlled by the following expression.

$$C(E_{K^+}E_{\min((K\cup K^+)^+,k_4)} + E_{K^+}E_{(K\cup K^+)(k_4-1)-1}).$$

We can make one small simplification to this expression. We know that $(K \cup K^+)^+$ is greater than K^+. If $(K \cup K^+)(k_4 - 1) - 1$ is no more than K^+, then the second term is controlled by the first term. On the other hand, if $(K \cup K^+)(k_4 - 1) - 1$ is greater than K^+, then it is equal to $K(k_4 - 1) - 1$. Therefore, the first term of equation 3 is bounded by the following expression.

$$C(E_{K^+}E_{\min((K\cup K^+)^+,k_4)} + E_{K^+}E_{K(k_4-1)-1}).$$

The second term of equation 3 is bounded by the following expression.

$$C(\sum_{i=2}^{K^-} E_i(E_{(K\cup i)(k_4-1)-1}) + E_{(K\cup i)^+}).$$

This expression simplifies considerably. Since i is at least 2, the term $(K \cup i)^+$ is equal to 1, which is always less than the other term. Since i is less than any number in K, $(K \cup i)(k_4 - 1)$ is equal to $K(k_4 - 1)$. Finally, the term E_i is always less than E_{K^-}. Therefore, the second term of equation 3 is bounded by the following expression.

$$CE_{K^-}E_{K(k_4-1)-1}.$$

Assembling our bounds for the two terms of equation 3, we get the follow inequality for γ.

$$|\gamma_K| < C(E_{K^+}E_{\min((K\cup K^+)^+,k_4)} + E_{K^+}E_{K(k_4-1)-1} + E_{K^-}E_{K(k_4-1)-1}). \qquad (4)$$

Next we let $\delta = \sum \delta_L dx^L$ be a primitive of a_3. In this equation, L is a set of $(k_3 - 1)$ numbers. By Theorem 7.1, we can choose δ to obey the following bounds.

$$|\delta_L| < C(E_{L^+} + E_{L^-}). \tag{5}$$

The integrand $P(Pa_4 \wedge a_1) \wedge Pa_3 \wedge a_2$ is equal to $\gamma \wedge \delta \wedge a_2$. Therefore, the norm of the integrand is bounded by the norm of $\gamma \wedge \delta$.

$$|\gamma \wedge \delta| < \max_{K,L \text{disjoint}} |\gamma_K||\delta_L|.$$

Expanding equations 4 and 5, we can bound this expression as follows.

$$< \max_{K,L \text{disjoint}} C$$

$$(E_{K^+} E_{\min((K \cup K^+)^+, k_4)} E_{L^+} + \tag{a}$$

$$E_{K^+} E_{\min((K \cup K^+)^+, k_4)} E_{L^-} + \tag{b}$$

$$E_{K^+} E_{K(k_4-1)-1} E_{L^+} + \tag{c}$$

$$E_{K^+} E_{K(k_4-1)-1} E_{L^-} + \tag{d}$$

$$E_{K^-} E_{K(k_4-1)-1} E_{L^+} + \tag{e}$$

$$E_{K^-} E_{K(k_4-1)-1} E_{L^-}). \tag{f}$$

Going term by term, we show that each monomial in the above expression is bounded by Z_6. We recall that Z_6 is the maximum of $E_{n-k_3+1} E_{n-k_3-k_2+2} E_{n-k_1-k_2-k_3+3}$ and $E_{n-k_4+1} E_{n-k_4-k_1+2} E_{n-k_4-k_1-k_2+3}$. We also recall that K is a set of $(k_1 + k_4 - 1)$ numbers and L is a set of $(k_3 - 1)$ numbers.

a. $E_{K^+} E_{\min((K \cup K^+)^+, k_4)} E_{L^+}$

Only one of K and L includes 1. If K includes 1, then this term is no more than $E_{k_4+k_1-1} E_{k_4} E_1$. This expression is equal to $E_{n-k_2-k_3+2} E_{n-k_2-k_3-k_1+3} E_1$, which is less than Z_6. If L includes 1 and 2, then this term is no more than $E_1 E_2 E_{k_3}$, which is less than Z_6. If L includes 1 but not 2, then this expression is less than $E_1 E_{k_4} E_{k_3}$, which is less than Z_6.

b. $E_{K^+} E_{\min((K \cup K^+)^+, k_4)} E_{L^-}$

This term is no more than $E_{k_1+k_4-1} E_{k_4} E_{n-k_3+1}$. After reordering the terms, this expression is equal to $E_{n-k_3+1} E_{n-k_2-k_3+2} E_{n-k_1-k_2-k_3+3}$, which is no more than Z_6.

c. $E_{K^+} E_{K(k_4-1)-1} E_{L^+}$

This term in no more than $E_{k_1+k_4-1} E_{n-k_4+1} E_{k_3}$, which is no more than Z_6.

d. $E_{K^+} E_{K(k_4-1)-1} E_{L^-}$

If K^+ is greater than k_1, then K must include each number from 1 to k_1. In this case $K(k_4-1)$ is equal to k_1, and the term above is bounded by $E_{k_1+k_4-1} E_{k_1-1} E_{n-k_3+1}$, which is less than Z_6. If K^+ is not greater than K_1, then this term is bounded by $E_{k_1} E_{n-k_3+1} E_{n-k_3-k_4+1}$, which is no more than Z_6.

e. $E_{K^-} E_{K(k_4-1)-1} E_{L^+}$

This term is no more than $E_{n-k_1-k_4+2} E_{n-k_4+1} E_{k_3}$. After reordering the terms, this expression is equal to $E_{n-k_4+1} E_{n-k_1-k_4+2} E_{n-k_1-k_2-k_4+3}$, which is no more than Z_6.

f. $E_{K^-} E_{K(k_4-1)-1} E_{L^-}$

This term is no more than $E_{n-k_3+1} E_{n-k_3-k_4+2} E_{n-k_3-k_4-k_1+3}$, which is no more than Z_6.

Therefore, $|\gamma \wedge \delta|$ is less than CZ_6. Therefore, for our choice of primitives, the integrand $P(Pa_4 \wedge a_1) \wedge Pa_3 \wedge a_2$ is less than CZ_6 pointwise. It follows that the integral $H_6(f) = \int_E P(Pa_4 \wedge a_1) \wedge Pa_3 \wedge a_2$ has norm less than $CZ_6 \text{volume}(E)$. Therefore, the invariant $H_6(f)$ has norm less than $CZ_6 E_1 ... E_n$. By analogous, somewhat easier, calculations, it follows that $|H_i(f)| < CZ_i E_1 ... E_n$ for each i.

Next we have to construct mappings which satisfy the lower bounds. We only need to do this construction when $E_1 > C$, which we assume from now on. The first step is to construct maps f_i which have large values of H_i. Again, we will use the example of H_6. By applying Theorem 7.2 to $X = S^{k_4} \vee S^{k_1} \vee S^{k_2} \vee S^{k_3}$, we get a 1-contracting map f from E to X with the following bound.

$\int_E Pa_4 \wedge Pa_1 \wedge Pa_2 \wedge a_3 >$

$cE_{n-k_4+1}E_{n-k_4-k_1+2}E_{n-k_4-k_1-k_2+3}E_1...E_n$.

We also know that $\int_E Pa_4 \wedge a_1 \wedge Pa_2 \wedge Pa_3 = 0$. Using this equality, it follows that, up to sign, the integral we just estimated is equal to $\int_E Pa_4 \wedge Pa_1 \wedge a_2 \wedge Pa_3$, which we call I_2. We also know that $\int_E a_4 \wedge Pa_1 \wedge P(a_2 \wedge Pa_3) = 0$. Using this equality, it follows that up to sign, the integral I_2 is equal to $\int_E Pa_4 \wedge a_1 \wedge P(a_2 \wedge Pa_3)$. Up to sign, this integral is equal to $\int_E P(a_1 \wedge Pa_4) \wedge a_2 \wedge Pa_3$, which is the definition of $H_6(f)$.

Similarly, if we apply Theorem 7.2 to $X = S^{k_3} \wedge S^{k_2} \wedge S^{k_1} \wedge S^{k_4}$, we get a 1-contracting map f' with $H_6(f') > cE_{n-k_3+1}E_{n-k_3-k_2+1}E_{n-k_1-k-2-k_3+3}E_1...E_n$. Taking the better choice of f' and f, we get a 1-contracting map from E to X with H_6 greater than $cZ_6E_1...E_n$. By analogous arguments, we can define 1-contracting functions f_i, so that $H_i(f_i) > cZ_iE_1...E_n$.

Finally, we will show that for any rational homotopy invariant $H = (\sum c_iH_i)$, one of the functions f_i gives a large value. This last part of the proof depends on the fact that the monomials Z_i are all different from each other. (The funny-looking basis H_i was chosen to ensure that the Z_i would be pairwise distinct.)

We consider the numbers $|c_i|Z_i$. First we deal with the case that one of these numbers is much larger than all of the others. Suppose for instance that $|c_3|Z_3 > 100(C/c)|c_i|Z_i$ for all the other values of i. In this case, $H(f_3)$ has a very large norm. As we proved above, $|c_3H_3(f_3)| > c|c_3|Z_3E_1...E_n$. On the other hand, since f_3 is a k_1-contracting map, every other term $|c_iH_i(f_3)|$ is less than $C|c_i|Z_iE_1...E_n$, which is less than $(1/100)c|c_3|Z_3E_1...E_n$. Therefore, the other terms are not sufficiently large to give any cancellation in the sum, and $H(f_3)$ is sufficiently large to satisfy our proposition.

Next we have to deal with the general case. It may happen that several of the numbers $|c_i|Z_i$ are roughly the same size as one another. This does happen for special ellipsoids, but it is rare in the space of all ellipsoids, because the monomials Z_i are all different. Therefore, we can always pick an ellipsoid E', with axes $E_i' > (1/C)E_i$, for which the numbers $|c_i|Z_i$ are widely different from one another. By the previous case, there is a map 1-contracting map f from E' to X with $H(f)$ as large as we need. But clearly there is a 1-contracting degree 1 map from E to E'. Composing these maps proves our estimate. □

A difficult invariant of order 4

The methods in this section can be used to prove analogous estimates for many invariants of all orders. When I first found them, I had hoped that they would allow a sharp estimate for all rational homotopy invariants of k_1-contracting maps provided that E_1 is sufficiently large. This hope has not been realized.

For example, consider the following fourth order rational homotopy invariant of maps from S^n to a wedge of 5 spheres $X = S^{k_1} \vee ... \vee S^{k_5}$.

$$H(f) = \int P(Pa_5 \wedge a_1) \wedge P(Pa_4 \wedge a_3) \wedge a_2 \pm P(Pa_3 \wedge a_2) \wedge Pa_4 \wedge Pa_5 \wedge a_1.$$

Suppose that $k_1 \leq k_2$ and that $k_2 = k_3 = k_4 = k_5$. For example, we could have $X = S^2 \vee S^3 \vee S^3 \vee S^3 \vee S^3$, and then H would give an invariant of maps from S^{10} to X.

The invariant H(f) is a sum of two terms. Using techniques analogous to those above, we can check that each term is bounded by a constant C times the following expression.

$E_{n-k_4+1}E_{n-k_4-k_5+2}E_{n-k_4-k_5-k_1+3}E_{n-k_4-k_5-k_1-k_2+4}E_1...E_n$.

The exact value of this expression is not so important. Let us refer to this expression as $P(E)$. When $E_1 > C$, an argument using Theorem 7.2 shows that this upper bound is sharp for each term. That is, Theorem 7.2 allows us to construct a 1-contracting map f_0 so that each of the two terms in the integral defining H has norm at least $cP(E)$. It turns out that if we take the correct sign in the definition of H, then $H(f_0)$ vanishes because the two terms cancel.

The following question then arises. Can we find a different k_1-contracting map for which $H(f)$ is at least $cP(E)$? In other words, is there a cancellation between the two terms of H whenever one of them is large, or did the cancellation occur only because the map f_0 was poorly chosen?

In some cases, it is possible to rewrite a sum of two monomials as a different monomial or sum of monomials. When that happens, one can prove in some cases that the sum satisfies a more stringent upper bound than either term individually. For example, consider the rational homotopy invariant for maps from an ellipsoid E to a wedge of three spheres $S^{k_1} \wedge S^{k_2} \wedge S^{k_3}$ (with $k_1 < k_2 < k_3$) given by $\int_E Pa_1 \wedge Pa_2 \wedge a_3 \pm \int_E Pa_1 \wedge a_2 \wedge Pa_3$. If $E_1 > c$, then each term in this sum can realize a value at least $cE_{n-k_1+1}E_{n-k_1-k_2+2}E_1...E_n$ for some 1-contracting map f. With an appropriate choice of sign, however, this sum is equal to $\int_E a_1 \wedge Pa_2 \wedge Pa_3$, which is bounded by $CE_{n-k_2+1}E_{n-k_1-k_2+2}E_1...E_n$ for all k_1-contracting maps f.

This algebraic strategy does not work for the fourth-order rational homotopy invariant defined above. I think that new geometric techniques will be needed to settle this point.

It may be, however, that the geometrical techniques in this chapter do suffice to give sharp estimates for a rational homotopy invariant H given by a monomial, provided that E_1 is sufficiently large. This question boils down to a combinatorial problem that I don't know how to solve.

7.7 Application to the k-dilation of diffeomorphisms

In the previous section, we proved a large number of estimates of the following form.

(STAR) For some rational homotopy invariant H of maps from S^n to some wedge of spheres X, of order d, we proved two bounds. The first bound says that if f is any k-contracting map from an ellipsoid E to X, then H(f) is less than $CP(E)$. The second bound says that if E_1 is at least C, then there is a 1-contracting map f from E to X with H(f) greater than $cP(E)$.

In the above equation, $P(E)$ and k depend on H. The number k was just the smallest degree of any form used to define H. The function $P(E)$ is a function depending on the principal axes of E, $E_1, ...E_n$. In simple cases, $P(E)$ is a monomial, and in more complicated cases, $P(E)$ is a maximum of a finite list of monomials. In all cases, if we rescale E by a factor λ, then $P(E)$ rescales by a factor λ^{n+d}.

Using these estimates, we can prove inequalities for the k-dilation of maps between ellipsoids and maps between rectangles.

Proposition 7.7.1. *Suppose that (STAR) holds for some triplet (H,k,P). Let f be any k-contracting map from an ellipsoid E to an ellipsoid F. Then the following inequality holds.*

$$(\mathit{deg}f)P(F) < CP(E).$$

Proof. At first, let us assume that F_1 is at least C. By hypothesis, we can construct a 1-contracting map g from F to a wedge of spheres X with $H(g)$ greater than $cP(F)$. The composition $g \circ f$ is a k-contracting map from E to X with $H(g \circ f) = (\deg f)H(g)$, which is greater than $c(\deg f)P(F)$. On the other hand, by hypothesis, any k-contracting map from E to X has H less than $CP(E)$. Therefore, $c(\deg f)P(F) < CP(E)$. Rearranging the constants gives the equation we want to prove.

Now we remove the assumption that F_1 is sufficiently large. Let F' be a rescaling of F by some very large factor λ, and let E' be the rescaling of E by the same factor. Then the rescaling of f is a k-contracting map f' from E' to F'. We can assume that F_1' is sufficiently large. By the preceding paragraph, we conclude that $(\deg f)P(F') < CP(E')$. But $P(F') = \lambda^{n+d}P(F)$, and $P(E') = \lambda^{n+d}P(E)$. Substituting into our last inequality, we get $(\deg f)P(F)\lambda^{n+d} < CP(E)\lambda^{n+d}$. Cancelling the factor of λ^{n+d} leaves the equation we want to prove. □

If we have a k-contracting diffeomorphism from a rectangle R to a rectangle S, then we can get a map between ellipsoids in two simple ways. First, we can restrict the diffeomorphism to the boundary of R. The restriction is a k-contracting degree 1 map from the boundary of R to the boundary of S. The boundary of any n-dimensional rectangle R is bilipshitz to the (n-1)-dimensional ellipsoid with principal axes $R_1 \leq ... \leq R_n$. Second, we can take the double of the diffeomorphism, which is a k-contracting degree 1 map from the double of R to the double of S. The double of a rectangle R is bilipshitz to the n-dimensional ellipsoid with principal axes $\epsilon \leq R_1 \leq$

... $\leq R_n$, for any number ϵ less than R_1. Therefore, using the estimates for rational homotopy invariants in the last section and the proposition above, we can prove inequalities for the k-dilation of diffeomorphisms between rectangles.

As far as I can see, the complete list of inequalities that can be proved in this way is a combinatorial mess. I don't know how to write the list in closed form. Therefore, let me instead record the inequalities that we can prove for small values of k and n. We assume throughout that there is a k-contracting diffeomorphism from R to S.

The three-dimensional case

k=2

We get one inequality by looking at the double.

$R_1 R_2^2 R_3 > c S_1 S_2^2 S_3$.

The four-dimensional case

k=2

By looking at the boundary we get one inequality.

$R_2 R_3^2 R_4 > c S_2 S_3^2 S_4$.

By looking at the double, we get two more inequalities.

$R_1 R_2 R_3^2 R_4 > c S_1 S_2 S_3^2 S_4$.

$R_1 R_2^2 R_3^2 R_4 > c S_1 S_2^2 S_3^2 S_4$.

k=3

We get no inequalities.

The five-dimensional case

k=2

By looking at the boundary, we get two inequalities.

$R_2 R_3 R_4^2 R_5 > c S_2 S_3 S_4^2 S_5$.

$R_2 R_3^2 R_4^2 R_5 > c S_2 S_3^2 S_4^2 S_5$.

By looking at the double we get four more inequalities.

$R_1 R_2 R_3^2 R_4 R_5 > c S_1 S_2 S_3^2 S_4 S_5$.

$R_1 R_2 R_3 R_4^2 R_5 > c S_1 S_2 S_3 S_4^2 S_5$.

$R_1 R_2 R_3^2 R_4^2 R_5 > c S_1 S_2 S_3^2 S_4^2 S_5$.

$R_1 R_2^2 R_3^2 R_4^2 R_5 > c S_1 S_2^2 S_3^2 S_4^2 S_5$.

k=3

By looking at the double we get one inequality.

$R_1 R_2 R_3^2 R_4 R_5 > c S_1 S_2 S_3^2 S_4 S_5$.

k=4

We get no inequalities.

Some of the inequalities listed above follow from the inequalities in Chapter 3, especially when n is small. For large n, however, I have no other proof for most of the inequalities in the list. When k is greater than $(n+1)/2$, then this method does not yield any inequalities. On the other hand, when n is large and k is small, this method yields a very long list of inequalities. As an example, we record the inequalities that we get by considering the boundary of R when n is 8 and k is 2. The resulting list has twelve different inequalities. As far as I can see, none of these inequalities can be deduced from those we have proven in earlier chapters together with the other eleven. They show that the 2-dilation of any diffeomorphism from R to S is at least a small constant times any of the following twelve expressions in $Q_i = S_i/R_i$.

$(Q_2 Q_3 Q_4 Q_5 Q_6 Q_7^2 Q_8)^{2/8}$

$(Q_2 Q_3 Q_4 Q_5 Q_6^2 Q_7 Q_8)^{2/8}$

$(Q_2 Q_3 Q_4 Q_5^2 Q_6 Q_7 Q_8)^{2/8}$

$(Q_2 Q_3 Q_4 Q_5 Q_6^2 Q_7^2 Q_8)^{2/9}$

$(Q_2 Q_3 Q_4 Q_5^2 Q_6 Q_7^2 Q_8)^{2/9}$

$(Q_2 Q_3 Q_4 Q_5^2 Q_6^2 Q_7 Q_8)^{2/9}$

$(Q_2 Q_3 Q_4 Q_5^2 Q_6^2 Q_7^2 Q_8)^{2/10}$

$(Q_2 Q_3 Q_4^2 Q_5 Q_6^2 Q_7^2 Q_8)^{2/10}$

$(Q_2 Q_3 Q_4^2 Q_5^2 Q_6 Q_7^2 Q_8)^{2/10}$

$(Q_2 Q_3 Q_4^2 Q_5^2 Q_6^2 Q_7 Q_8)^{2/10}$

$(Q_2 Q_3 Q_4^2 Q_5^2 Q_6^2 Q_7^2 Q_8)^{2/11}$

$(Q_2 Q_3^2 Q_4^2 Q_5^2 Q_6^2 Q_7^2 Q_8)^{2/12}$

7.8 The Hopf invariant of k-contracting maps between ellipsoids

In this chapter so far, we have given estimates for the Hopf invariant of a k-contracting map from an ellipsoid E to the unit sphere, as well as for more complicate rational homotopy invariants. In this section, we will give some estimates for the Hopf invariant of a k-contracting map between ellipsoids. For k-contracting maps from E^3 to F^2, we are able to give a good estimate for either $k = 1$ or $k = 2$. (In an appropriate sense, this estimate is sharp up to a constant factor.) For maps from E^7 to F^4, the problem looks quite difficult, but we will give several upper bounds for the Hopf invariant of k-contracting maps.

Maps from S^3 to S^2

For 2-contracting maps from a 3-dimensional ellipsoid E to a 2-dimensional ellipsoid F, the methods above give the following estimate.

$$(\text{Hopf} f) F_1^2 F_2^2 < C E_1 E_2^2 E_3.$$

Proof. First, notice that the problem is invariant if we rescale both E and F by the same factor. Therefore, we can assume that F_1 is much bigger than 1, which allows us to find a 1-contracting map g from F to the unit 2-sphere with degree roughly $F_1 F_2$. The Hopf invariant of $g \circ f$ is equal to $(\text{Hopf} f) F_1^2 F_2^2$. On the other hand, according to the results in section 6, any 2-contracting map from E to the unit 2-sphere has Hopf invariant less than $C E_1 E_2^2 E_3$. □

This inequality is not sharp for all pairs of ellipsoids. If $E_1 E_2$ is much smaller than $F_1 F_2$, then any 2-contracting map from E to F has Hopf invariant zero. This result follows from the following lemma, which was worked out jointly with Daniel Biss.

Lemma 7.2. *Let f be a map from S^3 to S^2. We can view f as a family of pointed maps from S^2 to S^2, parameterized by the unit circle, with the base point of the unit* *circle corresponding to the constant map. If every map in the family is not surjective, then f is null-homotopic.*

Proof. For some large number n, we can choose points y_i in S^2 so that y_i is not in the range of f_θ for $2\pi(i-2)/n < \theta < 2\pi(i+2)n$. (The index i varies from 1 to n, and the values of θ are taken modulo 2π.) At each value $2\pi i/n$, we cut open the circle and insert an additional segment. Along the additional segment inserted at $2\pi i/n$, the family of mappings contracts the to the basepoint while avoiding y_i, and then expands to its original value again avoiding y_i. The resulting family defines a new map which is homotopic to f. After reparamaterizing the circle, we can assume that $f_{2\pi i/n}$ is the constant map for each i, and that the region $2\pi i/n < \theta < 2\pi(i+1)/n$ in the original parametrization is mapped into the same region in the new parametrization. Since $f_{2\pi i/n}$ is the constant map, we can define f_i to be map from S^3 to S^2 corresponding to the family f_θ as θ varies from $2\pi i/n$ to $2\pi(i+1)/n$. The map f is homotopic to the sum of the homotopy classes of f_i. But each f_i is homotopic to the suspension of a map in $\pi_2(S^1)$ and hence null-homotopic. □

Now suppose that f is a 2-contracting map from E to F and that $E_1E_2 < cF_1F_2$. The ellipsoid E is swept out by 2-spheres S_θ with area less than CE_1E_2. We consider a 2-contracting map f from E to F as a family of maps f_θ parameterized by S^1, each map going from a set $S_\theta \subset E$ to F. Since each S_θ has area less than CE_1E_2, and $F_1F_2 > CE_1E_2$, each map f_θ is not surjective. Applying the above lemma, it follows that f is null-homotopic.

Now we have enough tools to estimate the Hopf invariant of 2-contracting maps.

Proposition 7.8.1. *Let f be a 2-contracting map from E to F. If $F_1F_2 > CE_1E_2$, then Hopf(f) is zero. In any case, the Hopf invariant of f has norm less than $CE_1E_2^2E_3/(F_1^2F_2^2)$. On the other hand, if $F_1F_2 < cE_1E_2$, then we will construct a 2-contracting map f with Hopf invariant greater than $cE_1E_2^2E_3/(F_1^2F_2^2)$.*

Proof. The two upper bounds were proven above. It remains only to do the construction. The ellipsoid E contains two thick linked tubes, one bilipshitz to $[0, E_1] \times [0, E_2] \times$

$S^1(E_3)$, and the other bilipshitz to $[0, E_2] \times [0, E_3] \times S^1(E_1)$. Since $E_1E_2 > CF_1F_2$, we can construct an area-contracting map from $[0, E_1] \times [0, E_2]$ to F with degree on the order of $(E_1E_2)/(F_1F_2)$. Similarly, we can construct an area-contracting map from $[0, E_2] \times [0, E_3]$ to F with degree on the order of $(E_2E_3)/(F_1F_2)$. Taking the Pontryagin-Thom collapses of either map or of both maps, we get three different area-contracting maps from E to F. One of these maps has Hopf invariant at least $c(E_1E_2^2E_3)/(F_1^2F_2^2)$. □

Proposition 7.8.2. *Let f be a 1-contracting map from E to F. If $F_1 > CE_1$ or $F_2 > CE_2$, then the Hopf invariant of f is zero. In any case, the Hopf invariant of f has norm less than $CE_1E_2^2E_3/(F_1^2F_2^2)$. On the other hand, if $F_1 < cE_1$ and $F_2 < cE_2$, then we will construct a 1-contracting map f with Hopf invariant greater than $cE_1E_2^2E_3/(F_1^2F_2^2)$.*

Proof. The ellipsoid E admits a map to the disk with fibers of diameter less than $4E_1$. Every ball in F of radius less than $(1/2)F_1$ is convex. According to a theorem of Gromov from [14], if $4E_1$ is smaller than $(1/2)F_1$, then any contracting map from E to F factors through the map from E to the disk and is null-homotopic.

The ellipsoid E is swept out by 2-spheres S_θ with diameter less than CE_2. If $F_2 > CE_2$, then each map f_θ from S_θ to F is not surjective, because the diameter of F is on the order of F_2. By the above lemma, f is null-homotopic.

For the final estimate, we embed in E, an ellipsoid minus its tip with principal axes E_1, E_2, cross S^1, linked with a rectangle $E_2 \times E_3$ crossed with an S^1. The two S^1's are linked once. Because $E_1 > CF_1$ and $E_2 > CF_2$, we get a map from the tipless ellipsoid to F of degree E_1E_2/F_1F_2. Similarly, since $E_2 > CF_2$, we easily get a map from the rectangle $E_2 \times E_3$ to F with degree E_2E_3/F_1F_2. Using these maps as in the proof of the last proposition, we get a 1-contracting map from E to F with Hopf invariant at least $c(E_1E_2^2E_3)/(F_1^2F_2^2)$. □

Maps from S^7 to S^4

The last two propositions give a pretty good accounting of the Hopf invariant for 1-contracting or 2-contracting maps from S^3 to S^2. My knowledge of the situation for maps from S^7 to S^4 is much less complete, but we can still prove some estimates.

Proposition 7.8.3. *If f is a 4-contracting map from E^7 to F^4, then*

$$(Hopf f)F_1^2F_2^2F_3^2F_4^2 < CE_1E_2E_3E_4^2E_5E_6E_7.$$

If f is a 2-contracting map from E^7 to F^4, then we can prove two additional, rather strange-looking, inequalities.

$$(Hopf f)F_1^2F_2^2F_3^4F_4^2 < CE_1E_2E_3^2E_4^2E_5^2E_6E_7.$$

$$(Hopf f)F_1^2F_2^4F_3^4F_4^2 < CE_1E_2^2E_3^2E_4^2E_5^2E_6^2E_7.$$

Proof. As usual, we can assume that F_1 is large because the problem is scale invariant. There is a 1-contracting map g from F to the unit sphere with degree on the order of $F_1F_2F_3F_4$. The Hopf invariant of $g \circ f$ is at least $c\text{Hopf(f)}F_1^2F_2^2F_3^2F_4^2$. But since $g \circ f$ is a 4-contracting map, according to Proposition 7.6.1, it has Hopf invariant at most $CE_1E_2E_3E_4^2E_5E_6E_7$. This proves the first inequality.

By Theorem 7.2, we can construct a 1-contracting map g from F to $X = S^2 \vee S^3$ with $\int_F Pa_3 \wedge a_2$ greater than $cF_1F_2F_3^2F_4$. Now the map $g \circ f$ from E to X is a 2-contracting map with $\int_E P(Pa_3 \wedge a_2) \wedge Pa_3 \wedge a_2$ greater than $c(\text{Hopf}f)(F_1F_2F_3^2F_4)^2$. Applying Proposition 7.6.3 to estimate this integral over E, we get the second inequality above.

Similarly, we can construct a 1-contracting map g from F to $X = S^2 \vee S^2 \vee S^2$ with $\int_F Pa_{2,1} \wedge Pa_{2,2} \wedge a_{2,3}$ greater than $cF_1F_2^2F_3^2F_4$. The composition $g \circ f$ is a 2-contracting map from E to X with $\int_E P(Pa_{2,1} \wedge Pa_{2,2} \wedge a_{2,3}) \wedge (Pa_{2,1} \wedge Pa_{2,2} \wedge a_{2,3})$ greater than $c(\text{Hopf}f)(F_1F_2^2F_3^2F_4)^2$. Using Proposition 7.4.1 to estimate the directional norms of the primitives, we get the third inequality above. □

I suspect that these inequalities are far from all the upper bounds that exist on the Hopf invariant of k-contracting maps from E^7 to F^4. It would be interesting to know whether there are 4-contracting maps from E to F with non-vanishing Hopf invariant even when $E_1E_2E_3E_4$ is much smaller than $F_1F_2F_3F_4$.

7.9 Homotopically non-trivial maps with small k-dilation

So far in this chapter, we have given bounds for the rational homotopy invariants of maps with small k-dilation. It is natural to ask whether we can also bound the torsion homotopy invariants of maps with small k-dilation. Gromov proved that maps with sufficiently small 2-dilation are null-homotopic.

Theorem. *(Gromov) Let M and N be compact simply connected Riemannian manifolds. There is a positive constant ϵ, depending on M and N, so that every map from M to N with 2-dilation less than ϵ is null-homotopic.*

A sketch of the proof appears on pages 229-230 of [10]. (For more information, see appendix B.)

We will show in this section that small 3-dilation has a much weaker effect on homotopy type. For example, we will construct maps from S^4 to S^3, homotopic to the suspension of the Hopf fibration, with arbitrarily small 3-dilation. More generally, we will construct maps homotopic to high order suspensions with small k-dilation.

Proposition 7.9.1. *Fix a homotopy class a in $\pi_m(S^n)$ and then consider its p-fold suspension $\Sigma^p a$ in $\pi_{m+p}(S^{n+p})$. If k is any integer greater than $n + (n/m)p$, then there are maps from S^{m+p} to S^{n+p} in the homotopy class $\Sigma^p a$, with arbitrarily small k-dilation.*

Proof. Inside of the unit sphere, we can quasi-isometrically embed a rectangle R with dimensions $[0,\epsilon]^m \times [0,\epsilon^{-m/p}]^p$. Let f be a map in the homotopy class a from $[0,1]^m$ to the unit n-sphere (taking the boundary of the n-cube to the base point of S^n). We can assume that the map f is Lipshitz with some Lipshitz constant L. Now we construct a map F from R to $S^n \times S^p$. The map F is simply a direct product of a map F_1 from $[0,\epsilon]^m$ to S^n and a map F_2 from $[0,\epsilon^{-m/p}]^p$ to S^p. The map F_1 is just a dilation from $[0,\epsilon]^m$ to the unit cube, composed with the map f. The map F_2 is a rescaling from $[0,\epsilon^{-m/p}]^p$ to $[0,1]^p$, followed by a standard degree 1 map from the p-cube to the p-sphere, taking the boundary of the cube to the basepoint. The k-dilation of F is $(L\epsilon^{-1})^n(\epsilon^{m/p})^{k-n}$. Expanding this expression gives $L^n\epsilon^{-n+(m/p)k-(m/p)n}$. The important part of the expression is the power of ϵ, which is equal to $(m/p)(k-n-(n/m)p)$. We have assumed that k is greater than $n+(n/m)p$, and so the exponent of ϵ is positive. For ϵ sufficiently small, the k-dilation of F is arbitrarily small. The map F takes the boundary of R to $S^n \vee S^p$. We compose F with a smash map, which is a degree 1 map from $S^n \times S^p$ to S^{n+p}, taking $S^n \vee S^p$ to the base point. The result is a k-contracting map from R to S^{n+p} which takes the boundary of R to the basepoint. We can easily extend this map to all of S^{m+p} by mapping the complement of R to the basepoint of S^{n+p}. The resulting map is homotopic to $\Sigma^p(a)$. □

For example, this proposition shows that the suspension of the Hopf map in $\pi_4(S^3)$ can be realized by maps with arbitrarily small 3-dilation. In this case, we are considering a 1-fold suspension of a map from S^3 to S^2. Therefore $p = 1$, $m = 3$, and $n = 2$. Since $3 > 2 + (2/3)1$, the result follows.

Using some fairly deep results in algebraic topology, we will show that many non-trivial homotopy classes can be realized with arbitrarily small 3-dilation.

Theorem 7.3. *For every N, there are infinitely many M so that there are homotopically non-trivial maps from S^M to S^N with arbitrarily small 3-dilation.*

Proof. If a is a homotopy class in $\pi_m(S^2)$, then the p-fold suspension $S^p a$ can be realized with arbitrarily small 3-dilation when $3 > (2/m)p+2$, that is when $p < m/2$. In order to prove Theorem 7.3, we have to find classes a so that $\Sigma^p a$ is not zero. In fact, there are infinitely many such classes. I would like to thank Haynes Miller, who helped me to find the relevant theorems in the literature.

When $i = 8j+1$, the homotopy group $\pi_i(SO)$ is equal to $\mathbb{Z}_2$. The J homomorphism maps $\pi_i(SO)$ to the stable i-stem of the homotopy groups of spheres. When $i = 8j+1$, the map J is injective, and its image is a copy of $\mathbb{Z}_2$. This image contains a non-trivial element in $\pi_{i+n}(S^n)$, for large n. It turns out that this non-trivial element is the suspension of a class in $\pi_{i+2}(S^2)$. This statement is made clearly in the introduction to the paper [6], and the proof appears in the older paper [5]. For each $i = 8j+1$, let

a_i be a homotopy class in $\pi_{i+2}(S^2)$ whose suspension is the non-trivial element in the image $J(\pi_i(SO))$. In particular, the p-fold suspension $\Sigma^p a_i$ is non-trivial for every p and every i.

Fix a number $N \geq 2$. For each $i = 8j + 1$ greater than $2N - 6$, the class $\Sigma^{N-2} a_i$ in $\pi_{i+N}(S^N)$ can be realized by maps with arbitrarily small 3-dilation. All of these homotopy classes are non-trivial. $\square$

It would be interesting to know whether the 3-dilation controls any torsion homotopy invariants. For example, the 100-fold suspension of the Hopf map is a non-trivial element in $\pi_{103}(S^{102})$. According to the above proposition, it can be realized by maps with arbitrarily small 69-dilation. It would be interesting to know whether it can be realized by maps with arbitrarily small 3-dilation.

Chapter 8

Partial Results On 2-Contracting Diffeomorphisms Between 4-Dimensional Rectangles

In Chapter 2, we estimated the smallest (n-1)-dilation of all diffeomorphisms between n-dimensional rectangles R and S, up to a constant factor. The next simplest case is to consider the 2-dilation of diffeomorphisms between 4-dimensional rectangles. In this case, I am not able to give a complete solution. In this chapter, we will give some partial results. In the first part we prove some lower bounds for the 2-dilation. In the second part, we construct a variety of non-linear diffeomorphisms with small 2-dilation, generalizing the snake map.

These partial results show that the case of 2-contracting diffeomorphisms between 4-dimensional rectangles is more complicated in some ways than the case of (n-1)-contracting maps. We'll say more about this point at the end of the chapter. Also, the partial results allow us to solve the problem in the special case that R is a cube.

Theorem 8.1. *If there is a 2-contracting diffeomorphism from the unit 4-cube to S then $S_2 S_3^3 S_4^2 < C$. On the other hand, if $S_2 S_3^3 S_4^2 < c$, then there is a 2-contracting diffeomorphism from the unit 4-cube to S.*

8.1 Lower bounds for the 2-dilation

A. Embedding inequalities

A 2-contracting diffeomorphism from R to S is a special case of a 2-expanding embedding of S into R. We solved the k-expanding embedding problem in Chapter 6. In particular, there is a 2-expanding embedding of S into R roughly if and only if the following inequalities hold.

$A1. R_1 R_2 > S_1 S_2$.

$A2. R_1 R_2 R_3 > S_1 S_2 S_3$.

$A3. R_1^2 R_2 R_3 > S_1^2 S_2 S_3$.

$A4. R_1 R_2 R_3 R_4 > S_1 S_2 S_3 S_4$.

$A5. R_1^3 R_2 R_3 R_4 > S_1^3 S_2 S_3 S_4$.

B. Boundary inequalities

A 2-contracting diffeomorphism from R to S restricts to a 2-contracting diffeomorphism from the boundary of R to the boundary of S. These two boundaries are bilipshitz to ellipsoids. We will now solve the problem of deciding whether there is a 2-contracting diffeomorphism between two ellipsoids up to a constant factor.

Proposition 8.1.1. *Let E be the ellipsoid with principal axes $R_1 \leq R_2 \leq R_3 \leq R_4$, and let F be the ellipsoid with principal axes $S_1 \leq S_2 \leq S_3 \leq S_4$. There is a 2-contracting diffeomorphism from E to F approximately if and only if the following inequalities hold.*

B1. $R_2 R_3 > S_2 S_3$.

B2. $R_2^2 R_3 R_4 > S_2^2 S_3 S_4$.

B3a. If $R_2^2 > S_2 S_3$, *then* $R_2 R_3 R_4 > S_2^{1/2} S_3^{3/2} S_4$.

B3b. If $R_2^2 < S_2 S_3$, *then* $R_3 R_4 > S_3 S_4$.

Proof. Inequalities B1 and B2 are quite similar to results we have proven about 3-dimensional rectangles.

The 2-width of the ellipsoid E is roughly $R_2 R_3$. It is not any bigger, because the function x_4 has level sets with volume less than $C R_2 R_3$. It is not smaller, because it contains a set quasi-isometric to the rectangle $R_2 \times R_3 \times R_4$ which has 2-width greater than cR_2R_3. If there is a 2-contracting diffeomorphism from E to F, then the 2-width of E is at least as big as the 2-width of F. This proves inequality B1.

The ellipsoid F contains N disjoint subsets each quasi-isometric to a cube with side length S_2, where N is at least cS_3S_4/S_2^2. The inverse image of each of these subsets in E has 2-width at least cS_2^2. According to the width-volume inequality for subsets of ellipsoids, the volume of each of these inverse images is at least $cR_2^{-1}S_2^4$. (In Chapter 3 we proved a width-volume inequality for subsets of rectangles. In section 5.1, we used a simple trick to show that the same inequality holds for subsets of ellipsoids.) It follows that the total volume of the inverse images of all the cubes is at least $cR_2^{-1}S_2^2S_3S_4$. Since these inverse images are disjoint subsets of E, their total volume is bounded by $CR_2R_3R_4$. This proves inequality B2.

The ellipsoid F also contains a subset quasi-isometric to the product of a rectangle with dimensions $S_3 \times S_4$ with a circle of length S_2. We are going to estimate the volume of the inverse image of this set. We view the set as a family of circles. The transverse area to this family of circles increases under inverse image, since our diffeomorphism is 2-contracting.

We will use the isoperimetric inequality to estimate the length of the inverse image of each circle. We restrict our attention to the circles which go through the central sub-rectangle of $S_3 \times S_4$ with dimensions $1/3S_3 \times 1/3S_4$. Each of these circles has filling area greater than cS_2S_3. Therefore, the inverse image of each circle in E has filling area greater than cS_2S_3.

We will estimate the length of the inverse image using the isoperimetric inequality in E. If c is a curve in E of length $L < cR_2$, then it sits inside of a ball that is bilipshitz to Euclidean. Therefore it has filling area less than CL^2. If $R_2^2 > S_2S_3$, then the inverse image of each circle has filling area greater than $c(S_2S_3)$. Therefore the inverse image of each circle must have length at least $c(S_2S_3)^{1/2}$. Since the transverse area measure increases, the inverse image of our solid torus has volume at least $cS_2^{1/2}S_3^{3/2}S_4$. Since this inverse image is a subset of E, its volume is bounded by $CR_2R_3R_4$. This proves inequality B3a.

On the other hand, if $R_2^2 < S_2S_3$, then the inverse image of each circle must have length at least cR_2. As we saw above, if the length is less than cR_2, then the filling area is less than $CL^2 < R_2^2 < S_2S_3$. Therefore, the inverse image of the solid torus must have volume at least $cR_2S_3S_4$. Since this inverse image is a subset of U, it follows that $R_3R_4 > cS_3S_4$. This proves inequality B3b.

This concludes the proof of the necessity of inequalities B1, B2, and B3. Next, if these inequalities hold, we construct a diffeomorphism from E to F with 2-dilation bounded by C.

Case a. $R_2^2 > S_2S_3$.

Let G be the ellipsoid with principal axes $((S_2S_3)^{1/2}, (S_2S_3)^{1/2}, S_2^{-1/2}S_3^{1/2}S_4)$. There is a diffeomorphism from G to F with 2-dilation less than C. Removing a line from G and expanding the metric we get the rectangle with dimensions $S_2^{1/2}S_3^{1/2} \times S_2^{1/2}S_3^{1/2} \times S_2^{-1/2}S_3^{1/2}S_4$. The volume of this rectangle is $S_2^{1/2}S_3^{3/2}S_4$. By inequality 3A., this is less than the volume of E. Also, by the assumption $R_2^2 > S_2S_3$, the rectangle is thin compared to E, and it admits an expanding embedding into E. The inverse map of this embedding is a 2-contracting diffeomorphism from a subset of E to the complement of a line in G. We extend this diffeomorphism by mapping the rest of E to (a tiny neighborhood of) this line. Since the line is 1-dimensional, this map can be made 2-contracting.

Case b. $R_2^2 < S_2S_3$.

In this case, we have the following three inequalities: $R_2R_3 > S_2S_3$, $R_2^2R_3R_4 > S_2^2S_3S_4$, and $R_3R_4 > S_3S_4$. According to Theorem 2.1, there is a diffeomorphism from the rectangle $R_2 \times R_3 \times R_4$ to the rectangles $S_2 \times S_3 \times S_4$ with 2-dilation less than C. The double of this diffeomorphism is a diffeomorphism from E to F, with 2-dilation less than C. □

Finally, we notice that inequality B3 gives the upper bound that we need for Theorem 8.1. Suppose that there is a 2-contracting diffeomorphism from the unit cube to S. We know from inequality B1 that $1 > cS_2S_3$. Therefore, when we apply inequality B3, we are in the first case. The inequality tells us that $1 > cS_2^{1/2}S_3^{3/2}S_4$. Therefore, $S_2S_3^3S_4^2 < C$.

C. Other inequalities

In Chapter 7, using rational homotopy invariants, we proved two more inequalities about 2-contracting diffeomorphisms. If there is a 2-contracting diffeomorphism from R to S, then we proved that $R_1R_2R_3^2R_4 > cS_1S_2S_3^2S_4$, and $R_1R_2^2R_3^2R_4 > cS_1S_2^2S_3^2S_4$. The second of these inequalities follows from the inequalities in parts A. and B., but the inequality $R_1R_2R_3^2R_4 > cS_1S_2S_3^2S_4$ does not.

The proof of this last inequality is based on the linear isoperimetric inequality. We might expect to get sharper information by using the exact isoperimetric inequality, and not only its linearization. Inequality B3 above follows from this line of thinking. Ideally, I would like to know the isoperimetric profile function $I_k(A)$ which is defined to be the largest filling (k+1)-volume of any relative k-cycle in the rectangle R. (There are a number of variations of this function: one can study real cycles or integral cycles, oriented cycles or non-oriented cycles, relative cycles or absolute cycles, and so on.) While I am not able to compute this function, even up to a constant factor, we will prove an isoperimetric estimate that gives some more information than the linear isoperimetric inequality from Chapter 7.

Proposition 8.1.2. *Let C be an oriented relative 2-cycle in the rectangle R with area less than $(1/2)R_1R_2$. Then C has an oriented filling with volume less than $CR_1R_2^2$.*

Proof. We are going to construct this filling C by the same method that we used in Chapter 7. First we approximate C by a rectilinear cycle C'. (Recall that a rectilinear cycle is made up of polyhedra each of which is a rectangle parallel to the coordinate axes.) If C' is sufficiently close, we can construct a 3-chain with boundary $C - C'$ and with arbitrarily small volume. Therefore, it suffices to find a filling of C' with volume less than $CR_1R_2^2$. If the approximation is sufficiently close, we can assume that the projection of C' to the (x_1, x_2)-plane has area at most $(2/3)R_1R_2$.

Now we consider the filling $F_p(C')$ defined in section 7.4. For a random point p, the average volume of this filling is less than C area$(C')R_3$, which is less than $CR_1R_2R_3$. Now let S be the subset of $[0, R_1] \times [0, R_2]$ disjoint from the projection of C'. Because the area of C is less than $(1/2)R_1R_2$, we can assume that the area of S is at least $(1/3)R_1R_2$. For a random point p in $S \times [0, R_3]$, the average volume of $F_p(C')$ is less than C area$(C')R_2$. In particular, we can choose p so that the volume of $F_p(C')$ is less than $CR_1R_2^2$. (The reason for this better estimate is as follows. The definition of $F_p(A)$ for a 2-face A in C' depends on whether (p_1, p_2) lies in the projection of A onto $[0, R_1] \times [0, R_2]$. If it does, the volume of $F_p(A)$ can be very large, and if we know it doesn't, we get a better estimate for the volume of $F_p(A)$.) □

Using this isoperimetric inequality, we can prove another estimate about the 2-dilation of diffeomorphisms.

Proposition 8.1.3. *Suppose there is a 2-contracting diffeomorphism from R to S. If $CR_1R_2^2 < S_1S_2S_3$, then $R_3R_4 > cS_3S_4$.*

Proof. For each pair (y_3, y_4) in $[0, S_3] \times [0, S_4]$, we consider the plane $[0, S_1] \times [0, S_2] \times \{y_3\} \times \{y_4\}$. This plane is a relative 2-cycle in S. If we choose (y_3, y_4) in the central sub-rectangle $[(1/3)S_3, (2/3)S_3] \times [(1/3)S_4, (2/3)S_4]$, then the filling volume of this cycle is at least $cS_1S_2S_3$.

We consider the inverse image of one of these 2-cycles in R. It must also have filling volume at least $cS_1S_2S_3$, which is greater than $CR_1R_2^2$. According to the last proposition, the 2-cycle itself must therefore have area at least cR_1R_2.

Let U be the inverse image of $[0, S_1] \times [0, S_2] \times [(1/3)S_3, (2/3)S_3] \times [(1/3)S_4, (2/3)S_4]$. This set is fibered by the inverse image of the 2-cycles above, each of which has area at least cR_1R_2. The area transverse to these fibers in U is larger than in the corresponding region of S, because our diffeomorphism is 2-contracting. Therefore, the volume of U is at least $c(R_1R_2)(S_3S_4)$. Since U is a subset of R, its volume is less than $R_1R_2R_3R_4$. Therefore, $R_3R_4 > cS_3S_4$. □

8.2 Variations of the snake map

In this section, we will describe five non-linear maps between 4-dimensional rectangles, generalizing the snake map in various ways. Each map takes the boundary of the domain to the boundary of the range, has degree 1, and has 2-dilation less than C. By slightly perturbing them, it is possible to construct diffeomorphisms with 2-dilation less than C.

1. Snake Map

This degree 1 map takes

$R_1 \times R_2 \times R_3 \times R_4$

onto

$R_1 \times R_2 \times \lambda^{-1} R_3 \times \lambda R_4$.

The constant λ is allowed to take any value between 1 and R_3/R_2.

Construction:

Let I be a quasi-isometric embedding of $\lambda^{-1} R_3 \times \lambda R_4$ into $R_3 \times R_4$. Let the set $U \subset R$ be the union of $\{0\} \times \{0\} \times R_3 \times R_4$ with $R_1 \times R_2 \times$ image(I).

Step 1. There is a Lipshitz retraction of R onto U. (This uses the fact that $R_2 < \lambda^{-1} R_3$.)

Step 2. There is a 2-contracting retraction of U onto $R_1 \times R_2 \times$ image(I).

Step 3. There is a quasi-isometric diffeomorphism from $R_1 \times R_2 \times$ image(I) onto $R_1 \times R_2 \times \lambda^{-1} R_3 \times \lambda R_4$.

2. Codimension 1 Snake Map

This degree 1 map takes

$R_1 \times R_2 \times R_3 \times R_4$

onto

$R_1 \times \lambda^{-1} R_2 \times \lambda R_3 \times \lambda^{-1} R_4$.

The constant λ is allowed to take any value which is bigger than 1, smaller than $(R_4/R_3)^{1/2}$, and smaller than (R_2/R_1).

Construction:

In section 2.1, we constructed a snake map S, which gives a 2-contracting diffeomorphism from $R_1 \times R_2 \times R_3$ to $R_1 \times \lambda^{-1} R_2 \times \lambda R_3$. (The constant λ must be at least 1 and at most R_2/R_1, which we have already assumed.) The Lipshitz constant of S is at most λ. Therefore, the map $F(x_1, x_2, x_3, x_4) = (S(x_1, x_2, x_3), \lambda^{-1} x_4)$ is also 2-contracting. The map F is a diffeomorphism from $R_1 \times R_2 \times R_3 \times R_4$ to $R_1 \times \lambda^{-1} R_2 \times \lambda R_3 \times \lambda^{-1} R_4$.

3. A map that stretches the short side of the domain

This degree 1 map takes

$R_1 \times R_2 \times R_3 \times R_4$

onto

$\lambda R_1 \times \lambda^{-3} R_2 \times \lambda R_3 \times \lambda^{-1} R_4$.

The constant λ is allowed to take any value which is at least 1, smaller than $(R_2/R_1)^{1/4}$, and smaller than $(R_4/R_3)^{1/2}$.

Construction:

Let I be a quasi-isometric embedding of $\lambda^{-2} R_2 \times \lambda^2 R_3 \times R_4$ into $R_2 \times R_3 \times R_4$. Let U be the union of $R_1 \times$ image(I) with $\{0\} \times R_2 \times R_3 \times R_4$.

Step 1. There is a Lipshitz retraction of R onto U.

Step 2. There is a retraction of U onto $R_1 \times$ image(I), which has 2-dilation bounded by λ^2. Moreover, at any point of U, either this map is 2-contracting, or else the derivative of the retraction maps into the 3-plane spanned by $\partial/\partial x_2, \partial/\partial x_3$, and $\partial/\partial x_4$.

Step 3. There is a quasi-isometric mapping from $R_1 \times$ image(I) to $R_1 \times \lambda^{-2} R_2 \times \lambda^2 R_3 \times R_4$, which preserves the coordinate x_1.

Step 4. Compose with the linear map from $R_1 \times \lambda^{-2} R_2 \times \lambda^2 R_3 \times R_4$ to $\lambda R_1 \times \lambda^{-3} R_2 \times \lambda R_3 \times \lambda^{-1} R_4$. This linear map is 2-contracting. Moreover, its restriction to the 3-plane $x_1 = 0$ has 2-dilation λ^{-2}. Therefore, the composition of all four maps is 2-contracting.

Technical remark: If we take I to be a quasi-isometric embedding of $\lambda^{-2} R_2 \times \lambda R_3 \times \lambda R_4$ into $R_2 \times R_3 \times R_4$, then the same argument gives us a map from R to $\lambda R_1 \times \lambda^{-3} R_2 \times R_3 \times R_4$. By an argument interpolating between these cases, we get a map to $\lambda R_1 \times \lambda^{-3} R_2 \times A \times B$, where $\lambda^4 < (R_2/R_1)$, $R_3 < A < (R_3 R_4)^{1/2}$ and $R_3 R_4 = AB$.

4. Pinching Map

This degree 1 map takes

$R_1 \times R_2 \times R_3 \times R_4$

onto

$A \times A \times A \times B$,

for any numbers A and B satisfying the following conditions.

1. $R_1 > A$.

2. $R_2 R_3 R_4 > A^2 B$.

Construction:

Let I be a quasi-isometric embedding of $A \times A \times B$ into $R_2 \times R_3 \times R_4$. The conditions above guarantee that such an embedding exists.

Let U be the union of $A \times \text{image(I)}$ with $\{0\} \times R_2 \times R_3 \times R_4$.

Step 1. There is a Lipshitz retraction of R onto U.

Step 2. There is a retraction of U onto $A \times \text{image(I)}$. This retraction has no bound on its 2-dilation, but it does have the following good property. For each point in U, either the retraction is 2-contracting or else that point is mapped into the plane $x_1 = 0$ and its tangent space is mapped into the tangent space of that plane.

Step 3. There is a quasi-isometric map from $A \times \text{image(I)}$ to $A \times A \times A \times B$, which preserves the coordinate x_1.

Step 4. There is a degree 1 Lipshitz pinch map from this rectangle onto itself, whose restriction to the plane $x_1 = 0$ is $p(0, x_2, x_3, x_4) = (0, 0, 0, x_4)$.

This pinching maps collapses all of the region with large 2-dilation to the line $x_1 = x_2 = x_3 = 0$. Since this line is one-dimensional, the 2-dilation of the composition is less than C everywhere.

This pinching map provides the non-linear diffeomorphism that we need to prove Theorem 8.1. Let A be equal to $S_2^{1/2} S_3^{1/2}$, and let B be equal to $S_2^{-1/2} S_3^{1/2} S_4$. There is a 2-contracting linear diffeomorphism from $A \times A \times A \times B$ to S. There is a 2-contracting pinching map from the unit cube to A provided that $1 > CA^2 = CS_2 S_3$, and that $1 > CA^2 B = CS_2^{1/2} S_3^{3/2} S_4$. Since the second inequality implies the first inequality, there is a 2-contracting diffeomorphism from the unit cube to S whenever $S_2 S_3^3 S_4 < c$. This finishes the proof of Theorem 8.1.

5. Double Pinching Map

This degree 1 map takes

$R_1 \times R_2 \times R_3 \times R_4$

onto

$R_1^2/A \times A \times A \times B$,

for any numbers A and B satisfying the following conditions.

1. $R_1 < A$.

2. $A^2 < R_2 R_3$.

3. $A^3 B < R_1 R_2 R_3 R_4$.

Construction:

Let I be a quasi-isometric embedding of $R_1 \times A^2/R_1 \times AB/R_1$ into $R_2 \times R_3 \times R_4$. The conditions 1, 2, and 3 above guarantee that such an embedding exists.

Let U be the union of $R_1 \times$ image(I) with $\{0\} \times R_2 \times R_3 \times R_4$.

Step 1. There is a Lipshitz retraction of R onto U.

Step 2. There is a retraction of U onto $R_1 \times$ image(I). This retraction has no bound on its 2-dilation, but it does have the following good property. For each point in U, either the retraction is 2-contracting or else that point is mapped into the plane $x_1 = 0$ and its tangent space is mapped into the tangent space of that plane.

Step 3. There is a quasi-isometric diffeomorphism from $R_1 \times$ image(I) to $R_1 \times R_1 \times A^2/R_1 \times AB/R_1$, which preserves the coordinate x_1.

Step 4. There is a degree 1 Lipshitz pinch map from this rectangle onto itself, whose restriction to the plane $x_1 = 0$ is $p(0, x_2, x_3, x_4) = (0, 0, x_3, x_4)$.

Step 5. Compose with the linear map 2-contracting map from this rectangle to $R_1^2/A \times A \times A \times B$. This map preserves the 2-plane $x_1 = 0, x_2 = 0$.

Step 6. There is a Lipshitz pinch map p from this rectangle to itself, whose restriction to the 2-plane $x_1 = 0, x_2 = 0$ is given by $p(0, 0, x_3, x_4) = (0, 0, 0, x_4)$.

The composition of all of these maps is a degree 1 map from $R_1 \times R_2 \times R_3 \times R_4$ to $R_1^2/A \times A \times A \times B$, with 2-dilation less than C.

Final Remarks

In the first part of this chapter, we gave lower bounds for the 2-dilation of diffeomorphisms from R to S. Let $L(R, S)$ be the lower bound for the 2-dilation of any diffeomorphism from R to S which we proved in section 1. In the second section of this chapter, we constructed several non-linear diffeomorphisms with small 2-dilation.

Let $U(R,S)$ be the smallest 2-dilation of any diffeomorphism from R to S that we constructed (including the composition of the diffeomorphisms we constructed with one another or with linear maps). It turns out that $U(R,S)/L(R,S)$ can be arbitrarily large. To be frank, I think that the techniques in this thesis are not even close to estimating the best 2-dilation of diffeomorphisms between 4-dimensional rectangles.

The partial results in this chapter do show that this problem is more complicated in certain ways than the problem of estimating the best (n-1)-dilation between n-dimensional rectangles. We recall from Chapter 2 that, up to a constant factor, this best (n-1)-dilation is equal to the maximum value of a list of monomials in the variables $Q_i = S_i/R_i$.

Our upper and lower bounds show that the best 2-dilation of diffeomorphisms from R to S cannot be determined, even up to a constant factor, by the quotients $Q_i = S_i/R_i$. We need to look at the actual side lengths R_i and S_i, not just the quotients S_i/R_i in order to calculate the smallest 2-dilation of any diffeomorphism from R to S. (In contrast, we proved in Chapter 6 that if we want to know whether there is a k-expanding embedding of S into R, then up to a constant factor, we only need to look at the quotients Q_i.)

Now we describe another way that the 2-dilation problem is more complicated than the (n-1)-dilation problem. Let $\rho_i = logR_i$, let $\sigma_i = logS_i$, and let $D_{k,n}(\rho,\sigma)$ be the logarithm of the best k-dilation of any diffeomorphism from R to S. The function $D_{n-1,n}$ is roughly equal to the maximum of a list of linear functions. Therefore, it is approximately equal to a concave function. The function $D_{2,4}$ is not even approximately a concave function. We can find two points with $D_{2,4}$ less than 0, and with $D_{2,4}$ taking arbitrarily large values on the line between these points. This fact has the following geometrical interpretation. There are several different strategies for constructing a 2-contracting diffeomorphism between 4-dimensional rectangles. Using two different strategies, we can show that $D_{2,4}$ is less than zero at two different points. The fact that $D_{2,4}$ takes arbitrarily large values on the line between these points shows that there is no way to "interpolate" between these two strategies.

Chapter 9

Applications to the Topology of 3-Manifolds

In this section, we use area-contracting maps in order to bound homotopy invariants of maps from 3-manifolds.

In particular, we will give bounds for a generalization of the Hopf invariant. Let M be a closed oriented 3-manifold and Σ be a closed oriented surface. Let α be a 2-form on Σ with $\int_\Sigma \alpha = 1$. Let f be a smooth map from M to Σ. If the form $f^*(\alpha)$ is not exact, then the Hopf invariant of f is not defined. If the form $f^*(\alpha)$ is exact, then the Hopf invariant of f is defined to be $\int_M Pf^*(\alpha) \wedge f^*(\alpha)$, where $Pf^*(\alpha)$ denotes any primitive of $f^*(\alpha)$. The Hopf invariant of f does not depend on the choice of the 2-form α nor on the choice of primitive, and it is a homotopy invariant. When M is the 3-sphere and Σ is the 2-sphere, then we recover the original Hopf invariant.

Theorem 9.1. *Let M be a closed oriented 3-manifold which can be triangulated with N simplices. Let f be any map from M to a closed oriented surface of genus 2. If the Hopf invariant of f is defined at all, then it is less than C^N. On the other hand, we will construct manifolds M and maps f with Hopf invariant greater than c^N (for some number $c > 1$).*

As another application of our techniques, we will give lower bounds for the number of simplices needed to build a homologically non-trivial singular cycle in an oriented hyperbolic 3-manifold with small injectivity radius.

Theorem 9.2. *Let M be a closed oriented hyperbolic 3-manifold with injectivity radius ϵ. Let X be a pseudo-manifold that can be triangulated with N simplices. Suppose that there is a map of non-zero degree from X to M. Then ϵ is greater than C^{-N}. On the other hand, we will give examples of closed oriented hyperbolic 3-manifolds with injectivity radius ϵ which can be triangulated with N simplices, where ϵ is less than c^{-N} (for some number $c > 1$).*

As a corollary, we will give a new proof of the following result of Teruhiko Soma, proven in [19].

Theorem. *(Soma) Let M be any closed oriented 3-manifold. Then the set of all closed oriented hyperbolic 3-manifolds which admit a map of non-zero degree from M is finite.*

The main motivation for the work in this chapter was to better understand this theorem of Soma.

9.1 Homotopies to 2-contracting maps

The bridge which relates topological problems to area-contracting maps is the following lemma.

Lemma 9.1. *Let X be a 3-dimensional simplicial complex. Let f be a continuous map from X to a complete hyperbolic manifold M (of any dimension). Put the standard metric on each 3-simplex of X, so that the metrics agree on all the boundaries. Then f can be homotoped to a smooth map with 2-dilation less than C.*

Proof. This lemma is a refinement of Thurston's simplex-straightening construction, which appears in [18]. We quickly review simplex straightening. We first homotope the map f to a map f_1, so that f_1 agrees with f on the 0-skeleton of X, and so that the restriction of f_1 to the 1-skeleton of X maps each 1-simplex to a geodesic in M. The boundary of each 2-simplex in X is mapped by a lift of f_1 to the boundary of a geodesic triangle in hyperbolic space. We homotope the lift of f_1 to map the 2-simplex to the interior of this geodesic triangle. In this way, we construct a homotopy from f_1 to f_2, where the lift of f_2 maps each 2-simplex of X to a geodesic 2-simplex in hyperbolic space. The boundary of each 3-simplex in X is mapped by the lift of f_2 to the boundary of a hyperbolic 3-simplex. We then homotope the lifted map to send the 3-simplex in X diffeomorphically to the interior of the hyperbolic 3-simplex. For more details, we refer the reader to [18].

At the end of this process, we have homotoped f to a map $\bar{f}$ which sends each p-simplex of X to the pushdown of a geodesic p-simplex in the universal cover of M. We let $\bar{X}$ be the simplicial complex X, where each simplex of X is equipped with the metric of the corresponding hyperbolic simplex in the universal cover of M. The map $\bar{f}$ is a 1-contracting map from $\bar{X}$ to M. The goal of the rest of this proof is to construct a map F from X to $\bar{X}$, homotopic to the identity, and with 2-dilation less than C.

It turns out to be easy to find a diffeomorphism from any simplex of X to the corresponding simplex of $\bar{X}$ with 2-dilation less than C. Because of this fact, I assumed for a while that it would be easy to construct F. When I tried to write down the map F, though, I found it surprisingly difficult to get these diffeomorphisms to agree on the boundaries of the simplices. With some labor, we will now construct the map F.

A general strategy for constructing 2-contracting maps is the following. Remove a 1-dimensional polyhedron from the range. Then, find a 1-expanding embedding of the rest of the range into the domain. Finally, extend the inverse of this embedding to the entire domain, in such a way that it maps the rest of the domain to the 1-dimensional polyhedron.

We define P to be a 1-dimensional polyhedron in $\bar{X}$, consisting of the 1-skeleton of $\bar{X}$, together with the 1-skeleton of a very fine triangulation of each 2-simplex in $\bar{X}$. The complement of P consists of the interiors of the 3-simplices of $\bar{X}$ together with many very thin tubes connecting them. We now construct a 1-expanding embedding of this complement into X.

Step 1. Embed each hyperbolic 3-simplex of $\bar{X}$ into the corresponding standard simplex of X. This step is not difficult. A hyperbolic 3-simplex consists of 1 or 2 thick regions which fit inside of a unit ball, together with up to 5 long tubes, whose width decreases exponentially in the distance from the thick part of the simplex. To fit it inside of a ball, we simply wind up the long thin tubes into a snake shape in the plane.

Step 2. The restriction of the complement of P to each 2-simplex of $\bar{X}$ consists of a large number of small triangles which are mostly not far from equilateral. Suppose that the complement of P in a 2-simplex $\bar{\Delta}$ of $\bar{X}$ consists of triangles T_i. We can choose P so that most of the triangles are roughly equilateral with side length δ, and so that every triangle has diameter less than 10δ. Then we pick disks D_i in the corresponding 2-simplex of Δ of X with radius r greater than $c\delta$. Because the area of a hyperbolic triangle is less than π, we can make the disks D_i disjoint. There is an obvious expanding map from $\bar{\Delta} - P$ to Δ that maps each triangle T_i to D_i.

Step 3. Let Δ be a 3-simplex in X and $\bar{\Delta}$ the corresponding simplex in $\bar{X}$. From step 1, we have an embedding of $\bar{\Delta}$ into Δ. The boundary of $\bar{\Delta}$ - P consists of a union of triangles T_i. For each T_i there is a corresponding disk D_i in the boundary of Δ. We now have to build thick tubes connecting the image of each T_i to the corresponding D_i.

We pick a cube C inside of the simplex Δ, but not in the image of $\bar{\Delta}$. The complicated part of the tube construction will take place inside of C. First, we build tubes going from each D_i to the top of C. Second, we build tubes going from the image of each T_i to the bottom of C. This process is not difficult, because the total cross-sectional area of the tubes is bounded. It remains only to build thick tubes that connect the corresponding disks at the top and the bottom of the cube C. The problem is that we have made no assumption about which disk on the bottom of C should be connected to a given disk on the top of C.

The problem of embedding thick networks of tubes into Euclidean space was studied by Kolmogorov and Barzdin in [17]. They found an algorithm for connecting the disks in this situation, which we will now explain.

After rescaling by a bounded factor, we can assume that the cube C is the unit cube. On the top face $z = 1$, there are roughly $(1/100)\epsilon^{-2}$ disks D_i of radius $(1/100)\epsilon$, and on the bottom face $z = 0$, there are the same number of disks D'_i, also of radius $(1/100)\epsilon$. We can assume that the disks are centered at the points of a lattice with side length 10ϵ.

We consider the 1-skeleton of the cubical lattice with side length ϵ. The 1-skeleton consists of some straight lines parallel to either the x, the y, or the z directions. We translate each line parallel to the x-axis by a small vector $v(x)$, (less than ϵ), and similarly for the y direction. For the z-direction, we replace each line by two slightly translated copies of that line. After these translations, no two lines in our list of lines intersect. In fact, the $(1/100)\epsilon$-neighborhoods of our lines are disjoint. We consider these neighborhoods as something like a system of highways. In a neighborhood of each lattice point, the highways come close to one another, and we add a system of entrance and exit ramps connecting them, all in a $(1/10)\epsilon$ neighborhood of the lattice point.

For each line in the z direction in the 1-skeleton of the lattice, we now have two slightly thick tubes. We call one of them the starting tube, and we call the other one the finishing tube. We slightly move the disks D_i on the top face $z = 1$ so that they coincide with the tops of the starting tubes, and we slightly move the disks D'_i on the bottom face so that they coincide with the bottoms of the finishing tubes.

We have numbered the disks so that D_i is supposed to connect to D'_i. Let $(x_i, y_i, 1)$ be the coordinates of the center of D_i, and let $(x'_i, y'_i, 0)$ be the coordinates of the center of D'_i. We have to construct a path P_i from $(x_i, y_i, 1)$ to $(x'_i, y'_i, 0)$. The path P_i follows the starting highway from $(x_i, y_i, 1)$ in the downward z-direction down to some height H_i which we will choose later. The height H_i will be a multiple of ϵ, so it corresponds to a juncture in the system of highways. At the juncture, P_i changes onto the highway in the x-direction, which it takes to the juncture closest to x'_i. At that juncture, P_i changes onto the highway in the y-direction, which it takes to the juncture closest to y'_i. At this juncture, the path P_i gets onto the finishing highway in the z-direction, which it takes down to the endpoint $(x'_i, y'_i, 0)$. The thick tube connecting D_i to D'_i will be the $(1/100)\epsilon$-neighborhood of P_i. We need to check that these neighborhoods are disjoint (for a smart choice of the heights H_i).

Two paths P_i and P_j can come too close to one another only if they take the same highway between the same two exits. We say that i and j are related if either $y_i = y_j$ or $x'_i = x'_j$. If i and j are not related, then the corresponding paths will stay far apart no matter how we choose the heights H_i. If i and j are related, the paths P_i and P_j stay far apart as long as H_i is different from H_j. The number of j related to a given i is less than $2(1/10)\epsilon^{-1}$. On the other hand, the number of choices for H_i is ϵ^{-1}. Therefore, we can choose H_i so that H_i is different from H_j whenever i and j are related. The resulting thick tubes are disjoint.

Finally, we can slightly tilt the paths P_i so that they are transverse to the foliation by (x,y)-planes.

We have found an expanding embedding of $\bar{X} - P$ into X, with some image U. The inverse map is a contracting diffeomorphism from U onto $\bar{X} - P$. Next, we should check that the expanding embedding is isotopic to the identity. This follows because the thin tubes constructed in step 3 are transverse to a foliation by 2-spheres of the space between the embedded 3-simplex and the boundary of the original 3-simplex. Since our embedding is isotopic to the identity, we can extend the inverse map from U to all of X, taking the complement of U onto P. Since P is one-dimensional, this extension is 2-contracting. Finally, we can smooth the resulting map from X to M with a negligible effect on the 2-dilation. This finishes the proof of the lemma. $\square$

Thurston used his straightening lemma to bound the degrees of maps to hyperbolic manifolds. We will use our lemma in an analogous way to bound the Hopf invariant of maps from 3-manifolds to hyperbolic surfaces.

It is natural to ask whether this lemma can be extended to higher-dimensional complexes X. It appears likely to me that it extends to complexes of all dimensions. This extension would say that any map from a simplicial complex X to a hyperbolic manifold can be homotoped to a map with 2-dilation less than C, for some constant C which depends only on the dimensions of the domain and the range. Unfortunately, I don't know how to prove such an extension.

9.2 The Hopf volume

We define the Hopf volume of a closed oriented Riemannian 3-manifold (M, g) as the supremal value of the integral $\int_M P\omega \wedge \omega$ as ω varies over all exact 2-forms which are bounded by 1 pointwise. We will abbreviate the Hopf volume of M by $HV(M)$.

Although we do not need any other definition, I think it is interesting to see the dual description in terms of the flow lines of a vector field. This description is due to Arnold in [3]. Any 2-form ω on (M, g) corresponds to a vector field v according to the definition $i_v dvol = \omega$. If x_i are coordinates which are oriented and orthonormal at p, and if the 2-form ω at p is equal to $adx_1 \wedge dx_2 + bdx_2 \wedge dx_3 + cdx_3 \wedge dx_1$, then the vector field v at p is equal to $a(\partial/\partial x_3) + b(\partial/\partial x_1) + c(\partial/\partial x_2)$. The 2-form ω is closed if and only if the vector field is divergence-free. In this case, the vector field defines a possible flow for an incompressible fluid. The 2-form ω is exact if, in addition, the flow lines of the flow are homologically trivial over long time periods. In this case, it is possible to define the asymptotic linking number of the flow lines over long time periods. Arnold refers to this as the helicity of the flow. Arnold has proved that the integral $\int_M P\omega \wedge \omega$ is equal to the asymptotic linking number of the flow lines of v. Therefore, the Hopf volume measures the largest asymptotic linking number of a slowly moving, incompressible flow on (M, g).

Next, we define a combinatorial Hopf volume for closed oriented triangulated 3-manifold by putting a standard metric on each simplex in the triangulation. We will abbreviate the combinatorial Hopf volume of X by $HV_{comb}(X)$. This definition applies not only to triangulated manifolds but also to triangulated pseudo-manifolds. As a technical point, the strict definition is as follows. Embed the pseudo-manifold X in $\mathbb{R}^n$ for a large n. Let $\tilde{X}$ be a small neighborhood of X. Put a standard metric on each simplex of X. Pick a retraction of $\tilde{X}$ onto X, and put a metric on $\tilde{X}$ which is only slightly bigger than the pullback of the metric on X by this retraction. The Hopf volume of X is defined to be the supremal value of the integral $\int_X P\omega \wedge \omega$. We will prove in the next section that the Hopf volume is finite for any closed oriented Riemannian manifold or for any pseudo-manifold.

Because of our homotopy lemma 9.1, the Hopf invariant of any map from a 3-manifold X to a hyperbolic surface is bounded by the Hopf volume of X.

Proposition 9.2.1. *Let f be a map from a triangulated 3-manifold X to a surface of genus 2. If the Hopf invariant of f is defined, then it is bounded by* $CHV_{comb}(X)$.

Proof. A surface of genus 2 has a hyperbolic metric. We can homotope f to a map $\bar{f}$ with 2-dilation less than C with respect to this metric. The form ω on the surface of genus 2 can be taken to be $1/(4\pi)$ times the area form, which has norm $1/(4\pi)$. The pullback $\bar{f}^*(\omega)$ has norm less than C. But the Hopf invariant of f is equal to $\int_X P(\bar{f}^*(\omega)) \wedge \bar{f}^*(\omega)$, which is bounded by $CHV_{comb}(X)$. □

Similarly, the Hopf volume controls the degrees of maps from a 3-manifold to a hyperbolic 3-manifold.

Proposition 9.2.2. *If X is a triangulated closed oriented 3-manifold and M is a closed oriented hyperbolic 3-manifold, and f is any map from X to M, then the degree of f is bounded as follows.*

$$HV_{comb}(X) > c\ deg(f) HV(M).$$

Proof. We choose an exact form ω on M with $\int_M P\omega \wedge \omega$ essentially equal to the Hopf volume of M and the norm of ω bounded by 1 pointwise. After applying a homotopy with Lemma 9.1, we can assume that f is a smooth map with 2-dilation bounded by C. Then the pullback $f^*(\omega)$ has norm bounded by C. The integral $\int_X (1/C) P f^*(\omega) \wedge (1/C) f^*(\omega)$ is less than the combinatorial Hopf volume of X. On the other hand, $f^*(P\omega)$ is a primitive of $f^*(\omega)$, so this integral is equal to $(1/C)^2 \int_X f^*(P\omega) \wedge f^*(\omega)$, which is equal to $(1/C)^2(\text{deg} f) \int_M P\omega \wedge \omega$. Therefore the combinatorial Hopf volume of X is greater than $(1/C)^2(\text{deg} f) HV(M)$. □

(This proposition continues to hold, with the same proof, if X is merely a pseudo-manifold.)

9.3 Estimates for the combinatorial Hopf volume

In this section, we will give a variety of estimates for the Hopf volumes of three-manifolds.

We begin with a standard estimate involving the eigenvalue of an operator. Let λ_1 be the smallest non-zero eigenvalue of the operator d taking coclosed 1-forms to 2-forms.

Proposition 9.3.1. *The Hopf volume of (M, g) is bounded by λ_1^{-1} Volume(M, g).*

Proof. Let ω be an exact 2-form bounded by 1 pointwise, with $\int_M P\omega \wedge \omega$ essentially equal to the Hopf volume of (M, g). We can always choose a primitive $P\omega$ with $|P\omega|_{L^2} \leq \lambda_1^{-1}|\omega|_{L^2}$. Therefore, the integral $\int_M P\omega \wedge \omega$ is bounded by $\lambda_1^{-1}|\omega|^2_{L^2}$. Since ω is bounded by 1 pointwise, this expression is bounded by λ_1^{-1}Volume(M, g). □

This proposition shows that every closed oriented Riemannian 3-manifold has finite Hopf volume.

In order to use the last proposition, it would be necessary to give estimates for the eigenvalue λ_1, which seems difficult to me. We now give a different estimate, using an isoperimetric constant in place of an eigenvalue.

Define Iso(M,g) to be the infimal constant so that every exact real 1-cycle in M with length L bounds a real 2-chain with area less than $IsoL$.

Proposition 9.3.2. *The Hopf volume of (M,g) is less than $Iso(M, g)$ Volume(M, g).*

Proof. As usual, let ω be an exact 2-form with norm bounded by 1 pointwise and $\int_M P\omega \wedge \omega$ essentially equal to the Hopf volume of (M,g). According to the Sullivan trick (see section 7.2), we can choose a bounded measurable primitive $P\omega$, with norm bounded by $Iso(M, g)$ almost everywhere. We also checked in section 7.2 that this measurable primitive gives the same value in our integral as a smooth primitive would. The integrand is bounded by Iso and the integral is bounded by $Iso(M, g)$Volume(M, g). □

Using this proposition, we can prove an upper bound for the combinatorial Hopf volume of a triangulated pseudo-manifold in terms of the number of the simplices.

Proposition 9.3.3. *If X is a triangulated pseudo-manifold with N simplices then the combinatorial Hopf volume of X is bounded by C^N.*

Proof. We need to estimate Iso(X) or else $Iso(\tilde{X})$ if X is only a pseudomanifold. Let c be an exact real 1-cycle in $\tilde{X}$, with length L. We can easily push c into X itself, with a negligible increase in length, through an annulus of area ϵL. Next, we can push c to the 1-skeleton of the triangulation of X without increasing its length by more than a factor C, through a surface of area less than CL. This pushing is accomplished by the Federer-Fleming method of pushing outward from a random point to get to the 2-skeleton and then again pushing outward from a random point to get to the 1-skeleton. Then we rewrite c as $\sum a_i e_i$, where e_i is an oriented edge of the 1-skeleton of X and $\sum |a_i|$ is less than CL.

We pick an orientation for each 1-simplex and each 2-simplex in the triangulation of X. There is a boundary map from the free vector space over 2-simplices to the free vector space over 1-simplices, and by assumption the cycle $\sum a_i e_i$ is in the image of this boundary map. The boundary map is given by a matrix M with entries ± 1 or 0, and with dimensions bounded by CN. Because each 2-simplex has only three 1-simplices in its boundary, each column of the matrix has at most three non-zero entries. Recall that the Hilbert-Schmidt norm of a matrix with entries M_{ij} is defined to be $(\sum |M_{ij}|^2)^{1/2}$. The Hilbert-Schmidt norm of the boundary matrix is bounded by $(CN)^{1/2}$.

Let r be the rank of the boundary matrix M. We can choose orthogonal matrices O_1 and O_2 so that $O_1 M O_2$ is zero outside of the top left-hand $r \times r$ sub-matrix, which we call M'. Because the rank of M is r, one of the $r \times r$ sub-determinants of M is not zero. Because the entries of M are all integers, this sub-determinant is an integer, and so its norm is at least 1. The norm of the determinant of M' is at least as large as the norm of any $r \times r$ sub-determinant of M. Therefore, the determinant of M' has norm at least 1. The Hilbert-Schmidt norm of M' is equal to that of M, which is less than $(CN)^{1/2}$. Any $(r-1) \times (r-1)$ sub-matrix of M' has smaller Hilbert-Schmidt norm. The determinant of an arbitrary matrix M_{ij} is bounded by $\Pi_i(\sum_j |a_{ij}|^2)^{1/2}$. It follows from this estimate and our bound on the Hilbert-Schmidt norm that every $(r-1) \times (r-1)$ sub-matrix of M' has determinant bounded by C^N. Since the determinant of M' is at least 1, the norm of every entry in the inverse of M' is bounded by C^N. The norm of the inverse of M' is bounded by CNC^N.

Therefore, the exact 1-chain $\sum a_i e_i$ is equal to the boundary of a 2-chain $\sum b_i \Delta_i$, with $\sum |b_i|$ bounded by $CN^2C^N \sum |a_i|$. After increasing the constant C, we can say that $\sum |b_i|$ is bounded by $C^N \sum |a_i|$. Therefore, the 1-cycle $\sum a_i e_i$ bounds a real 2-chain with mass less than $C^N L$. In other words, $Iso(X)$ is bounded by C^N.

According to the last proposition, the combinatorial Hopf volume of X is at most $Iso(X)$ times the volume of X, which is less than CNC^N. After increasing the constant C, the Hopf volume of X is bounded by C^N. □

This exponential upper bound seemed very high to me at first. For comparison, the sphere of radius $N^{1/3}$ admits a triangulation by N simplices that are roughly standard geometrically, and the combinatorial Hopf volume of this triangulation is roughly $N^{4/3}$. We will show, however, that the Hopf volume really can grow exponentially in N for some triangulated manifolds with N simplices.

Proposition 9.3.4. *There exists a triangulation of S^3 with N simplices and combinatorial Hopf volume greater than $exp(cN)$.*

Proof. We start with a torus T^2, equipped with a triangulation T and with a choice of basis for $H_1(T^2)$, called a and b, with intersection pairing of a and b equal to 1. Next we take the product $T^2 \times [0,1]$, and we Dehn fill both boundary components to get S^3. More precisely, we Dehn fill the component $T^2 \times \{0\}$ in such a way that the homology class a bounds a disk in the new solid torus, and we Dehn fill the component $T^2 \times \{1\}$ in such a way that the homology class b bounds a disk in the new solid torus. We pick triangulations of the two solid tori extending T. (So far we have only constructed the simplest Heegaard splitting of S^3.)

Next we will construct a triangulation of the central cylinder $T^2 \times [0, 1]$ which restricts to the triangulation T on each boundary component. We pick an Anosov diffeomorphism Ψ of T^2. For instance, Ψ might act on homology by $\Psi_*(a) = 2a+b$ and $\Psi_*(b) = a + b$. We pick a triangulation of the mapping cylinder of Ψ which restricts to the triangulation T on each boundary component. Our triangulation of $T^2 \times [0, 1]$ consists of N copies of this triangulated mapping cylinder laid end to end, followed by N copies of the mirror image of the mapping cylinder. The whole triangulation involves less than CN simplices.

We will show that this triangulated 3-sphere contains two thick tubes with linking number at least $exp(cN)$. The first tube T_1 is localized near the middle of the long cylinder. Using the local coordinates for T^2, it is homologous to a. The second tube T_2 is the core of the Dehn filling of $T^2 \times \{0\}$. The linking number of T_1 with T_2 is equal to the pairing of $\Psi_*^N(b)$ with a, which grows exponentially with N.

For each thick tube, we have a Lipshitz map from S^3 to S^2, given by the Pontryagin-Thom collapse of a suitable framing of the thick tube. (What we have to check is simply that we can find a framing which does not wind around very much compared to the triangulation near the tube.) Let us call these maps f_1 and f_2. Let a_1 and a_2 be the pullback of the area form of S^2 by f_1 and f_2 respectively. The forms a_1 and a_2 are bounded pointwise by C. We define another map f_3, which is the composition of a map from S^3 to $S^2 \vee S^2$ using the Pontryagin-Thom collapse on both thick tubes with a degree (1,1) map from $S^2 \vee S^2$ to S^2. The pullback of the area form of S^2 by f_3 is equal to $a_1 + a_2$.

Now a simple calculation shows that $\text{Hopf}(f_3) - \text{Hopf}(f_2) - \text{Hopf}(f_1) = 2 \int_{S^3} Pa_1 \wedge a_2$, which is proportional to the linking number of T_1 with T_2, which is greater than $exp(cN)$. Therefore, the Hopf volume of this triangulation is at least $exp(cN)$. □

Of course, the manifold S^3 admits much simpler triangulations than the one given above. If we use the triangulation above to guide a surgical operation, we can construct manifolds M^3 which can be triangulated by N simplices, but so that every triangulation of M has Hopf volume on the order of $exp(cN)$.

Proposition 9.3.5. *We will construct closed oriented 3-manifolds M which can be triangulated by N simplices, and which admit maps to a surface of genus 2 with Hopf invariant at least $exp(cN)$.*

Proof. The manifold M is constructed by doing a kind of surgery on the above triangulation X near the two tubes T_1 and T_2. For each tube, we do the following procedure. Identify T with $D^2 \times S^1$, and refine the triangulation X so that the boundary of $D^2 \times S^1$ lies in the 2-skeleton. We can do this operation without increasing the number of simplices by more than a constant number, and moreover, the number of simplices in the triangulation of the boundary of $D^2 \times S^1$ can be bounded by C. Finally, the triangulation is well behaved with respect to the basis for H_1 of the boundary given by the curves $\partial D^2 \times \{q\}$ and $\{p\} \times S^1$. Next, we let Σ' be a surface of genus two with one boundary component. We cut out $D^2 \times S^1$ from our manifold and glue in $\Sigma' \times S^1$. The gluing map is the diffeomorphism from $\partial\Sigma' \times S^1$ to $\partial D^2 \times S^1$ given by the product of a diffeomorphism from $\partial\Sigma'$ to ∂D^2 and the identity map on S^1. Finally, we extend our triangulation to $\Sigma' \times S^1$, again adding only finitely many simplices. Applying this procedure to both T_1 and T_2, we get our 3-manifold M.

Now we construct some maps from M to a closed surface of genus 2. Inside of M, we have two copies of $\Sigma' \times S^1$, which we added with our two surgeries. For each copy, we can construct a map in the following way. We map $\Sigma' \times S^1$ to Σ' by projecting to the first factor, and then compose with a degree 1 map from Σ' to a surface of genus 2, taking the boundary of Σ' to the base point of the target surface. The resulting map takes the boundary of $\Sigma' \times S^1$ to the base point of Σ, and so we can extend the map to all of M, mapping the rest of M to the basepoint of Σ. In this way, we construct two maps, called f_1 and f_2. We can also construct a third map f_3 by applying the above construction to both copies of $\Sigma' \times S^1$. As in the previous proposition, $\mathrm{H}opf(f_3) - \mathrm{H}opf(f_1) - \mathrm{H}opf(f_2)$ is equal to twice the linking number of T_1 with T_2, which is greater than $exp(cN)$. ☐

We are now ready to prove Theorem 9.1. First we recall the statement.

Theorem 9.1. Let X be any pseudomanifold that can be triangulated by N simplices. Let f be a map from X to a closed surface of genus 2. If the Hopf invariant of f is defined, then it is bounded by C^N. There are examples of 3-manifolds M which can be triangulated by N simplices and maps from M to a surface of genus 2 with Hopf invariant greater than c^N.

Proof. We give the surface of genus 2 a hyperbolic metric. We define $\alpha = 1/(4\pi) darea$. The integral of α over the surface is 1, and so $\int_X Pf^*(\alpha) \wedge f^*(\alpha)$ is the Hopf invariant of f. Using the homotopy lemma 9.1, we can assume that the 2-dilation of f is at most C. Since the norm of α is bounded by $1/(4\pi)$, the norm of $f^*(\alpha)$ is bounded by C. Therefore, the integral $\int_X Pf^*(\alpha) \wedge \alpha$ is bounded by a constant times the combinatorial Hopf volume of X. By proposition 9.3.3, the combinatorial Hopf volume is less than C^N.

The examples were constructed in the last proposition. □

Given a 3-manifold X, it is interesting to know the minimal combinatorial Hopf volume of any triangulation of X. If X can be triangulated by N simplices, then this minimum is less than C^N. According to proposition 9.3.5, this estimate is fairly sharp for some 3-manifolds. For many other 3-manifolds, the estimate is not at all sharp. In particular, we will now give much stronger estimates for circle bundles and for a large variety of mapping tori. These estimates are not necessary to prove Theorems 9.1 or 9.2. Using these estimates, we could strengthen Theorems 9.1 and 9.2 when the domain is in the list above.

Proposition 9.3.6. *Let M be the total space of the circle bundle over a closed surface Σ of genus G with Euler number 1. Then M admits a triangulation with Hopf Volume bounded by CG^2.*

Proof. Let U be equal to the product of the surface Σ minus a small ball with S^1. We can triangulate U with CG simplices. The manifold M is a Dehn filling of U along the torus boundary, and we can extend the triangulation to M with only C additional simplices.

We will prove that for this triangulation, $Iso(M)$ is bounded by CG. To see this, let c be a closed null-homologous curve in M, with length L. By a homotopy through an area less than CL, we can move c to a curve that lies in U. Since c is null-homologous in M, the number of times c goes around the fiber S^1 is equal to the number of times it goes around the small ball removed from Σ with opposite signs. We can find a 2-chain in M - U with area bounded by CL whose boundary lies in the boundary of U, and which is homologous to the curve c in $H_1(U)$. It suffices to find a filling of the difference c' of c with the boundary of this 2-chain.

We divide Σ minus a ball into G regions, each of which is equal to a torus with one or two boundary components. This division leads to an analogous division of U into G regions, each of which is the product of one of the regions above with S^1. With a slight addition of length, we can construct a small area cobordism of our curve to a union of curves, one supported in each region. Each curve in the union is homologous to a curve in the boundary of the union through a cobordism with area less than C L. The total length of all the curves is less than C L. These regions in the boundary can be taken to be either the boundary circles of the regions in Σ or fibers of U. Now each boundary circle of the regions in Σ can be filled by area G. The fibers can all be moved to the fiber over a given single point by a cobordism with area CGL. Finally, since the curve we are working with is null-homologous in U, the fibers over this point cancel. Therefore, Iso is less than CG.

According to proposition 9.3.2, the combinatorial Hopf Volume of M is less than Iso times the volume, which is bounded by CG^2. □

(As a topological application, we see that any map from M to a closed surface of genus 2 has Hopf invariant less than CG^2. On the other hand, the projection from M to Σ, followed by the degree $G-1$ map to from Σ to the surface of genus 2, has Hopf invariant $(G-1)^2$. Therefore, our estimate was sharp up to a constant factor. On the other hand, there is a completely elementary proof that any map from M to a surface of genus 2 has Hopf invariant at most $(G-1)^2$. The fundamental group of M is generated by a_i, b_i, where i go from 1 to g, with the relation that the element $\Pi_i[a_i, b_i]$ commutes with all other elements. For any map f from M to a surface of genus 2, consider the image of this central element in π_1 of the target surface. If this image is the identity, then the map f factors through the projection from M to Σ. The Hopf invariant of the resulting map is equal to the square of the degree of the induced map from Σ to the surface of genus 2. This degree is at most G-1 according to Kneser's theorem (see Chapter 5 of [15]), and so the Hopf invariant is at most $(G-1)^2$ in this case. If the above image is not the identity, however, then its commutator subgroup is cyclic, and so the fundamental group of M is mapped to a cyclic subgroup of the fundamental group of the surface of genus 2. In this case, the mapping factors through a circle and the Hopf invariant is zero.)

Next, we bound the Hopf volume of certain mapping tori.

Proposition 9.3.7. *Let Σ be a closed oriented surface of genus G, and let γ_i be a set of homologically trivial curves in Σ. Then there is a constant C, which depends on G and γ_i, so that the mapping torus of any sequence of N Dehn twists around curves in the list γ_i admits a triangulation with combinatorial Hopf Volume bounded by CN^2.*

Proof. Pick a triangulation of Σ. Then pick a triangulation of the mapping cylinder of the Dehn twist around each γ_i, which restricts to the first triangulation on each boundary component. Now, for any sequence of Dehn twists, we can assemble these pieces to get a triangulation of the mapping torus.

We show that for the standard metric on this triangulation, Iso is bounded by CN. (Throughout this paragraph, C will denote a constant which depends only on G and the curves γ_i.) Let c be a null-homologous curve in our triangulated mapping torus. Let t be a function from the mapping cylinder to $\mathbb{R}$ modulo N, which goes from i to i+1 along the mapping cylinder of the ith Dehn twist. Mark the points of c at half-integer values of t, and straighten c between the marked points. This produces a cobordism with area controlled by a multiple of L and the final curve of length less than L. Then cut the straightened c at integer values of t. Since c winds zero times around M in total, each cut gives rise to a 0-cycle on a copy of Σ. We fill the zero cycle at a cost of zero area and increasing the length of the curve by a factor of G. The new curve is divided into components, each component supported on the mapping cylinder of a single Dehn twist around some γ_i. We then straighten each curve to lie on a constant t slice with cost of area CL, and length of the final curve less than CL. Now because each curve γ_i is null-homologous, each Dehn twist acts trivially on the homology of Σ. Using this fact, we can construct a cobordism moving each curve to lie on a single copy of Σ at t=0, without lengthening them, and at a cost of area CLN, where this constant depends on the curves γ_i. The resulting 1-cycle is null-homologous in Σ, and so we can bound it by a 2-chain with area CL. In total, we have constructed a filling for c with area CLN, which shows that Iso is bounded by CN.

According to Proposition 9.3.2, the Hopf volume of our triangulation is bounded by Iso times the volume. Since the volume of the triangulation is clearly bounded by CN, the Hopf volume is bounded by CN^2. □

Proposition 9.3.8. *Let Ψ be a diffeomorphism of a closed oriented surface Σ, so that the action of Ψ_* on the first homology group of Σ has no eigenvalues on the unit circle. Then the mapping torus of Ψ^N admits a triangulation with Hopf volume bounded by CN, where the constant C depends on Ψ.*

Proof. Pick a triangulation of Σ, and a triangulation of the mapping cylinder of Ψ which restricts to the first triangulation on both boundary components. Composing these cylinders we get a triangulation of the mapping torus of Ψ^N with CN simplices.

We will show that for the standard metric on this triangulation, Iso is bounded by a constant C independent of N. (This constant does depend on Ψ though.) Let c be a null-homologous curve in our triangulated mapping torus, of length L. We let t be a function from the mapping torus to $\mathbb{R}$ modulo N, which goes from i to i+1 on the ith copy of the mapping cylinder of Ψ. Mark the points of c at half-integer values of t, and straighten c between the marked points. This produces a cobordism with area controlled by a multiple of L and the final curve of length less than L. Then cut the straightened c at integer values of t. Since c winds zero times around the mapping torus in total, each cut gives rise to a 0-cycle on the surface Σ. We fill the zero cycle at a cost of zero area and increasing the length of the curve by a factor of G. The new curve is divided into components, each supported on one copy of the mapping cylinder of Ψ. At the cost of area CL, we can straighten each component to lie on a copy of Σ and to be equal to a sum of standard generators of the homology of Σ. Therefore, it suffices to bound the isoperimetric constant of a real chain with boundary equal to one of these standard generators.

Let a be one standard generator for the homology of Σ. We have assumed that the eigenvalues of Ψ_* are bounded away from the unit circle, and so the eigenvalues of $Id - \Psi_*^N$ are bounded away from 0. Therefore, we can find an element b in $H_1(\Sigma, \mathbb{R})$ with $b - \Psi_*^N(b) = a$, and with the norm of b bounded by C, independent of N. The element b is represented by a real 1-cycle in Σ with mass bounded by C. The cycle $b - \Psi_*^N(b)$ can be filled by a real 2-chain with area bounded by $C \sum_{i=1}^N |\Psi_*^i(b)|$. Since a and b are bounded, the equation $\Psi_*^N(b) = b - a$ shows that $\Psi_*^N(b)$ is also bounded. Since the eigenvalues of Ψ_* are bounded away from the unit circle, the size of $\Psi_*^i(b)$ is bounded by $Cexp[-\min(i, N - i)]$. Therefore, the above sum is bounded by a constant independent of N.

Since Iso is bounded by a constant independent of N, the Hopf volume of this triangulation is bounded by CN. $\square$

9.4 The Hopf volume of hyperbolic 3-manifolds

In this section, we will show that a closed oriented hyperbolic 3-manifold with small injectivity radius is complicated in one of four ways. Either it has large volume, or it requires surfaces of large genus to span H_2, or it has a torsion element in H_1 of high order, or it has large Hopf volume. Any of these four features leads to bounds on the degrees of maps into the manifold.

Let M be a closed oriented hyperbolic 3-manifold with injectivity radius ϵ. Then M contains a closed embedded geodesic γ of length on the order of ϵ. If γ is torsion in homology with a fairly small order, then we will prove that either M has large Hopf volume or else M has large volume. This estimate is the main idea in the proof of Theorem 9.2.

Proposition 9.4.1. *Let M be a closed oriented hyperbolic 3-manifold with a closed geodesic γ of length ϵ which is torsion in homology. Then either the volume of M is* *at least $c\epsilon^{-1/6}$, or the order of γ is at least $c\epsilon^{-1/6}$, or the Hopf volume of M is at least $c\epsilon^{-1/4}$.*

Proof. Let T be the Margulis tube around γ. (The thin part of M is defined to be the subset of M where the injectivity radius is less than a certain constant, and the Margulis tube is the connected component of the thin region containing γ.) The universal cover of γ is a geodesic in hyperbolic 3-space. We can parameterize hyperbolic space by the upper half-space model so that the universal cover of γ is equal to the vertical line through the origin (i.e. the line $x = 0, y = 0, z > 0$). The group of covering transformations of M include a loxodromic isometry that fixes this vertical line. This isometry is given by multiplying the three coordinates by a constant on the order of $(1 + \epsilon)$, and rotating in the (x,y)-plane by an angle θ. Consider the quotient of hyperbolic 3-space by this isometry, and let $\tilde{T}$ be the Margulis tube around the core geodesic of this quotient. Taking the quotient by the other covering transformations of M gives a map from $\tilde{T}$ into M, which takes the core geodesic of $\tilde{T}$ onto γ. According to the Margulis Lemma, the map is an embedding of $\tilde{T}$ into the Margulis tube around γ.

The Margulis tube $\tilde{T}$ has a simple form. For some number R depending on ϵ and θ, the tube $\tilde{T}$ is equal to the quotient of the region $x^2 + y^2 < z^2R^2$ by the action of the loxodromic isometry corresponding to γ. A fundamental region for this action is given by the intersection of the region above with the region $1 \leq z \leq 1 + \epsilon$. The boundary of the fundamental region includes two disks, the first given by $z = 1$ and $x^2 + y^2 \leq R^2$, and the second given by $z = 1 + \epsilon$ and $x^2 + y^2 \leq (1 + \epsilon)^2R^2$. The loxodromic isometry corresponding to γ takes the first disk onto the second disk. In several constructions, we will use the radial curves in $\tilde{T}$, which are the straight rays through the origin in the Euclidean metric on the (x,y,z)-space. These radial lines are not geodesics in hyperbolic space. On the other hand, their images foliate the quotient $\tilde{T}$.

We will give two estimates for R in terms of ϵ and θ. The first estimate says that if ϵ is small, then R is large. More precisely, the radius R is at least $c\epsilon^{-1/2}$. Beginning at a point p on the edge of $\tilde{T}$, we can follow a radial curve with length $N\epsilon R$ going

N times around the edge of $\tilde{T}$. This curve hits the circle $z = 1$, $x^2 + y^2 = R^2$ in N points. Connecting the closest two of these N points by an arc of the circle, we get a homotopically non-trivial closed curve with length less than $N\epsilon R + R/N$. Since this curve lies on the edge of the Margulis tube, its length is greater than a constant on the order 1. We can make this construction for any number N. In particular, if $N = \epsilon^{-1/2}$, and we get the inequality $\epsilon^{1/2}R \geq c$. This proves our lower bound $R > c\epsilon^{-1/2}$. On the other hand, we can assume that R is not too big. The volume of $\tilde{T}$ is roughly ϵR^2. If the volume of M is at least $c\epsilon^{-1/6}$ then we are done, so we may assume that $R < C\epsilon^{-7/12}$.

Our second estimate says that if $\theta/(2\pi)$ is well-approximated by rational numbers, then R is large. Again, we begin at a point at radius R from the the center of the horosphere $z = 1$, and we trace a radial curve that goes vertically q times around the tube, and then connect the endpoint of this curve to the starting point within the horosphere $z = 1$. The total length of this curve is roughly $q\epsilon R + |q\theta/(2\pi) - p|R$, where p is the integer that makes this expression smallest. Since this curve is homotopically non-trivial and lies on the edge of the Margulis tube, it must have length at least on the order of 1. Therefore, $|\theta/(2\pi) - p/q| > c(1/q)(1/R) - \epsilon$.

We choose generators for the homology of the boundary of $\tilde{T}$, given by l, the longitude, and m the meridian. (The choice of longitude is related to the choice of θ in the following way. Take a radial curve on the edge of our fundamental domain, going once around the tube from $z = 1$ to $z = 1 + \epsilon$. This line connects to a point on the base horosphere $z = 1$. From that point, follow an arc of the circle with directed length $-\theta$ to the initial point of the radial curve. The resulting closed curve is homologous to the longitude.) Let $pm + ql$ be the homology class of the primitive curve which bounds in $M - \tilde{T}$. (The numbers p and q are relatively prime integers. Since the proposition assumes that γ is torsion in homology, q is not zero. The absolute value of q is the order of γ.) If q is at least $c\epsilon^{-1/6}$ we are done, so we will assume that q is smaller than $c\epsilon^{-1/6}$. The inequality at the end of the last paragraph tells us that $|\theta/(2\pi) - p/q| > c\epsilon^{3/4} - \epsilon$. Since the volume of M is at least on the order of 1, we can assume that ϵ is quite small, and therefore $|\theta/(2\pi) - p/q| > c\epsilon^{3/4}$.

Finally, we estimate the Hopf volume of M. We claim that the Hopf volume of M is at least $c|\theta/(2\pi) - p/q|R^2$, which is greater than $c\epsilon^{-1/4}$. To see this, we build a collection of tubes in $\tilde{T}$. Each tube goes around the tube $\tilde{T}$ N times, and goes around the core of $\tilde{T}$ M times, where (M/N) is very close to $\theta/2\pi$. If $\theta/(2\pi)$ is rational with denominator N, the tubes are just small neighborhoods of the radial curves. If $\theta/(2\pi)$ is irrational, but very close to M/N, then the tube $\tilde{T}$ is bilipshitz to the corresponding tube with a loxodromic rotation through an angle M/N, and we pull back the tubes from there. We divide the tubes into inner tubes, which stay in the central part of $\tilde{T}$ given by $x^2+y^2 < 1/4z^2R^2$, and outer tubes, which stay on the outside of this region. Geometrically, each tube is bilipshitz to a Euclidean product of a small disk and a long circle. The total cross-sectional area of the outer tubes is on the order of (R/N), and the total cross-sectional area of the inner tubes is also on the order of (R/N). The linking number of an outer tube with an inner tube is given by $((M/N) - (p/q))N^2$. For large values of N, this quantity converges to $(\theta/(2\pi) - p/q)N^2$.

Now we use the bilipshitz map from each tube to the product of a Euclidean disk with a long circle to construct some exact 2-forms. We start with a compactly supported 2-form on the disk with integral roughly equal to the area of the disk. We pull this 2-form back to the product and then back to the tube in $\tilde{T}$, which gives us a closed 2-form with norm bounded by 1 pointwise. Adding together all the 2-forms on the outer tubes, we get a closed 2-form a, and adding together all the 2-forms on the inner tubes we get a closed 2-form b. Since all the tubes are torsion in homology, a and b are exact 2-forms on M. Now the Hopf volume of M is at least $(1/3)|\int_M P(a+b) \wedge (a+b) - \int_M Pa \wedge a - \int_M Pb \wedge b|$. This expression, however, is equal to the linking number of each outer tube with each inner tube, times the total cross-sectional area of the outer tubes, times the total cross-sectional area of the inner tubes, which is at least $c|\theta/(2\pi) - p/q|N^2(R/N)(R/N)$. This expression is equal to $c|\theta/(2\pi) - p/q|R^2$, which is at least $c\epsilon^{3/4}R^2$, which is at least $c\epsilon^{-1/4}$. □

(To understand this result better, it might help to see how it plays out for the hyperbolic Dehn fillings of a compact oriented hyperbolic manifold with a single cusp. Let X_0 be such a manifold, and let a and b be a basis for the homology of the boundary torus so that the homology class a bounds in X_0, and so that the intersection number of a and b is 1. Let $X(m,n)$ be the manifold formed by Dehn filling the boundary torus along the curve homologous to $ma+nb$, for relatively prime numbers m and n. According to Thurston's theory, the manifold $X(m,n)$ admits a hyperbolic structure for all but finitely many choices of (m,n). These hyperbolic manifolds have uniformly bounded volume. They contain a short core geodesic with length on the order of $(m^2+n^2)^{-1}$. Except for the single case when n is zero, the core geodesic is torsion with order n. On the other hand, the Hopf volume of $X(m,n)$ is at least $c(m/n)$ when m or n is large. The reason is that uniformly in m and n, the manifold $X(m,n)$ contains two disjoint thick tubes homologous to b, which lie in the thick part of X_0. Neither tube is null-homologous in X_0, so their linking number is not defined in X_0. But in $X(m,n)$, the tubes are (rationally) null-homologous, and a short calculation shows that their linking number is equal to (m/n). Therefore, we see that each hyperbolic Dehn filling with a short core geodesic either has a torsion element in H_1 with large order or else has large Hopf volume.)

To complement this proposition, we prove a different kind of inequality for hyperbolic 3-manifolds with short geodesics that are non-torsion in homology.

Proposition 9.4.2. *If a closed oriented hyperbolic manifold has a closed geodesic which is non-trivial in $H_1(M,\mathbb{Q})$ of length ϵ, then any surface with non-zero intersection number with γ must have both area and genus at least $c\epsilon^{-1/2}$.*

Proof. We consider as above the Margulis tube around the short geodesic γ, which has radius at least $R=c\epsilon^{-1/2}$. It contains thin tubes, each going around N times, with total cross-sectional area at least R/N. Any surface with a non-zero intersection number with γ meets each of these tubes at least N times, and so it has total area at least $N(R/N)=R>c\epsilon^{-1/2}$. By the Thurston simplex straightening argument, the genus of such a surface must be at least $c\epsilon^{-1/2}$ as well. □

We define the spanning genus of a 3-dimensional simplicial complex X to be the smallest genus G so that $H_2(X,\mathbb{Q})$ is spanned by surfaces of genus at most G. We can rephrase the last proposition in this language: if M is a closed oriented hyperbolic

3-manifold with a closed geodesic of length ϵ which is not torsion in $H_1(M)$, then the spanning genus of M is at least $c\epsilon^{-1/2}$.

In the last section, we proved that the Hopf volume of a complex with N simplices is less than C^N. Similarly, we will prove that the spanning genus of a complex with N simplices is less than C^N.

Proposition 9.4.3. *Let X be a simplicial 3-complex with N simplices. Then the spanning genus of X is less than C^N.*

Proof. We consider the boundary map from the free vector space of 2-simplices of X to the free vector space of 1-simplices of X. The boundary map is given by a matrix M, with dimensions less than CN, and with each entry equal to ± 1 or 0. Moreover, since the boundary of each 2-simplex is only three 1-simplices, each column of the matrix has only three non-zero entries. Let r be the rank of the matrix M. Pick r 2-simplices Δ_i in X, so that $M(\Delta_i)$ gives a basis for the image of M. After renumbering the 2-simplices, we suppose that these are Δ_1 through Δ_r.

Now for every other 2-simplex Δ_j, we can express $M(\Delta_j)$ as a sum $a_{i,j}M(\Delta_i)$. (In this formula, the index j runs from $r+1$ to the last 2-simplex, and the index i runs from 1 to r.) Let I be the restriction of M to the span of the first r simplices. The linear map I is an isomorphism from $\mathbb{R}^r$ onto the image of M. This isomorphism has Hilbert-Schmidt norm less than $CN^{1/2}$, and it has determinant at least 1. Therefore, its inverse has norm at most C^N. For a fixed j, the sum $a_{i,j}\Delta_i$ is simply $I^{-1}(M(\Delta_j))$. Since $M(\Delta_j)$ has length $\sqrt{3}$, the vector $a_{i,j}$ has length less than C^N. In particular, each coefficient $a_{i,j}$ is less than C^N.

The coefficients $a_{i,j}$ don't have to be integers, but they are rational numbers with controlled denominators. The vectors $M(\Delta_j)$ are integral vectors in the image of M. The vectors $M(\Delta_i)$ span a lattice L of integral vectors in the image of M. Let D be the index of the lattice L inside all of the integral vectors in the image of M. In particular, $DM(\Delta_j)$ lies in the integral span of the $M(\Delta_i)$, and so $Da_{i,j}$ is an integer for every i and j. Since the determinant of I is bounded by C^N, D is also bounded by C^N.

Now the integral 2-cycles $D\Delta_j - \sum_i Da_{i,j}\Delta_i$ span the kernel of the matrix M. Each of these 2-cycles has less than NC^NC^N 2-simplices, and so it lies in the span of some surfaces with genus less than C^N. □

Over the course of the last proof, we have also bounded the size of the torsion subgroup of $H_1(X, \mathbb{Z})$. This subgroup is equal to the integral vectors in the image of M modulo the lattice L, and so D is the order of the torsion subgroup of $H_1(X, \mathbb{Z})$. As we showed above, this order is less than C^N. By an analogous argument, the torsion subgroup of $H^2(X, \mathbb{Z})$ has order less than C^N.

We can now give the proof of Theorem 9.2. First we recall the statement.

Theorem 9.2 Let M be a closed oriented hyperbolic 3-manifold with injectivity radius ϵ. Let X be a 3-dimensional pseudo-manifold which can be triangulated by N simplices. If there is a map of non-zero degree from X to M, then ϵ is at least C^{-N}. On the other hand, we will give examples of closed oriented hyperbolic 3-manifolds which can be triangulated with N simplices and have injectivity radius less than c^{-N}, for a constant $c > 1$.

Proof. Let f be a map of non-zero degree from X to M. According to Lemma 9.1, we can homotope f to a map with 2-dilation less than C. Since M has injectivity radius ϵ, it must contain a closed geodesic of length less than 2ϵ. According to Propositions 9.4.1 and 9.4.2, either M has volume at least $c\epsilon^{-1/6}$, or $H_1(M, \mathbb{Z})$ has a torsion element with order at least $c\epsilon^{-1/6}$, or M has Hopf volume at least $c\epsilon^{-1/4}$, or M has spanning genus at least $c\epsilon^{-1/2}$.

If M has volume at least $c\epsilon^{-1/6}$, then according to Thurston's straightening theorem, N is greater than $c\epsilon^{-1/6}$, and we are done.

If $H_1(M, \mathbb{Z})$ has a torsion element of order T, greater than $c\epsilon^{-1/6}$, then we proceed as follows. Let α be this torsion class, and let α^* be the Poincare dual class in $H^2(M, \mathbb{Z})$. The element $f^*(\alpha^*)$ in $H^2(X, \mathbb{Z})$ must also be torsion with some order S dividing T. We claim that T divides the product of S with the degree of f. Let α_T be the image of α in $H_1(M, \mathbb{Z}_T)$. Let α_T^* be the Poincare dual class in $H^2(M, \mathbb{Z}_T)$. The pullback $Sf^*(\alpha_T^*)$ also vanishes. Since $H_0(M, \mathbb{Z})$ is free, $H^1(M, \mathbb{Z}_T)$ is equal to

$Hom(H_1(M,\mathbb{Z}),\mathbb{Z}_T)$. In particular, we can choose a class a in $H^1(M,\mathbb{Z}_T)$ so that $a(\alpha) = a(\alpha_T) = 1$ modulo T. Therefore, the cup product $a \cup \alpha_T^*$ is equal to the fundamental cohomology class O_T in $H^3(M,\mathbb{Z}_T)$. We let $[X]_T$ denote the image of the fundamental homology class of X in $H_3(X,\mathbb{Z}_T)$. The degree of f modulo T is given by the pairing $f^*(O_T)([X]_T)$. Finally, S times the degree of f modulo T is given by $f^*(SO_T)([X]_T)$. But $f^*(SO_T)$ is equal to $f^*(a) \cup Sf^*(\alpha_T^*)$, which vanishes. Since the volume of M is at least on the order of 1, the degree of f is at most CN by Thurston's simplex straightening argument. Also, the number S is bounded by the order of the torsion subgroup of $H^2(X,\mathbb{Z})$, which is bounded by C^N. Therefore C^N must be greater than $c\epsilon^{-1/6}$, and we are done.

If the Hopf volume of M is at least $c\epsilon^{-1/4}$, we proceed as follows. Since the map f has 2-dilation less than C and degree non-zero, the combinatorial Hopf volume of X must be at least $c\epsilon^{-1/4}$. According to Proposition 9.3.3, the combinatorial Hopf volume of X is less than C^N. Therefore, C^N must be greater than $c\epsilon^{-1/4}$, and we are done.

If the spanning genus of M is at least $c\epsilon^{-1/2}$, we proceed as follows. Since the degree of f is non-zero, and since M obeys Poincare duality, the map f_* from $H_2(X,\mathbb{Q})$ to $H_2(M,\mathbb{Q})$ must be surjective. Therefore, the spanning genus of X must be at least $c\epsilon^{-1/2}$. According to proposition 9.4.3, the spanning genus of X is less than C^N. Therefore, C^N must be greater than $c\epsilon^{-1/2}$ and we are done.

We have now proven that if X admits a map of non-zero degree to M, then ϵ is at least C^{-N}. On the other hand, we will construct closed oriented hyperbolic manifolds M which can be triangulated by N simplices and which contains closed geodesics of length less than c^{-N}. Begin with a non-compact finite volume hyperbolic 3-manifold M_0 with a single cusp. We view M_0 as a manifold with boundary and we triangulate it. Let T be the restriction of our triangulation to the torus boundary. Next, we pick an Anosov diffeomorphism Ψ of the torus and a triangulation of the mapping cylinder of this diffeomorphism, which restricts to the triangulation T on each boundary component. Finally, we pick a curve c_0 in the boundary torus and a triangulation of a solid torus which restricts to T on the boundary and which Dehn fills the curve c_0. To form the 3-manifold M, we glue together M_0, N copies of the mapping cylinder, and one copy of the solid torus. (At each gluing, we glue one torus with triangulation T to another torus with the same triangulation by using the identity map.) The resulting 3-manifold can be triangulated with less than CN simplices. It is diffeomorphic to the Dehn filling of M_0 along the curve $\Psi^{-N}(c_0)$. If we fix any basis of $H_1(T^2)$, then for most choices of c_0, the coefficients of the homology class $\Psi_*^{-N}([c_0])$ grow exponentially with N. By Thurston's theory, almost all of these Dehn fillings are closed hyperbolic 3-manifolds, with uniformly bounded volume.

If the short geodesic in the core of the Dehn filling has length ϵ, then the radius of the horosphere cross-section of the Margulis tube must be roughly $\epsilon^{-1/2}$, and so there is a homologically non-trivial simple curve in the boundary of the Margulis tube of length less than $\epsilon^{-1/2}$ which bounds a disk in the Margulis tube. For large values of N, we have proved above that this length must be at least c^N, for a constant $c > 1$. Therefore, the length of the shortest geodesic is less than c^{-N}. □

Corollary. *Given any closed oriented 3-manifold X, there are at most finitely many closed oriented hyperbolic 3-manifolds M with maps of non-zero degree from X to M.*

Proof. If X can be triangulated with N simplices, then according to Theorem 9.2, M must have injectivity radius at least C^{-N}. According to the Thurston straightening argument, M must have volume at most CN. By Cheeger finiteness theorem, there are only finitely many hyperbolic manifolds obeying these bounds. □

This result is due to Teruhiko Soma. Trying to understand Soma's result was the main motivation for the work in this section. The proof above essentially shows the following finiteness result for closed oriented hyperbolic 3-manifolds.

Proposition 9.4.4. *For any number B, the set of closed, oriented, hyperbolic 3-manifolds with volume less than B, the order of the torsion subgroup of H_1 less than B, Hopf volume less than B, and spanning genus less than B is finite.*

Proof. Let M be a closed oriented hyperbolic 3-manifold obeying the above bounds. Because the Hopf volume of M is less than B, the volume of M is less than B, the spanning genus of M is less than B, and the order of the torsion subgroup of H_1 is less than B, the injectivity radius of M is at least C^{-B}. By Cheeger finiteness, the set of hyperbolic 3-manifolds with volume less than B and injectivity radius at least C^{-B} contains only finitely many diffeomorphism types. □

In particular, there are only finitely many hyperbolic homology three-spheres with volume less than B and Hopf volume less than B. It seems likely that there are only finitely many hyperbolic homology three-spheres with Hopf volume less than B.

9.5 The situation for non-orientable manifolds

Our proof of Soma's theorem does not extend to non-orientable manifolds, because intersection numbers, linking numbers, and the Hopf invariant are only defined modulo 2 in non-orientable manifolds. As Wang discovered, the analogue of Soma's theorem is false for non-orientable manifolds.

Theorem. *(Wang) There is a closed non-orientable 3-manifold X which admits degree 1 maps to infinitely many different closed oriented hyperbolic 3-manifolds M. (The degree is only defined modulo 2.)*

In the paper [4], Boileau and Wang published an incorrect proof that there is a closed oriented 3-manifold which admits maps of non-zero degree to infinitely many different closed oriented hyperbolic 3-manifolds. According to Soma's theorem, there is no such closed oriented 3-manifold. Wang told me that by carefully reading this incorrect proof, one gets a correct proof of the theorem above. I had some trouble adapting the argument from [4], but I eventually found a proof along similar lines. For the reader's reference, here is a proof of Wang's theorem.

Proof. The main idea of the proof is a clever choice of the domain which is due to Wang. We begin with a finite volume hyperbolic manifold X_0 with a single cusp, which we view as a manifold with a torus boundary. Let us pick a basis for $H_1(\partial X_0)$, with elements a and b, so that the intersection product of a and b equal to 1. We then perform a filling of X_0 roughly analogous to a Dehn filling but using a Mobius band in place of a disk. Let us make a precise construction. Let B denote the Mobius band. The boundary of $B \times S^1$ is a torus. We pick a diffeomorphism ϕ of the boundary of X_0 with the boundary of $B \times S^1$, so that $\phi(\partial B)$ is homologous to a and $\phi(S^1)$ is homologous to b. We define X to be the result of gluing X_0 to $B \times S^1$ with this diffeomorphism.

We will construct degree 1 maps from X to half of all the Dehn fillings of X_0. All but finitely many of these Dehn fillings admit hyperbolic metrics, and they include infinitely many different manifolds. For every pair (m, n) of relatively prime integers, let $X(m, n)$ denote the Dehn filling of X_0 along a curve homologous to $ma + nb$. The manifold $X(m, n)$ consists of the union of X_0 and a solid torus, glued along their boundaries. The map that we will construct from X to $X(m, n)$ takes X_0 to X_0 identically. The restriction of the identity map to the boundary gives a map from the boundary of $B \times S^1$ to the boundary of $D^2 \times S^1$.

We now investigate in what cases we can extend this map to a map from $B \times S^1$ to $D^2 \times S^1$. Since maps to D^2 automatically extend, it suffices to extend our map from the boundary of $B \times S^1$ to its interior.

The homotopy classes of maps from the boundary of $B \times S^1$ to S^1 are simply $H^1(T^2)$. The cohomology class of our map is $n(\partial B)^* - m(S^1)^*$, where $(\partial B)^*$ is the cocycle that takes the value 1 on ∂B and the value 0 on S^1. Now B is homotopic to S^1. Let c denote a circle in B which is a deformation retract of B. Then $H^1(B \times S^1)$ is generated by c^* and $(S^1)^*$. The inclusion map of $\partial B \times S^1$ into $B \times S^1$ induces a map on cohomology taking c^* to $2(\partial B)^*$ and $(S^1)^*$ to $(S^1)^*$. A given map from the boundary to S^1 extends to the interior if and only if the corresponding class in $H^1(\partial B \times S^1)$ lies in the image of $H^1(B \times S^1)$. This condition is met if n is even. Therefore, for every even n, there is a degree 1 map from the non-orientable 3-manifold X to the Dehn filling $X(m, n)$. □

Appendix A

Other Stretching Conditions

In this thesis, we have been studying k-contracting mappings. We might ask, though, whether there are other conditions which limit how much a mapping can stretch the underlying space. For simplicity, we restrict attention to maps from $\mathbb{R}^n$ to $\mathbb{R}^n$.

The way in which a linear map distorts the Euclidean metric is described by its singular values. We recall that any $n \times n$ matrix M can be written as a product $O_1 D O_2$, where O_1 and O_2 are orthogonal matrices and D is a non-negative diagonal matrix. The diagonal entries of D, arranged in order, are $0 \leq s_1 \leq ... \leq s_n$. They are called the singular values of M.

We will call a set S of $n \times n$ matrices a stretching condition if, whenever M is a matrix in S, and whenever N is another matrix with $s_i(N) \leq s_i(M)$ for each i, then N is also in S. In particular, membership in S only depends on the singular values of a matrix. The set S codifies some restriction on how much a matrix is allowed to stretch space.

A smooth map f from a domain U in $\mathbb{R}^n$ to $\mathbb{R}^n$ is said to obey the condition S if the derivative df lies in S at each point of U. Since the condition S depends only on the singular values of a matrix, it also makes sense to say that a smooth map f between Riemannian n-manifolds obeys the condition S.

The k-contracting condition is an example of a stretching condition. A linear map is k-contracting if its singular values obey $s_{n-k+1} ... s_n \leq 1$. A smooth map is k-contracting if its derivative is k-contracting at each point.

There are many other stretching conditions besides k-contracting maps. For instance, for maps from $\mathbb{R}^3$ to $\mathbb{R}^3$, we can consider the condition $s_2^{1/2}s_3 \leq 1$. I believe that a linear map obeys $s_2^{1/2}s_3 \leq 1$ if and only if it decreases the 3/2-Hausdorff measure of all subsets - but I'm not sure because of technicalities with Hausdorff measure. In any case, this inequality defines a stretching condition. The reader can see that the set of all stretching conditions is quite large.

In this section, we show that mappings obeying a certain condition S are often able to approximate any mapping obeying a much more general condition S'.

Before we state a general result, let's look at the simplest example. Suppose that n is 2, and consider the family of conditions $S(\alpha)$ given by $s_1^{\alpha}s_2 \leq 1$, for α in the range $[0, 2]$. From the point of view of linear algebra, we have a continuous family of different conditions. But when we consider the space of mappings that obey these inequalities, then a different, discrete situation appears. If α is in the range $(0, 1)$, and if L is a linear map that obeys $s_1s_2 < 1$, then L can be C^0 approximated by smooth maps that obey $S(\alpha)$. If α is bigger than 1, and if L is any linear map at all, then L can be C^0 approximated by smooth maps that obey $S(\alpha)$. From this geometric point of view, the conditions $S(\alpha)$ do not vary continuously with α. Instead, they exhibit exactly three behaviors. If $\alpha = 0$, then $S(\alpha)$ describes Lipshitz maps. If α lies in the range $(0, 1]$, then $S(\alpha)$ describes 2-contracting maps. If α is bigger than 1, then $S(\alpha)$ describes the class of all maps.

Let us sketch the construction of these C^0 approximations. Suppose that α is in the range $(0, 1]$. It follows that $S(\alpha)$ is satisfied when $s_1 = s_2 = 1$, and also when $s_1 = 0$. Let L be a linear map obeying $s_1s_2 \leq 1$. If $s_1(L) = 0$, then L is already in $S(\alpha)$, so we have nothing to prove. If not, then L is a diffeomorphism from the unit square to a rectangle with area less than 1. We make a fine lattice in the rectangle, and we consider the open squares in the lattice. If we take the lattice fine enough, then these open squares can be isometrically embedded in the unit cube, as some subset V of U. This isometric embedding gives us a local isometry from V to the union of these squares, and this local isometry certainly satisfies $S(\alpha)$. Moreover, if we make the lattice quite fine, we can assume that this mapping from V to R is

C^0 close to the linear map L. We extend the mapping we have constructed so far by mapping the complement of A to the 1-skeleton of the fine lattice in R. Since the 1-skeleton is only 1-dimensional, the resulting map has $s_1 = 0$, and so it obeys $S(\alpha)$. We can make the extension so that the resulting map is C^0 close to L. (The map we have just constructed is not smooth along the boundary of A, but with a little more work, we can construct a smooth map.)

A similar construction can be made when α is greater than 1 in order to C^0-approximate any linear map L. We can assume that L maps the unit square onto some rectangle R with area A. Again we divide the rectangle into squares of side-length ϵ. Let Q be a rectangle with dimensions $A^{-\frac{\alpha}{\alpha-1}}\epsilon \times A^{\frac{1}{\alpha-1}}\epsilon$. There is a linear map from Q to the square with side length ϵ, with singular values $s_1 = A^{-\frac{1}{\alpha-1}}$ and $s_2 = A^{\frac{\alpha}{\alpha-1}}$. In particular $s_1^\alpha s_2 = 1$, and so this linear map obeys $S(\alpha)$. The area of Q is $(1/A)\epsilon^2$. There are $A\epsilon^{-2}$ small squares in R, and for each of them we pick a rectangle isometric to Q in U. Because the area of Q is sufficiently small, we can find this many disjoint rectangles in U. Now we define the union of these rectangles to be a set V. As before, we map V onto the union of the open squares in R and the complement of V to the 1-skeleton. This map obeys $S(\alpha)$, and it can be made arbitrarily C^0 close to L by taking ϵ small.

The main result of this appendix is a general C^0-approximation theorem that applies to all open stretching conditions. First we make a few definitions. For an open stretching condition S, define the k-dilation of S, $D_k(S)$ to be the supremal k-dilation of any matrix in S. The number $D_k(S)$ can take any value in $(0, \infty]$. We define the condition $D(S)$ to be the set of matrices M so that the k-dilation of M is less than $D_k(M)$ for every k. The set $D(S)$ is clearly a stretching condition, and $S \subset D(S)$.

Theorem A.1. *If f is any smooth map on the closed unit ball which obeys $cD(S)$, then f can be C^0 approximated by Lipshitz maps f_i, whose derivatives obey S almost everywhere.*

As usual, c denotes a constant that depends only on the dimension n.

The main ingredient of the proof is Gromov's theory of convex integration, which is a general, efficient tool for constructing C^0 approximations. In the first section, we review the relevant part of convex integration. The tools of convex integration reduce the proof to an algebraic problem about the space of matrices, which we carry out in the second section.

1. Review of Convex Integration

Let R be any subset of the space of $n \times n$ matrices. The set R is called principally convex if, whenever M_1 and M_2 belong to R and the rank of $M_1 - M_2$ is equal to 1, then the line segment from M_1 to M_2 also belongs to R. The principal convex hull of a set R is defined to be the intersection of all principally convex sets containing R. The principal convex hull of R is denoted $\mathrm{Conv_P}(\mathrm{R})$.

Theorem. *(Gromov) Let R be a bounded open set in the space of matrices. Let f be a smooth function on the closed ball which obeys the condition* $\mathrm{Conv_P}(\mathrm{R})$. *Then f can be C^0 approximated by Lipshitz functions f_i, whose derivatives almost everywhere obey R.*

This theorem is a simple case of the theorem on page 218 of Gromov's book *Partial Differential Relations* [13].

The idea behind this theorem is the following. Suppose that the rank of $M_1 - M_2$ is equal to 1. We are going to construct a piecewise linear approximation to the linear map $tM_1 + (1-t)M_2$, for a number t between 0 and 1. Since the rank of $M_1 - M_2$ is equal to 1, it follows that there is an (n-1)-plane W in $\mathbb{R}^n$ so that M_1 and M_2 agree on W. We can view W as the vanishing set of a linear function L. Next, we define a piecewise linear function F on $\mathbb{R}^n$, whose derivative is equal to M_1 if the value of L modulo 1 lies in $[0, t]$ and is equal to M_2 if the value of L modulo 1 lies in $[t, 1]$. Because M_1 and M_2 agree on W, these derivatives fit together on the boundaries to give a PL map. (To specify the map F completely, we also set $F(0) = 0$.) We can calculate $F(x)$ by integrating the derivative of F along the line from 0 to x. If $|L(x)|$ is large, then this line spends about $t/1$ of its length in the region where $dF = M_1$ and about $(1-t)/1$ of its length in the region where $dF = M_2$. Therefore, $F(x)$

is approximately equal to $tM_1(x) + (1-t)M_2(x)$. Next we define piecewise linear functions F_ϵ, whose derivative is equal to M_1 if the value of L modulo ϵ lies in $[0, t\epsilon]$ and is equal to M_2 if the value of L modulo ϵ lies in $[t\epsilon, \epsilon]$. As ϵ goes to zero, F_ϵ converges in C^0 to the linear map $tM_1 + (1-t)M_2$.

In order to prove Gromov's theorem, it is necessary to repeat this argument at many locations and on many scales.

2. The Principal Convex Hull of a Stretching Condition

Given Gromov's theorem in the last section, the next step is to compute the principal convex hull of a stretching condition. We are not able to compute it exactly, but we are able to approximate it up to a constant factor.

Proposition A.0.1. *Let S be an open stretching condition. Then the principal convex hull of S contains $cD(S)$. On the other hand, the principal convex hull of S is contained in $D(S)$.*

This proposition is a problem in linear algebra, but the proof is fairly long. We begin by proving the easy direction that the principal convex hull of S is contained in $D(S)$. It suffices to show that the set of matrices with $s_i ... s_n < B$ is principally convex for any number B. If not, pick matrices M_1 and M_2 in this set, with rank $M_1 - M_2 = 1$, and with $M = tM_1 - (1-t)M_2$ not in this set. After changing basis in the domain and the range, we can assume that M is diagonal with $M_{i,i} = s_i(M)$. Therefore, the lower right (n-i+1) by (n-i+1) sub-matrix of M has determinant bigger than B. Call this submatrix M', and similarly, define M_1' and M_2' to be the lower right (n-i+1) by (n-i+1) sub-matrices of M_1 and M_2. By assumption the determinants of M_1' and M_2' are less than B. But the rank of $M_1' - M_2'$ is at most 1, and so $detM' = tdetM_1' + (1-t)detM_2'$. Therefore, the determinant of M' is less than B, and so the set $s_i ... s_n < B$ is principally convex. It follows that the set $D(S)$ is principally convex, and so the principal convex hull of S lies in $D(S)$.

Now we begin the proof in the hard direction. The first step of the proof deals with a special case for the set S.

Lemma A.1. *Suppose that S is equal to the set of matrices with $s_i \leq B_i$, for some* *numbers* $0 < B_1 \leq ... \leq B_n$. *Then the principal convex hull of* S *is equal to* $D(S)$.

Proof. The key to this proof is to break the construction into very small steps.

We define $S(j)$ to be the set of matrices obeying the following inequalities.

1. $s_i \leq B_i$ for i from 1 to j-1.

2. $s_i ... s_n \leq B_i ... B_n$.

The set $S(n)$ is equal to S, and the set $S(1)$ is equal to $D(S)$. Therefore, it suffices to prove that the principal convex hull of $S(j+1)$ includes $S(j)$.

We will prove this inclusion in a series of even smaller steps.

We define $S(j,k)$ to be the set of matrices obeying the following conditions.

1. M is in $S(j)$.

2. For some number l in the range $j < l \leq k$, we have the equality $s_l ... s_n = B_l ... B_n$.

This defines $S(j,k)$ when k is in the range $j+1 \leq k \leq n$. As k increases, the set $S(j,k)$ becomes a larger subset of $S(j)$. We also define $S(j, n+1)$ to be equal to all of $S(j)$.

The set $S(j, j+1)$ actually lies in $S(j+1)$. If M lies in $S(j, j+1)$, then $s_j ... s_n \leq B_j ... B_n$, but $s_{j+1} ... s_n = B_{j+1} ... B_n$. Therefore, $s_j \leq B_j$. Since M lies in $S(j)$, it automatically obeys all the other conditions of $S(j+1)$.

The main inductive step of our proof is to show that the principal convex hull of the union of $S(j+1)$ and $S(j,k)$ includes $S(j,k+1)$. Once we have this inductive step, it follows immediately that the principal convex hull of $S(j+1)$ includes $S(j)$, and therefore that the principal convex hull of S contains $D(S)$. We now prove the inductive step.

Suppose that M lies in $S(j,k+1)$. Since all of our conditions depend only on the singular values of M, we can assume that M is a diagonal matrix with $M_{i,i} = s_i(M)$. We define $M(t)$ to be the matrix that agrees with M except that $M(t)_{jk} = t$. We will prove that for some choice of t, $M(t)$ lies in the union of $S(j+1)$ and $S(j,k)$.

We can assume that M does not itself lie in $S(j+1)$, and therefore $s_j(M) > B_j$. On the other hand, since M does lie in $S(j,k+1)$ which lies in $S(j)$, $s_{j-1}(M) \leq B_{j-1}$.

We can also assume that M does not itself lie in $S(j,k)$. Since M lies in $S(j,k+1)$ but not $S(j,k)$, we know that $s_{k+1} ... s_n = B_{k+1} ... B_n$. Since M does lie in $S(j)$, we know that $s_k ... s_n \leq B_k ... B_n$. Therefore, $s_k \leq B_k$. Since M lies in $S(j)$, we also know that $s_{k+2} ... s_n \leq B_{k+2} ... B_n$. Therefore, $s_{k+1} \geq B_{k+1}$.

Now, we consider the singular values of $M(t)$. As t increases, $s_j(M(t))$ decreases and $s_k(M(t))$ increases, while keeping the product $s_j s_k$ constant. All of the other singular values remain the same. We continue to increase t until either $s_j(M(t)) = B_j$ or $s_k(M(t)) = B_k$.

To make sure that this happens for some value of t, we need to check two things. The first thing is to make sure that $s_k(M(t))$ remains less than or equal to $s_{k+1}(M(t))$. We know that $s_{k+1}(M) \geq B_{k+1} \geq B_k$, so this inequality holds. Similarly, we have to make sure that $s_j(M(t))$ remains greater than or equal to $s_{j-1}(M(t))$. We know that $s_{j-1}(M(t)) \leq B_{j-1} \leq B_j$, so this inequality holds also. The second thing is to check that as we increase t very strongly, the singular value $s_k(M(t))$ will indeed get high enough. This part is obvious, because $s_n(M(t))$ is obviously at least t.

If $s_j(M(t)) \leq B_j$, then $M(t)$ lies in $S(j+1)$. On the other hand if $s_k(M(t)) = B_k$, then $M(t)$ lies in $S(j,k)$. The singular values of $M(-t)$ are the same as the singular values of $M(t)$. Therefore, $M(-t)$ also lies in the union of $S(j+1)$ and $S(j,k)$. The rank of $M(t) - M(-t)$ is equal to 1, so our principal convex hull includes $(1/2)M(t) + (1/2)M(-t)$. But this average is just M. This finishes the proof of the inductive step. Our induction then shows that the principal convex hull of S contains $D(S)$. On the other hand, we have already proven that S is contained in $D(S)$. $\square$

Now that we have proven the lemma, we return to the general case that S is any open stretching condition. We need to prove that the convex hull of S contains $cD(S)$. Let $D(r)$ denote the matrices in $D(S)$ with rank at most r. It suffices to prove the following inductive step.

We will prove that the principal convex hull of S together with $(const)D(r)$ includes $c(const)D(r+1)$.

Let M be a matrix in $D(r+1)$, with singular values $s_i(M)$.

Since M lies in $D(S)$, there is some matrix M_k in S with $s_{n-k+1}(M_k) ... s_n(M_k) \geq s_{n-k+1}(M) ... s_n(M)$. Because S is a stretching condition, it includes any matrix with singular values $s_i \leq s_i(M_k)$. According to our lemma, the principal convex hull of S includes any matrix with singular values obeying $s_i...s_n \leq s_i(M_k)...s_n(M_k)$ for each i. In particular, the principal convex hull of S contains any matrix with $s_1 = ... = s_{n-k} = 0$ and $s_{n-k+1} \leq ... \leq s_n \leq (s_{n-k+1}(M)...s_n(M))^{1/k}$.

Now we define $M(p)$ to be the space of all matrices with singular values obeying the following inequalities.

1. For i from 1 to p, $s_i \leq s_i(M)$.
2. For all other i, $s_i \leq (s_{p+1}...s_n)^{\frac{1}{n-p}}$.

Since the rank of M is equal to $r+1$, the space $M(n-r-2)$ is contained in the principal convex hull of S, according to the results in the last paragraph, taking $k = r+1$. On the other hand, the space $M(n-1)$ contains the matrix M. Using these space M(p), we can break our task into steps. It suffices to prove the following smaller inductive step.

We will prove that the principal convex hull of the union of S, $cD(r)$, and $cM(p)$ contains $3^{-n}cM(p+1)$.

This inductive step, we break into yet smaller pieces. We define $M(p,q)$ to be the set of all matrices whose singular values obey the following conditions.

1. For i from 1 to p, $s_i \leq s_i(M)$.
2. For $i = p+1$, $s_i \leq s_{p+1}^{\frac{q+1}{n-p}}(s_{p+2}...s_n)^{\frac{n-p-1-q}{(n-p)(n-p-1)}}$.
3. For i from $p+2$ to $n-q$, $s_i \leq (s_{p+1}...s_n)^{\frac{1}{n-p}}$.
4. For i from $n-q+1$ to n, $s_i \leq (s_{p+2}...s_n)^{\frac{1}{n-p-1}}$.

These spaces are defined for p from n-r-2 to n-2 and q from 0 to $n-p-1$. The space $M(p,0)$ is equal to $M(p)$, and the space $M(p,n-p-1)$ is equal to $M(p+1)$. Therefore, it suffices to prove an even smaller inductive step.

We will prove that the principal convex hull of the union of S, $cD(r)$ and $cM(p,q)$ contains $(1/3)cM(p,q+1)$.

Let N be a matrix in $M(p,q+1)$. We have to prove that $(1/3)cN$ lies in the principal convex hull of the union of S, $cD(r)$ and $cM(p,q)$. To see this, it suffices to check that N lies in the principal convex hull of the union of S, $3D(r)$ and $3M(p,q)$. We may assume without loss of generality that N does not itself lie in $M(p,q)$, and therefore that $s_{n-q}(N)$ is at least $(s_{p+1}...s_n)^{\frac{1}{n-p}}$.

Since all the spaces involved depend only on the singular values of the matrices involved, we can select a diagonal representative for N, with $N_{i,i}$ equal to $s_i(N)$.

Finally, we will actually do a principal convex interpolation.

We define two matrices $N(1)$ and $N(2)$, by modifying N. The matrix $N(1)$ agrees with N except for three entries. First, $N(1)_{n-q,n-q} = 0$. Second, $N(1)_{p+1,n-q}$ is a number t, on the order of $N_{n-q,n-q}$, which we will choose later. Third, $N(1)_{p+1,p+1}$ is equal to $N_{p+1,p+1}N_{n-q,n-q}(s_{p+1}(M)...s_n(M))^{-\frac{1}{n-p}}$. Because N lies in $M(p, q+1)$, the term $N_{p+1,p+1}$ obeys an inequality written above, which implies that the last expression is bounded by $N_{n-q,n-q}$. The matrix $N(1)$ has rank r. Since N lies in $D(r+1)$ and since the number t will be less than $2N_{n-q,n-q}$, it easy to check that $N(1)$ lies in $3D(r)$.

The matrix $N(2)$ agrees with N except for two entries. First, $N(2)_{n-q,n-q}$ is equal to $2(s_{p+1}...s_n)^{\frac{1}{n-p}}$. Second, $N(2)_{p+1,p+1}$ is equal to $N(1)_{p+1,p+1}$. Since N lies in $M(p, q+1)$, a simple but boring calculation shows that $N(2)$ lies in $2M(p, q)$.

The rank of $N(1)-N(2)$ is equal to 1. Therefore, the average $(1/2)N(1)+1/2N(2)$ is contained in the principal convex hull of the union of $3D(r)$ and $2M(p, q)$.

We claim that for an appropriate choice of t, less than $2N_{n-q,n-q}$, the singular values of this average will be equal to the singular values of N. This average matrix has only one off-diagonal term. We should think of it as a 2×2 submatrix at the $p+1$ and $n-q$ coordinates, plus a diagonal matrix. The singular values of the diagonal matrix are just the (n-2) singular values of N, leaving out $s_{p+1}(N)$ and $s_{n-q}(N)$. Therefore, it suffices to understand the singular values of the 2×2 submatrix. The sub-matrix is upper triangular. Its determinant is equal to $s_{p+1}(N)s_{n-q}(N)$. When t=0, the submatrix is diagonal. Its singular values are the diagonal terms. Their product is $s_{p+1}(N)s_{n-q}(N)$, but they are closer together than the singular values of N, because we assumed that $s_{n-q}(N)$ is at least $(s_{n-p+1}...s_n)^{\frac{1}{n-p}}$, which is the larger singular value of the diagonal submatrix. The upper right-hand term of the sub-matrix is $t/2$. As t increases, its singular values move farther apart, maintaining the same product. When $t/2$ reaches $s_{n-q}(N)$, then the larger singular value will be at least $s_{n-q}(N)$. Therefore, for some intermediate value of t, the singular values of the average are equal to those of N.

Since all the spaces involved depend only on the singular values, the matrix N itself is also in the principal convex hull of the union of S, $3D(r)$ and $3M(p,q)$. This statement completes the inductive step of our proof. Applying the induction, we see that the principal convex hull of S contains $cD(S)$. This finishes the proof of the proposition.

Given Gromov's theorem and this tedious calculation of principal convex hulls, the proof of our main theorem is very quick. Suppose that S is a bounded open stretching conditions, and that f is a smooth map on the unit ball that obeys $cD(S)$. According to our last calculations, $cD(S)$ lies in the principal convex hull of S. Therefore, according to Gromov's theorem, the map f can be approximated in C^0 by Lipshitz maps f_i that obey S almost everywhere. Now suppose that S is only an open stretching condition. Suppose that f is a smooth map on the closed unit ball that obeys $cD(S)$. Since the closed ball is compact, f actually obeys a compact subset of $cD(S)$. Therefore, f obeys $cD(S')$ for some bounded open subset S' of S which is also a stretching condition. According to Gromov's theorem, the map f can be approximated in C^0 by Lipshitz maps f_i that obey S' almost everywhere. Since S' is a subset of S, the maps f_i obey S almost everywhere.

3. The S-Dilation of Diffeomorphisms

We have had in mind a generalization of our first problem about k-contracting maps to other stretching conditions.

Let S be an open stretching condition. Since S contains an open ball around 0, any matrix lies in λS for some sufficiently large number λ. We define the S-dilation of a matrix M to be the infimal number λ so that M lies in λS. For example, if S is the set of k-contracting matrices, then the S-dilation of M is equal to the k-dilation of M raised to the power $(1/k)$. The S-dilation of a smooth map f is defined to be the supremal value of the S-dilation of df at any point in the domain of f.

Problem. *Let R and R' be n-dimensional rectangles, and let S be a stretching condition. Estimate the smallest S-dilation of any diffeomorphism from R to R'.*

Since S is contained in $D(S)$, the S-dilation of any map is at least as great as its $D(S)$-dilation. On the other hand, our theorem provides some evidence for the following conjecture.

Conjecture. *If there is a diffeomorphism from R to R' with $D(S)$-dilation B, then there is a nearby diffeomorphism from R to R' with S-dilation less than CB.*

Our theorem constructs a degree 1 Lipshitz map from R to R' (taking the boundary of R to the boundary of R'), with S-dilation less than CB almost everywhere.

If the conjecture is true, then, up to a constant factor, our problem reduces to the special case that $S = D(S)$. In other words, S is given by the condition k-dilation$(M) < \Lambda_k$ for some sequence of numbers Λ_k. These considerations highlight the following special case of our problem.

Problem. *Given n-dimensional rectangles R and R', estimate the largest number λ so that there is a diffeomorphism from R to $\lambda R'$ with k-dilation less than Λ_k for each k.*

Appendix B

Literature on Area-Contracting Maps

In this appendix, we will survey all of the results that I know in the mathematical literature which relate to area-contracting maps. I would be very interested to learn of more. One goal of this appendix is to give some context for the results in the thesis.

All of the results we have mentioned in the thesis come from remarks in two papers by Gromov.

In appendix 1 of the paper "Filling Riemannian Manifolds" [12], Gromov discusses the k-width. Using Almgren's Morse theory (see B below), he gives the sharp value for the k-width of the unit n-sphere, and also estimates the k-width of (M, g) in terms of sectional curvature and injectivity radius. Then he gives his elementary estimate for the k-width of the unit sphere, by repeatedly using the isoperimetric inequality. (We repeat this argument in Proposition 3.1.1.) This argument shows in particular that the unit ball in the n-dimensional space $l^\infty(n)$ has k-width at least $c(n)$. Gromov raised the problem whether $c(n)$ is bounded below as n goes to infinity (for a fixed k). Such a bound would imply that every closed Riemannian k-manifold with the volume of every 1-ball less than $\epsilon(k)$ has filling radius less than 1. (I believe this problem is still open.)

There is a scattered discussion of k-dilation in the long essay "Carnot-Caratheodory spaces as seen from within" [10]. This essay is mostly about generalizations of

Riemannian geometry, but certain results about Riemannian geometry appear in it as starting points for investigations on Carnot-Caratheodory spaces. In particular, the essay includes several results about the homotopy types of maps with small k-dilations. First Gromov gives an estimate for the rational homotopy invariants of a map with k-dilation less than Λ. He goes on to show that rational homotopy invariants can be controlled by the L^q norm of $|\Lambda^k df|$ for appropriate choices of q. Even the L^1 norm of $|\Lambda^k df|$ has an effect on homotopy. Namely, Gromov shows that if the L^1 norm of $|\Lambda^k df|$ is sufficiently small, then f can be homotoped to a map that takes the k-skeleton of the domain to the (k-1)-skeleton of the range. If $k = 2$, this homotopy shows that the map $\pi_1(M) \to \pi_1(N)$ factors through a free group.

Finally, Gromov proves that a map from a compact simply connected manifold with sufficiently small 2-dilation is null-homotopic. A sketch of the proof appears on pages 229-230 of [10]. (The proof is not very detailed. Here is my understanding of the main idea. For example, suppose we have a map f from S^n to a complex X in $\mathbb{R}^N$ with very small 2-dilation. View S^n as a family of 2-spheres, parameterized by the (n-2)-ball. Each 2-sphere is mapped to a 2-sphere of small area. By Riemann mapping theorem, we can change coordinates so that the mapping on each 2-sphere has bounded Dirichlet energy $\int_{S^2} |df|^2$. By the Sobolev inequality, the map on each 2-sphere then has small BMO norm. We homotope the map by convolving it with a smoothing kernel, so that at the end of the homotopy it maps each 2-sphere to a point. At any stage in the middle of the homotopy, a given point is mapped to a weighted average of the value of f on a ball around that point in S^2. Since our map has small BMO norm, it actually comes close to taking that average value, and so the average value lies inside of a small neighborhood of X. If the 2-dilation of f is sufficiently small, this neighborhood retracts to X. Therefore, the map from S^n to X is homotopic to a map which collapses each 2-sphere in our family to a point. Hence the map from S^n factors through the map to the (n-2)-ball and is null-homotopic.)

Area-contracting maps also appear very briefly in Gromov's book *Partial Differential Relations* [13]. In an exercise, Gromov states that the k-contracting condition is C^0-closed. This means that if f_i are C^1 k-contracting maps that converge in C^0 to a C^1 map f, then f is also k-contracting. (Since Gromov does not include the proof, we provide it here. If f is not k-contracting, then its derivative at a certain point p is not k-contracting. Therefore, it maps a small (k-1)-dimensional sphere, centered at p, bounding a k-ball of volume V, to a (k-1)-dimensional ellipsoid with filling volume more than V. A very good C^0 approximation of f maps that (k-1)-dimensional sphere to a cycle that lies very close to the ellipsoid, and so also has filling volume more than V. Therefore, a good C^0 approximation of f is not k-contracting.)

To my knowledge, these are the only results in the literature that deal with area-contracting maps for their own sake. I have found a number of other papers which relate area-contracting maps to other topics. Usually, these papers do not mention area-contracting maps by name. They are listed here, with short descriptions, according to topic.

A. Scalar curvature

Gromov and Lawson mention area-contracting maps in their famous work on scalar curvature [11]. In particular, they prove the following estimate, using the Atiyah-Singer index theorem.

Theorem. *(Gromov and Lawson) If (M, g) is a complete spin Riemannian n-manifold with scalar curvature greater than a constant C, then any compactly supported 2-contracting map from M to the unit n-sphere has degree zero.*

As far as I know, the paper by Gromov and Lawson is the first paper to refer to area-contracting maps. In section 5 of [11], they prove the above theorem for 1-contracting maps. In section 6, they generalize it to get the result above for 2-contracting maps. The generalization to 2-contracting maps allows them to understand positive scalar curvature on certain non-compact manifolds. For instance, they prove that the manifold $T^{n-1} \times \mathbb{R}$ does not admit a complete metric with positive scalar curvature.

B. Families of cycles

The sharpest estimates for k-width come from Almgren's work on Morse theory on the space of cycles. Almgren studied the space of k-cycles on a Riemannian manifold, equipped with the topology coming from the flat norm. The k-volume is a function on the space of k-cycles, and its critical points are minimal surfaces. There is a unique empty cycle with volume 0, and every other cycle has positive volume. If the k-volume were a smooth function V on a finite-dimensional manifold, with a unique minimum with value 0, and if C were the first non-zero critical value of V, then the set $V^{-1}((0,C))$ would be contractible. Almgren proved an analog of this result on the space of cycles.

Theorem. *(Almgren) Let f be a continuous map from S^m into the space of k-cycles on a Riemannian manifold M. Let V_0 be the smallest k-volume of a non-trivial minimal k-cycle in M. If the image of f consists of k-cycles with volume less than V_0, then f is contractible.*

(I have had some trouble finding a reference for this theorem. I believe that it appears in Almgren's paper "The theory of varifolds - a variational calculus in the large for the k-dimensional area integrated". Almgren's paper was never published, and I have not been able to find a copy. According to Jean Taylor's introduction to [1], there is a copy in the Mathematics and Physics Library at Princeton University.)

In an earlier paper [2], Almgren determined the homotopy groups of the spaces of cycles.

Theorem. *(Almgren) Let $C(k)$ denote the space of integral k-cycles in a Riemannian manifold M. Then $\pi_i(C(k)) = H_{i+k}(M, \mathbb{Z})$.*

I think that Almgren's work provides the most natural definition of k-width. If we pick an orientation for M, we can define α to be the homotopy class in $\pi_{n-k}(C(k))$ which corresponds to the generator of $H_n(M, \mathbb{Z})$. Then, we define the Almgren k-width of M to be less than W if there exists a map f from S^{n-k} to $C(k)$ in the class α so that every k-cycle in the image of f has k-volume less than W. (The Almgren k-width does not depend on the orientation that we chose.) The Almgren k-width is the smallest k-volume of a family of k-cycles which sweep out (M, g) with degree 1. Because the fibers of a generic PL map form a family of flat cycles in this homotopy class, the Almgren k-width is less than or equal to the k-width we have used in this thesis.

According to the first theorem we cited above, the Almgren k-width of the unit n-sphere is at least the smallest k-volume of a minimal surface in the unit n-sphere. This smallest volume turns out to be the volume of the unit k-sphere. On the other hand, if we take the fibers of a linear projection from S^n to $\mathbb{R}^{n-k}$, they form a family of flat cycles in the homotopy class α, each with k-volume at most the k-volume of the unit k-sphere. Therefore, the Almgren k-width of the unit n-sphere (and also our k-width of the unit n-sphere) is exactly equal to the volume of the unit k-sphere.

C. Magnetic relaxation

The first example of an interesting area-contracting map was given by Zel'dovich and Gehring, working independently on quite different problems, around 1970. They each found an open set U inside an arbitrarily small ball in $\mathbb{R}^3$ which admits an area-contracting diffeomorphism onto $\mathbb{D}^2(1) \times S^1(\delta)$, for some very small number δ. Let us explain the construction. Begin with ϵ^{-2} ordinary circular tubes, of length ϵ and cross-sectional area on the order of ϵ^2. All together the tubes have volume on the order of ϵ, and so we can fit them inside a ball of radius $\epsilon^{1/3}$. Let c be a curve through this ball, which passes once through each tube, as though the tubes were beads which we were stringing together. The set U is simply the ball with a small neighborhood of the curve c cut out. The area-contracting diffeomorphism is constructed as follows. We map each small tube quasi-isometrically to a tube in the image of the form $B^2(\epsilon) \times S^1(\epsilon)$. Then we extend the diffeomorphism in any way to the rest of U. After that, we squeeze the image tubes closer together and squeeze in the S^1-direction as much as we like. The complement of the tubes is mapped to a region which is essentially 1-dimensional, and so the map is area-contracting.

(A more careful analysis shows that we can take $\delta > c\epsilon^2$. The separation between the beads is on the order of ϵ, and so the Lipshitz constant of the map we have constructed to the disk need be no more than ϵ^{-1}, which occurs when two adjacent beads are mapped to opposite ends of the disk. Therefore, after squeezing the beads together as much as we like, we only need to shrink the S^1-direction by a factor of ϵ. This estimate is sharp up to a constant factor. Any circle near the center of $D^2(1) \times S^1(\delta)$ has filling area on the order of δ and so its image must have length at least on the order of $\delta^{1/2}$ in $\mathbb{R}^3$. Since the area transverse to the circles increases, the set U must have volume at least $c\delta^{1/2}$. But U lies in a ball with volume ϵ. Therefore, δ is at most $C\epsilon^2$.)

The Russian physicist Zel'dovich discovered this construction while he and Sakharov were studying the magnetic fields of neutron stars. A neutron star is a liquid made from a material which carries a magnetic field. The magnetic field is naturally described by a closed 2-form B. If the fluid in the star flows from one position to another according to a diffeomorphism ϕ, then the magnetic field transforms according to the usual pushforward for a 2-form. (This last rule is only an approximation, which disregards a friction effect.) The energy of the magnetic field is the integral of $|B|^2$. Sakharov and Zel'dovich wondered whether the fluid can flow into another position so as to reduce the energy of the magnetic field, which would be released in the form of heat. This process is called magnetic relaxation.

If the initial magnetic field is given by the equation $B(x, y, z) = (-y, x, 0)$, so that the field lines simply circle around the z-axis, then Zel'dovich found volume-preserving diffeomorphisms that decrease the energy of the magnetic field as close as one likes to zero. Zel'dovich's map takes the complement of the z-axis to the set U described above (rescaled to have volume 1). Each small bead in U has, as its preimage, a long thin tube running parallel to the magnetic field. Zel'dovich's map stretches the long thin tube perpendicular to the magnetic field lines of B while shrinking it parallel to the magnetic field lines. Therefore, the resulting magnetic field in U is very small on each bead. The preimages of the beads in U can be chosen to fill almost all of the volume of the star. With a little care, one can control the energy of the magnetic field on the complement of the beads.

Zel'dovich raised the problem whether a star can dissipate all of its magnetic energy in this way for other initial states of the magnetic field. He explained the problem to Arnold. Arnold wrote a paper [3], which contains a good exposition of the problem. In that paper, he also proved an estimate for a slightly idealized version of the problem, in which the magnetic fluid fills a homology 3-sphere. He defined the generalized Hopf invariant of the magnetic field B to be the integral $\int_M B \wedge PB$, where PB denotes any primitive of the exact 2-form B. Arnold explained that this integral measures the average linking of the flow lines of B. If λ_1 denotes the first eigenvalue of the Laplacian acting from coclosed 1-forms to 2-forms, then the Hopf invariant of B is at most $\lambda_1^{-1}|B|_2^2$. In other words, the energy of B is at least $\lambda_1|\mathrm{Hopf}(B)|$. Since the Hopf invariant of B is diffeomorphism invariant, a magnetic fluid with non-zero Hopf invariant cannot release all of its energy by magnetic relaxation.

D. Quasi-conformal geometry

Gehring discovered the set U while working on conformal geometry. If g is a metric on $D^2 \times S^1$, then Gehring defined a function $m_1(g)$ by the following formula.

$$m_1(g) = \mathrm{inf}_\rho \mathrm{sup}_\gamma (\int_{D^2 \times S^1} \rho^3 dvol)(\int_\gamma \rho)^{-3}.$$

In this formula, ρ represents any positive function, and γ represents any circle homotopic to S^1. The resulting number $m_1(g)$ is a conformal invariant, because a conformal change in the metric g by a factor λ is equivalent to multiplying ρ by the function $\lambda^{1/2}$. For example, if g is the product metric of a disk with area A and a circle with length L, then $m_1(g)$ is equal to AL^{-2}.

Gehring studied the possible values of $m_1(U)$ for a solid torus embedded in $\mathbb{R}^3$. It is easy to find solid tori with arbitrarily small m_1 by taking long thin tubes. It is less obvious how to find a solid torus with very large m_1. The set U described in the last section does the job. Recall that we constructed a subset U contained in a ball of volume ϵ, equipped with an area-contracting diffeomorphism Ψ to $D^2(1) \times S^1(\delta)$. We select γ to be one of the curves $\Psi^{-1}(\{x\} \times S^1)$. If we choose γ randomly with respect to $darea(x)$, the average value of $\int_\gamma \rho$ will be less than $(1/pi) \int_U \rho$. Therefore, the value of $m_1(U)$ is at least $inf_\rho(\int \rho^3 dvol)(\int \rho dvol)^{-3}$. This quantity is at least $(\mathrm{volume}(U))^{-2}$, which is greater than ϵ^{-2}. Therefore, a solid torus in $\mathbb{R}^3$ can take an arbitrarily large value of m_1. By a continuity argument, a solid torus in $\mathbb{R}^3$ can have m_1 equal to any value in $(0, \infty)$.

In the paper [8], Gehring then proved that if U and V are two linked solid tori, then one of them must have m_1 at least c. As a special case, it follows that if U and V are linked and have small volume, then they do not both admit area-contracting diffeomorphisms to $D^2(1) \times S^1(\delta)$.

E. Knotted flow lines

In the paper [7], Freedman and He improved Arnold's lower bounds for the energy of a magnetic field. Let B be a magnetic field on a domain in $\mathbb{R}^3$, which we can think of as either a divergence-free vector field or a closed 2-form. Suppose that B leaves invariant a solid torus U. (By this I mean that the flow of the vector field leaves U invariant. Equivalently, we can say that the restriction of the 2-form to the boundary of U is zero.) Suppose that the flux of the vector field through the solid torus is equal to F. (By this I mean that the integral of the 2-form B over a fundamental relative 2-cycle of the solid torus is equal to F.) It the tube U is unknotted, then the energy of the magnetic field may be arbitrarily small. If the tube U is tied in a knot with genus G, then Freedman and He proved the following estimate about the magnetic field.

Theorem. *(Freedman and He) Let B be a divergence free vector field on a domain D in $\mathbb{R}^3$. Suppose that B leaves invariant a tube U, which is tied in a knot of genus G, and suppose that the flux of B across this tube is equal to F. The the following inequality holds.*

$$\int_D |B|^{3/2} \geq (\frac{16}{\pi})^{1/4} |F|^{3/2} (2G-1)^{3/4}.$$

As a very special case, we can suppose that there is an area-contracting diffeomorphism from U to the product $D^2(1) \times S^1(\delta)$. We can define B to the pullback of the area 2-form on $D^2(1)$, and we can set the domain D to be equal to U. The above inequality shows that the volume of U is at least $2\pi^{5/4}(2G-1)^{3/4}$. (Freedman and He also gave estimates for the conformal modulus $m_1(U)$, defined in the last section.)

F. Networks in three dimensional space

In the early 60's, Kolmogorov and Barzdin wrote a geometry paper [17], which originated in trying to understand one feature of the structure of the brain. The neurons on the brain lie in a thin layer along its outer edge (the grey matter), while the inside of the brain is taken up by axons joining them. This arrangement allows many fewer neurons than if the neurons were placed throughout the whole brain. It is a challenge to explain this small number of neurons.

They modeled the brain as a graph, with each neuron corresponding to a vertex and each axon corresponding to an edge. The graph has bounded valence. Realizing the graph in the brain means finding a thick embedding in three space, where each vertex corresponds to a ball of radius 1 and each edge to the 1-neighborhood of a curve. They proved that every graph of valence 4 and cardinality N can be embedded into a cube of side length $C\sqrt{N}$. On the other hand, they proved that with high probability, a random graph cannot be embedded into a cube of side length less than $c\sqrt{N}$.

The embedding of any graph into a cube of side length $C\sqrt{N}$ is achieved by putting the vertices of the graph along the boundary of the cube and all the edges of the graph on the inside of the cube. For most graphs, this is a fairly efficient arrangement.

The interesting point is to explain why most graphs require so much more space than cubical lattices. In modern terminology, the key fact from graph theory is that most graphs are $1/100$ - expanders. A graph is called a $1/100$ - expander if it obeys the following isoperimetric property: for each subset A of the vertices which contains less than half of all vertices, the number of edges joining A to its complement is at least $|A|/100$. Now the cube of side length S can be swept out by surfaces of area S^2, and one of these surfaces must divide the vertices of Γ into two equal halves. (Technical remark: any surface can easily be homotoped, without much increasing its area, to avoid all the 1-balls corresponding to vertices of Γ.) Such a surface must cut $N/200$ edges, meeting each in a surface with area π, and so the total area $S^2 > \pi N/200$. Therefore the side length of the cube is at least $\sqrt{N}/20$.

This argument gives a lower bound for the size of a cube containing a thick expanding graph. Kolmogorov and Barzdin also proved a lower bound for the volume of the graph, on the order of $N^{3/2}$. The argument in the preceding paragraph really shows that the width of the embedded network is at least cN. According to the width volume inequality, the volume of the embedded network is at least $cN^{3/2}$. (Kolmogorov and Barzdin have a different proof of the volume bound. I thought of the width-volume inequality while I was trying to understand their proof.)

The published paper includes the theorems on embedding networks in Euclidean space with no mention of the biological problem. (Although Barzdin mentions this problem in the notes accompanying the paper in Kolmogorov's collected works.) The idea about the structure of the brain seems very interesting to me, but it is not clear that the argument above explains the small number of neurons in the brain in a convincing way. The brain need not be a random graph, especially since random graphs fit poorly into three-dimensional space. Here is a question that would clarify things. Let F be a random boolean function from $\{0,1\}^m$ to $\{0,1\}^n$, for large numbers m and n. We try to find a circuit with standard gates which computes the function F along with an embedding of a thickened version of this circuit into three-dimensional space, so that it fits into a ball which is as small as possible. One technique is to find a circuit with a minimal number of gates, without paying attention to the geometry of the graph, and then to embed the circuit with all gates on the boundary of a ball and the wires on the inside. It would be interesting to know whether a circuit can be found that substantially beats this technique, for random functions F. If this technique turns out to be (roughly) optimal, it would reinforce Kolmogorov's point of view about the brain.

Appendix C

Open Questions

Since so little is known about area-contracting maps, there are naturally many open questions. I have picked out four questions that seem important to me.

1. The sponge problem

If U is an open set in $\mathbb{R}^n$ with very small volume, say less than $\epsilon(n)$, then is there a 1-expanding embedding of U into the unit ball?

The width-volume inequality gives a very weak result in this direction. With a very small extra effort, the width-volume inequality shows that for each k there is a k-expanding embedding of U into $B^k \times \mathbb{R}^{n-k}$, where B^k denotes the unit ball in $\mathbb{R}^k$. The sponge conjecture asks for an embedding obeying a stronger condition into a smaller region.

This question is related to estimates for the Lipshitz constant of a diffeomorphism from the unit n-sphere to (S^n, g). Suppose that the target sphere contains disjoint metric balls $B(x_i, R_i)$. Then any diffeomorphism from the unit n-sphere to the target sphere must have Lipshitz constant at least $c(\sum R_i^n)^{1/n}$. Let L be largest lower bound obtained from this estimate after considering all sets of disjoint balls in (S^n, g). Then is there a diffeomorphism from the unit sphere to the target sphere with Lipshitz constant less than CL? An affirmative answer to this question would imply an affirmative answer to the sponge problem, as well as other corollaries.

2. The extension problem for k-contracting maps

Given a map F from the unit sphere in $\mathbb{R}^n$ to $\mathbb{R}^N$ with k-dilation 1, can F be extended to a map from the unit ball with k-dilation less than some constant C?

This question is a special case of a problem that Gromov raised in [10], on page 220.

When k is equal to 1, the answer is trivially yes, and the sharpest constant C is also known by the Kirszbraun theorem.

Theorem. *(Kirszbraun) Let S be a subset of $\mathbb{R}^n$, and let f be a 1-Lipshitz map from S to $\mathbb{R}^N$ (using the extrinsic metric on S). Then S extends to a 1-Lipshitz map from $\mathbb{R}^n$.*

This theorem is explained in [15], at the end of Chapter 1. In particular, when k=1, our question has an affirmative answer and the best constant C is equal to $\pi/2$.

The Kirszbraun theorem can be interpreted as follows. If the map f can be extended to any single curve with endpoints in S, then it can be extended to all of $\mathbb{R}^n$. A possible generalization of the Kirszbraun theorem to $k > 1$ would be the following.

Question: Let S be a compact submanifold of $\mathbb{R}^n$, and let f be a map from S to $\mathbb{R}^N$. Suppose that for any k-chain C with boundary in S, there is a map F from C to $\mathbb{R}^N$, which agrees with f on the boundary of C, and with the property that the k-volume of F(C) is less than the k-volume of C. Does it follow that f extends to a k-contracting map from $\mathbb{R}^n$ to $\mathbb{R}^N$?

An affirmative answer to this question easily implies that any k-contracting map from the boundary of the unit ball extends to a map on the unit ball with k-dilation less than C.

3. Variations of the isoperimetric inequality

In this thesis, we have touched on many ways of describing the size of domains in Euclidean space, or more generally of Riemannian manifolds. I am very curious to know if there are isoperimetric inequalities involving some of these sizes. Here are some examples. Throughout, we assume that U is a bounded open set in $\mathbb{R}^n$ with smooth boundary.

a. Suppose that there is a k-contracting diffeomorphism from U to the unit ball. Is there a (k-1)-contracting diffeomorphism from the boundary of U to the unit sphere?

(We can also consider degree 1 maps in place of diffeomorphisms.)

I have proven a weak result in this direction, [16]. In 3-dimensional space, if there is a 2-contracting diffeomorphism from U to the unit ball, then there is a degree 1 map from the boundary of U to the unit sphere with Lipshitz constant less than 400.

What if there is a k-contracting diffeomorphism from U to some rectangle R?

b. The extension problem can be seen as a kind of isoperimetric inequality. Suppose that there is a k-contracting diffeomorphism from the unit sphere to the boundary of U. Is there a diffeomorphism from the unit ball to U with k-dilation less than C?

c. In Chapter 4, we proved that the boundary of U contains disjoint sets S_i, so that the volume of U is less than $C \sum$ (n-2)-width$)(S_i)^{n/n-2}$. Suppose that U contains disjoint sets U_i with $\sum$k-width$(U_i)^q$ greater than 1. What can we conclude about the boundary of U? Suppose that U has k-width at least 1 around some simplicial complex X. (The k-width of an open set around a complex is defined in Chapter 6.) What can we conclude about the boundary of U?

d. Suppose that there is a k-contracting map from $(U, \partial U)$ to a wedge of spheres X with a large rational homotopy invariant. What can we conclude about the boundary of U?

We wish to compare the size of an n-dimensional object to the size of an (n-1)-dimensional object. If the objects had the same dimension, then it would be possible to compare them more "directly", for instance, by mapping one onto the other. If U were a bounded open subset of Hilbert space, then the boundary of U would also be infinite dimensional, and we could ask if there is a distance-expanding map from U into its own boundary. In finite dimensions this is impossible. Still, we can ask for some geometric control of a map from U into its boundary, as in the following question.

e. Let $V(U)$ denote the volume of U. Is there a map F from U to the boundary of U so that, for each p-dimensional manifold M with p-volume V(M) in the boundary of U, the (p+1)-volume of $F^{-1}(M)$ is less than $CV(U)^{1/n}V(M)$?

f. Can we find any Sobolev inequalities in the same spirit, bounding some size of a compactly supported C^1 function f by some other size of df?

4. A filtration of homotopy groups

We can use the k-dilation to define a filtration on the homotopy groups of a simplicial complex X. Let X be a finite simplicial complex. Put the standard metric on each simplex of X. (But don't worry: any metric on X would be bilipshitz to this one and define the same filtration.) Next, we say that a class $a \in \pi_n(X)$ lies in $V_k\pi_n(X)$ if it is represented by pointed maps f_i from $\mathbb{S}^n$ to X with k-dilation of f_i tending to zero.

These spaces form a filtration: $0 = V_1\pi_n \subset V_2\pi_n \subset ... \subset V_n\pi_n \subset \pi_n$.

Here is what I know about this filtration.

a. The identity class in $\pi_n(S^n)$ is not in $V_n\pi_n$, because the volume of the image of any map in a non-trivial homotopy class is the whole target sphere.

b. More generally, Gromov showed that any rational homotopy invariant defined using forms of degree at least k vanishes on V_k. In particular, the homotopy class of the Hopf fibration does not lie in $V_2\pi_3(S^2)$.

c. If a homotopy class factors through a space Y of dimension m, then that homotopy class lies in V_{m+1}.

d. The suspension operation maps $V_k\pi_n(X)$ to $V_{k+1}\pi_{n+1}(SX)$.

e. Gromov proved in [10], on pages 229-230, that $V_2\pi_n(X) = 0$ for any complex X and any n greater than 1.

f. In section 7.9, we showed that high suspensions admit maps with small k-dilation. For example, the homotopy class of the suspension of the Hopf fibration lies in $V_3\pi_4(S^3)$.

More generally, we proved that if a is a homotopy class in $\pi_n(S^m)$, then the p-fold suspension of a lies in $V_k\pi_{n+p}(S^{m+p})$ for any $k > (m/n)p + n$. According to a fairly deep result in homotopy theory, $\pi_n(S^2)$ contains stably homotopically non-trivial elements for infinitely many n. Therefore, for each N, $V_3\pi_M(S^N)$ is non-empty for infinitely many M.

The 100-fold suspension of the Hopf map defines a class in $\pi_{103}(S^{102})$. According to f, this class lies in V_{69}. According to e, it does not lie in V_2. What is the smallest k so that V_k contains this class?

It would be interesting to know if this filtration coincides with a filtration defined purely in terms of homotopy theory.

How does the filtration interact with the long exact sequence of a fibration? If F is the fiber of the fibration and B is the base, is it true that the boundary map from $\pi_n(B)$ to $\pi_{n-1}(F)$ takes $V_k\pi_n(B)$ to $V_k\pi_n(F)$?

Bibliography

[1] Frederick J. Almgren. *Plateau's Problem, revised edition.* American Mathematical Society, United States, 2001.

[2] Frederick Justin Almgren. The homotopy groups of the integral cycle groups. *Topology*, 1:257–299, 1962.

[3] V. I. Arnold. The asymptotic hopf invariant and its applications. *Sel. Math. Sov.*, 5:327–45, 1986.

[4] Michel Boileau and Shicheng Wang. Non-zero degree maps and surface bundles over s^1. *Journal of Differential Geometry*, 43(4):789–806, 1996.

[5] E. B. Curtis. Some non-zero homotopy groups of spheres. *Bulletin of the American Mathematical Society*, 75:541–544, 1969.

[6] Donald M. Davis and Mark Mahowald. The image of the stable j-homomorphism. *Topology*, 28(1):39–58, 1989.

[7] Michael Freedman and Zheng-Xu He. Divergence-free vector fields: energy and asymptotic crossing number. *Annals of Math*, 134(1):189–229, 1991.

[8] F. Gehring. Inequalities for condensers, hyperbolic capacity, and extremal metrics. *Michigan Math. Journal*, 18:1–20, 1971.

[9] Philip Griffiths and John Morgan. *Rational Homotopy Theory and Differential Forms*, volume 16 of *Progress in Mathematics*. Birkhauser, Boston, 1981.

[10] Mikhael Gromov. Carnot-caratheodory spaces as seen from within. In Andre Bellaiche and Jean-Jacques Risler, editors, *Sub-Riemannian Geometry*, number 144 in Progress in Mathematics, pages 85–325. Birkhauser, Basel, 1996.

[11] Mikhael Gromov and Blaine Lawson. Positive scalar curvature and the dirac operator on complete riemannian manifolds. *Publ. Math. I.H.E.S.*, 58:83–196, 1983.

[12] Mikhail Gromov. Filling riemannian manifolds. *Journal of Differential Geometry*, 18:1–147, 1983.

[13] Mikhail Gromov. *Partial Differential Relations*. Springer-Verlag, Berlin, 1986.

[14] Mikhail Gromov. Width and related invariants of riemannian manifolds. *Asterisque*, 163:93–109, 1988.

[15] Misha Gromov. *Metric Structures on Riemannian and Non-Riemannian Spaces*, volume 152 of *Progress in Mathematics*. Birkhauser, Boston, 1999.

[16] Larry Guth. Lipshitz maps from surfaces. to appear in Geometry and Functional Analysis.

[17] A. N. Kolmogorov and Ya. M. Barzdin. On the realization of networks in three-dimensional space. In A. N. Shiryayev, editor, *Selected Works of Kolmogorov, Volume 3*. Kluwer Academic Publishers, Dordrecht, 1993.

[18] John Milnor and William Thurston. Characteristic numbers of 3-manifolds. *Enseignement Math.*, 23:249–54, 1977.

[19] Teruhiko Soma. Non-zero degree maps to hyperbolic 3-manifolds. *Journal of Differential Geometry*, 49(3):517–46, 1998.

[20] Hassler Whitney. *Geometric Integration Theory*, volume 21 of *Princeton Mathematical Series*. Princeton University Press, Princeton, New Jersey, 1957.