THE UNIVERSITY OF CHICAGO

DISCRETE METRIC SPACES AND COARSE CHARACTERISTIC CLASSES

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF MATHEMATICS

BY

KEVIN WHYTE

CHICAGO, ILLINOIS

JUNE 1998

To my family

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

I would like to thank my advisor, Shmuel Weinberger, to whom I am greatly indebted for the knowledge and guidance he gave me.

I am indebted to my colleagues Jerome Benveniste, Bryan Clair, Alex Gurevich, Tom Haines, Eugenie Hunsicker, Heather Johnston, Anastasios Politis, Todd Rowland, and Jennifer Taback for their help and friendship.

I also owe a great deal to Jonathon Block, Kevin Corlette, Alex Eskin, Benson Farb, Steve Gersten, Bruce Kleiner and Mel Rothenberg for many useful discussions and insights.

Finally I would like to thanks Paul Halmos without whose encouragement I might not ever have entered mathematics.

CHAPTER 1

INTRODUCTION AND STATEMENT OF RESULTS

The geometry of discrete metric spaces has been a very active research area recently. The motivation comes, primarily, from two (not completely separate) areas of mathematics. One is the study of non-compact manifolds equipped with metrics well defined up to uniformly bounded distortion. These arise naturally in the study of compact manifolds; for example as universal covers or leaves of foliations. To focus attention purely on the large scale structure, the local topology is thrown away by passing to a discrete subset, called a net, which evenly fills out the space. A typical example is $\mathbb{Z}^n$ in $\mathbb{R}^n$. Two nets in a space are bilipschitz equivalent on the large scale, but not the small. This leads to the notion of a quasi-isometry of metric spaces, and gives a functor from manifolds of bounded geometry, up to bounded distortion, to discrete metric spaces up to quasi- isometry.

The second source of motivation and examples is geometric group theory. A finitely generated group can be given a metric by measuring the distance from an element to the identity by the minimal length of a word in the generators which representing the element. This extends to a metric by translation invariance. This metric, called the word metric, is well-defined up to bilipschitz equivalence (this can be seen simply by expanding one set of generators in another). The geometric group theorist studies the group via the geometry of this space.

The coarse geometry of manifolds and geometric group theory are related, most obviously by the fundamental group. If M is a compact manifold then $\bar{M}$ has obvious nets given by the orbits of the action of $\pi_1(M)$. These nets are canonically identified, at least as sets, with $\pi_1(M)$. Conversely, if Γ is any finitely presented group there is a compact manifold with $\pi_1 = \Gamma$. It is a fundamental observation in the subject that the quasi-isometry type of Γ as a net agrees with the word metric. More generally, the Milnor-Svarc theorem says that any path metric space on which Γ acts cocompactly by isometries is quasi-isometric to Γ with its word metric. Thus, for instance, any two compact hyperbolic manifolds of the same dimension have quasi-isometric fundamental groups. The same holds for other locally symmetric spaces. This is sense in which we expect the word metric to capture the "geometric type" of a group.

The reader may have noticed a small distinction between the geometric and algebraic viewpoints: nets are well defined only up to quasi-isometry while the word metric gives a canonical bi-lipschitz type. This discrepancy is real - one instance where it causes problems is free products. It is not hard to see that the free products of bi-lipschitz equivalent groups are bilipschitz equivalent. One might therefore expect that quasi-isometric groups have quasi-isometric free products. This is not true, as the case of finite groups shows.

Our first result is a complete description of the difference between quasi-isometries and bi-lipschitz maps. UDBG stands for uniformly discrete of bounded geometry. See the next section for definitions.

Theorem 1.0.1. If $f : X \to Y$ is a quasi-isometry between UDBG-spaces then there is a bilipschitz map at bounded distance from f iff $f_*([X]) = [Y]$ where $[X]$ and $[Y]$ are the fundamental classes in H_0^{uf}.

This theorem also gives, more or less directly, a classification of the bilipschitz classes of spaces within a given quasi-isometry type.

Theorem 1.0.2. For any UDBG-space X, $S_{bilip}(X) = H_0^{uf+}(X)$.

Results of [3] show that $H_0^{uf}(X)$ is closely related to amenability. Amenability was originally defined, for groups, in terms of invariant means (finitely additive probability measures). This was shown by Folner to be equivalent to the existence of subsets whose boundaries are arbitrarily small compared to the entire subset. This definition makes sense for any discrete space. One of the early conjectures about amenability, called the Von Neumann conjecture (although Pierre de la Harpe informs me that it was, in fact, due to Day) states that a group is either amenable or contains a non-abelian free group (wlog on two generators). This is now known to be false [17], at least for finitely generated groups, but one does not know whether there is a finitely presented counterexample. We formulate and prove a geometric version of this conjecture:

Theorem 1.0.3. Let Z be a UDBG-space. Z is non-amenable iff Z has a partition whose pieces are uniformly lipschitz embedded copies of the free group on two generators with its word metric.

As a consequence we show that the vanishing of 0-dimension l^p-homology for $p > 1$ is equivalent to non-amenability.

Theorem 1.0.4. Let X be a UDBG-space, and $p \in (1, \infty]$. X is non-amenable iff $H_0^{(p)}(X) = 0$.

The case $p = \infty$ of this theorem is a theorem of [3].

As discussed above, one case where it is useful to know if a quasi-isometry can be replaced by a bi-lipschitz equivalence is in forming free products. Theorem 1.0.1 thus gives a way of constructing groups which are quasi-isometric, which do not obviously act isometrically on the same space (the only other source of such examples seems to be [8]).

Theorem 1.0.5. If Γ_1 and Γ_2 are quasi-isometric, finitely generated non-amenable groups then for any finitely generated group G, $\Gamma_1 * G$ and $\Gamma_2 * G$ are quasi-isometric.

This allows us to construct many new quasi-isometries. In particular we show that the sign of the Euler characteristic is not a quasi-isometry invariant, which was a question of Gromov's [12]. We also work out the quasi-isometric and commensurability classifications of free products of surface groups, in the course of which we prove a general coarse unique factorization theorem for free products:

Theorem 1.0.6. If $G_1 * G_2 * ... * G_n$ is quasi-isometric to $H_1 * H_2 * ... * H_m$ where the G_i and H_i are one ended then each G_i is quasi-isometric to some H_j.

For manifolds, uniformly finite homology is Poincare dual to the cohomology of the complex of bounded forms. If the manifold has bounded geometry, Chern-Weil theory gives characteristic classes living in these groups. For this reason, uniformly finite homology is important for manifolds as well as for groups. One of our results on H_0^{uf} gives:

Theorem 1.0.7. Let M be a manifold of bounded geometry. For any n there are constants $C_0, C_1, ..., C_n$ and r so that if $\omega \in \Omega^m_{(\infty)}(M)$ then ω is d of a bounded form iff for some C

$$\int_X w \leq Cvol(\partial X)$$

for all $(X, \partial X)$ compact, codimension 0 submanifold with $\nabla^i II_{\partial X} \leq C_i$ for $i = 0, 1, ..., n$.

We next study unformly finite characteristic classes, in particular the $\hat{A}$ genus and its relationship to scalar curvature.

The Atiyah-Patodi-Singer index theorem computes the index of an operator on a manifold with boundary in terms of the integral of a characteristic form and a spectral invariant (the η invariant) of the boundary. Results of Cheeger and Gromov ([5]) show that for the signature operator the η invariants is bounded linearly in the volume of the boundary. This is extended to many other geometric operators, and the Dirac operator on spinors in particular, in [18]. This shows that the only term in the index theorem which contributes to the uniformly finite characteristic class is the index.

For closed spin manifolds, one knows from Bochner type formulae that these indices vanish under appropriate positivity assumptions on the curvature. When the manifold has boundary, this method does not work directly, as there is an error term picked up when integrating by parts. The Bochner argument gives the positivity of this term, which can be computed from the spectrum decomposition on the boundary. Thus we get a lower bound for the largest eigenvalue which shows up in the boundary values of a harmonic spinor. The boundary condition one imposes on D to get a reasonable index, following [1], is that only negative eigenvalues occur in the boundary values of the admissible spinors. Thus, we can bound the index by the dimension of a piece of the spectrum of the operator on the boundary. The uniformly finite charcteristic class therefore vanishes if we have good estimates for the spectrum. For the case of greatest interest to us, the $\hat{A}$ class, we prove the spectral bound:

Theorem 1.0.8. If N^n is a compact spin manifold, then for each λ there is a C_λ depending only on the curvature and injectivity radius of N, for which $N_{D^2}(\lambda) \leq C_\lambda vol(N)$

By the argument outlined above this gives a vanishing theorem for manifolds of positive scalar curvature:

Theorem 1.0.9. If M has non-negative scalar curvature then $\hat{A} = 0$ in $H_0^{uf}(M)$.

Thus we have an obstruction to finding such metrics within a bounded distortion class of metrics, and we give several new examples of manifolds which cannot admit metrics of non-negative scalar curvature. In particular we provide a converse to a theorem of Block and Weinberger which produces metrics of (uniformly) positive scalar curvature on certain infinite connected sums. We show:

Theorem 1.0.10. Let M be a spin manifold of bounded geometry and uniformly positive scalar curvature, and N a compact spin manifold of the same dimension with $\hat{A}(N) \neq 0$. If $S \subset N$ is uniformly discrete, then the manifold $M \#_S N$ admits a metric of uniformly positive scalar curvature and bounded geometry in the canonical bounded distortion class iff $[S] = 0$ in $H_0^{uf}(M)$.

Our obstruction to positive scalar curvature does not seem to fit with the standard C^* algebra approaches, as the group the invariant lives in, H_0^{uf} is not Hausdorff in its natural topology. This raises some interesting questions about higher order invariants, which we discuss.

CHAPTER 2

NOTATION AND DEFINITIONS

As geometric group theory and coarse geometry are relatively young fields, not all the notation and terminology is uniform, below are our choices.

Definition 2.0.1. A space X is *uniformly discrete* if $\exists \varepsilon > 0$ such that $d(x_1, x_2) < \varepsilon \ \Rightarrow x_1 = x_2$.

Definition 2.0.2. If X is a metric space, a subset $S \subset X$ is *C-dense* if $d(x, S) \leq C$ for all x in X.

When we are not concerned with the value of C, we will simply say S is dense.

Definition 2.0.3. If X a metric space, $S \subset X$ is a *net* if S is uniformly discrete and dense.

The study of non-compact manifolds via their nets only really makes sense if the local geometry is controlled. For us that usually means the following:

Definition 2.0.4. A manifold M has *bounded geometry* if its section curvature is bounded in absolute value and its injectivity radius is bounded away from 0.

This local control survives in limited form in any net, so we make the following parallel definition:

Definition 2.0.5. A discrete metric space X has *bounded geometry* if it is uniformly discrete and, for any $r > 0$ there is a bound N_r on the cardinality of any r-ball in X.

While in theory it could be confusing to use the same term for both concepts it is always clear from context. To avoid repeating "uniformly discrete metric space of bounded geometry" we use the abbreviation UDBG-space.

2.1. Amenability

Let X be a UDBG-space, and $S \subset X$.

Definition 2.1.1.

$$N_r(S) = \{x \in X \ s.t. \ d(x,S) \leq r\}$$

and,

$$\partial_r(S) = N_r(S) - S = \{x \in X \ s.t. \ 0 < d(x,S) \leq r\}$$

The following definition of amenability is a direct generalization of a definition for groups due to Folner.

Definition 2.1.2. X a UDBG-space is *amenable* if

$$\exists \{S_i \subset X\} \ s.t. \ \forall r > 0 \lim_{i \to \infty} \frac{|\partial_r(S_i)|}{|S_i|} = 0$$

Such a sequence is called a regular sequence.

Remark 1. labelreg It is more common to require that $S_i \subset S_{i+1}$ and $\cup_i S_i = X$. This is equivalent [3]. A similar refinement we will have use for is that the S_i can be taken disjoint (so long as X is infinite). This is achieved by taking any regular sequence R_i, passing to a subsequence R_{n_i} with $|R_{n_{i+1}}| >> |R_{n_i}|$ and then taking inductively $S_i = R_{n_i} \cap S_{i-1}^c$. The S_i are disjoint by construction, and are a regular sequence since the R_{n_i} are, and the S_i differ from them by a set whose size goes to 0 as a fraction of $|S_i|$.

Lemma 2.1.1. Let X be a UDBG-space. The following are equivalent:

1. X is non-amenable.

2. For all $C > 1 \ \exists r \ s.t. \ \forall S \subset X \quad C|S| \leq |N_r(S)|$

3. For some $C > 1 \ \exists r \ s.t. \ \forall S \subset X \quad C|S| \leq |N_r(S)|$

4. For all $C > 0 \ \exists r \ s.t. \ \forall S \subset X \quad |S| \leq C|\partial_r(S)|$

5. For some $C > 0 \ \exists r \ s.t. \ \forall S \subset X \quad |S| \leq C|\partial_r(S)|$

Proof. It is immediate from the definition that all the conditions imply non-amenability. The last two conditions are just rephrasings of the second and third, and (2) $\Rightarrow$ (3) is trivial. We prove are (3) $\Rightarrow$ (2) and (1) $\Rightarrow$ (5).

(3) $\Rightarrow$ (2)

For any S we have:

$$N_s(N_t(S)) \subset N_{s+t}(S)$$

so if (3) holds for C and r then:

$$|N_{kr}(S)| \geq |N_r(N_{(k-1)r}(S))| \geq 1/C|N_{(k-1)r}(S)|$$

so,

$$|S| \leq C^k |N_{kr}(S)|$$

(1) $\Rightarrow$ (5)

If (5) does not hold, then there is a sequence $\{S_i\}$ of sets in X for which the condition fails with $r = i$ and $C = \frac{1}{i}$. This is then a regular sequence, which contradicts (1). $\square$

2.2. Homology

To a UDBG-space X we associate a family of simplicial complexes as follows:

Definition 2.2.1. For any $r > 0$, $R_r(X)$, the *r-Rips complex of X*, is the simplicial complex with X as its 0-simplices and in which a set $\{x_0, x_1, ..., x_k\}$ are the vertices of a simplex if and only if $d(x_i, x_j) \leq r$ for all i and j.

We have inclusions $R_r(X) \subset R_{r'}(X)$ for $r \leq r'$. If $f : X \to Y$ is coarse lipschitz then for any $r > 0$ there is r' so that f extends canonically to a simplicial map from $R_r(X)$ to $R_{r'}(Y)$. Thus we can define invariants of a UDBG-space as limits of simplicial invariants of its Rips complexes. In particular this lets us define homologies which are functorial in the coarse lipschitz category. We will use the following:

Definition 2.2.2. For X a UDBG space, the l^p homology groups $H_*^{(p)}(X)$ are defined as the limit as $r \to \infty$ of the l^p homology groups of $R_r(X)$.

If $f : X \to Y$ is coarse lipschitz then, as above, f induces a map on the Rips complexes and thereby of $H_*^{(p)}(X)$ to $H_*^{(p)}(Y)$. If f is at bounded distance from the identity then it is homotopic to the identity on the Rips complexes for r large enough, and so is the identity on homology. Thus a quasi-isometry induces an isomorphism of homology.

The groups $H_*^{(2)}$ have nice duality properties, and have been studied by many authors. The group $H_1^{(1)}(X)$ is related to coarse hyperbolicity in [9], settling a conjecture of Weinberger. The case of $p = \infty$ is of particular importance to us. These groups were introduced in [3] under the name uniformly finite homology, written H_*^{uf}, and we use that terminology here as well. Notice that for these groups $\mathbb{Z}$-coefficients make sense, while for $p < \infty$ we need to use $\mathbb{R}$ for a sensible theory.

Any subset S of X defines a uniformly finite 0-cycle. We denote the class of this cycle by $[S]$.

CHAPTER 3

BILIPSCHITZ STRUCTURES

3.1. Relation with H_0^{uf}

The goal of this chapter it to understand the relationship between bilipschitz equivalence and quasi-isometry.

Definition 3.1.1. Let X be a UDBG-space. We define the *bilipschitz structures* on X, $S_{bilip}(X)$, as the set of pairs (Y, f), where Y is a UDBG-space and $f : Y \to X$ is a quasi-isometry, up to the relation that (Y_1, f_1) is equivalent to (Y_2, f_2) iff there is a bilipschitz map $Y_1 \to Y_2$ which commutes (up to bounded distance) with the f_i.

Proposition 3.1.1. $S_{bilip}(pt) = \mathbb{N}$

Proof. Cardinality is clearly a complete invariant. □

Proposition 3.1.2. $S_{bilip}(\mathbb{Z})$ is uncountable.

Proof. Let t be any positive real number, and define $f_t : \mathbb{Z} \to \mathbb{Z}$ by $n \mapsto [tn]$ where $[x]$ is the greatest integer less or equal to x. Each f_t is a quasi-isometry, so $(\mathbb{Z}, f_t)$ give elements x_t of $S_{bilip}(\mathbb{Z})$.

If x_s and x_t are equivalent then there is a bilipschitz map $\phi : \mathbb{Z} \to \mathbb{Z}$ such that $|t\phi(n) - sn|$ is bounded independently of n. Equivalently $|\phi(n) - \frac{s}{t}n|$ is bounded independently of n.

Lemma 3.1.3. If $n \mapsto [\alpha n]$, $\alpha \in \mathbb{R}^+$, is at bounded distance from a bijection then $\alpha = 1$.

Proof. Let ψ be the bijection. $\psi^{-1}\{0, 1, ..., n\}$ is at bounded distance from $\{0, [\frac{1}{\alpha}], ..., [\frac{n}{\alpha}]\}$ and hence has cardinality $\frac{n}{\alpha}$ up to a term bounded independently of n. But ψ is a bijection, so the cardinality is exactly n. Letting $n \to \infty$ proves $\alpha = 1$. □

This proves that the x_t are all distinct, and hence that $S_{bilip}(\mathbb{Z})$ is uncountable. $\square$

We now generalize to arbitrary X.

Definition 3.1.2. Let X be a UDBG-space.

$\chi : S_{bilip}(X) \to H_0^{uf}(X)$ is defined by $(Y, f) \mapsto f_*([Y])$

This is well defined because a bilipschitz map $Y_1 \to Y_2$ carries $[Y_1]$ to $[Y_2]$, even on the level of chains.

We start by asking what the image of χ is. If $f : Y \to X$ is a quasi-isometry then $f_*([Y])$, as a chain, is non-negative and is strictly positive on a dense set.

Definition 3.1.3. A chain $c \in C_0^{uf}$ is positive if it is everywhere non-negative and positive on a dense set. A class in H_0^{uf} is positive if it is represented by a positive chain. The cones of positive chains and homology classes will be denoted C_0^{uf+} and H_0^{uf+} respectively.

Proposition 3.1.4. The image of χ is exactly H_0^{uf+}

Proof. We saw above that the image is contained in H_0^{uf+}. Fix $c \in C_0^{uf+}$, and define:

$$Y_c = \{(x, n) : x \in X, 1 \leq n \leq c_x\}$$

with the metric it inherits from $X \times \mathbb{N}$.

Y_c has a canonical map π to X given by projection, and this map is clearly a quasi-isometry. As $\pi_*([Y_c]) = c$, we are done. $\square$

Next we want to know if χ is injective. If $z_1 = (Y_1, f_1)$ and $z_2 = (Y_2, f_2)$ are elements of $S_{bilip}(X)$, then there is a unique (up to bounded distance) quasi-isometry from Y_1 to Y_2 which commutes up to bounded distance with the f_i, namely $g_2 f_1$ where g_2 is a coarse inverse for f_2. For z_1 and z_2 to be equivalent we need to find a bilipschitz map within bounded distance of this quasi-isometry. If we assume $\chi(z_1) = \chi(z_2)$ then we know that $(g_2 f_1)_*([Y_1]) = [Y_2]$. That this is sufficient to find a bilipschitz map is the point of the next section.

3.2. Constructing Bi-Lipschitz Maps

Theorem 3.2.1. If $f : X \to Y$ is a quasi-isometry between UDBG-spaces with $f_*([X]) = [Y]$ then there is a bilipschitz map at bounded distance from f.

Proof. Observe that a quasi-isometry between UDBG-spaces is bilipschitz if and only if it is bijective. Since any map at bounded distance from a quasi-isometry is itself a quasi-isometry, we need to see that there is a bijection at bounded distance from f.

We accomplish this in two stages. First we find an injection at bounded distance from f. If g is a coarse inverse for f then we also get an injection at bounded distance from g. The proof of the Schroeder-Berstein theorem then assembles these to form a bijection.

Lemma 3.2.2. There is an injective map at bounded distance from f.

Proof. Let $r > 0$ be fixed. To find an injection within r of f is to select, for each x in X an element of $B_r(f(x))$ without repetition. This is exactly the sort of question addressed by Hall's selection theorem.

Theorem 3.2.3. (Hall's Selection Theorem [11]) If we have $\alpha : A \to Fin(B)$ (the finite subsets of B), then there's an injective map $\phi : A \to B$ with $\phi(a) \in \alpha(a) \forall a \in A$ iff $\forall S \subset A$ we have that

$$|S| \leq |\cup_{s\in S}\, \alpha(s)|$$

In the case of the lemma, we need $\forall S \subset X$ $|S| \leq |N_r(f(S))|$. The LHS is increased if we replace S by $f^{-1}(f(S))$. Therefore we need,

$$(*) \forall S \subset Y \quad |f^{-1}(S)| \leq |N_r(S)|$$

We have

$$|f^{-1}(S)| = |\Sigma_S f_*([X])|$$

where $\Sigma_S c$, for $c \in C_0^{uf}$, is just the sum of the coefficients of vertices in S

So $(*)$ becomes

$$\Sigma_S f_*([X]) \leq |N_r(S)|$$

equivalently,

$$\Sigma_S(f_*([X]) - [Y]) \leq |\partial_r(S)|$$

Let c be the chain $f_*([X]) - [Y]$. It is, by construction, everywhere at least -1, and is non-negative on a dense set.

Lemma 3.2.4. Let Z be a UDBG-space, $c \in C_0^{uf}(X)$ everywhere at least -1 and non-negative on a C-dense set. If $c = \partial b$ for $b \in C_1^{uf}$ then there is some r such that for any $S \subset Z$ we have:

$$\Sigma_S c \leq |\partial_r(S)|$$

Proof. We need some constants. Let M be a bound on the absolute values of the coefficients of b, and α such that $b \in C_0^{uf}(R_\alpha(Z))$.

Consider the graph Γ which is the subgraph of the 1-skeleton of R_α consisting of those edges in the support of b. We orient the edges such that all coefficients of b are positive. Let V be the maximal valence of any vertex of Γ. We can assume Γ has no oriented loops as we can subtract such a loop from b without increasing its norm or changing ∂b.

Fix S.

We next construct a family of oriented edge paths in Γ, with the following properties:

1. Every path ends at a point of S, and the number of paths ending at s is $max(c_s, 0)$.

2. Any edge e is traversed at most b_e times.

3. Every path either has no initial point or starts at a point where c is -1.

4. No two paths have a common initial point.

Starting with any $s \in S$ with $c_s > 0$, follow edges of Γ backwards, choosing at random if there is more than one. This process only gets stuck if it reaches a point without incoming edges. As we used an outgoing edge to reach this point, that implies c is negative at this point, and hence is -1.

This gives one path with the desired properties. Now subtract this path from b and repeat. Subtracting a path changes ∂b only at the endpoints, giving all but the second property above. This property also follows since each time an edge e is traversed, b_e decreases by 1, and when it reaches 0 that edge is removed from Γ.

Let r be any real greater than C.

We divide our family of paths into four sets:

P_0 is paths starting in S.
P_1 is paths starting in $\partial_C(S)$.
P_2 is paths starting in $\partial_r(S) - \partial_C(S)$.
P_3 is paths starting outside $\partial_r(S)$ (including the infinite paths)

Proposition 3.2.5. (Stoke's Theorem)

$$\Sigma_S c = |P_1| + |P_2| + |P_3|$$

Proof. The boundaries of paths in P_0 contribute nothing to $\Sigma_S c$, the boundaries of the others each contribute 1, so: $\square$

We estimate these terms individually.

First, we see $|P_1| \leq |\{z \in \partial_C(S) \text{ with } c_z = -1\}|$. We call this set $\partial_C^-(S)$, and the complement in $\partial_C(S)$ we call $\partial_C^+(S)$.

Lemma 3.2.6. There is $A > 1$, independent of r, such that

$$A|P_2| \leq |\partial_{r+C}(S)| + |\partial_C^+(S)|$$

Proof. To each path in P_2 we associate two points: its initial point, which by construction has $c = -1$, and a randomly chosen point which is within C of the initial point and has $c \geq 0$. There is such a second point by the definition of C, and any point with $c \geq 0$ can be associated to at most the maximal size of a C-ball (call this K) of points. Thus $A = 1 + \frac{1}{K}$ works. $\square$

Lemma 3.2.7. There is $B > 0$, independent of r, such that:

$$|P_3| \leq \frac{B(|\partial_r(S)| - |\partial_C(S)|)}{r - C}$$

Proof. Since each path in P_3 leaves $N_r(S)$, and each edge is at most length α, there must be at least $\frac{r-C}{\alpha}$ edges in the path in $\partial_r(S) - \partial_C(S)$. And as each edge can be on at most M edges, there are at least $|P_3|\frac{r-C}{M}$ edges, and therefore at least $|P_3|\frac{r-C}{VM\alpha}$ points, in this region. $\square$

Now, if $\frac{|P_2|+|P_3|}{|P_2|} \leq A$ then by 3.2.6 and 3.2.5 we have

$$\Sigma_S c \leq |\partial_{r+C}(S)|$$

If $\frac{|P_2|+|P_3|}{|P_2|} \geq A$ then by 3.2.7

$$\Sigma_S c \leq |\partial_C(S)| + \frac{AB}{r-C}|\partial_r(S) - \partial_C(S)|$$

If $r > AB + C$ then the RHS is less than $|\partial_r(S)|$.

Thus $r = AB + 2C$ works. $\square$

This bound is exactly what was needed to produce an injection, as the class $f_*([X]) - [Y]$ is a boundary by assumption.

$\square$

Applying this to f and a coarse inverse for f, we get injections $h_1 : X \to Y$ and $h_2 : Y \to X$, so that h_1 is at bounded distance from f, as is h_2^{-1} wherever its defined. The Schroeder-Berstein theorem (as proven in [15], for example) produces a bijection $\phi : X \to Y$ which everywhere agrees with h_1 or h_2^{-1} and hence is at bounded distance from f. $\square$

Corollary 3.2.8. For any UDBG-space X, $S_{bilip}(X) = H_0^{uf+}(X)$.

Remark 2. The bilipschitz map constructed is in no way canonical, and so enjoys no nice functorial properties. It is natural to ask about parametrized and equivariant analogues. For example, given two quasi-isometries from X to itself which commute up to bounded distance, when is there a bilipschitz action of $\mathbb{Z}^2$ on X or on something quasi-isometric to X. It is not hard to define higher order obstructions to this which live in higher uniformly finite homology groups, but one does not know either whether those obstruction are everything or whether they are ever non-zero.

One point of view of all this is that we have an exact sequence:

$1 \to Bilip_1(X) \to Bilip(X) \to QI(X) \to H_0^{uf+}(X) \to 1$

Here $Bilip_1(X)$ is the group of bilipschitz maps at bounded distance from the identity.

What we are asking about is the lifting problem from $QI(X)$ to $Bilip(X)$. The geometry of this sequence is somewhat mysterious. $Bilip_1(X)$ has the natural metric of distance in X. Is it bounded generated in this metric? This would be a start on a uniqueness theorem to go with 3.2.1. It is not easy, although true, for $X = \mathbb{Z}$. Even there the word metric from these generators and the natural metric are quadratically distorted.

CHAPTER 4

APPLICATIONS

We first recall the following:

Theorem 4.0.9. ([3]) If X is a non-amenable UDBG-space then $H_0^{uf}(X) = 0$ iff X is non-amenable.

We give a new proof in a later section (see 4.3.1).

4.1. Free Products

If G_1 and G_2 are finitely generated groups, then so is their free product. It is not hard to see that the bilipschitz class of the free product is determined by the bilipschitz classes of the factors: the obvious induced map on normal forms is a bilipschitz equivalence.

Theorem 4.1.1. If Γ_1 and Γ_2 are quasi-isometric, finitely generated non-amenable groups then for any finitely generated group G, $\Gamma_1 * G$ and $\Gamma_2 * G$ are quasi-isometric.

Proof. By 4.3.1 or [3], $H_0^{uf}(\Gamma_1) = 0$. By 3.2.1 Γ_1 and Γ_2 are bilipschitz equivalent, and thus so are $\Gamma_1 * G$ and $\Gamma_2 * G$. $\square$

This is not always true for amenable groups: $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ is infinite while $\{e\} * \mathbb{Z}/2\mathbb{Z}$ is finite, and $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/3\mathbb{Z}$ is not quasi-isometric to either.

This produces quasi-isometries which are algebraically surprising. As an example, we work out the following case:

Corollary 4.1.2. There are exactly five quasi-isometry classes of free products of surface groups: $\mathbb{Z}^2$, Σ_2, $\mathbb{Z}^2 * \mathbb{Z}^2$, $\mathbb{Z}^2 * \Sigma_2$, and $\Sigma_2 * \Sigma_2$.

Proof. All the surface groups of higher genus are of finite index in Σ_2 and therefore quasi-isometric to it. As these are non-amenable they are bilipschitz equivalent. Thus we need only consider free products of $\mathbb{Z}^2$ and Σ_2.

$\Sigma_2 * \Sigma_2$ has a subgroup of index k isomorphic to $\Sigma_{k+1} * \Sigma_2 * \Sigma_2 * ... \Sigma_2$ (k-times). Thus the free product two or more surface groups of higher genus is quasi-isometric to $\Sigma_2 * \Sigma_2$.

Likewise, the free product of at least one $\mathbb{Z}^2$ and at least one Σ_2 is quasi-isometric to $\mathbb{Z}^2 * \Sigma_2$.

While $\mathbb{Z}^2$ is amenable, $\mathbb{Z}^2 * \mathbb{Z}^2$ is non-amenable and so again by passing to subgroups of finite index we see that any free product of at least two copies of $\mathbb{Z}^2$ is quasi-isometric to $\mathbb{Z}^2 * \mathbb{Z}^2$.

That reduces us to the given list. Being a Poincare duality group is a quasi-isometry invariant [3], as is amenability. Thus we can distinguish $\mathbb{Z}^2$ and Σ_2 from each other and the rest of the list.

Proposition 4.1.3. If $G_1 * G_2 * ... * G_n$ is quasi-isometric to $H_1 * H_2 * ... * H_m$ where the G_i and H_i are one ended then each G_i is quasi-isometric to some H_j.

Proof. Since G_i is one ended, it has no coarse cut points. This means its image is within bounded distance of a unique conjugate of one of the H_j. The same holds for this H_j under the inverse quasi-isometry, and G_i is within bounded distance of the image of H_j which means the quasi-isometry restricts to give a quasi-isometry between of G_i and H_j.

□

The proposition distinguishes between the remaining groups listed. □

This should be contrasted with:

Proposition 4.1.4. $\Gamma_1 = \Sigma_{a_1} * \Sigma_{a_2} * ... * \Sigma_{a_m}$ and $\Gamma_2 = \Sigma_{b_1} * \Sigma_{b_2} * ... * \Sigma_{b_n}$ are commensurable iff

$$\frac{\chi(\Gamma_1)}{m-1} = \frac{\chi(\Gamma_2)}{n-1}$$

Proof. This is essentially a special case of an explicit version of the Kurosh subgroup theorem, as in [20].

The subgroups of finite index in such a free product are of the form $\Sigma_{g_1} * \Sigma_{g_2} * ... * \Sigma_{g_r} * F_l$ where F_l is the free group of rank l. If this is of index I in Γ_1 then the set $S = \{g_1, ..., g_r\}$ can be partitioned into $S_1, ..., S_m$ where S_i are the groups contained in conjugates of Σ_{a_i}. Thus, for each $s \in S_i$ we have that $a_i - 1$ divides $s - 1$ and that $\Sigma_{S_i} \frac{s-1}{a_i-1} = I$. The rank of the free factor is then determined by calculating Euler characteristics,

$$l = I\chi(\Gamma_1) + r - 1 - \Sigma_S(2 - 2s)$$

but we can calculate $\Sigma_S(2-2s)$ as $\Sigma_1^m \Sigma_{S_i}(2-2s)$ which is equal to $\Sigma_1^m I(2 - 2a_i) = I(\chi(\Gamma_1) + m - 1)$. Thus

$$l = r - 1 - I(m - 1)$$

Conversely, given such partition we can produce a corresponding graph of groups which is a finite cover of the graph for Γ_1. This graph can be chosen connected exactly when the rank of the free factor predicted by the Euler characteristic formula is non-negative.

So, to determine when Γ_1 and Γ_2 are commensurable we need to decide when there is a set S with partitions for both Γ_1 and Γ_2 which give the same (non-negative) rank for the free factor. By our earlier formula this says that $I_1(m-1) = I_2(n-1)$. But equating Euler characteristics gives:

$$I_1\chi(\Gamma_1) = I_2\chi(\Gamma_2)$$

so $\frac{\chi(\Gamma_1)}{m-1} = \frac{\chi(\Gamma_2)}{n-1}$, which shows that the condition of the theorem is necessary.

Going the other way, if these are equal we can find S with partitions $\{A_1, ..., A_m\}$ and $\{B_1, ..., B_n\}$ so that $\frac{\Sigma_{A_i} 2-2s}{2-2a_i}$ is independent of i, and likewise for the $\{B_i\}$. Replacing every s with $N!(s+1) - 1$ for N large does not change this property, but does give the divisibility and connectedness properties needed, so the condition of the theorem is also sufficient.

□

This shows that in general the quasi-isometric and commensurable classifications of groups are very different. We have enough examples to show that many commensurability invariants are not quasi-isometry invariants. The next corollary answers a question of Gromov.

Corollary 4.1.5. The sign of the Euler characteristic is not a quasi-isometry invariant.

Proof. Let $G = \Sigma_2 \times \Sigma_2$ and $H = \Sigma_2$.

Take G_k a subgroup of G of index k, and H_l a subgroup of H of index l.

All the groups are non-amenable, so all groups of the for $G_k * H_l$ are quasi-isometric, but

$$\chi(G_k * H_l) = k\chi(G) + l\chi(h) - 1 = 4k - 2l - 1$$

which can clearly be of either sign. □

This shows that the quasi-isometries produced by 4.1.1 are not "algebraic", as the sign of the Euler characteristic is a commensurability invariant, and is the same for different uniform lattices in the same Lie group.

Remark 3. All that was really used about the Euler characteristic here is that it is multiplicative under finite coverings and additive for free products. This is a property of many other invariants, L^2-betti numbers for example. Exactly the same construction shows no such invariants, for example ratios of L^2-betti numbers, can be quasi-isometry invariants. Note that the non-amenable condition will not get in the way as the L^2-betti numbers of any amenable group are all zero ([6]).

4.2. The Von Neumann Conjecture

The Von Neumann conjecture is that a finitely generated group is either amenable or contains a non-abelian free subgroup. While this is know to be false ([17]), it is a good intuitive guide and holds under some restrictions, for example for linear groups there is the Tits alternative [21]. In this section we give a geometric version of the conjecture, and prove it using our earlier results.

To give a geometric version of the conjecture, one needs a geometric analogue of having a free non-abelian subgroup. All free groups contain the free group on two generators, so there is no loss of generality in asking for that as a subgroup. A subgroup is clearly more than just a subset - at the least it partitions the space into cosets, each isomorphic to itself. It is that decomposition we mimic in our version of the conjecture. Our theorem also has an interpretation in terms of group actions which we discuss in the remark at the end of this section.

Theorem 4.2.1. (Geometric Von Neumann Conjecture) Let Z be a UDBG-space. Z is non-amenable iff Z has a partition whose pieces are uniformly lipschitz embedded copies of the tri-valent tree.

Remark 4. The free group on two generators with its word metric is the 4-valent tree. We could just as easily as asked for those, as by 3.2.1 the 3 and 4 valent trees are bilipschitz equivalent. We chose the 3-valent tree as it seems more natural geometrically.

Proof. If Z admits such a decomposition then it is clearly non-amenable, simply apply any of the criteria of 2.1.1 piecewise.

If Z is non-amenable, we produce a decomposition with pieces quasi-isometric to the tri-valent tree, and then use 3.2.1 to improve them to lipschitz. From now on we assume Z is non-amenable.

Lemma 4.2.2. There is a bilipschitz map f from $Z \times \{0, 1\}$ to Z at bounded distance from, but never equal to, the projection onto the first factor.

Proof. We simply copy the proof of 3.2.1, but as Z is non-amenable it is much easier, even with the extra restriction. If f is to be within r of the projection, the set of possible values for $f(z, \epsilon)$ is $B_r(z) - \{z\}$. So, for any S in Z, the set of possible values for $f(S \times \{0, 1\})$ is at least $\partial_r(S)$. By 2.1.1, $|\partial_r(S)|$ can be made larger than $|S \times \{0, 1\}| = 2|S|$ by taking r large enough. Thus, as in 3.2.1, we can produce an injection with the desired properties. The other way we have the inclusion $z \mapsto (z, 0)$, which can be moved to an injection so that z goes to neither $(z, 0)$ nor $(z, 1)$ in exactly the same way. The Schroeder-Berstein produces f.

□

Given f, let Γ be the graph with vertex set Z and directed edges from z to $f(z, 0))$ and $(z, f(z, 1))$. Thus every vertex has valence three (the obvious two "outgoing" edges, and the "incoming" edge from $\pi_Z f^{-1}(z)$. The fact that f is at bounded distance from π_Z means we have a bound on the lengths of the edges is Γ.

Proposition 4.2.3. Each connected component of Γ has at most one loop.

Proof. Consider a path in Γ. Since each vertex has only one incoming edge, any path starting with an oriented edge must be oriented. Thus any path is either oriented, or there is some vertex with both of its outgoing edges on the path. In this case the path is the union of two oriented paths leaving that point. These two branches can never meet again as that would require a point with two incoming edges. hus any loop in Γ is oriented,

If there are two loops in a component, connect them with a path. Whichever type of path it is, it has one endpoint with the end of the path oriented towards it. This is impossible as this point is also on a oriented loop, and would therefore need two incoming edges. □

Removing one edge from each loop from Γ to get a graph Γ'. Each component is a tree, which is either the 3-valent tree or a tree with one vertex of valence 2, and the rest of valence 3. This latter tree is quasi-isometric to the 3-valent tree simply by collapsing one edge adjacent to the vertex of valence 2, and therefore by 3.2.1 it is bilipschitz to the 3-valent tree. $\square$

Remark 5. We can reformulate 4.2.1 using the following idea:

Definition 4.2.1. If X is a UDBG space, and Γ a group then we call an action of Γ on X translation-like if:

1. For each $\gamma \in \Gamma$ the map $x \mapsto \gamma x$ is at bounded distance from the identity.

2. Γ acts freely.

Any group acting on itself from the right (if we use the left word metric) is translation-like. So is the action of any subgroup.

Theorem 4.2.4. (Geometric Von Neumann conjecture reformulated) A UDBG-space X is non-amenable iff it admits a translation-like action of a non-abelian free group.

It is reasonable to ask, given a group, which spaces it can act on in a translation-like manner. If the spaces are non-amenable (which they must be if the group is) then this is a quasi-isometry invariant. It is not clear what the answer is, even for the case of $\Gamma = \mathbb{Z}$. The argument of 4.2.1 show that any coarsely path connected space admits a bilipschitz self map which is fixed point free and at bounded distance from the identity, but it is unclear how to prevent periodic orbits.

4.3. Amenability and 0-dimensional Homology

Theorem 4.3.1. Let X be a UDBG-space, and $p \in (1, \infty]$. X is non-amenable iff $H_0^{(p)}(X) = 0$.

Proof. If X is amenable we construct a non-trivial 0-cycle as follows:

Let S_i be a regular sequence with the S_i disjoint (see the remark preceding 2.1.1). Then for a sequence a_i of positive reals, $c = \Sigma_i a_i S_i$ is an l^p chain iff $\Sigma_i a_i^p |S_i| < \infty$. Now, if $c = \partial b$ then, for each i, the total flow through $\partial_r S_i$ is at least $a_i |S_i|$ by Stokes theorem. Then, by Holder's inequality:

$$||b||^p \geq \Sigma_i (\frac{a_i |S_i|}{|\partial_r S_i})^p |\partial_r S_i|$$

$$= \Sigma_i a_i^p |S_i| (\frac{|S_i|}{|\partial_r S_i|})^{p-1}$$

$(\frac{|S_i|}{|\partial_r S_i})^{p-1} \to \infty$ since, by assumption, $\frac{|S_i|}{|\partial_r S_i|}$ does. Thus it is easy to produce a sequence a_i which gives an l^p chain but cannot be an l^p boundary.

For the other direction, suppose X is non-amenable and c is an l^p 0-chain on X. Write X as a disjoint union of trivalent trees, and choose one edge in each tree and cut it to get X as a union of rooted binary trees. Then c restricts to each tree to give an l^p chain.

Lemma 4.3.2. Let T be the infinite rooted binary tree. There is, for each p, a bounded operator $I : C_0^{l^p}(T) \to C_1^{l^p}(T)$ such that $\partial I = Id$.

Proof. For q in T, define Iq to be the 1-chain supported on the tree of points which connect to the base point through q, where each edge has weight $2^{-d(-,q)}$. Then $\partial I q = q$.

We define Ic for any 0-chain c by linearity. The sums are finite as each edge has non-zero weight in Iq for finitely many q.

We need to see that the I so defined takes l^p to l^p, and is bounded. For this we use a special case of the Marcinkiewicz interpolation theorem ([7]).

Theorem 4.3.3. (Marcinkiewicz Interpolation) If T is bounded from $l^\infty \to l^\infty$ and is weak type $(1,1)$ then T is bounded from l^p to l^p for all $p > 1$.

Recall that T weak type $(1,1)$ means that there is C such that for any α and any f in l^1:

$$\lambda_f(\alpha)|\{v : Tf \geq \alpha\}| \leq C||f||_1/\alpha$$

.

In our case, $||I||_\infty$ is 1.

Lemma 4.3.4. I is weak type $(1,1)$.

Proof. Fix α.

We look for f of l^1-norm 1 such that If is above α on as many edges as possible. Since If can clearly be above α only on finitely many edges, and f of the vertices below them does not contribute to them, we can take f to have finite support. Let A_n be the set of f with l^1-norm 1 which are supported in the ball of radius n around the root.

We show, by induction on n, that

$$sup_{f\in A_n} \lambda_{If}(\alpha) \leq 2[\frac{1}{\alpha} - 1]$$

$[x]$ is the greatest integer $\leq x$.

The case $n = 0$ is obvious.

Take $f \in A_{n+1}$

Call the components of $T - \{root\}$ the left and right subtrees, with roots $\{root_l\}$ and $\{root_r\}$.

Let $g_l(resp.\ g_r)$ be f restricted to the left (resp. right) subtree.

Consider $f_l = g_l + \frac{f(\{root\})}{2}\delta_{\{root_l\}}$, and similarly for f_r. Then If_r and If_l agree with the restrictions of If, and each is in A_n. Conversely, these determine f. By induction, letting t be the norm of a candidate f_l, we have:

$$\lambda_f(\alpha) \leq sup_t([\frac{1}{t\alpha}] + [\frac{1}{(1-t)\alpha}] + 2\chi_{[\alpha,1-\alpha]}(t))$$

The last term comes from the edges coming from the root.

The desired bound now follows from the fact that $\lambda_g(\alpha) = 0$ for $\alpha > \frac{1}{2}$ for any g of l^1-norm 1. $\square$

$\square$

Now combining the I from each tree gives a splitting of $C_1^{(p)}(X) \to C_0^{(p)}(X)$ $\square$

The case $p = \infty$ is a theorem of ([3]) (recall that $H_*^{uf} = H_*^{(\infty)}$).

We can also give a new proof of the following theorem asserted in ([2]) which we will use in the next chapter.

Theorem 4.3.5. If Z a UDBG-space, $c \in C_0^{uf}(Z)$, then $[c] = 0$ iff there are r and C such that for all $S \subset Z$:

$$|\Sigma_S c| \leq C|\partial_r S|$$

Proof. Let $n \in \mathbb{N}$ be large enough that $c + n[X]$ is positive.

Consider the elements x_n and y_n of $S_{bilip}(Z)$ corresponding to $c + n[Z]$ and $n[Z]$ respectively, with representatives (X_n, f_n) and (Y_n, g_n).

Clearly $[c] = 0$ iff $f_*([X_n]) = g_*([Y_n])$.

By 3.2.8, this is true iff there is a bilipschitz map $h : X_n \to Y_n$, commuting up to bounded distance with f_n and g_n. Will try to build h as in the proof of 3.2.1.

Fix r.

A map $h : X_n \to Y_n$ commutes within r with f and g iff

$$h(x) \in \{y : \ d(f(x), g(y)) \leq r\}$$

for all $x \in X_n$.

Thus, for any set S in X_n, the possible range is $g^{-1}(N_r(f(S))$. To find an injective map the Marriage lemma requires that this be no smaller than S. This is hardest when $S = f^{-1}(f(S))$. We can then write things in terms of $A = f(S)$:

$$\forall A \subset Z \ \ |f^{-1}(A)| \leq |g^{-1}(N_r(A))|$$

In terms of c, this says:

$$\forall A \subset Z \ \ \Sigma_A c + n[Z] \leq \Sigma_{N_r A} n[Z]$$

or,

$$\forall A \subset Z \ \ \Sigma_A c \leq \Sigma_{\partial_r A} n[Z]$$

or

$$\forall A \subset Z \ \ \Sigma_A c \leq n|\partial_r A|$$

This holds once n is sufficiently large iff

$$\exists C, r \ s.t. \ \forall S \subset Z \ \ \Sigma_S c \leq C|\partial_r S|$$

To build the injection the other way, we need

$$\forall A \subset Z \ \ \Sigma_A(n[Z]) \leq \Sigma_{N_r A}(c + n[Z])$$

equivalently,

$$\forall A \subset Z \ \ \Sigma_A(-c) \leq \Sigma_{\partial_r A}(c + n[Z])$$

which holds for n sufficiently large iff:

$$\exists C, r \ s.t. \ \forall S \subset Z \ \ \Sigma_S(-c) \leq C|\partial_r S|$$

The two conditions together give the Theorem.

□

CHAPTER 5

MANIFOLDS OF BOUNDED GEOMETRY

In this chapter we show that the preceding theory does have applications to the topology and geometry of manifolds as well as UDBG-spaces.

All manifolds are assumed complete, connected, smooth riemannian manifolds of bounded geometry, and all maps are smooth and of bounded distortion.

5.1. Index Theory and H_0^{uf}

We are interested in understanding geometrically defined elliptic operators on a manifold M. In particular we are interested in the Dirac operator on spinors.

For compact manifolds, one has the famous Atiyah-Singer Index theorem, which says that the index of the operator can be computed as a characteristic number of the manifold. In the case of the Dirac operator on spinors, the characteristic form is the $\hat{A}$ genus.

While the operators make sense in our non-compact setting, it is not a priori clear that anything else does. It is certainly not true that the kernel and cokernel of an elliptic operator D are finite dimensional, so the index as it is normally defined doesn't carry over.

Chern-Weil theory says that for compact manifolds, any characteristic number can be computed as the integral over M of a polynomial in the curvature. We can still form these polynomials in the curvature in the non-compact setting, and as the curvature is bounded, we get a bounded top dimensional form. Changing out choice of metric within its bounded distortion class changes this form by d of a bounded form. So we are naturally led to consider the cohomology groups of the complex of bounded forms. By Poincare duality these groups are isomorphic to the l^∞ homology groups ([3]). This ties in with the last chapter because of:

Lemma 5.1.1. $H_0^{(\infty)}(M) = H_0^{uf}(X)$ for any net X in M.

Proof. As all nets give the same answer, we can use any one we like. We take the vertices of a uniformly finite triangulation (which one can get as the nerve of a covering by geodesically convex balls, for example). Then we have a map from M to its Rips complex for r large enough. By chosing minimizing geodesics we can build a homotopy inverse on the 1-skeleton of the Rips complex. This is sufficient to get an isomorphisn for the 0 dimensional groups. □

Thus the characteristic numbers naturally take values in $H_0^{uf}(X)$. It is easier in this setting to work cohomologically with the forms, so we restate 4.3.5 in those terms:

Theorem 5.1.2. Given $\omega \in \Omega^m(M)$ bounded, ω is d of a bounded form iff there are C, r such that for all compact $S \subset M$:

$$\int_S w \leq Cvol(\partial_r S)$$

Cheeger and Gromov ([5]) prove a chopping theorem for manifolds of bounded geometry which says that for any n there are constants $C_0, C_1, ..., C_n$ and r so that for any S there is a codimension 0 manifold with boundary $(X, \partial X)$ such that:

1. $S \subset X \subset N_r(S)$
2. For $i = 0, 1, ..., n$ $\nabla^i II_{\partial X} \leq C_i$ where $II_{\partial X}$ is the second fundamental form, and ∇^i is the i^{th} covariant derivative.
3. $\frac{Vol(\partial X)}{Vol(N_r(S))}$ is bounded above and below independent of S.

This is important to us because it means:

Theorem 5.1.3. For any n there are constants $C_0, C_1, ..., C_n$ and r so that if $\omega \in \Omega^m(M)$ then ω is d of a bounded form iff for some C

$$\int_X w \leq Cvol(\partial X)$$

for all $(X, \partial X)$ compact, codimension 0 submanifold with $\nabla^i II_{\partial X} \leq C_i$ for $i = 0, 1, ..., n$.

Proof. The condition is necessary by Stokes theorem. To see that it is sufficient we use the previous criterion. Given $S \subset M$, approximate it as above. By the third condition we have that we need only see

$\int_S \omega \leq CVol(\partial X)$

But by (1) $\int_S \omega$ and $\int_X \omega$ differ by at most $||\omega|| Vol(\partial_r S)$ which, again by (3), just means taking a larger C. □

Our elliptic operator D restricts to one on each such X. If the operator is one of the classical Dirac operators, or sufficiently related to one (these are the geometric operators of [18] or [10]) then with suitable boundary conditions, it has an index which is related to characteristic forms via the Atiyah-Patodi-Singer index theorem ([1]).

Theorem 5.1.4. (Atiyah-Patodi-Singer) For any geometric operator D, there is a characteristic form ω, a polynomial P, and $n \in \mathbb{N}$, so that for any $(X, \partial X)$ compact we have:

$$index(D) = \int_X w + \int_{\partial X} P(II_{\partial X}, \nabla II_{\partial X}, ..., \nabla^n II_{\partial X}) + \eta(\partial X)$$

Where $\eta(N) = \lim_{s \to 0^+} \Sigma_{\lambda \in Spec(D_\partial)} \lambda^{-s}$

Given the bounds on $\nabla^i II_{\partial X}$, that term of 5.1.4 is bounded by a constant multiple of $Vol(\partial X)$.

Likewise, the η are bounded linearly in $Vol(\partial X)$. This was shown for the signature operator in [5] and for a wide range of geometric operators including the Dirac operator on spinors in [18]. In view of 5.1.3, we can interpret 5.1.4 as saying that the characteristic form in H_0^{uf} is an “index” of our operator.

5.2. The $\hat{A}$-Genus and Scalar Curvature

We now focus on the case of the Dirac operator on spinors, so for this section M is spin, and D is the Dirac operator on spinors. X will be a submanifolds with boundary as in 5.1.3.

We are interested in using the $\hat{A}$ genus to help understand scalar curvature. These are related by:

Theorem 5.2.1. (Lichnerowicz formula, [16])

$$D^2 = \nabla^* \nabla + \frac{\kappa}{4}$$

Where ∇ is the canonical spinor connection, and κ the scalar curvature.

So, if s is a harmonic spinor (meaning $Ds = 0$) on X, we have:

$$0 = < \nabla^* \nabla s, s > + \frac{\kappa}{4} ||s||^2$$

Remark 6. If X were closed, we could integrate over X, and the first term on the RHS would be $||\nabla s||^2$.

Then, if $\kappa > 0$ both terms on the RHS would be ≥ 0, and therefore 0. This would mean $s = 0$. i.e. that there are no harmonic spinors, so that the index would be zero. Then by the Atiyah-Singer index theorem the $\hat{A}$ genus would be 0.

It is this argument we try to extend to our setting.

We can still integrate the Lichnerowicz formula over X when it has boundary. The difference is that now ([16])

$$\int_X < \nabla^* \nabla s, s > = ||\nabla s||^2 - \int_{\partial X} < \nabla_\nu s, s >$$

(ν is the unit normal vector to ∂X):

The second term on the RHS is introduced when we integrate by parts.

So, when we have boundary, the Lichnerowicz formula becomes:

Lemma 5.2.2. If s is a harmonic spinor on X, we have:

$$0 = ||\nabla s||^2 + \int_X \frac{\kappa}{4}||s||^2 - \int_{\partial X} < \nabla_\nu s, s >$$

If we assume the scalar curvature is ≥ 0 then the first two terms are non-negative. This can only happen if:

$$\int_{\partial X} < \nabla_\nu s, s > \geq 0$$

We can expand D in normal coordinates around the boundary:

Lemma 5.2.3. ([10]) Along ∂X we have:

$$D = G(D_{\partial X} - \nabla_\nu - \frac{1}{2}tr(II))$$

Where G is the bundle automorphism induced by clifford multiplication by the normal vector, and $D_{\partial X}$ is the Dirac operator intrinsic to ∂_X.

$Ds = 0$ gives

$$< \nabla_\nu s, s > = < D_{\partial X} s, s > - \frac{1}{2}tr(II)||s||^2$$

Thus we have proven:

Proposition 5.2.4. If X has non-negative scalar curvature, and s is a harmonic spinor, then:

$$\int_{\partial X} < D_{\partial X} s, s > \geq \int_{\partial X} \frac{1}{2}tr(II)||s||^2$$

The boundary conditions for harmonic spinors in the Atiyah-Patodi-Singer index theorem are that when $s|_{\partial X}$ is expanded in eigenfunctions of $D_{\partial X}$, only negative eigenvalues are used.

Writing s as $\Sigma_\lambda a_\lambda s_\lambda$, where the s_λ are the eigenvectors of $D_{\partial X}$, 5.2.4 becomes:

$$\Sigma_\lambda \lambda ||a_\lambda||^2 \geq \frac{1}{2}\int_{\partial X} tr(II)||s||^2$$

Since all the λ must be negative, we must have some $\lambda \geq \lambda_0 = inf(\frac{1}{2}tr(II))$. By projecting onto these eigenspaces between λ_0 and 0, we get:

Theorem 5.2.5. If $(X, \partial X)$ is spin and has non-negative scalar curvature then there is Λ, depending only on the second fundamental form, such that $dim(H) \leq N_{D^2_{\partial X}}(\Lambda)$, where H is the space of harmonic spinors on X with A-P-S boundary values, and $N_{D^2_{\partial X}}(\Lambda)$ is the dimension of the space of eigenfuctions of $D^2_{\partial X}$ below Λ.

As the index of D with APS boundary conditions is bounded by $dim(H)$ we next bound $dim(H)$.

Theorem 5.2.6. If N^n is a compact spin manifold, then for each λ there is a C_λ depending only on the curvature and injectivity radius of N, for which $N_{D^2}(\lambda) \leq C_\lambda vol(N)$

Proof. By ([4],Prop 4.20(ii)) there are A and B so that $\lambda_n(D^2) \geq \lambda_{An-B}(\Delta)$, where Δ is the laplacian on functions. Thus the bound we need follows immediately from the same statement (with different C_λ) for the laplacian on functions.

Theorem 5.2.7. ([13]) There is a constant K depending only on curvature bounds and dimension for which for any compact manifold N we have the bound:

$$\lambda_{V(\varepsilon)} \geq K\varepsilon^{-2}$$

For any $\varepsilon \leq inj\ rad$. Here $V(\varepsilon)$ is the minimal number of ε balls which cover N.

Since we have bounds on the curvature, the Heintz-Karcher comparison theorem ([13]) gives an L such that:

$$Vol(B(\varepsilon)) \geq L\varepsilon^n$$

for $B(\varepsilon)$ any such ε ball in N.

Choose a maximal family of disjoint $\frac{\varepsilon}{2}$ ball in N. By maximality the concentric balls of radius ε cover. But by disjointness there are at most

$$\frac{2^n Vol(N)\varepsilon^{-n}}{L}$$

balls. Therefore we have:

Proposition 5.2.8. There is a constant C depending only on the curvature such that, for any $\varepsilon \leq inj\ rad(N)$

$$\lambda_{Cvol(N)\varepsilon^{-n}} \geq \epsilon^{-2}$$

Turning this around to an upper bound on the spectral counting function gives:

Proposition 5.2.9. There is a constant C depending only on the curvature, and a λ_0 depending only on curvature and injectivity radius, such that for any $\lambda \geq \lambda_0$

$$N(\lambda) \leq CVol(N)\lambda^{\frac{n}{2}}$$

This gives the C_λ we needed for 5.2.6

$C_\lambda = C\lambda^{\frac{n}{2}}$ works for $\lambda \geq \lambda_0$.

and

$C_\lambda = C\lambda_0^{\frac{n}{2}}$

works for $\lambda \leq \lambda_0$, as

$N(\lambda) \leq N(\lambda_0)$

$\square$

Since we have now bounded all the terms in 5.1.4 linearly in the volume, 5.1.3 proves:

Theorem 5.2.10. If M has non-negative scalar curvature then $\hat{A} = 0$ in $H_0^{uf}(M)$.

Corollary 5.2.11. If $\hat{A} \neq 0$ in $H_0^{uf}(M)$ then M does not admit a metric of non-negative scalar curvature and bounded geometry within its bounded distortion class.

Remark 7. This does not quite fit into the standard C^* algebra framework. The difficulty comes from the fact that $H_0^{uf}(M)$ is not Hausdorff in the topology it gets from the L^∞ norms on cycles. One can form reduced groups $\bar{H}_*^{uf}$ by quotienting out by the closure of 0 (equivalently one can define the groups as cycles mod the closure of the boundaries). The analogue of 5.2.10 for the reduced groups can also be proven by the methods of [3]. The reduced $\hat{A}$ genus is more compatable with C^* algebra methods, which is useful if one wants to generalize 5.2.10 to higher $\hat{A}$ classes and to families. Unfortunately, however, we will see examples in the next section which show that this reduction does lose information.

5.3. Infinite Connected Sums

We use 5.2.11 to give a converse to a theorem of Block and Weinberger, which itself grew out of a converse to a theorem of Roe [19]. Given M of bounded geometry, $S \subset M$ uniformly discrete, and N compact, we form the manifold $M\#_S N$ by connected summing at each point of S with a copy of N. This manifold is quasi-isometric to M, but the local topology has become more complicated.

Theorem 5.3.1. ([2]) Let M be a spin manifold of bounded geometry and uniformly positive scalar curvature, and N a compact spin manifold of the same dimension with $\hat{A}(N) \neq 0$. If $S \subset N$ is uniformly discrete with $[S] = 0$ in $H_0^{uf}(M)$ then the manifold $M\#_S N$ admits a metric of uniformly positive scalar curvature and bounded geometry.

Clearly $\hat{A}(M\#_S N) = \hat{A}(M) + (\hat{A}(N))[S]$.

According to 5.2.10 $\hat{A}(M) = 0$,

Thus if $M\#_S N$ admits a metric of non-negative scalar curvature then $[S] = 0$.

So we have the converse to 5.3.1:

Theorem 5.3.2. Let M and N be as above, and $S \subset M$ uniformly discrete. The manifold $M\#_S N$ admits a metric of uniformly positive scalar curvature iff $[S] = 0$ in $H_0^{uf}(M)$.

Theorem 5.2.10 gives a stronger statement: when $[S] \neq 0$ there is not even a metric of non-negative scalar curvature. It would be nice to have an extension of 5.3.1 to the non-negative case. To explain the difficulties in achieving this we first explain the methods used to prove 5.3.1. The notation is as above.

Lemma 5.3.3. If M is non-compact, spin, and has scalar curvature bounded from below by a positive number then $M\#N$ admits a metric of positive scalar curvature.

Proof. Choose a sequence of points $\{p_i \in M\}$ heading to infinity with $d(p_i, p_{i+1}) = 1$. Let N' be N with the opposite orientation.

Consider the manifold $X = M\#_{\{p_{even}\}}N\#_{\{p_{odd}\}}N'$. By grouping the $\{p_i\}$ as (p_0, p_1), (p_2, p_3), ... and observing $N\#N'$ is spin null cobordant we see that X is bounded geometry spin cobordant to M. However, by grouping the $\{p_i\}$ as p_0,(p_1, p_2),(p_3, p_4), ... we can likewise see that X is bounded geometry spin null cobordant to $M\#N$. Thus M and $M\#N$ are bounded geometry spin cobordant. So what we need is a way to propagate positive scalar curvature across cobordisms, or what is equivalent, across surgeries.

There is a surgery theory for manifolds of positive scalar curvature, due to Gromov anf Lawson [14]. The idea to keep in mind for us is that to keep positive scalar curvature while doing surgery, the surgery must be done using a small tubular neighborhood of the sphere to be surgered. This is so that the positive curvature from the linking spheres is large enough to overcome any negative curvature present on the sphere being surgered. How small this may need to be depends on how small the scalar curvature of M is nearby. Thus given the unform lower bound in the lemma there is no trouble surgering M to $M\#N$ while keeping bounded geometry and positive scalar curvature.

□

One can now prove 5.3.1 by constructing "tails" running to infinity from each point of S, so long as only a uniformly bounded number pass through any ball of given radius. Such tails exist iff $[S] = 0$ in H_0^{uf} ([3]).

It is now clear what goes wrong if the scalar curvature of M is not uniformly positive: the surgeries going off to infinity must be done with smaller and smaller tubular neighborhoods of the spheres and thus distort the geometry by larger and larger amounts.

Remark 8. These examples show that one really needs H_0^{uf} not just $\bar{H}_0^{uf}$ (see the remark at the end of the last section). For example, if $M = \mathbb{R} \times S^3$, $N = K3$, and $S = \{n^2, n \in \mathbb{N}\} \times \{pt\}$ then $\hat{A}(M \#_S N) = 0$ in $\bar{H}_0^{uf}$ but not in H_0^{uf}. This means these invariants do not work nicely with the C^* algebra machinery which causes real difficulty, for example, if one wants to know whether $(M \#_S N) \times \mathbb{R}^i$ carries a metric of positive scalar curvature.

REFERENCES

[1] M. Atiyah, V. Patodi, and I. Singer, *Spectral asymmetry and riemannian geometry 1*, Math. Proc. Cambridge Philosophical Society, 77 (1975).

[2] J. Block and S. Weinberger, *Aperiodic tilings, positive scalar curvature, and amenability*, Journal of the AMS, 5 (1992).

[3] ——, *Large scale homology theories and geometry*, AMS/IP Studies in Advanced Mathematics, 2 (1997).

[4] M. Bordoni, *Spectral estimates for schroedinger and dirac-type operators on riemannian manifolds*, Math. Ann., 298 (1994).

[5] J. Cheeger and M. Gromov, *On the characteristic numbers of complete manifolds of bounded curvature and finite volume*, Differential Geometry and Complex Analysis, Springer, (1985).

[6] ——, *l_2-cohomology and group cohomology*, Topology, 25 (1986).

[7] G. Folland, *Real Analysis. Modern Techniques and their Applications*, Wiley, 1984.

[8] S. Gersten, *Bounded cocycles and combings of groups*, International J. of Algebra and Computation, 2 (1992).

[9] ——, *A homological characterization of hyperbolic groups*, Preprint, (1998).

[10] P. Gilkey, *On the index of geometrical operators for riemannian manifolds with boundary*, Adv. Math., 102 (1993).

[11] J. Graver and M. Watkins, *Combinatorics with Emphasis on the Theory of Graphs*, Springer-Verlag, 1977.

[12] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric Group Theory, LMS lecture notes, 182 (1993).

[13] ——, *Metric Structures on Riemannian and Non-Riemannian Spaces*, preprint, 1998.

[14] M. Gromov and H. Lawson, *The classification of simply connected manifolds of positive scalar curvature*, Annals of Math(2), 111 (1980).

[15] P. Halmos, *Naive Set Theory*, Springer-Verlag, 1973.

[16] H. Lawson and M. Michelsohn, *Spin Geometry*, Princeton University Press, 1989.

[17] V. Olshanskii, *On the question of the existence of an invariant mean on a group*, Uzpekhi Mat. Nauk., 35 (1980).

[18] M. Ramachandran, *Von neumann index theorems for manifolds with boundary*, Journal of the Differential Geometry, 38 (1993).

[19] J. Roe, *An index theorem on open manifolds. i,ii.*, Journal of the Differential Geometry, 27 (1988).

[20] J. Serre, *Trees*, Springer-Verlag, 1980.

[21] J. Tits, *Free subgroups in linear groups*, J. Algebra, 20 (1972).