The Pennsylvania State University
The Graduate School

HOMOLOGICAL CALCULATIONS WITH THE ANALYTIC STRUCTURE GROUP

A Thesis in
Department of Mathematics
by
Paul Siegel

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

December 2012

The thesis of Paul Siegel was reviewed and approved* by the following:

John Roe
Professor of Mathematics
Thesis Advisor, Chair of Committee

Dmitri Burago
Professor of Computer Science

Sean Hallgren
Professor of Mathematics

Nigel Higson
Professor of Mathematics

*Signatures are on file in the Graduate School.

Abstract

We build a Mayer-Vietoris sequence for the analytic structure group defined by Higson and Roe and use it to give a new proof of and generalizations of Roe's partitioned manifold index theorem. We give applications of the generalized partitioned manifold index theorem to the theory of positive scalar curvature invariants. Finally, we construct an analogue of the Kasparov product for the analytic structure group and examine how positive scalar curvature invariants behave with respect to this product.

Table of Contents

Acknowledgments

My most substantial debt of gratitude is to my parents and my sister who have trusted and supported me in innumerable ways for my entire life. I would never have made it this far without them. I owe a similar debt to my many friends over the years, among whom I would like to single out Nate Walton, my best friend of almost 15 years. All of my friends have provided me with seemingly limitless patience, generosity, empathy, and good humor, but Nate has suffered the burden longer and more enthusiastically than anyone else.

I must also thank my many brilliant teachers over the years; it would be impossible to sustain my curiosity if it were not channelled in productive directions by more knowledgeable people. This includes my fourth grade teacher, Sharon Viol, who was the first person to show me that mathematics can be an enjoyable experience. It also includes my high school math teacher, Ceci Anderson, and my high school physics teacher, John Plough, who responded to my endless stream of half-baked questions with wisdom and encouragement. Among my college professors I have to single out Brian Conrad, who I could always count on to give me detailed explanations of just about anything mathematical and whose guidance has been enormously helpful throughout the years.

I have many teachers to thank at Penn State. Some of my best teachers have been fellow students: Rufus Willett, Tyrone Crisp, Shillin Yu, Alan Yashinski, and Yanli Song, to name a few. In spite of the fact that he does not work in my area (or perhaps because of it), Vaughn Climenhaga's penetrating questions, crazy ideas, and unique insight have been very energizing for me over the years. Among the faculty, I must single out Dmitri Burago and Nigel Higson for nurturing my abilities and sharing their knowledge - I have learned as much from the two of them as I have from nearly anyone else.

I wrote "nearly" anyone else because there is a single person who has taught me more about mathematics than anybody: my thesis adviser, John Roe. My officemate, Shawn Ryan, often expressed jealousy over the fact that I would often return from my meetings with my adviser with an expression of contented enlightenment

on my face. I would be hard pressed to identify a part of my mathematical knowledge that he hasn't heavily influenced, and there are entire areas of mathematics that I would probably never have encountered but for his guidance. His mentorship extended beyond just mathematical understanding: he heavily influenced my development as a teacher and as a communicator.

There are certainly many others who belong on this list, so for them I offer one last sincere expression of gratitude: thank you all.

Chapter 1

Introduction

Let S be a closed surface. When describing the geometry of S, it is useful to distinguish between two types of invariants: *global* and *local*. For closed surfaces there is essentially only one global invariant: the genus, or "number of holes" in S. The genus is global in the sense that it doesn't change if one stretches or deforms the surface. The typical example of a local invariant, on the other hand, is the Gaussian curvature function $K\colon S \to \mathbb{R}$. This is a smooth function which at any given point $p \in S$ measures the "fatness" or "thinness" of an infinitesimal triangle along the surface near p as compared to a Euclidean triangle with the same side lengths. It is local in the sense that $K(p)$ depends only on the geometry of S in an arbitrarily small neighborhood of p.

One of the most basic problems in geometry is to relate the local invariants of a geometric object to its global invariants, and in the context of the differential geometry of surfaces this is achieved using the classical Gauss-Bonnet formula:

$$\chi_S = \frac{1}{2\pi} \int_S K$$

where χ_S is the *Euler characteristic* of S, given by two minus twice the genus. Other kinds of geometry have similar formulas which achieve a similar objective; for instance algebraic geometry has the Riemann-Roch formula and algebraic topology has the Hirzebruch signature theorem.

In the first half of the twentieth century Hodge discovered that many global invariants, including the Euler characteristic, count the number of solutions to certain systems of partial differential equations which arise naturally in various geometric settings. To be more specific, begin with a self-adjoint differential operator D on a smooth compact manifold M. If D is *elliptic*, meaning it is “approximately invertible” in a certain sense, then its kernel and cokernel are finite dimensional spaces. An operator with this property is said to be *Fredholm*, and associated to any Fredholm operator is an important invariant called the *Fredholm index*:

$$\mathrm{Index}(D) = \dim\ker(D) - \dim\mathrm{coker}(D)$$

The index of an elliptic operator on a compact manifold is a very stable invariant: it is constant along continuous paths of elliptic operators, for instance. Consequently indices of elliptic operators which arise naturally in geometry are often global invariants of the underlying manifold, and it is not obvious that they can be calculated in terms of purely local data. Atiyah and Singer did just that in the 1960's ([3]), arriving at the following general formula:

$$\mathrm{Index}(D) = \int_{T^*M} ch(D)Todd(TM \otimes \mathbb{C})$$

Here $ch(D)$ and $Todd(TM \otimes \mathbb{C})$ are characteristic classes, the first of which is constructed using the ellipticity of D in a crucial way. Thus the integrand can be expressed as an explicit differential form.

Not long after the Atiyah-Singer index theorem was proved, a considerable body of literature emerged with the aim of extending, applying, and generalizing it. One important challenge was to extend the tools of index theory to non-compact manifolds; aside from standard examples such as Euclidean space, there was considerable interest in equivariant elliptic operators on spaces equipped with cocompact group actions and in elliptic operators acting on the leaves of foliated manifolds. One difficulty with elliptic operators on non-compact manifolds is that they often fail to be Fredholm and thus it is not obvious how to even pose an intelligent index problem. One systematic strategy for overcoming this difficulty was developed by Roe in the 1980's (beginning with [19] using large-scale geometry.

The following is a brief sketch of Roe's idea. A classical theorem of Atkinson (see [6], for instance) asserts that a bounded operator on Hilbert space is Fredholm if and only if it is invertible modulo compact operators. Roe observed that if D is a suitable elliptic operator on a complete Riemannian manifold M then D determines an operator on the Hilbert space $L^2(M)$ which is invertible modulo a well-chosen algebra of bounded operators which contains the compact operators. This algebra, called the *coarse algebra* of M and denoted by $C^*(M)$, has the structure of a C*-algebra and it depends only on the geometry of M "at infinity" in a sense determined by the metric. Associated to every C*-algebra is a system of algebraic invariants called its *K-theory* groups, and Roe showed that the K-theory groups of the coarse C*-algebra provide a natural home for a *coarse index* of D. If M happens to be compact then its coarse C*-algebra is simply the C*-algebra of compact operators whose K_0 group is isomorphic to $\mathbb{Z}$, and the coarse index of D is simply the ordinary Fredholm index.

The coarse C*-algebra associated to a complete Riemannian manifold M fits into a very convenient general framework for doing index theory. It is an ideal in another C*-algebra called the *structure algebra* of M, denoted by $D^*(M)$. $D^*(M)$ depends on both the large- and small-scale structure of M, and it turns out that the quotient $Q^*(M) = D^*(M)/C^*(M)$ depends only on small-scale geometry. Indeed, the K-theory of $Q^*(M)$ is a model for the *K-homology* of M (up to a customary shift in degree), i.e. the generalized homology theory which is naturally dual to Atiyah and Hirzebruch's topological K-theory. This yields a long exact sequence joining the K-theory of $C^*(M)$, the K-theory of $D^*(M)$, and the K-homology of M:

$$\ldots \to K_{p+1}(D^*(M)) \to K_p(M) \xrightarrow{\partial} K_p(C^*(M)) \to \ldots \tag{1.0.1}$$

In chapter 3 we will show that elliptic operators determine elements of the K-homology groups of M, and in Chapter 4 we will explain that the map $\partial\colon K_p(M) \to K_p(C^*(M))$ sends the K-homology class of an elliptic operator to its coarse index. This map, which passes from invariants which are built out of the local structure of M on the left-hand side to invariants which are built of the global structure of M on the right-hand side, is an algebraic model for the passage from local invariants to global invariants in the Atiyah-Singer index theorem.

It is useful to consider an elaboration on the long exact sequence (1.0.1). Suppose that M is a compact Riemannian manifold, G is a countable discrete group, and $\widetilde{M}$ is a locally isometric G-cover of M. Then any elliptic operator on M lifts to a G-equivariant elliptic operator on $\widetilde{M}$ (which may be non-compact!), and the lifted operator is invertible modulo an equivariant counterpart of the coarse algebra denoted by $C^*_G(\widetilde{M})$. This C*-algebra is an ideal in an equivariant counterpart of the structure algebra, denoted by $D^*_G(\widetilde{M})$, and it turns out that, once again, the K-theory of $D^*_G(\widetilde{M})/C^*_G(\widetilde{M})$ is a model for the K-homology of M.

Since M is compact, $\widetilde{M}$ has the same large-scale geometric structure as the group G (equipped with a word metric) so one may expect that $C^*_G(\widetilde{M})$ depends only on G. At least at the level of K-theory we have an isomorphism:

$$K_p(C^*_G(\widetilde{M})) \cong K_p(C^*_r(G))$$

where $C^*_r(G)$ is a C*-algebra familiar to operator algebraists and representation theorists called the *reduced group C*-algebra* of G. Thus the equivariant analogue of the long-exact sequence (1.0.1) takes the form:

$$\ldots \to S_p(\widetilde{M}, G) \to K_p(M) \to K_p(C^*_r(G)) \to \ldots \tag{1.0.2}$$

Here $S_p(\widetilde{M}, G)$, called the *analytic structure group* of the pair (M, G), is given by the K-theory of $D^*_G(\widetilde{M})$ (with the same customary degree shift as above). In the case where $M = BG$ is a classifying space for G and $\widetilde{M} = EG$ is its universal cover, the resulting map $K_p(BG) \to K_p(C^*_r(G))$ is called the *Baum-Connes assembly map* and the *Baum-Connes conjecture* ([16]) asserts that this map is an isomorphism. This conjecture, which is known to be true for many groups G, has important implications in geometry, topology, and functional analysis. The long exact sequence (1.0.2) itself has an important topological interpretation: Higson and Roe showed in [10], [11], and [12] that it is closely related to the topologists' *surgery exact sequence*. For our purposes it is sufficient to view $K_p(M) \to K_p(C^*_r(G)$ as a generalized index map which passes from local structure to global structure, mediated by the analytic structure group.

With this algebraic machinery in hand we take up an important index theorem for *partitioned manifolds* due to Roe in chapter 5. We say that a manifold M is *partitioned* by a hypersurface N if M is the union of two submanifolds-with-boundary M^+ and M^- whose boundaries are both N. Roe proved in [20] that if N is compact and D is an elliptic operator on M which has a favorable local structure in a neighborhood of N then the coarse index of D is given by the ordinary index of its restriction to N. The main result of chapter 5 is a generalization of this theorem to partitioned manifolds with non-compact hypersurfaces and to a partitioned manifold index theorem for equivariant indices.

The main result is appealing due both to its greater generality and the conceptual nature of our proof. Our strategy is to us the *Mayer-Vietoris sequence*, a tool originating in algebraic topology which relates the homology groups of a space to the homology groups of smaller pieces. Specifically, if X is a space and $X = Y_1 \cup Y_2$ where Y_1 and Y_2 are appropriate subspaces then the Mayer-Vietoris sequence takes the form:

$$\ldots \to H_n(Y_1 \cap Y_2) \to H_n(Y_1) \oplus H_n(Y_2) \to H_n(Y) \to H_{n-1}(Y_1 \cap Y_2) \to \ldots$$

In chapter 5 we will use an abstract construction in K-theory for C*-algebras to build Mayer-Vietoris sequences for the K-theory of the coarse algebra, the K-theory of the structure algebra, and K-homology. We will carry out this construction equivariantly with respect to a free and proper group action, and we will fit the three Mayer-Vietoris sequences together using the analytic surgery exact sequence (1.0.2). The result is the following braid diagram:

$$\begin{tikzcd}
K_p(Y_1) \oplus K_p(Y_2) \arrow[r, bend left] \arrow[dr] & K_p(C^*_G(Y_1)) \oplus K_p(C^*_G(Y_2)) \arrow[r, bend left] \arrow[dr] & K_p(C^*_G(X)) \arrow[r, bend left] \arrow[dr] & S_{p-1}(\widetilde{X}, G) \\
K_p(C^*_G(Y_1 \cap Y_2)) \arrow[ur] \arrow[dr] & K_p(X) \arrow[ur] \arrow[dr] & S_{p-1}(\widetilde{Y}_1, G) \oplus S_{p-1}(\widetilde{Y_2}, G) \arrow[ur] \arrow[dr] & K_{p-1}(C^*_G(Y_1 \cap Y_2)) \\
S_p(\widetilde{X}, G) \arrow[ur] \arrow[r, bend right] & S_{p-1}(\widetilde{Y}_1 \cap \widetilde{Y}_2, G) \arrow[ur] \arrow[r, bend right] & K_{p-1}(Y_1 \cap Y_2) \arrow[ur] \arrow[r, bend right] & K_{p-1}(Y_1) \oplus K_{p-1}(Y_2)
\end{tikzcd}$$

The diagram interweaves analytic surgery exact sequences with Mayer-Vietoris sequences; it is exact by the naturality of long exact sequences in K-theory.

A partitioned manifold comes naturally equipped with a decomposition suitable for this diagram, and our proof of the partitioned manifold index theorem uses the sub-diagram

$$\begin{tikzcd} K_p(X) \arrow[r] \arrow[d] & K_p(C^*_G(X)) \arrow[d] \\ K_{p-1}(Y_1 \cap Y_2) \arrow[r] & K_{p-1}(C^*_G(Y_1 \cap Y_2)) \end{tikzcd}$$

together with some calculations with the Mayer-Vietoris boundary maps. Iterating our proof of the partitioned manifold index theorem yields an index theorem for *k-partitioned manifolds* for which the submanifold N has codimension k instead of just 1.

Finally, in chapter 6 we discuss applications of our results to the theory of positive scalar curvature invariants in Riemannian geometry. If M is a Riemannian manifold which has a higher orientation structure called a *spin structure* then there is a specific elliptic operator on M, called the *spinor Dirac operator*, which according to a theorem of Lichnerowicz (see the book [14]) has a close relationship to the scalar curvature function of M. In particular, if the scalar curvature is positive then the index of this operator vanishes and, according to the long exact sequence (1.0.2), it determines a class in the analytic structure group of M called the *structure invariant* of the positive scalar curvature metric on M. Using the structure invariant and the index theorem for k-partitioned manifolds we give a new proof of a theorem of Gromov and Lawson ([7]):

Theorem 1.0.1. *If M is a compact manifold which admits a Riemannian metric of nonpositive sectional curvature then it admits no Riemannian metric of positive scalar curvature.*

We conclude the thesis with some preliminary calculations involving structure invariants and the Mayer-Vietoris sequence for the structure group. If M is a Riemannian manifold with positive scalar curvature then its suspension SM also has positive scalar curvature, and it is natural to expect that if M is spin then the Mayer-Vietoris boundary map for the structure group sends the positive scalar curvature invariant for SM to the positive scalar curvature invariant for M.

This is analogous to the fact that the Mayer-Vietoris boundary map for K-homology sends the K-homology fundamental class of the suspension a spin^c manifold M to the K-homology fundamental class of M. One proof of this fact makes crucial use of the Kasparov product in K-homology, and in an attempt to imitate this argument we construct a counterpart of the Kasparov product for the analytic structure group. We then prove that this product structure is compatible with positive scalar curvature invariants in an appropriate sense; if we could prove that it is also compatible with the Mayer-Vietoris sequence then we could establish the desired relationship between positive scalar curvature invariants and the suspension construction. The full details of this calculation elude us, but we hope to complete them in future work.

Chapter 2

Preliminaries

2.1 Background in C*-algebra Theory

The goal of this thesis, broadly speaking, is to use functional analysis to solve problems in geometry and topology. The basic functional-analytic structure which we will use to accomplish this is a C*-algebra, an object which encapsulates both algebras of bounded operators on Hilbert space and algebras of continuous functions on topological spaces. C*-algebras were first introduced as algebras of Hilbert space operators by Von Neumann in the 1930's in order to provide a rigorous framework for Heisenberg's algebraic formulation of quantum mechanics. Later, in the 1940's, Gelfand and Naimark provided the theory with an axiomatic foundation and began to study C*-algebras as purely abstract objects.

The celebrated Gelfand-Naimark theorem asserts that every commutative C*-algebra has the form $C(X)$ where X is a locally compact Hausdorff topological space. This set the stage for Connes' noncommutative geometry program in the 1980's which aims to extend the tools of geometry and topology from commutative C*-algebras to noncommutative C*-algebras. Connes' program relies heavily on connections between C*-algebra theory and the Atiyah-Singer index theorem, and these connections are the direct motivation for introducing C*-algebra theory here.

2.1.1 C*-algebras and Representations

We begin with Gelfand and Naimark's abstract definition of a C*-algebra. Much of what follows is described in detail in [6].

Definition 2.1.1. *Let A be an algebra over the complex numbers equipped with a norm $\|\cdot\|$.*

- *A is a* Banach Algebra *if $\|xy\| \leq \|x\| \|y\|$ for all x and y in A and if A is complete with respect to its norm.*
- *A is a* C*-Algebra *if it is a Banach algebra equipped with a linear involution $* : A \to A$ which satisfies the* C*-identity $\|a^*a\| = \|a\|^2$ *for every a in A.*
- *A is* unital *if it has a multiplicative identity element* 1.

A $*$-*homomorphism* between two C*-algebras is an algebra homomorphism which preserves the involution. There is a category whose objects are C*-algebras and whose morphisms are $*$-homomorphisms. Every $*$-homomorphism automatically preserves the norm and hence is injective.

A *C*-Subalgebra* of a C*-algebra A is a norm closed subalgebra of A which is closed under the involution $*$ of A. A *C*-ideal* is defined to be an ideal in A which has the structure of a C*-subalgebra; we will rarely need to consider ideals which are not C*-ideals, so from now on we will simply refer to them as ideals. If A is a C*-algebra and J is an ideal in A then A/J carries the structure of a C*-algebra.

If A is (possibly non-unital) C*-algebra then $A \oplus \mathbb{C}$ naturally carries the structure of a unital C*-algebra called the unitalization of A, often denoted $\widetilde{A}$. It can be regarded as the C*-algebra generated by A and a multiplicative identity element which commutes with A. It contains A as a ideal.

Example 2.1.2. *Let $A = \mathbb{B}(H)$ equipped with the operator norm; it is elementary that A is a Banach algebra. A has an involution $*$ which sends an operator to its adjoint, and the C*-identity is a classical result in Hilbert space theory. So $\mathbb{B}(H)$ - as well as any C*-subalgebra of $\mathbb{B}(H)$ - is a C*-algebra. The space $\mathbb{K}(H)$ of compact operators on H is an ideal in $\mathbb{B}(H)$, and the quotient $\mathbb{B}(H)/\mathbb{K}(H)$ is a C*-algebra called the* Calkin algebra *of H.*

Example 2.1.3. *Let X be a locally compact Hausdorff space and let $A = C_0(X)$ be the algebra of complex valued continuous functions on X equipped with the uniform norm. A is a Banach algebra, and it carries an involution $*$ which sends a function to its complex conjugate. The C*-identity is trivial for A, so A is a C*-algebra. Note that multiplication of continuous functions is commutative, so $C_0(X)$ is a commutative C*-algebra.*

According to the following important results of Gelfand and Naimark, these two examples are each universal in an appropriate sense.

Theorem 2.1.4 (Gelfand-Naimark Theorem 1)**.** *Every C*-algebra is isometrically $*$-isomorphic to a C*-subalgebra of $\mathbb{B}(H)$ for some Hilbert space H.*

Theorem 2.1.5 (Gelfand-Naimark Theorem 2)**.** *Every commutative C*-algebra A is isometrically $*$-isomorphic to $C_0(X)$ for a unique locally compact Hausdorff space X. A is unital if and only if X is compact.*

Proofs of both of these theorems appear in [6]; we will not comment further on the proofs here. Theorem 2.1.4 justifies importing language from Hilbert space theory into C*-algebra theory. For example, if A is a unital one defines the *spectrum* of an element a of a C*-algebra to be the set of all complex numbers λ with the property that $a - \lambda 1$ is not invertible. An element $a \in A$ is said to be *positive* (written $a \geq 0$ if its spectrum consists only of real numbers.

If A is unital a is a normal element of A, meaning a commutes with its adjoint, then the C*-subalgebra $\langle a \rangle$ of A generated by a, a^*, and 1 is a commutative unital C*-algebra. Thus it has the form $C(X)$ for some compact Hausdorff space X; one can check that X is in fact the spectrum of a. This is the content of the spectral theorem for normal operators. The element of $\langle a \rangle$ which corresponds to a continuous function f on the spectrum of a is denoted by $f(a)$, and the isometrically $*$-isomorphic assignment $f \mapsto f(a)$ is called the *functional calculus* for a.

We will need a few basic facts about the representation theory of C*-algebras, so we will record them here.

Definition 2.1.6. *Let A be a C*-algebra and let H be a Hilbert space. A* representation *of A on H is $*$-homomorphism*

$$\rho : A \to \mathbb{B}(H)$$

ρ is said to be nondegenerate *if $\rho(A)H$ is dense in H.*

Every $*$-homomorphism between C*-algebras is isometric, so in particular the same is true for any representation. Every C*-algebra A has a representation ρ by Theorem 2.1.4, and by restricting H to the subspace $\rho(\bar{A})H$ we can always construct a nondegenerate representation. The existence of representations of C*-algebras allows us to construct a C*-algebraic notion of tensor product. Let A and B be C*-algebras, and assume they are represented on Hilbert spaces H_A and H_B, respectively. The algebraic tensor product of H_A and H_B can be completed with respect to the inner product $\langle v_A \otimes v_B, w_A \otimes w_B \rangle = \langle v_A, w_A \rangle \langle v_B, w_B \rangle$ to form a Hilbert space $H_A \otimes H_B$, and $T_A \otimes T_B$ is a bounded operator on $H_A \otimes H_B$ whenever $T_A \in \mathbb{B}(H_A)$ and $T_B \in \mathbb{B}(H_B)$. Thus the algebraic tensor product of A and B admits a representation on $\mathbb{B}(H_A \otimes H_B)$.

Definition 2.1.7. *The* minimal *tensor product of A and B is the C*-algebra obtained as the completion of the algebraic tensor product of A and B regarded as a subalgebra of $\mathbb{B}(H_A \otimes H_B)$.*

Remark 2.1.8. *The minimal tensor product is independent of the representations used to define it. There are many other notions of tensor products of C*-algebras, but we will only be concerned with the minimal one.*

Representations of commutative C*-algebras arise in a particularly natural way thanks to Theorem 2.1.5. Every separable locally compact Hausdorff space X admits a Borel measure μ, and the map $C_0(X) \to L^2(X, \mu)$ given by pointwise multiplication operators is a representation. If μ is chosen so that no open set has measure 0 then this representation is nondegenerate. Every representation of $C_0(X)$ is unitarily equivalent to a direct sum of representations of the form $L^2(X, \mu)$, and consequently every representation of $C_0(X)$ extends to a representation of the space of bounded Borel functions on X.

We conclude this section with a technical lemma concerning the behavior of the restriction of a representation ρ of a C*-algebra A on a Hilbert space H to a closed subspace H_1 of H. Let H_2 denote its orthogonal complement, so that ρ is a $*$-homomorphism $\rho : A \to \mathbb{B}(H_1 \oplus H_2)$. Fix an operator $T \in \mathbb{B}(H_1)$ and consider the C*-subalgebra $\mathbb{B}(T)$ of $\mathbb{B}(H_1)$ consisting of those operators S such that S and S^* commute with T modulo compact operators. Let $\mathbb{K}(T)$ denote the set of all operators $S \in \mathbb{B}(T)$ such that ST and S^*T are both compact operators; it is easy to see that $\mathbb{K}(T)$ is an ideal in $\mathbb{B}(T)$. For example if T is the identity then $\mathbb{B}(T) = \mathbb{B}(H_1)$ and $\mathbb{K}(T) = \mathbb{K}(H_1)$.

Lemma 2.1.9. *Express $\rho : A \to \mathbb{B}(H_1 \oplus H_2)$ as:*

$$\rho = \begin{pmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{pmatrix}$$

If ρ_{11} is a $*$-homomorphism modulo $\mathbb{K}(T)$, meaning $\rho_{11}(aa') - \rho_{11}(a)\rho_{11}(a') \in \mathbb{K}(T)$ for every $a, a' \in A$, then $T\rho_{12}(a) : H_2 \to H_1$ and $\rho_{21}(a)T : H_1 \to H_2$ are compact operators for every $a \in A$.

Proof. We can dilate H_1 and H_2 so that they are both isomorphic as Hilbert spaces to some larger Hilbert space, and we can extend ρ and T to this larger space. So we can assume without loss of generality that $H_1 = H_2$. Since ρ is a $*$-homomorphism we have the identity $\rho(aa^*) = \rho(a)\rho(a)^*$; viewing both sides of this equation as 2×2 matrices and comparing the upper left entries we get $\rho_{11}(aa^*) = \rho_{11}(a)\rho_{11}(a)^* + \rho_{12}(a)\rho_{12}(a)^*$. Since ρ_{11} is a $*$-homomorphism modulo $K(T)$ we have that $\rho_{12}(a)\rho_{12}(a)^* \in K(T)$ and thus $\rho_{12}(a)\rho_{12}(a)^*$ projects to 0 in the quotient C*-algebra $\mathbb{B}(T)/\mathbb{K}(T)$. By the C*-identity in this algebra it follows that $\rho_{12}(a) \in K(T)$ and thus $T\rho_{12}(a)$ is compact. The fact that $\rho_{21}(a)T$ is compact follows from the same argument applied to the identity $\rho(a^*a) = \rho(a)^*\rho(a)$. $\square$

2.2 Completely Positive Maps

Let X be a compact Hausdorff space, $Y \subseteq X$ a closed subspace. The inclusion $Y \hookrightarrow X$ induces a surjective $*$-homomorphism $C(X) \to C(Y)$ which sends a continuous function on X to its restriction to Y. This map fits into a short exact sequence

$$0 \to C_0(X - Y) \to C(X) \to C(Y) \to 0$$

There are in general topological obstructions to splitting this short exact sequence by a $*$-homomorphism $C(Y) \to C(X)$. The same phenomenon occurs in many other contexts in C*-algebra theory, notably in Brown, Douglas, and Filmore's theory of extensions, and it complicates many natural algebraic constructions. To overcome these difficulties we introduce a kind of map between C*-algebras more general than $*$-homomorphisms; these maps figure prominently in two deep results in C*-algebra theory which will become essential later on.

Definition 2.2.1. *A bounded, unital, linear map $\sigma : A \to B$ between unital C*-algebras is* completely positive *if for every collection of elements $a_1, \ldots, a_n \in A$ and $b_1, \ldots, b_n$ we have:*

$$\sum_{i,j} b_i^* \sigma(a_i^* a_j) b_j \geq 0$$

Remark 2.2.2. *Every $*$-homomorphism is completely positive, but the most important examples of completely positive maps are not.*

The language of completely positive maps is in fact overkill for our humble goal of splitting short exact sequences of commutative C*-algebras, but there is a substantial amount of theory built around the concept which we wish to tap into. Our first order of business is to prove that a restriction map $C(X) \to C(Y)$ described above admits a completely positive section.

Let $A = C(X)$ where X is a compact Hausdorff space and let B be any unital C*-algebra. Given a finite set $x_1, \ldots, x_k \in X$ a collection $P_1, \ldots, P_k$ of positive elements of B whose sum is 1, we can define a positive linear map $A \to B$ by $f \mapsto \sum_{i=1}^k f(x_i) P_i$. We will refer to such maps as *positive discrete maps.*

Lemma 2.2.3. *Any positive unital linear map $\sigma : A \to B$ is the point-norm limit of a sequence of positive discrete maps $\sigma_m : A \to B$, meaning $\|\sigma(f) - \sigma_m(f)\| \to 0$ as $m \to \infty$.*

Proof. For each m choose a finite cover of X by open sets of diameter no larger than $\frac{1}{m}$ and choose a collection of points $x_1, \ldots, x_{k_m}$, one from each set in the cover. Let $h_1, \ldots, h_{k_m}$ be a partition of unity subordinate to the cover. Set $P_i = \sigma(h_i)$ for each i from 1 to k_m; P_i is positive since σ is positive. Thus the map $\sigma_m : A \to B$

given by $\sigma_m(f) = \sum_{i=1}^{k_m} f(x_i)P_i$ is a positive discrete map. We have:

$$\begin{aligned}\|\sigma(f) - \sigma_m(f)\| &= \left\|\sigma(f) - \sum_i f(x_i)\sigma(f_i)\right\| \\ &= \left\|\sum_i (f_i\sigma(f) - f(x_i)\sigma(f_i))\right\| \\ &= \left\|\sigma(\sum_i (f - f(x_i))f_i)\right\|\end{aligned}$$

Since the diameter of the open cover tends to 0 as m tends to infinity, $\sum_{i=1}^{k_m}(f - f(x_i))f_i$ tends to 0 uniformly. Thus $\|\sigma(f) - \sigma_m(f)\| \to 0$ since σ is continuous. □

The restriction map $C(X) \to C(Y)$ has a positive section by a variation on the Tietze extension theorem, so we will prove that completely positive sections exist by, in effect, showing that every positive unital linear map $C(Y) \to C(X)$ is completely positive.

Proposition 2.2.4. *If X is a compact Hausdorff space and $Y \subseteq X$ is a closed subspace then the restriction map $C(X) \to C(Y)$ admits a completely positive section.*

Proof. Choose any positive unital linear map $C(Y) \to C(X)$ which extends a continuous function on Y to a continuous function on X. By the previous lemma, this map is the point-norm limit of positive discrete maps; complete positivity is preserved under point-norm limits, so it suffices to show that any positive discrete map $\sigma : C(Y) \to C(X)$ is completely positive. Assume $\sigma(g) = \sum_{i=1}^{k} g(y_i)f_i$ where $y_1, \ldots, y_k$ is any finite set of points in Y and $f_1, \ldots, f_k$ is a collection of positive functions on X which sum to 1. We must show that $\sum_{j,k} b_j^*\sigma(a_j^*a_k)b_k \geq 0$ where $a_1, \ldots, a_n \in C(Y)$ and $b_1, \ldots, b_n \in C(X)$. By the definition of σ we have:

$$\sum_{j,k} b_j^*\sigma(a_j^*a_k)b_k = \sum_{i,j,k} \overline{b}_j b_k \overline{a}_j(y_i)a_k(y_i)f_i$$

For any fixed pair j, k and any i we calculate:

$$\overline{b}_j b_j \overline{a}_j(y_i)a_j(y_i) + \overline{b}_j b_k \overline{a}_j(y_i)a_k(y_i) + \overline{b}_k b_j \overline{a}_k(y_i)a_j(y_i) + \overline{b}_k b_k \overline{a}_k(y_i)a_k(y_i)$$

$$= b_j a_j(y_i)\overline{b_j a_j(y_i)} + 2\mathrm{Re}(b_j a_j(y_i)\overline{b_k a_k(y_i)}) + b_k a_k(y_i)\overline{b_k a_k(y_i)}$$
$$= |b_j a_j(y_i) + b_k a_k(y_i)|^2 \geq 0$$

Summing over all pairs j, k and all i, the desired positivity follows. □

Remark 2.2.5. *Passing to the one-point compactification, Proposition 2.2.4 also applies to restriction maps $C_0(X) \to C_0(Y)$ where Y is a closed subspace of a locally compact Hausdorff space X.*

The previous proposition together with the observation that $*$-homomorphisms are completely positive maps provide all of the examples of completely positive maps that we will need. Thus we turn to the results about completely positive maps that we will need to exploit later.

The first result, due to Stinespring, actually predates the definition of a completely positive map and most likely motivated it. It asserts that every completely positive map into the C*-algebra of bounded operators on a Hilbert space can be extended to a representation on a larger Hilbert space.

Theorem 2.2.6 (Stinespring's Theorem)**.** *Let A be a unital C*-algebra, H a Hilbert space, and $\sigma : A \to \mathbb{B}(H)$ a completely positive map. Then there exists a Hilbert space H', an isometry $V : H \to H'$, and a nondegenerate representation $\rho : A \to \mathbb{B}(H')$ such that $\sigma(a) = V^*\rho(a)V$ for every $a \in A$.*

Proof. See Chapter 3 of [9]. □

The second result is a deep theorem of Voiculescu which was originally discovered in order to organize extensions of C*-algebras into a group. This group turned out to be a model for the *K-homology* of a C*-algebra (which we will develop in the next chapter), and Voiculescu's result often makes a crucial appearance in other models of K-homology as well.

Theorem 2.2.7 (Voiculescu's Theorem)**.** *Let A be a unital separable C*-algebra, let H be a separable Hilbert space, and let $\rho : A \to \mathbb{B}(H)$ be a nondegenerate representation. Suppose H' is another separable Hilbert space and $\sigma : A \to \mathbb{B}(H')$ is a completely positive map with the property that $\sigma(a) = 0$ whenever $\rho(a)$ is compact for $a \in A$. Then there exists an isometry $V : H' \to H$ with the property that $\sigma(a) \sim V^*\rho(a)V$ for all $a \in A$.*

Proof. See Chapter 3 of [9] □

2.3 K-Theory

For much of this thesis our interest in C*-algebra theory will be focused on *K-theory*, a system of algebraic invariants which houses a plethora of algebraic and analytic information. K-theory was first developed by Grothendieck in the 1950's in order to generalize the Riemann-Roch formula of classical algebraic geometry. Not long after it was adapted to algebraic topology by Atiyah and Hirzebruch; thanks to Bott's famous periodicity theorem, topological K-theory satisfies Eilenberg and Steenrod's axioms for a generalized cohomology theory. Atiyah discovered a new proof of the Bott periodicity theorem which used techniques belonging to Banach algebra theory, and this quickly inspired operator algebraists to develop K-theory for C*-algebras.

Our motivation for introducing K-theory originates in Atiyah and Singer's first published proof of their index theorem (see [3]) in which they used topological K-theory as a home for index invariants. We will be interested in index invariants which are adapted to more complicated geometric contexts, such as the large-scale geometry of a space or a group action. C*-algebra K-theory provides a convenient unified framework for working in a variety of these different contexts at once. Specific applications of the theory will appear in later chapters; for now we simply review some basic facts about K-theory for abstract C*-algebras. This material is standard in textbooks on K-theory for C*-algebras; see [9] or [4], for example.

Definition 2.3.1. *Let A be a unital C*-algebra. A* projection over A *is an element $\boldsymbol{p}$ of a matrix algebra $M_n(A)$ over A with the property that $\boldsymbol{p}^2 = \boldsymbol{p}^* = \boldsymbol{p}$. A* unitary *over A is an element $\boldsymbol{u}$ of a matrix algebra $M_n(A)$ with the property that $\boldsymbol{u}^*\boldsymbol{u} = \boldsymbol{u}\boldsymbol{u}^* = 1$*

There is a natural notion of direct sum for matrices over A: if $x \in M_n(A)$ and $y \in M_{n'}(A)$ then we define $x \oplus y = \begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix}$, an element of $M_{n+n'}(A)$. Note that the direct sum of two projections over A is a projection over A.

The group $K_0(A)$ of a unital C*-algebra A is an abelian group generated by projections over A whose group law corresponds to direct sum of projections. To motivate this construction, consider the case where X is a compact Hausdorff space and $A = C(X)$. A projection over A is in general the same thing as a projective module over A, and by the Serre-Swan theorem a projective module over $C(X)$ is the same thing as a vector bundle over X. Thus the K_0 group of the C*-algebra $C(X)$ recovers Atiyah and Hirzebruch's definition of topological K-theory as an abelian group generated by vector bundles over a space. With this geometric motivation in mind, we give the precise definition of K_0.

Definition 2.3.2. *Let A be a unital C*-algebra. Define $K_0(A)$ to be the abelian group generated by projections over A with the following relations:*

- *The additive identity element is represented by the zero matrix in $M_n(A)$ for each n.*
- $[\boldsymbol{p} \oplus \boldsymbol{q}] = [\boldsymbol{p}] + [\boldsymbol{q}]$ *for any pair of projections $\boldsymbol{p}, \boldsymbol{q}$ over A.*
- *If $\boldsymbol{p}_t$ is a continuous path of projections in $M_n(A)$ for $t \in [0,1]$ then $[\boldsymbol{p}_0] = [\boldsymbol{p}_1]$. In this case we say that $\boldsymbol{p}_0$ and $\boldsymbol{p}_1$ are* homotopic *as projections.*

Remark 2.3.3. *By the first two axioms for $K_0(A)$, every element of $K_0(A)$ is a formal difference $[\boldsymbol{p}] - [\boldsymbol{q}]$ of projections in some matrix algebra $M_n(A)$.*

Remark 2.3.4. *By standard results in C*-algebra theory, two homotopic projections $\boldsymbol{p}$ and $\boldsymbol{q}$ are* unitarily equivalent *in the sense that there is a unitary $\boldsymbol{u} \in A$ such that $\boldsymbol{p} = \boldsymbol{u}\boldsymbol{q}\boldsymbol{u}^*$. Moreover there is a partial converse: if $\boldsymbol{p}$ and $\boldsymbol{q}$ are unitarily equivalent projections then $\boldsymbol{p} \oplus \boldsymbol{0}$ is homotopic to $\boldsymbol{q} \oplus \boldsymbol{0}$ for some zero matrix $\boldsymbol{0}$.*

Example 2.3.5. *Let us compute $K_0(\mathbb{C})$. To projections in $M_n(\mathbb{C})$ are homotopic if and only if they have the same rank, so the map which sends a projection over $\mathbb{C}$ to its rank determines an isomorphism $K_0(\mathbb{C}) \cong \mathbb{Z}$.*

K_0 is easily seen to be covariantly functorial for unital $*$-homomorphisms between unital C*-algebras. If $\phi : A \to B$ is a unital $*$-homomorphism then there are induced maps $\phi_n : M_n(A) \to M_n(B)$ for every n, and if $\mathbf{p}$ is a projection in $M_n(A)$ we have that $\phi_n(p)$ is a projection in $M_n(B)$. This assignment is compatible with the relations defining K-theory, so it determines a map $\phi_* : K_0(A) \to K_0(B)$.

It is immediate from the definitions that if $\phi_t : A \to B$ is a continuous family of $*$-homomorphisms for $t \in [0,1]$ then $(\phi_0)_* = (\phi_1)_*$ as homomorphisms $K_0(A) \to K_0(B)$.

If $A = C(X)$ for some compact Hausdorff space X then we will often write $K^0(X)$ instead of $K_0(C(X))$. The assignment $X \mapsto C(X)$ is a contravariant functor (in fact an equivalence of categories by Theorem 2.1.5) from the category of compact Hausdorff spaces with continuous maps to the category of unital commutative C*-algebras with unital $*$-homomorphisms, so $K^0(X)$ is contravariantly functorial for continuous maps. As discussed above, this agrees with Atiyah and Hirzebruch's topological definition of K-theory.

K-theory can also be defined for non-unital C*-algebras. The definition uses the short exact sequence of C*-algebras

$$0 \to A \to \widetilde{A} \to \mathbb{C} \to 0$$

where A is any C*-algebra and $\widetilde{A}$ is its unitalization.

Definition 2.3.6. *Let A be any C*-algebra. Define $K_0(A)$ to be the kernel of the induced homomorphism $K_0(\widetilde{A}) \to K_0(\mathbb{C}) \cong \mathbb{Z}$.*

Note that if A is unital then $\widetilde{A} = A \oplus \mathbb{C}$, so this definition recovers our definition above. If $A = C_0(X)$ where X is a locally compact Hausdorff space then $K^0(X)$ defined this way corresponds to the relative K-theory of X in its one point compactification, and this precisely Atiyah and Hirzebruch's definition of $K^0(X)$.

As we mentioned at the beginning of this section, topological K-theory gives rise to a generalized cohomology theory. The higher K-theory groups $K^p(X)$ of a locally compact Hausdorff space X are obtained by forming K_0 of the suspension of X, and we will now imitate this construction to define higher K-theory groups for a C*-algebra.

Definition 2.3.7. *Let A be a C*-algebra. The* pth suspension *of A is the C*-algebra $S^p(A) = C_0(\mathbb{R}^p) \otimes A$. The K-theory group of A in degree p is defined to be $K_p(A) = K_0(S^p(A))$.*

Remark 2.3.8. *$S^p(A)$ can also be described as $C_0(\mathbb{R}^p, A)$. It is often convenient to describe $S(A)$ as the set of all continuous functions $f : [0,1] \to A$ such that*

$f(0) = f(1) = 0$.

It is natural to expect that a short exact sequence of C*-algebras

$$0 \to J \to A \to A/J \to 0$$

gives rise to a long exact sequence in K-theory:

$$\ldots \to K_{p+1}(A/J) \to K_p(J) \to K_p(A) \to K_p(A/J) \to \ldots \tag{2.3.1}$$

The key to this construction is defining the boundary map $\partial : K_{p+1}(A/J) \to K_p(J)$. This is obtained via a mapping cone construction.

Definition 2.3.9. *Let A be a C*-algebra and let $J \subseteq A$ be an ideal. The* mapping cone *of $\pi : A \to A/J$ is the C*-algebra $C(A, A/J)$ consisting of all pairs (a, f) where $a \in A$ and $f : [0,1] \to A/J$ is a continuous path with $f(0) = 0$ and $f(1) = \pi(a)$.*

A standard homotopy theoretic calculation allows us to calculate the K-theory of the mapping cone of π; we formulate the result but omit the proof:

Proposition 2.3.10. *The map excision map $J \to C(A, A/J)$ given by $j \mapsto (j, 0)$ induces an isomorphism on K-theory.*

Proof. See Chapter 4 of [9]. □

We have an embedding $S(A/J) \to C(A, A/J)$ which sends a map $f : [0,1] \to A/J$ satisfying $f(0) = f(1) = 0$ to the element $(0, f) \in C(A, A/J)$ and we have a surjection $C(A, A/J) \to A$ given by $(a, f) \mapsto a$ whose kernel is the image of $S(A/J)$. Thus we have a short exact sequence

$$0 \to S(A/J) \to C(A, A/J) \to A \to 0$$

The map $C(A, A/J) \to A$ induces the same map $K_0(J) \to K_0(A)$ as the inclusion $J \to A$, so we define the K-theoretic boundary map $\partial : K_1(A/J) \to K_0(J)$ to be the map induced by $S(A/J) \to C(A, A/J)$. The entire construction can be repeated replacing A with $S^p(A)$, so the long exact sequence ((2.3.1)) follows.

The long exact sequence (2.3.1) is simpler than it first appears due to the celebrated Bott periodicity theorem:

Theorem 2.3.11 (Bott Periodicity)**.** *Let A be a C^*-algebra. Then $K_0(A)$ is naturally isomorphic to $K_2(A)$.*

Proof. See Chapter 4 of [9] or section 9 of [4]. $\square$

Thus $K_p(A)$ is isomorphic to either $K_0(A)$ or $K_1(A)$, depending on the parity of p. Because of this it is common to write (2.3.1) as the *six term exact sequence*

$$\begin{tikzcd} K_0(J) \arrow[r] & K_0(A) \arrow[r] & K_0(A/J) \arrow[d] \\ K_1(A/J) \arrow[u] & K_1(A) \arrow[l] & K_1(J) \arrow[l] \end{tikzcd}$$

Given the privileged role of $K_1(A)$, it is convenient for many purposes to have available a more concrete description of its generators and relations. Assume for this discussion that A is unital. By definition $K_1(A) = K_0(S(A))$ and $K_0(S(A))$ is defined to be the kernel of the map $K_0(\widetilde{S(A)}) \to Z$ induced by the unitalization map $\widetilde{S(A)} \to \mathbb{C}$. Since A is unital $\widetilde{S(A)}$ can be described as the set of all maps $f : [0,1] \to A$ such that $f(0) = f(1) \in \mathbb{C}$, and thus a projection over $\widetilde{S(A)}$ is a projection-valued loop $\mathbf{p} : [0,1] \to M_n(A)$ such that $\mathbf{p}(0) = \mathbf{p}(1) \in M_n(\mathbb{C})$. $K_1(A)$ is generated by formal differences of such loops.

Just as $K_0(A)$ is generated by (formal differences of) projections over A, it turns out that $K_1(A)$ is generated by unitaries over A. To identify classes in $K_1(A)$ with unitaries, we build an auxiliary group $K_u(A)$ generated by unitaries and exhibit an isomorphism between $K_u(A)$ and $K_1(A)$.

Definition 2.3.12. *Let A be a unital C^*-algebra. Define $K_u(A)$ to be the abelian group generated by unitaries over A with the following relations:*

- *The additive identity element is represented by the identity matrix in $M_n(A)$ for each n.*
- $[\boldsymbol{u} \oplus \boldsymbol{v}] = [\boldsymbol{u}] + [\boldsymbol{v}]$ *for any pair of unitaries $\boldsymbol{u}, \boldsymbol{v}$ over A.*
- *If $\boldsymbol{u}_t$ is a continuous path of unitaries in $M_n(A)$ for $t \in [0,1]$ then $[\boldsymbol{u}_0] = [\boldsymbol{u}_1]$. In this case we say that $\boldsymbol{u}_0$ and $\boldsymbol{u}_1$ are* homotopic *as unitaries.*

Remark 2.3.13. *One can show that the additive inverse of a unitary $\boldsymbol{u}$ is given by $\boldsymbol{u}^*$, so every class in $K_u(A)$ is represented by a unitary over A (formal differences are not needed).*

$K_u(A)$ is isomorphic to $K_1(A)$, but we will not give a complete proof of this fact. We will, however, need an explicit identification $K_u(A) \to K_1(A)$. Given a unitary $\mathbf{u} \in M_n(A)$, let $\mathbf{U}(t)$ be any path of unitaries in $M_{2n}(A)$ with the property that

$$\mathbf{U}(0) = \begin{pmatrix} \mathbf{u} & 0 \\ 0 & \mathbf{u}^* \end{pmatrix}, \mathbf{U}(1) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Such a path always exists. Define $\mathbf{p(t)} = \mathbf{U}(t)(\mathbf{1}_n \oplus \mathbf{0}_n)\mathbf{U}(t)^* \in M_{2n}(A)$ and $\mathbf{q}(t) = \mathbf{1}_n \oplus \mathbf{0}_n$. Note that $\mathbf{p}$ and $\mathbf{q}$ are both loops of projections based at a projection in $M_{2n}(\mathbb{C})$, so their formal difference $[\mathbf{p}] - [\mathbf{q}]$ defines a class in $K_1(A)$.

Proposition 2.3.14. *The assignment $\boldsymbol{u} \mapsto [\boldsymbol{p}] - [\boldsymbol{q}]$ gives rise to a well-defined isomorphism of groups $K_u(A) \to K_1(A)$ independent of the path $\boldsymbol{U}(t)$ used to define it.*

We conclude our overview of K-theory for C*-algebras by formulating a few lemmas concerning inner automorphisms. Given a unital C*-algebra A and a unitary $\mathbf{u} \in A$, the inner automorphism associated to $\mathbf{u}$ is given by $\mathrm{Ad}_{\mathbf{u}}(a) = \mathbf{u}a\mathbf{u}^*$. By Remark 2.3.4, such automorphisms act trivially on K-theory. There is an analogous statement for non-unital algebras.

Lemma 2.3.15. *Let A be any C*-algebra and assume A is an ideal in a unital C*-algebra B. If $\mathbf{u} \in B$ is a unitary then $Ad_{\mathbf{u}}$ induces the identity map on $K_p(A)$.*

Proof. Form the C*-algebra $D = \{b_1 \oplus b_2 \in B \oplus B : \ b_1 - b_2 \in A\}$. The inclusion $i : A \to D$ given by $a \mapsto a \oplus 0$ and the map $D \to B$ given by $b_1 \oplus b_2 \mapsto b_2$ fit into a short exact sequence

$$0 \to A \to D \to B \to 0$$

This short exact sequence is split by the map $b \mapsto b \oplus b$, so the long exact sequence in K-theory splits into short exact sequences:

$$0 \to K_p(A) \to K_p(D) \to K_p(B) \to 0$$

Thus i_* is injective. Let $\mathbf{v}$ be the unitary $\mathbf{u} \oplus \mathbf{u}$ and observe that $i \circ \mathrm{Ad}_{\mathbf{u}} = \mathrm{Ad}_{\mathbf{v}} \circ i$. Since i_* is injective and $(\mathrm{Ad}_{\mathbf{v}})_* : K_p(D) \to K_p(D)$ is the identity (D is unital), $(\mathrm{Ad}_{\mathbf{u}})_*$ is also the identity. $\square$

A consequence of this is an important uniqueness statement about maps between K-theory groups that we will use several times.

Lemma 2.3.16. *Let A and B be C^*-subalgebras of $\mathbb{B}(H_A)$ and $\mathbb{B}(H_B)$ where H_A and H_B are Hilbert spaces. Suppose $V_1, V_2 : H_B \to H_A$ are isometries such that Ad_{V_1} and Ad_{V_2} map B into A and the operators $V_i V_j^*$ lie in A for each $i, j \in \{1, 2\}$. Then Ad_{V_1} and Ad_{V_2} induce the same homomorphism $K_p(B) \to K_p(A)$.*

Proof. It suffices to show that the maps $B \to M_2(A)$ given by

$$T \mapsto \begin{pmatrix} V_1 T V_1^* & 0 \\ 0 & 0 \end{pmatrix} \qquad T \mapsto \begin{pmatrix} 0 & 0 \\ 0 & V_2 T V_2^* \end{pmatrix}$$

induce the same map on K-theory. The second map can be obtained from the first by conjugating with the unitary

$$\begin{pmatrix} 1 - V_1 V_1^* & V_1 V_2^* \\ V_2 V_1^* & 1 - V_2 V_2^* \end{pmatrix}$$

in $M_2(\widetilde{A})$, so we are done since inner automorphisms act trivially on K-theory. $\square$

Chapter 3

K-Homology and Elliptic Operators

3.1 Introduction

As explained in Chapter 1, the goal of this thesis is to calculate a "generalized index" of an appropriate elliptic operator on a partitioned manifold. In this chapter we describe a framework for posing and analyzing index problems which is suitable for our needs. The first hints at this framework can be found in the original papers of Atiyah and Singer on the index of elliptic operators (beginning with [3]), so we shall take a moment to review the basic features of their approach.

Let M be a smooth compact manifold and let D be a first order differential operator acting on smooth sections of a vector bundle $S \to M$. PDE theorists have found it useful to describe the behavior of D in terms of the *principal symbol* σ_D of D, a smoothly varying family of polynomials on the cotangent spaces of M which captures the highest order behavior of D. In particular if σ_D is invertible away from 0 on each cotangent space (this is the condition that D is *elliptic*) then D itself is invertible modulo compact operators and therefore has a Fredholm index. Atiyah and Singer realized that the principal symbol of an elliptic operator determines a class in $K^0(T^*M)$, the topological K-theory group of the total space of the cotangent bundle of M, and that the Fredholm index of D depends only on this K-theory class. Indeed, there is a group homomorphism

$$\text{Index}\colon K^0(T^*M) \to \mathbb{Z}$$

which sends the symbol class of an elliptic operator to its index.

Atiyah and Singer formulated their original index theorem as a topological statement about this homomorphism, which they called the *analytic index*. They realized very quickly that if D is an elliptic operator on a smooth vector bundle $S \to M$ and $E \to M$ is another smooth vector bundle then $D_E := D_E \otimes 1$ is an elliptic operator on the vector bundle $S \otimes E \to M$ and that the pairing $(D, E) \mapsto \text{Index}(D_E)$ respects the relations in the second variable which define the K-theory group of M. Indeed, one of Atiyah and Singer's great accomplishments was to show that the analytic index map determines a bilinear pairing

$$K^0(T^*M) \otimes K^0(M) \to \mathbb{Z}$$

Atiyah recognized some formal similarities between this map and the Poincare duality pairing between ordinary homology and cohomology. K-theory defines a generalized cohomology theory in the sense of the Eilenberg-Steenrod axioms, and Atiyah proposed in [1] that there is a model for the corresponding homology theory such that the homology groups of a manifold are generated by (generalized) elliptic operators. Kasparov implemented Atiyah's proposal by specifying the relations among operators which correspond to the relations satisfied by their symbols in $K^0(T^*M)$. This yielded an analytic model for *K-homology*, the generalized homology theory which corresponds to K-theory.

Working in the setting of C*-algebra theory rather than algebraic topology, Kasparov actually achieved much more: he developed a bivariant functor $KK(\cdot, \cdot)$ on pairs of C*-algebras which generalizes K-theory and K-homology as well as various maps and pairings between K-groups. Included in Kasparov's theory is a sophisticated product structure which captures a variety of important phenomena in algebraic topology and analysis. We will not need the full power of Kasparov's bivariant theory in this thesis, but we will use it to build a model of K-homology and we will make use of the Kasparov product.

We will begin the chapter by reviewing some of the basic functional-analytic properties of differential operators which motivate our definition K-homology. We will then define K-homology explore some of its basic properties, culminating in a construction of long exact sequences. Our proof of the partitioned manifold index theorem later on hinges on calculating certain boundary maps, and to this end we will conclude the chapter with some computations involving the suspension map in K-homology and the Kasparov product. Our approach to the subject is mostly borrowed from [9]; many of our proofs are adapted from arguments in chapters 5, 8, 9, and 10.

3.2 Elliptic Operators on Manifolds

K-homology theory, as we will use it in this thesis, is based on an interaction between geometry and functional analysis which is mediated by index theory for differential operators on manifolds. Let M be a smooth manifold and let $S \to M$ be a smooth Hermitian vector bundle. By a (linear, first order) differential operator on S we mean a linear map $D\colon C^\infty(M;S) \to C^\infty(M;S)$ has the form

$$D = \sum_i A_i \partial_{x_i} + B \tag{3.2.1}$$

in a coordinate neighborhood which trivializes S, where the A_i's and B are smooth sections of $\mathrm{End}(S)$. Sometimes we will suppress S and refer to D as a differential operator on M. Note that differential operators are *local* in the sense that if s_1 and s_2 are smooth sections of S which agree on an open set $U \subseteq M$ then Ds_1 and Ds_2 agree on U.

Recall that a differential operator D is *symmetric* if $\langle Ds_1, s_2\rangle = \langle s_1, Ds_2\rangle$ for every $s_1, s_2 \in C^\infty(M;S)$ where $\langle \cdot,\cdot\rangle$ is the standard inner product on smooth sections of S defined relative to a smooth measure on M and the given Hermitian structure on S. A symmetric differential operator is *essentially self-adjoint* if it has a unique extension to a self-adjoint unbounded operator on $L^2(M;S)$. Essentially self-adjoint operators are convenient because they are accessible to the techniques of Hilbert space theory via the spectral theorem:

Theorem 3.2.1 (Spectral Theorem). *Let D be an essentially self-adjoint first order differential operator on $S \to M$. Then there is a sequence of Borel measures*

μ_n on $\mathbb{R}$ and a unitary isomorphism

$$L^2(M;S) \cong \bigoplus_n L^2(\mathbb{R}, \mu_n)$$

such that D decomposes as

$$D = \bigoplus_n M_n$$

where M_n is the multiplication operator on $L^2(\mathbb{R}, \mu_n)$ given by $M_n f(x) = xf(x)$

Proof. See [21], for instance. $\square$

The spectral theorem allows us to build new operators from D:

Definition 3.2.2. *Let D be an essentially self-adjoint first order differential operator on $S \to M$ and let H be the Hilbert space $L^2(M;S)$. Let $\mathcal{B}(\mathbb{R})$ denote the C^*-algebra of bounded Borel functions on $\mathbb{R}$ and let $H \cong \bigoplus_n L^2(\mathbb{R}, \mu_n)$ be the decomposition guaranteed by the spectral theorem. Given $\varphi \in \mathcal{B}(\mathbb{R})$, let $\varphi(D)$ denote the bounded operator on H which acts on the summand $L^2(\mathbb{R}, \mu_n)$ as*

$$\varphi(D)f(x) = \varphi(x)f(x)$$

The functional calculus map *for D is the isometric $*$-homomorphism*

$$\mathcal{B}(\mathbb{R}) \to \mathbb{B}(H)$$

given by $\varphi \mapsto \varphi(D)$.

Our approach to index theory for differential operators on manifolds uses the spectral theorem and the functional calculus extensively, and hence it is important that we give conditions which guarantee that a differential operator is essentially self-adjoint. The conditions we will describe depend only on the top order behavior of D (and the geometry of M), so we begin by defining an object which captures this behavior.

Definition 3.2.3. *Let U be an open subset of $\mathbb{R}^n$, let $S \to U$ be a trivial vector bundle, and let*

$$D = \sum_i A_i \partial_{x_i} + B$$

be a first order differential operator on $S \to U$. The symbol *of D is the bundle map $\sigma_D : T^*U \to End(S)$ given by*

$$\sigma_D(x, \xi) = \sum_i A_i \xi_i$$

Thus the symbol of D is the Fourier multiplier obtained by "freezing the coefficients" of the top order part of D. We have the following global formula for the symbol:

$$\sigma_D(x, df)s(x) = ([D, M_f]s)(x)$$

where s is a smooth section of S, f is a smooth function on M, and $M_f : C^\infty(M; S) \to C^\infty(M; S)$ is multiplication by f. This shows that σ_D is coordinate independent and allows us to define the symbol of a differential operator on a manifold.

Definition 3.2.4. *Let D be a first order differential operator on a vector bundle $S \to M$. The* symbol *of D is the bundle map $\sigma_D : T^*M \to End(S)$ given by*

$$\sigma_D(x, df)s(x) = ([D, M_f]s)(x)$$

Together with a Riemannian structure on M, the symbol of D allows one to quantify how D disturbs the support of a smooth section of S.

Definition 3.2.5. *Let M be a Riemannian manifold and let D be a differential operator acting on smooth sections of a Hermitian vector bundle $S \to M$. The* propagation speed *of D is defined to be*

$$c_D = \sup\{\|\sigma_D(x, \xi)\| : x \in M, \|\xi\| = 1\}$$

where the Hermitian structure on S is used to define $\|\sigma_D(x, \xi)\|$ and the Riemannian structure on M is used to define $\|\xi\|$.

This allows us to formulate our main result about essentially self-adjoint operators.

Proposition 3.2.6. *Let M be a complete Riemannian manifold and let D be a symmetric differential operator acting on smooth sections of a Hermitian vector* *bundle $S \to M$. If the propagation speed of D is finite then D is essentially self-adjoint.*

Proof. See Chapter 10 of [9]. □

The following example illustrates some of the subtleties of this proposition.

Example 3.2.7. *Consider the differential operator $D = i\frac{d}{dx}$ acting on smooth sections of the trivial Hermitian line bundle over $\mathbb{R}$ equipped with the standard Riemannian metric. We have $\sigma_D(x,\xi) = i\xi$ for every $(x,\xi) \in T^*\mathbb{R}$, so the propagation speed of D is precisely 1. It follows that D is essentially self-adjoint since $\mathbb{R}$ is complete.*

Now consider the same operator $D = i\frac{d}{dx}$ acting instead on smooth sections of the trivial Hermitian line bundle over the interval $(0,1)$ with the standard Riemannian metric. The propagation speed is the same, but D is not essentially self-adjoint: it has a variety of self-adjoint extensions corresponding to different choices of boundary conditions.

Of course these subtleties disappear in the compact case:

Corollary 3.2.8. *Every symmetric differential operator on a compact manifold is essentially self-adjoint.*

Proof. Let D be a symmetric differential operator acting on a Hermitian vector bundle S over a compact manifold M. For any Riemannian metric on M the unit sphere bundle of T^*M is compact so the propagation speed of D is finite. Moreover every compact Riemannian manifold is automatically complete, so D is essentially self-adjoint by Proposition 3.2.6 □

To do index theory we need to impose an additional constraint on the symbol of a differential operator.

Definition 3.2.9. *A differential operator D on $S \to M$ is* elliptic *over an open set $U \subseteq M$ if its symbol $\sigma_D(x,\xi)$ is an invertible endomorphism of S_x for every nonzero $\xi \in T^*_xM$ and every $x \in U$.*

The operator $i\frac{d}{dx}$ appearing in Example 3.2.7 is a canonical example of a first order elliptic operator. The ellipticity condition ensures that operators obtained from D via the functional calculus are compatible with the representation

$$\rho\colon C_0(M) \to \mathbb{B}(L^2(M;S))$$

of the C*-algebra $C_0(M)$ by multiplication operators. Specifically we have the following:

Proposition 3.2.10. *Let M be a smooth manifold and let D be an essentially self-adjoint first order differential operator on $S \to M$. If D is elliptic over an open set $U \subseteq M$ then for every $\varphi \in C_0(\mathbb{R})$ and every $f \in C_0(U)$ the operator $\rho(f)\varphi(D) \in \mathbb{B}(L^2(M;S))$ is compact.*

Proof. See Chapter 10 of [9] $\square$

A bounded operator T on $L^2(M;S)$ with the property that $\rho(f)T$ is compact for every $f \in C_0(U)$ is said to be *locally compact for* U (or simply *locally compact* if $U = M$). Thus the proposition asserts that if D is essentially self-adjoint and elliptic over U then $\varphi(D)$ is locally compact for U for any $\varphi \in C_0(\mathbb{R})$. This result is important because it encapsulates the functional-analytic properties of elliptic operators which make them suitable for index theory, as the following corollary illustrates:

Corollary 3.2.11. *Let M be a compact smooth manifold and let D be an essentially self-adjoint first order differential operator on $S \to M$ which is elliptic over all of M. Then D is Fredholm.*

Proof. Let $K_D = \{\varphi(D)\colon\ \varphi \in C_0(\mathbb{R})\}$, a commutative C*-algebra bounded operators on $L^2(M;S)$. Every operator in K_D is compact by Proposition 3.2.10, so the operators in K_D can be simultaneously diagonalized by the spectral theorem for compact operators. D is diagonalized by any orthonormal basis which simultaneously diagonalizes K_D; since the operator e^{-D^2} (for instance) is compact it follows that $e^{-\lambda_n^2} \to 0$ where $\{\lambda_n\}$ is the spectrum of D, and hence $|\lambda_n| \to \infty$. In particular the 0-eigenspace of D is finite dimensional; since D is essentially self-adjoint this implies that the kernel and cokernel of D are finite dimensional, as desired. $\square$

Note that

$$\text{Index}(D) = \dim \ker(D) - \dim \ker(D^*)$$

so the index of an essentially self-adjoint Fredholm operator D is automatically 0. At first glance this makes index theory for elliptic operators on manifolds appear vacuous. However, many of the operators which arise naturally in geometry come equipped with additional algebraic structure that allows for a more interesting notion of index. Recall that a vector bundle $S \to M$ is *graded* if it has a direct sum decomposition $S = S^+ \oplus S^-$ and an operator D acting on smooth sections of S is *odd* if it decomposes as:

$$D = \begin{pmatrix} 0 & D^- \\ D^+ & 0 \end{pmatrix}$$

The overall operator D must have index 0, but the proof of Corollary 3.2.11 can easily be adapted to show that D^+ and D^- have finite dimensional kernels, and this can be used to produce a nontrivial invariant:

Definition 3.2.12. *Let D be a graded operator acting on smooth sections of a graded vector bundle $S \to M$. The* graded index *of D is defined to be* $\dim \ker D^+ - \dim \ker D^-$.

Elliptic operators on non-compact manifolds are not in general Fredholm, and for this reason it is useful to view the conclusion of Proposition 3.2.10 as a substitute for the Fredholm property. Later on we will use Proposition 3.2.10 in tandem with some geometry to define a more refined notion of index which applies to the non-compact case. This suggests that we isolate the following class of operators:

Definition 3.2.13. *Let M be a Riemannian manifold and let $S \to M$ be a Hermitian vector bundle. A* Dirac-type *operator on $S \to M$ is a symmetric first order elliptic differential operator D which has finite propagation speed.*

Remark 3.2.14. *The notion of a Dirac-type operator has a number of different definitions in the literature, many of which are more specific than the definition given here. The literature also has a variety of different notions of a* Dirac operator*, which is generally a Dirac-type operator equipped with a specific algebraic structure.*

Every Dirac-type operator on a complete Riemannian manifold is essentially self-adjoint and satisfies the hypotheses of Proposition 3.2.10; thus Dirac-type operators are particularly suitable for index theory. However, it can be difficult to work directly with a Dirac-type operator D since differential operators are by nature unbounded. It is more convenient to use the functional calculus to replace D with a bounded proxy, and for that purpose we introduce a special kind of function on $\mathbb{R}$.

Definition 3.2.15. *A continuous function $\chi\colon \mathbb{R} \to [-1,1]$ is a* normalizing function *if:*

- *χ is an odd function*
- *$\chi(\lambda) > 0$ for $\lambda > 0$*
- *$\chi(\lambda) \to \pm 1$ as $\lambda \to \pm\infty$.*

Lemma 3.2.16. *Let M be a compact smooth manifold and let D be a graded essentially self-adjoint first order differential operator on $S \to M$. Then $\chi(D)$ is a graded-Fredholm operator for any normalizing function χ with the same graded index as D.*

Proof. We begin by proving that $\chi(D)$ is odd. Let γ denote the grading operator on $L^2(M;S)$, so that an operator on $L^2(M;S)$ is odd if and only if it anti-commutes with γ. Since D is odd we have $\gamma(i+D)^{-1} = (i-D)^{-1}\gamma$ and $\gamma(i-D)^{-1} = (i+D)^{-1}\gamma$, and thus the operator $(i+D)^{-1} - (i-D)^{-1}$ anti-commutes with γ. By the Stone-Weierstrass theorem any odd function in $C_0(\mathbb{R})$ is the uniform limit of functions in the $*$-subalgebra of $C_0(\mathbb{R})$ generated by $(i+x)^{-1} - (i-x)^{-1}$, so $\varphi(D)$ is odd for any odd function $\varphi \in C_0(\mathbb{R})$. But χ is the pointwise limit of functions in $C_0(\mathbb{R})$, so $\chi(D)$ is the strong limit of odd operators and hence is itself odd.

Now we prove that $\chi(D)$ has the same graded kernel and cokernel as D. As in the proof of Corollary 3.2.11, $L^2(M;S)$ admits an orthonormal basis of simultaneous eigenvectors for D and $\chi(D)$; if $\{\lambda_n\}$ is the set of eigenvalues for D then $\{\chi(\lambda_n)\}$ is the set of eigenvalues for $\chi(D)$. Since χ vanishes only at 0 it follows that $\chi(D)$ has the same kernel as D and hence the same graded kernel and cokernel. $\square$

Remark 3.2.17. *If χ_1 and χ_2 are two different normalizing functions then $\chi_1 - \chi_2 \in C_0(\mathbb{R})$ and hence $\chi_1(D) - \chi_2(D) = (\chi_1(D) - \chi_2(D))(\rho(1))$ is compact by Proposition 3.2.10, so their index is the same since the index is stable under compact perturbations. More generally if D is a Dirac-type operator on a complete Riemannian manifold then $\chi_1(D)$ and $\chi_2(D)$ differ by a locally compact operator.*

This allows us to replace D with $\chi(D)$ for the purposes of index theory. For this reason it makes sense to use the functional-analytic properties of operators of the form $\chi(D)$ as a guide when we define K-homology, which is supposed to be an algebraic abstraction of index theory. It turns out that the right abstract property of $\chi(D)$ is that it is *pseudolocal*, meaning the commutator $[\chi(D), \rho(f)]$ is compact for every $f \in C_0(M)$. Our proof of this fact uses an alternative description of the functional calculus based on the wave equation.

The wave equation for D is the partial differential equation $\frac{\partial s}{\partial t} = iDs$. If D is essentially self-adjoint then standard PDE theory shows that this equation has a unique solution for any L^2 initial data and that the solution operator e^{itD} is a unitary operator in $\mathbb{B}(L^2(M; S))$. If D is a Dirac-type operator on a complete Riemannian manifold then finite propagation speed arguments yield the following result about the wave operator e^{itD}:

Proposition 3.2.18. *Let D be a Dirac-type operator on a complete Riemannian manifold M and let f, g be bounded continuous functions on M such that the support of g is compact and disjoint from the support of f. Then there exists $\varepsilon > 0$ such that $\rho(f)e^{itD}\rho(g) = 0$ whenever $|t| < \varepsilon$.*

Proof. See Chapter 10 of [9]. □

The importance of the wave operator stems from the fact that it determines the functional calculus for D via some modest distribution theory.

Proposition 3.2.19. *Let D be a Dirac-type operator on a complete Riemannian manifold M. If $\varphi\colon \mathbb{R} \to \mathbb{R}$ is a bouned Borel function with compactly supported Fourier transform then:*

$$\varphi(D) = \frac{1}{2\pi} \int_{\mathbb{R}} \hat{\varphi}(t) e^{itD} dt$$

in the sense of distributions.

Proof. See Chapter 10 of [9]. □

Because our main interest in the functional calculus involves normalizing functions, we must show that there are normalizing functions compatible with Proposition 3.2.19.

Proposition 3.2.20. *For every $a > 0$ there is a normalizing function χ with the property that $\hat{\chi}$ is supported in $(-a, a)$ and $s\hat{\chi}(s)$ is smooth where $\hat{\chi}$ is the (distributional) Fourier transform of χ.*

Proof. Let g be any smoooth, even, compactly supported, real valued function on $\mathbb{R}$ and let $f = g * g$ be the convolution of g with itself. Choose g so that $f(0) \neq 0$ and scale g so that $f(0) = \frac{1}{\pi}$. Define

$$\chi(t) = \int_{\mathbb{R}} \frac{e^{itx} - 1}{ix} f(x) dx$$

We claim that χ has the required properties. Straightforward substitutions show that f and consequently χ are odd functions. $\chi(0)$ is 0 since f is odd and compactly supported, and $\chi'(t)$ is positive since it is the inverse Fourier transform of f which in turn is the square of the (real valued) inverse Fourier transform of g. Thus $\chi(t) > 0$ for $t > 0$. A similar argument with inverse Fourier transforms together with the scaling $f(0) = \frac{1}{\pi}$ implies that $\chi(t) \to 1$ as $t \to \infty$, so χ is a normalizing function.

Finally, the fact that $\chi'(t)$ is the inverse Fourier transform of f and the fact that f is compactly supported together imply that $s\hat{\chi}(s)$ is smooth and compactly supported. At the possible cost of rescaling $g(x)$ in x, we can arrange for $\hat{\chi}(s)$ to have support in $(-a, a)$, as desired. □

We are now ready to prove the main result of this section:

Proposition 3.2.21. *If D is a Dirac-type operator on a complete Riemannian manifold M then $[\chi(D), \rho(f)]$ is compact for any normalizing function χ and any $f \in C_0(M)$.*

Proof. We work over the one-point compactification $\widetilde{M}$ of M, noting that ρ extends to a representation of $C_0(\widetilde{M})$ on $L^2(M; S)$. By Kasparov's lemma (reference) it suffices to show that $\rho(f)\chi(D)\rho(g)$ is compact for any pair of continuous functions on $\widetilde{M}$ with disjoint supports. At least one of f or g must have compact support inside M, and at the possible cost of replacing $\rho(f)\chi(D)\rho(g)$ with its adjoint we may assume that g has this property. By Proposition 3.2.18, there exists $\varepsilon > 0$ such that $\rho(f)e^{itD}\rho(g) = 0$ for every $t \in (-\varepsilon, \varepsilon)$.

By Proposition 3.2.20 there is a normalizing function χ_ε with the property that its distributional Fourier transform $\widehat{\chi_\varepsilon}$ is supported in $(-\varepsilon, \varepsilon)$. By Proposition 3.2.19 we have:

$$\rho(f)\chi_\varepsilon(D)\rho(g) = \frac{1}{2\pi}\int_{-\varepsilon}^{\varepsilon} \widehat{\chi_\varepsilon}(t)\rho(f)e^{itD}\rho(g)dt = 0$$

Since $(\chi(D) - \chi_\varepsilon(D))\rho(g)$ is compact, we conclude that $\rho(f)\chi(D)\rho(g)$ is compact as desired. $\square$

3.3 Dual Algebras and K-Homology

In this section we define and develop some of the basic properties of K-homology, following chapter 5 of [9]. Our definition, due to Paschke in [17], is motivated by Proposition 3.2.10 and Proposition 3.2.21 from the previous section. These two results help to characterize the index theoretic properties of an elliptic operator on a smooth vector bundle S over a smooth manifold M purely in terms of the Hilbert space $L^2(M; S)$ and the representation of $C_0(M)$ as bounded multiplication operators on this Hilbert space.

From this perspective it is not important that M is a manifold, or even an ordinary topological space: the entire theory can be expressed in terms of an abstract C*-algebra A equipped with a representation on some Hilbert space H. However we will almost exclusively be interested in the case where A is a commutative C*-algebra, so that $A = C_0(X)$ for some locally compact Hausdorff space X. Similarly, when we introduce an ideal J in A and discuss relative K-homology the example of interest is $J = C_0(U)$ where U is an open subset of X. In spite of these remarks it is worthwhile to express the main results in the language of abstract C*-algebras because many of the proofs involve sophisticated results in C*-algebra theory that lack an obvious geometric interpretation.

3.3.1 Dual Algebras

Following Paschke, we will define the K-homology groups of a C*-algebra A in terms of the K-theory groups of an auxiliary C*-algebra associated to A called the *dual algebra*. An element of the dual algebra is meant to be a "generalized elliptic operator" in a sense suggested by Proposition 3.2.21. For what follows we will use the symbol "$\sim$" to denote equality up to compact operators.

Definition 3.3.1. *Let A be a C*-algebra and let $\rho\colon A \to \mathbb{B}(H)$ be a representation of A on a Hilbert space H.*

- *An operator $T \in \mathbb{B}(H)$ is* pseudolocal *for A if $[T, \rho(a)] \sim 0$ for every $a \in A$.*
- *The* dual algebra *of A is the C*-subalgebra $\mathfrak{D}^*_\rho(A) \subseteq \mathbb{B}(H)$ of all pseudolocal operators for A.*

In elementary index theory one considers bounded operators which are invertible modulo the ideal of compact operators. In K-homology we replace the compact operators with the ideal in $\mathfrak{D}^*(A)$ consisting of locally operators, following Proposition 3.2.10.

Definition 3.3.2. *Let A be a separable C*-algebra and let $\rho\colon A \to \mathbb{B}(H)$ be an ample representation of A on a separable Hilbert space H.*

- *An operator $T \in \mathfrak{D}^*_\rho(A)$ is* locally compact *for A if $\rho(a)T \sim T\rho(a) \sim 0$ for every $a \in A$.*
- *The* locally compact algebra *of A is the ideal $\mathfrak{C}^*(A)$ in $\mathfrak{D}^*(A)$ of all locally compact operators for A.*

Remark 3.3.3. *In this thesis the notation C (for spaces of continuous functions), $\mathfrak{C}^*$ (for the algebra of locally compact controlled operators), and C^* (for the coarse C*-algebra, to be defined later) will all be used regularly. The reader is asked to take care to distinguish between them.*

Before we define K-homology, we introduce a specific class of well-behaved representations which share some key features with the representation $C_0(\mathbb{R}) \to \mathbb{B}(L^2(\mathbb{R}))$ by multiplication operators.

Definition 3.3.4. *Let A be a C^*-algebra and let H be a Hilbert space. A representation $\rho\colon A \to \mathbb{B}(H)$ is said to be* ample *if it extends to a representation $\widetilde{\rho}\colon \widetilde{A} \to \mathbb{B}(H)$ of the unitalization of A which has the following properties:*

- *$\widetilde{\rho}$ is nondegenerate, meaning $\widetilde{\rho}(\widetilde{A})H$ is dense in H*
- *$\widetilde{\rho}(a)$ is compact for $a \in \widetilde{A}$ if and only if $a = 0$*

Remark 3.3.5. *Some authors define a representation ρ to be ample if the two conditions above are satisfied by ρ itself rather than by the unital extension $\widetilde{\rho}$.*

Remark 3.3.6. *In fact $\widetilde{\rho}$ is nondegenerate in the sense described above if and only if it is unital, but we will not need this fact and the condition given in the definition is sometimes easier to check.*

We are now ready to define K-homology.

Definition 3.3.7. *Let A be a separable C^*-algebra and let $\rho\colon A \to \mathbb{B}(H)$ be an ample representation of A on a separable Hilbert space H. The pth* (unreduced) K-homology group *of A, $p \in \mathbb{Z}$, is defined by $K^p(A) = K_{1-p}(\mathfrak{D}^*_\rho(A)/\mathfrak{C}^*_\rho(A))$.*

Remark 3.3.8. *Though the C^*-algebra $\mathfrak{D}^*_\rho(A)/\mathfrak{C}^*_\rho(A)$ depends on the representation ρ, we will show that its K-theory groups do not so long as ρ is ample. The proof uses the separability of A and H as well as the fact that ρ is ample.*

Because of this fact we will often suppress the representation ρ from the notation and simply write $\mathfrak{D}^*(A)$ and $\mathfrak{C}^*(A)$, with the understanding that the implied representation is ample.

Example 3.3.9. *Let $A = \mathbb{C}$ and let H be any separable infinite dimensional Hilbert space. The representation $\rho : \mathbb{C} \to \mathbb{B}(H)$ given by $\rho(\lambda) = \lambda I$ where I is the identity operator on H is ample, and we have*

$$\mathfrak{D}^*_\rho(\mathbb{C}) = \mathbb{B}(H)$$
$$\mathfrak{C}^*_\rho(\mathbb{C}) = \mathbb{K}(H)$$

Thus $K^p(\mathbb{C})$ is isomorphic to $\mathbb{Z}$ for p even and 0 for p odd.

For geometric applications, we define the analytic K-homology groups $K_*(X)$ of a second countable locally compact Hausdorff space to be the K-homology groups of the C*-algebra $C_0(X)$, noting that a locally compact Hausdorff space is second countable if and only if $C_0(X)$ is separable. Suppose that X has the structure of a smooth manifold and let D be an essentially self-adjoint first order elliptic operator acting on smooth sections of a vector bundle S over X. Represent $C_0(X)$ as multiplication operators on the Hilbert space $L^2(X;S)$; this is an ample representation, so we can use it to define the dual algebra $\mathfrak{D}^*(X)$. If χ is any normalizing function then Proposition 3.2.21 implies that $\chi(D) \in \mathfrak{D}^*(X)$. We would like to show that $\chi(D)$ determines a K-homology class.

- Suppose first that D is an ungraded operator. Set $P = \frac{1}{2}(\chi(D)+1)$ so that $P^2 = \frac{1}{4}(\chi(D)^2 + 2\chi(D) + 1)$. Note that the function $\chi^2 - 1$ belongs to $C_0(\mathbb{R})$, so according to Proposition 3.2.10 we have that $\rho(f)(\chi(D)^2 - 1)$ is compact for any $f \in C_0(X)$. It follows that P^2 differs from P by a locally compact operator and hence P maps to a projection in the quotient algebra $\mathfrak{D}^*_\rho(X)/\mathfrak{C}^*_\rho(X)$. This in turn determines a class in the K-homology group $K_1(X)$.

- Suppose instead that D is a graded operator so that D has the form

$$D = \begin{pmatrix} 0 & D^- \\ D^+ & 0 \end{pmatrix}$$

  Following the argument in Lemma 3.2.16 we have that $\chi(D)$ is an odd self-adjoint operator for any normalizing function χ. Write

$$\chi(D) = \begin{pmatrix} 0 & U^* \\ U & 0 \end{pmatrix}$$

  As above we have that $\rho(f)(\chi(D)^2 - 1)$ is compact for any $f \in C_0(X)$, and hence $\rho(f)(UU^* - 1)$ and $\rho(f)(U^*U - 1)$ are compact operators. Consider the sequence of Hilbert spaces H_n, $n \in \mathbb{Z}$, given by:

$$H_n = \begin{cases} L^2(X; S^-) & -\infty < n \leq 0 \\ L^2(X; S^+) & 1 \leq n < \infty \end{cases}$$

Define H to be the Hilbert space $\bigoplus_n H_n$, and note that H carries an ample representation ρ' of $C_0(X)$. Let U' be the bounded operator on H which acts as $U\colon L^2(X; S^+) \to L^2(X; S^-)$ on H_1 and as the identity on H_n for $n \neq 1$. We have shown that U' is a unitary in $\mathfrak{D}^*_{\rho'}(X)/\mathfrak{C}^*_{\rho'}(X)$ and hence it defines a class in the K-homology group $K_0(X)$.

In either case we will use the notation $[D]$ to signify the K-homology class determined by D.

3.3.2 Functoriality

Implicit in our claim at the beginning of this chapter that K-homology satisfies the axioms of a generalized homology theory is the assertion that K-homology is covariantly functorial, at least for continuous maps between compact spaces. Since a continuous map $X \to Y$ between compact spaces is equivalent to a $*$-homomorphism $C(Y) \to C(X)$ it will suffice to establish the contravariant functoriality of K-homology for $*$-homomorphisms between unital C*-algebras. However, out interest in K-homology is not limited to compact spaces, so we will need to address functoriality for $*$-homomorphisms between non-unital C*-algebras as well. Certain technical complications arise in the most general cases which can be avoided by restricting our attention to separable C*-algebras (correspondingly, second countable spaces), so we shall do so for the remainder of this section.

We will implement the functiality of K-homology by associating to an appropriate $*$-homomorphism $A \to B$ between separable C*-algebras a $*$-homomorphism $\mathfrak{D}^*_{\rho_B}(B)/\mathfrak{C}^*_{\rho_B}(B) \to \mathfrak{D}^*_{\rho_A}(A)/\mathfrak{C}^*_{\rho_A}(A)$. This construction will depend on a number of choices, including the (ample) representations ρ_A and ρ_B, but we shall see that the ambiguity disappears at the level of K-theory. The key to the construction is the following definition:

Definition 3.3.10. *Let A and B be separable C^*-algebras equipped with ample representations $\rho_A\colon A \to \mathbb{B}(H_A)$ and $\rho_B\colon B \to \mathbb{B}(H_B)$ on separable Hilbert spaces, and let $\phi\colon A \to B$ be a $*$-homomorphism. We say that an isometry $V\colon H_B \to H_A$* covers *ϕ if $V^*\rho_A(a)V \sim \rho_B(\phi(a))$ for every $a \in A$.*

Proposition 3.3.11. *Let A and B be separable C^*-algebras, let $\rho_A\colon A \to \mathbb{B}(H_A)$ be an ample representation on a separable Hilbert space, and let $\rho_B\colon B \to \mathbb{B}(H_B)$ be a representation. Then any $*$-homomorphism $\phi\colon A \to B$ is covered by an isometry $V\colon H_B \to H_A$.*

Proof. Let ρ_A and ρ_B denote the ample representations of A and B, respectively. Let $\widetilde{\rho}_A$ and $\widetilde{\rho}_B$ denote the respective extensions of ρ_A and ρ_B to the unitalizations $\widetilde{A}$ and $\widetilde{B}$ guaranteed by the definition of ample. Finally let $\widetilde{\phi}\colon \widetilde{A} \to \widetilde{B}$ denote the (unique) extension of ϕ to a unital $*$-homomorphism between $\widetilde{A}$ and $\widetilde{B}$. Our strategy is to apply Voiculescu's theorem (Theorem 2.2.7) to the nondegenerate representation $\widetilde{\rho}_A$ and the completely positive map $\widetilde{\rho}_B \circ \widetilde{\phi}$ (recall that any $*$-homomorphism is completely positive). To do so we must check that $\widetilde{\rho}_B\circ\widetilde{\phi}(a) = 0$ if $\widetilde{\rho}_A(a)$ is compact; $\widetilde{\rho}_A(a)$ is compact if and only if $a = 0$ since ρ_A is ample, so this follows immediately. Voiculescu's theorem thus guarantees the existence of an isometry V such that

$$V^*\widetilde{\rho}_A(a)V \sim \widetilde{\rho}_B(\widetilde{\phi}(a))$$

for every $a \in \widetilde{A}$; since $\widetilde{\rho}_A$, $\widetilde{\rho}_B$, and $\widetilde{\phi}$ restrict to ρ_A, ρ_B, and ϕ, respectively, it follows that V covers ϕ. □

Now, given an isometry $V\colon H_B \to H_A$ which covers a $*$-homomorphism $\phi\colon A \to B$, we can define a $*$-homomorphism $\mathrm{Ad}_V\colon \mathbb{B}(H_B) \to \mathbb{B}(H_A)$ by $\mathrm{Ad}_V(T) = VTV^*$. We shall prove that Ad_V restricts to a map $\mathfrak{D}^*_{\rho_B}(B) \to \mathfrak{D}^*_{\rho_A}(A)$ which sends the ideal $\mathfrak{C}^*_{\rho_B}(B)$ to $\mathfrak{C}^*_{\rho_A}(A)$. First, we need a technical lemma:

Lemma 3.3.12. *Let V be an isometry which covers a $*$-homomorphism $\phi\colon A \to B$. Then the projection $P = VV^*$ is an element of $\mathfrak{D}^*_{\rho_A}(A)$.*

Proof. Decompose H_A as the orthogonal direct sum $PH_A \oplus (1-P)H_A$ and express ρ_A relative to this decomposition as:

$$\rho_A = \begin{pmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{pmatrix}$$

Observe that ρ_{11} is a $*$-homomorphism modulo compact operators since V covers ϕ. So by Lemma 2.1.9 applied to the identity operator we have that $\rho_{12}(a)\colon PH_A \to (1-P)H_A$ and $\rho_{21}(a)\colon (1-P)H_A \to PH_A$ are compact operators for every $a \in A$. But $\rho_{12}(a)P = \rho_A(a)P - P\rho_A(a)P$ and $\rho_{21}(a)(1-P) = P\rho_A(a) - P\rho_A(a)P$, so $[P, \rho_A(a)]$ is compact. $\square$

Proposition 3.3.13. *Ad_V maps $\mathfrak{D}^*_{\rho_B}(B)$ into $\mathfrak{D}^*_{\rho_A}(A)$ and $\mathfrak{C}^*_{\rho_A}(A)$ into $\mathfrak{C}^*_{\rho_B}(B)$.*

Proof. Take $T \in \mathfrak{D}^*_{\rho_B}(B)$. Since T is pseudolocal for ρ_B, we have

$$T\rho_B(\phi(a)) - \rho_B(\phi(a))T \sim 0$$

By the definition of covering isometry, this gives

$$TV^*\rho_A(a)V - V^*\rho_A(a)VT \sim 0$$

and hence

$$VTV^*\rho_A(a)VV^* - VV^*\rho_A(a)VTV^* \sim 0$$

By Lemma 3.3.12 this becomes

$$VTV^*VV^*\rho_A(a) - \rho_A(a)VV^*VTV^* \sim 0$$

Finally, since V is an isometry,

$$VTV^*\rho_A(a) - \rho_A(a)VTV^* \sim 0$$

This means that $\mathrm{Ad}_V(T)$ is pseudolocal for ρ_A, as desired.

The proof for locally compact algebras is similar. Take $S \in \mathfrak{C}^*_{\rho_B}(B)$, so that $S\rho_B(\phi(a)) \sim \rho_B(\phi(a))S \sim 0$ for every $a \in A$. By the definition of covering isometry,

$$SV^*\rho_A(a)V \sim 0$$

and hence

$$VSV^*\rho_A(a)VV^* \sim 0$$

By Lemma 3.3.12 this becomes

$$VSV^*VV^*\rho_A(a) \sim 0$$

Since V is an isometry it follows that $VSV^*\rho_A(a) \sim 0$. By a nearly identical argument (or by appealing to the fact that VSV^* is pseudolocal) we have that $\rho_A(a)VSV^* \sim 0$ as well, so $\mathrm{Ad}_V(S)$ is locally compact. $\square$

If $\phi\colon A \to B$ and $\psi\colon B \to C$ are $*$-homomorphisms covered by isometries $V\colon H_B \to H_A$ and $W\colon H_C \to H_B$, respectively, then it is clear that $VW\colon H_C \to H_A$ covers $\psi\circ\phi$ and that $\mathrm{Ad}_{VW} = \mathrm{Ad}_V \circ \mathrm{Ad}_W$. However, the dual algebra construction is not quite functorial for $*$-homomorphisms between C*-algebras equipped with ample representations since a choice of covering isometry must be made. The next result shows that different choices yield the same result at the level of K-homology.

Lemma 3.3.14. *Suppose $\phi\colon A \to B$ is a $*$-homomorphism covered by two isometries $V_1, V_2\colon H_B \to H_A$. Then Ad_{V_1} and Ad_{V_2} induce the same map*

$$K_p(\mathfrak{D}^*_{\rho_B}(B)/\mathfrak{C}^*_{\rho_B}(B)) \to K_p(\mathfrak{D}^*_{\rho_A}(A)/\mathfrak{C}^*_{\rho_A}(A))$$

Proof. According to Lemma 2.3.16 and Proposition 3.3.13, it suffices to show that $V_iV_j^* \in \mathfrak{D}^*_{\rho_A}(A)$ for every pair $i, j \in \{1, 2\}$. If $i = j$ then this follows from Lemma 3.3.12, so we need only show that $V_1V_2^*$ and $V_2V_1^*$ are in $\mathfrak{D}^*_{\rho_A}(A)$. Since V_1 and V_2 cover the same map ϕ we have

$$V_1V_1^*\rho_A(a)V_1V_2^* \sim V_1V_2^*\rho_A(a)V_2V_2^*$$

It follows from Lemma 3.3.12 that

$$\rho_A(a)V_1V_1^*V_1V_2^* \sim V_1V_2^*V_2V_2^*\rho_A(a)$$

Since V_1 and V_2 are isometries we conclude that

$$\rho_A(a)V_1V_2^* \sim V_1V_2^*\rho_A(a)$$

which means that $V_1V_2^* \in \mathfrak{D}^*_{\rho_A}(A)$. By the same argument with the roles of V_1 and V_2 reversed, $V_2V_1^* \in \mathfrak{D}^*_{\rho_{A(a)}}(A)$ as well. $\square$

It follows easily from this that $K_p(\mathfrak{D}^*_{\rho_1}(A)) \cong K_p(\mathfrak{D}^*_{\rho_2}(A))$ if ρ_1 and ρ_2 are two ample representations of the same C*-algebra A.

Corollary 3.3.15. *Let A be a separable C*-algebra and let $\rho_1\colon A \to \mathbb{B}(H_1)$ and $\rho_2\colon A \to \mathbb{B}(H_2)$ be ample representations of A on separable Hilbert spaces. Then the identity map $1_A\colon A \to A$ induces a canonical isomorphism $K_p(\mathfrak{D}^*_{\rho_1}(A)/\mathfrak{C}^*_{\rho_1}(A)) \cong K_p(\mathfrak{D}^*_{\rho_2}(A)/\mathfrak{C}^*_{\rho_2}(A))$.*

Proof. The identity map is a $*$ - homomorphism which extends to $\widetilde{A}$, so by Proposition 3.3.11 there are isometries $V\colon H_2 \to H_1$ and $W\colon H_1 \to H_2$ which cover 1_A. Clearly $VW\colon H_1 \to H_1$ is also an isometry which covers 1_A. The identity map $H_1 \to H_1$ is yet another isometry which covers 1_A, and it induces the identity map $K_p(\mathfrak{D}^*_{\rho_1}(A)/\mathfrak{C}^*_{\rho_1}(A)) \to K_p(\mathfrak{D}^*_{\rho_1}(A)/\mathfrak{C}^*_{\rho_1}(A))$. But by the previous lemma it induces the same map on K-theory as Ad_{VW}, so $(\mathrm{Ad}_{VW})_* = (\mathrm{Ad}_V)_* \circ (\mathrm{Ad}_W)_*$ is the identity. $(\mathrm{Ad}_W)_* \circ (\mathrm{Ad}_V)_*$ is the identity by the same argument, so $(\mathrm{Ad}_V)_*\colon K_p(\mathfrak{D}^*_{\rho_2}(A)/\mathfrak{C}^*_{\rho_2}(A)) \to K_p(\mathfrak{D}^*_{\rho_1}(A)/\mathfrak{C}^*_{\rho_1}(A))$ is an isompor-phism. $\square$

Remark 3.3.16. *The isomorphism on K-theory induced by the identity map is canonical in the sense that it depends only on A, ρ_1, and ρ_2 and not the choices of covering isometries.*

If $\phi\colon A \to B$ is a $*$-homomorphism between separable C*-algebras we will use the notation $\phi^*\colon K^p(B) \to K^p(A)$ for the map $(\mathrm{Ad}_V)_*\colon K_p(\mathfrak{D}^*(B)/\mathfrak{C}^*(B)) \to K_p(\mathfrak{D}^*(A)/\mathfrak{C}^*(A))$ where V is any isometry which covers ϕ. The results of this section can be summarized as follows:

Proposition 3.3.17. *The assignment $\{\phi\colon A \to B\} \mapsto \{\phi^*\colon K^p(A) \to K^p(B)\}$ is a contravariant functor from the category of separable C*-algebras to the category of abelian groups.*

We conclude with some remarks about how the ideas in this section interact with topology. If X and Y are compact second countable Hausdorff spaces and $\alpha\colon X \to Y$ is a continuous map then α induces a $*$-homomorphism $\alpha^*\colon C(Y) \to$

$C(X)$ given by pulling back to X: $\alpha^*(f) = f \circ \alpha$. Both $C(Y)$ and $C(X)$ are separable, so if we choose ample representations of $C(Y)$ and $C(X)$ on separable Hilbert spaces (such representations always exist by the GNS construction) then we can apply Proposition 3.3.11 to obtain an isometry which covers α^*.

Suppose now that X and Y are second countable locally compact Hausdorff spaces and $\alpha\colon X \to Y$ is a continuous map between them. The pullback construction no longer produces an induced $*$-homomorphism $C_0(Y) \to C_0(X)$ in general; consider the case where X is non-compact and Y is a single point, for instance. Thus the category of locally compact Hausdorff spaces with continuous maps is ill-suited for K-homology theory. However, if α is a continuous *proper* map then the pullback $f \mapsto f \circ \alpha$ does define a $*$-homomorphism $\alpha^*\colon C(Y) \to C(X)$.

Definition 3.3.18. *Let $\alpha\colon X \to Y$ be a proper map between locally compact Hausdorff spaces and let H_X and H_Y be Hilbert spaces carrying ample representations of $C_0(X)$ and $C_0(Y)$, respectively. An isometry $V\colon H_X \to H_Y$* topologically covers *α if it covers the pullback map $\alpha^*\colon C_0(Y) \to C_0(X)$ in the sense of Definition 3.3.10*

In fact if $\alpha\colon U \to Y$ is a continuous proper map defined on an open subset $U \subseteq X$ then α still induces a map $\alpha^*\colon C_0(Y) \to C_0(X)$: extend α to a map $\widetilde{\alpha}\colon \widetilde{X} \to \widetilde{Y}$ between one-point compactifications by sending the complement of U in X to the point at infinity in Y and restrict $\widetilde{\alpha}^*\colon C(\widetilde{Y}) \to C(\widetilde{U}) \subseteq C(X)$ to $C_0(Y)$. Proper maps defined on open subsets of X constitute the full collection of morphisms for which K-homology is functorial, though we will only need functoriality for proper maps defined on all of X.

3.3.3 Relative K-Homology

In order to build long exact sequences and boundary maps in K-homology it useful (and perhaps necessary) to develop a relative version of the theory. In geometry this means extending the K-homology functor K_p to a functor of pairs $K_p(X, Y)$, where X is a locally compact Hausdorff space and Y is a closed subspace of X, which fits into a long exact sequence

$$\ldots \to K_p(Y) \to K_p(X) \to K_p(X, Y) \to K_{p-1}(Y) \to \ldots$$

In the context of C*-algebra theory the group $K_p(X,Y)$ corresponds to a group $K_p(A,A/J)$ where A is a C*-algebra and $J \subseteq A$ is an ideal. In this generality the construction of long exact sequences is a very delicate matter which, for our purposes, requires the additional assumption that the quotient map $A \to A/J$ is semisplit (i.e. it has a completely positive section). This assumption is always satisfied for commutative C*-algebras by Proposition 2.2.4, so the required long exact sequence exists in all of the cases of interest in this thesis.

For what follows let A be a separable C*-algebra and let J be an ideal in A. As usual ρ_A and $\rho_{A/J}$ will refer to ample reprsentations of A and A/J, respectively, on separable Hilbert spaces.

Definition 3.3.19.

- *An operator $T \in \mathfrak{D}^*(A)$ is* locally compact for J *if $T\rho_A(j) \sim \rho_A(j)T \sim 0$ for every $j \in J$.*
- *Denote by $\mathfrak{D}^*(A//J)$ the ideal in $\mathfrak{D}^*(A)$ of all operators which are locally compact for J.*

Remark 3.3.20. *The ideal $\mathfrak{D}^*(A//A)$ is the same as the ideal $\mathfrak{C}^*(A)$ from Definition 3.3.2.*

As we have shown, there is an isometry $V\colon H_{A/J} \to H_A$ which covers the quotient map $\pi\colon A \to A/J$ and hence a map $\mathrm{Ad}_V\colon \mathfrak{D}^*(A/J) \to \mathfrak{D}^*(A)$. In fact, we can say more:

Lemma 3.3.21. *Ad_V maps $\mathfrak{D}^*(A/J)$ into $\mathfrak{D}^*(A//J)$*

Proof. For $T \in \mathfrak{D}^*(A/J)$ and $j \in J$, we have that $\mathrm{Ad}_V(T)\rho_A(j) = VTV^*\rho_A(j)$. Since V is an isometry this is equal to $VTV^*VV^*\rho_A(j)$, and by Lemma 3.3.12 this agrees with $VTV^*\rho_A(j)VV^*$ up to compact operators. Since V is a covering isometry we conclude that $\mathrm{Ad}_V(T)\rho_A(j) \sim VT\rho_{A/J}(\pi(j))V^* = 0$. A similar argument shows that $\rho_A(j)\mathrm{Ad}_V(T)$ is compact, so $\mathrm{Ad}_V(T)$ is locally compact for J. $\square$

We would like to prove that Ad_V induces an isomorphism on K-theory. This is not a trivial result, and it and we shall only prove it under the technical hypothesis that the quotient map $A \to A/J$ is semisplit.

Theorem 3.3.22. *If $\pi\colon A \to A/J$ admits a completely positive section then the map*

$$(Ad_V)_*\colon K_p(\mathfrak{D}^*(A/J)) \to K_p(\mathfrak{D}^*(A/\!/J))$$

is an isomorphism.

Proof. Let $\rho_A\colon A \to \mathbb{B}(H_A)$ and $\rho_{A/J}\colon A/J \to \mathbb{B}(H_{A/J})$ denote the ample representations of A and A/J, respectively, used to define the dual algebras appearing in the statement of the theorem. According to our definition of ample these representations extend to unital representations $\widetilde{\rho}_A\colon \widetilde{A} \to \mathbb{B}(H_A)$ and $\widetilde{\rho}_{A/J}\colon \widetilde{A}/J \to \mathbb{B}(H_{A/J})$, so we can assume without loss of generality that A, ρ_A, and $\rho_{A/J}$ are unital.

Let $\sigma\colon A/J \to A$ be a completely positive section of π. By Stinespring's theorem (Theorem 2.2.6), there are a Hilbert space H and a representation

$$\rho'_{A/J}\colon A/J \to \mathbb{B}(H_A \oplus H)$$

of the form

$$\rho'_{A/J} = \begin{pmatrix} \rho_A\sigma & \rho_{21} \\ \rho_{12} & \rho_{22} \end{pmatrix}$$

Let $H'_{A/J} = H_A \oplus H$ and let $W\colon H_A \to H'_{A/J}$ be the isometry $v \mapsto (v, 0)$. We claim that Ad_W maps $\mathfrak{D}^*_{\rho_A}(A/\!/J)$ into $\mathfrak{D}^*_{\rho'_{A/J}}(A/J)$. Given $T \in \mathfrak{D}^*_{\rho_A}(A/\!/J)$ and any $a \in A$ we have

$$[\mathrm{Ad}_W(T), \rho'_{A/J}(\pi(a))] = \begin{pmatrix} [T, \rho_A(\sigma\pi(a))] & T\rho_{12}(\pi(a)) \\ -\rho_{21}(\pi(a))T & 0 \end{pmatrix}$$

The upper left entry of this matrix is compact since T is pseudolocal. In the notation of Lemma 2.1.9 we have that $\rho_A(J) \subseteq \mathbb{K}(T)$ since T is locally compact for J, and thus $\rho_A\sigma\pi$ is a $*$-homomorphism modulo $\mathbb{K}(T)$ since $\sigma\pi$ is a $*$-homomorphism modulo J. It follows from Lemma 2.1.9 that $T\rho_{12}(\pi(a))$ and $\rho_{21}(\pi(a))T$ are compact, and hence $\mathrm{Ad}_W(T)$ is pseudolocal.

Now we show that the composition $WV\colon H_{A/J} \to H'_{A/J}$ covers the identity map $A/J \to A/J$, i.e. $\rho_{A/J}(\pi(a)) \sim V^*W^*\rho'_{A/J}(\pi(a))WV$. By our matrix expression for $\rho'_{A/J}$ we have $W^*\rho'_{A/J}(\pi(a))W = \rho_A(\sigma\pi(a))$, and since V covers π we have $V^*\rho_A(\sigma\pi(a))V \sim \rho_{A/J}(\pi\sigma\pi(a)) = \rho_{A/J}(\pi(a))$, as desired. But functoriality

$\mathrm{Ad}_W\mathrm{Ad}_V$ induces an isomorphism on K-theory and hence $(\mathrm{Ad}_V)_*$ is injective with left inverse $(\mathrm{Ad}_W)_*$.

To complete the proof we will use a version of Atiyah's rotation trick to show that $(\mathrm{Ad}_W)_*$ is injective and hence that $(\mathrm{Ad}_V)_*$ and $(\mathrm{Ad}_W)_*$ are inverses. Define $H'_A = H_A \oplus H'_{A/J}$ and define a representation $\rho_A\colon A \to \mathbb{B}(H'_A)$ by $\rho'_A = \rho_A \oplus \rho'_{A/J}\pi$. Let $X\colon H'_{A/J} \to H'_A$ be the isometry $v \mapsto (0, v)$ and observe that:

$$\begin{aligned} X^*\rho'_A(a)X &= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \rho_A(a) & 0 & 0 \\ 0 & \rho_A(\sigma\pi(a)) & \rho_{12}(\pi(a)) \\ 0 & \rho_{21}(\pi(a)) & \rho_{22}(\pi(a)) \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \\ &= \rho'_{A/J}(\pi(a)) \end{aligned}$$

Thus X covers π. The composition $XW\colon H_A \to H'_A = H_A \oplus H_A \oplus H$ includes H_A as the second summand, and a standard rotation yields a homotopy between the $*$-homomorphisms $\mathrm{Ad}_X\mathrm{Ad}_W$ and Ad_Y where Y is the inclusion of H_A as the first summand. But ρ'_A acts as ρ_A on the first summand, so Y covers the identity map $A \to A$. By homotopy invariance of K-theory, $(\mathrm{Ad}_X)_*(\mathrm{Ad}_W)_*$ is an isomorphism and hence $(\mathrm{Ad}_W)_*$ is injective. □

Corollary 3.3.23. *For any separable C^*-algebra A, $K^p(A) \cong K_{1-p}(\mathfrak{D}^*(A))$.*

Proof. It suffices to show that $\mathfrak{C}^*(A)$ has trivial K-theory. Since $\mathfrak{C}^*(A)$ can be expressed as $\mathfrak{D}^*(A//A)$, Theorem 3.3.22 implies that $K_p(\mathfrak{C}^*(A)) \cong K_p(\mathfrak{D}^*(\{0\}))$. The unitalization of $\{0\}$ is $\mathbb{C}$, so an ample representation of $\{0\}$ is simply the zero map $\{0\} \to \mathbb{B}(H)$ where H is an infinite dimensional Hilbert space. Thus $\mathfrak{D}^*(\{0\}) = \mathbb{B}(H)$ and hence $K_p(\mathfrak{D}^*(\{0\})) = 0$ for every p. □

We are now ready to define the relative K-homology groups $K_p(A, A/J)$ and fit them into a long exact sequence with $K_p(A)$ and $K_p(A/J)$.

Definition 3.3.24. *The* relative K-homology group *for the pair $(A, A/J)$ in degree p is defined to be*

$$K^p(A, A/J) = K_{1-p}(\mathfrak{D}^*(A)/\mathfrak{D}^*(A//J))$$

(where the dual algebra is as usual defined using an ample representation).

Thus if A is a separable C*-algebra, J is an ideal in A, and the quotient map $A \to A/J$ is semisplit then the short exact sequence $0 \to \mathfrak{D}^*(A//J) \to \mathfrak{D}^*(A) \to \mathfrak{D}^*(A)/\mathfrak{D}^*(A//J) \to 0$ gives rise to a long exact sequence in K-homology:

$$\ldots \to K^p(A/J) \to K^p(A) \to K^p(A, A/J) \to K^{p-1}(A/J) \to \ldots$$

Just as in K-theory for C*-algebras, the relative term $K^p(A, A/J)$ can be replaced with $K^p(J)$ via the following nontrivial excision result.

Theorem 3.3.25 (Excision Theorem for K-Homology)**.** *Let A be a separable C*-algebra equipped with a relative representation ρ on a separable Hilbert space H and let J be a C*-ideal in A equipped with the representation obtained by restricting ρ to J. Then the inclusion $\mathfrak{D}^*(A) \hookrightarrow \mathfrak{D}^*(J)$ induces an isomormphism $\mathfrak{D}^*(A)/\mathfrak{D}^*(A//J) \cong \mathfrak{D}^*(J)/\mathfrak{C}^*(J)$ and hence an isomoprhism $K^p(A, A/J) \cong K^p(J)$.*

Proof. From the trivial identity $\mathfrak{D}^*(A//J) = \mathfrak{D}^*(A) \cap \mathfrak{C}^*(J)$ we deduce that the inclusion $\mathfrak{D}^*(A) \hookrightarrow \mathfrak{D}^*(J)$ induces a well-defined injective map on quotient algebras. To prove that this map is surjective we must show that $\mathfrak{D}^*(J) = \mathfrak{D}^*(A) + \mathfrak{C}^*(J)$. Our strategy is to use Kasparov's technical theorem, Theorem B.0.6. To streamline the notation, identify A with its image $\rho(A)$ in $\mathbb{B}(H)$ in what follows.

Fix $T \in \mathfrak{D}^*(J)$. Let E_1 denote the C*-subalgebra of $\mathbb{B}(H)$ generated by all commutators $[a, T]$ for $a \in A$, let E_2 denote the C*-subalgebra $J \subseteq \mathbb{B}(H)$, and let Δ be the linear subspace $A \subseteq \mathbb{B}(H)$. E_1, E_2, and Δ are each separable, and it is clear from the definitions that Δ derives E_1. To apply Kasparov's technical theorem we need only check that $E_1 E_2 \subseteq \mathbb{K}(H)$, i.e. that $j[a, T]$ is compact for every $j \in J$, $a \in A$. Indeed:

$$\begin{aligned} j[a, T] &= jaT - jTa \\ &= jaT - Tja + Tja - jTa \\ &= [ja, T] + [T, j]a \end{aligned}$$

This operator is compact since T is pseudolocal for J.

Let L be the operator guaranteed by Kasparov's technical theorem, so that $(1 - L)[a, T]$, Lj, and $[L, a]$ are each compact operators for every $a \in A$, $j \in J$.

We shall prove that $(1-L)T \in \mathfrak{D}^*(A)$ and $LT \in \mathfrak{C}^*(J)$, implying that $T = (1-L)T + LT \in \mathfrak{D}^*(A) + \mathfrak{C}^*(J)$.

To prove that $(1-L)T \in \mathfrak{D}^*(A)$, write:

$$\begin{aligned}[(1-L)T, a] &= (1-L)Ta - a(1-L)T \\ &= (1-L)Ta - (1-L)aT + (1-L)aT - a(1-L)T \\ &= (1-L)[T,a] + [(1-L),a]T \sim 0\end{aligned}$$

Finally, $LT \in \mathfrak{C}^*(J)$ since $LTj \sim LjT \sim 0$ and $jLT \sim LjT \sim 0$. □

Thus the long exact sequence above becomes:

$$\ldots \to K^p(A/J) \to K^p(A) \to K^p(J) \to K^{p-1}(A/J) \to \ldots$$

A key fact about long exact sequences in K-homology is that they are natural with respect to $*$-homomorphisms.

Lemma 3.3.26. *Let A and B be separable C^*-algebras with ideals $J_A \subseteq A$ and $J_B \subseteq B$. Let $\phi\colon A \to B$ be a $*$-homomorphism which carries J_A into J_B. Then the following diagram commutes:*

$$\begin{tikzcd}
{} \arrow[r] & K^p(B/J_B) \arrow[r] \arrow[d] & K^p(B) \arrow[r] \arrow[d] & K^p(J_B) \arrow[r] \arrow[d] & K^{p-1}(B/J_B) \arrow[r] \arrow[d] & {} \\
{} \arrow[r] & K^p(A/J_A) \arrow[r] & K^p(A) \arrow[r] & K^p(J_A) \arrow[r] & K^{p-1}(A/J_A) \arrow[r] & {}
\end{tikzcd}$$

where the rows are long exact sequences in K-homology.

Proof. Fix ample representations ρ_A and ρ_B of A and B, respectively, on separable Hilbert spaces and let V be an isometry which covers ϕ. We begin by checking that Ad_V carries the ideal $\mathfrak{D}^*(B//J_B)$ into $\mathfrak{D}^*(A//J_A)$. Given $T \in \mathfrak{D}^*(B//J_B)$ we have already shown that $\mathrm{Ad}_V(T) \in \mathfrak{D}^*(A)$. For any $a \in J_A$ we have $T\rho_B(\phi(a) \sim \rho_B(\phi(a))T \sim 0$ since $\phi(a) \in J_B$ and T is locally compact for J_B. It thus follows from the same calculation which proves that Ad_V maps $\mathfrak{C}^*(B)$ into $\mathfrak{C}^*(A)$ (see Proposition 3.3.13) that $\mathrm{Ad}_V(T)$ is locally compact for J_A.

Consequently the following diagram commutes:

$$\begin{tikzcd}
0 \arrow[r] & \mathfrak{D}^*(B//J_B) \arrow[r] \arrow[d, "\mathrm{Ad}_V"] & \mathfrak{D}^*(B) \arrow[r] \arrow[d, "\mathrm{Ad}_V"] & \mathfrak{D}^*(B)/\mathfrak{D}^*(B//J_B) \arrow[r] \arrow[d, "\mathrm{Ad}_V"] & 0 \\
0 \arrow[r] & \mathfrak{D}^*(A//J_A) \arrow[r] & \mathfrak{D}^*(A) \arrow[r] & \mathfrak{D}^*(A)/\mathfrak{D}^*(A//J_A) \arrow[r] & 0
\end{tikzcd}$$

The result now follows from the naturality of long exact sequences in K-theory for $*$-homomorphisms. $\square$

We conclude this section by interpreting our results in the case $A = C_0(X)$ where X is a second countable locally compact Hausdorff space. The ideals in A are precisely the subalgebras of the form $C_0(X - Y)$ where Y is a closed subspace of X. This ideal fits into the short exact sequence

$$0 \to C_0(X - Y) \to C_0(X) \to C_0(Y) \to 0$$

where the surjection $C_0(X) \to C_0(Y)$ is given by restriction of functions. Let is introduce some additional notation in this setting:

- $\mathfrak{D}^*(Y \subseteq X) := \mathfrak{D}^*(C_0(X)//C_0(X - Y))$
- $K_p(Y \subseteq X) := K_{1-p}(\mathfrak{D}^*(C_0(X)//C_0(X - Y)))$
- $K_p(X, Y) := K_{1-p}(C_0(X), C_0(Y))$

According to Proposition 2.2.4, the map $C_0(X) \to C_0(Y)$ always has a completely positive section and hence Theorem 3.3.22 implies that the inclusion $Y \hookrightarrow X$ induces an isomorphism $K_p(Y) \cong K_p(Y \subseteq X)$. The excision theorem corresponds to the *strong excision axiom* $K_p(X, Y) \cong K_p(X - Y)$, and thus there is a long exact sequence in K-homology given by:

$$\ldots \to K_p(Y) \to K_p(X) \to K_p(X - Y) \to K_{p-1}(Y) \to \ldots$$

There are more attractive proofs of the excision theorem available in the setting of commutative C*-algebras; we will encounter one such argument when we discuss Mayer-Vietoris sequences.

3.4 Kasparov's K-Homology

So far we have considered Paschke's model for the K-homology groups of a C*-algebra A which requires fixing an ample representation ρ of A and passing to the K-theory of the dual algebra $\mathfrak{D}^*_\rho(A)$. However, it is convenient for certain constructions to have the freedom to vary the representation and to consider representations which are more degenerate than those allowed by Paschke. In this section we introduce Kasparov's model of K-homology which includes this extra flexibility at some cost to ease of computation. In the end Kasparov's and Paschke's models are equivalent for separable C*-algebras, and we will use the notation $K^p(A)$ for both of them. Their equivalence is mediated by Voiculescu's theorem, which implies that any ample representation absorbs any nondegenerate representation up to compact operators; more detail will be provided in an appendix.

There is a product structure in K-homology

$$K^{p_1}(A_1) \times K^{p_2}(A_2) \to K^{p_1+p_2}(A_1 \otimes A_2)$$

called the *Kasparov product* which is much easier (though still very difficult!) to construct in Kasparov's model. The Kasparov product is an extremely powerful tool which organizes many important results in K-homology theory. After defining Kasparov's model of K-homology we will carefully construct the Kasparov product and use it to compute boundary maps in the long exact sequence for K-homology. Our exposition follows chapters 8 and 9 of [9].

3.4.1 Basic Definitions

We begin by defining the objects which will become generators in Kasparov's model of K-homology.

Definition 3.4.1. *Let A be a separable C^*-algebra. A* Fredholm module *over A is a triple (ρ, H, F) where H is a separable Hilbert space $\rho\colon A \to \mathbb{B}(H)$ is a representation, and F is a bounded operator on H which satisfies:*

- $(F^2 - 1)\rho(a) \sim 0$
- $(F - F^*)\rho(a) \sim 0$

- $[F, \rho(a)] \sim 0$

for every $a \in A$.

A Fredholm module (ρ, H, F) over A is said to be degenerate *if $(F^2 - 1)\rho(a)$, $(F - F^*)\rho(a)$, and $[F, \rho(a)]$ are all exactly 0.*

As usual the symbol "$\sim$" means "equals up to a compact operator".

Example 3.4.2. *Let M be a smooth manifold, $S \to M$ a smooth vector bundle over M, and D an essentially self-adjoint first order elliptic operator acting on smooth sections of S. Let $H = L^2(M; S)$ equipped with the representation $\rho\colon C_0(M) \to \mathbb{B}(H)$ by multiplication operators, and let χ be any normalizing function. Then according to Proposition 3.2.10 and Proposition 3.2.21 the triple $(\rho, H, \chi(D))$ is a Fredholm module over $C_0(M)$.*

Just as it was important for the purpose of index theory to consider differential operators with additional grading structure, it is important to develop a notion of graded Fredholm module. Recall that a Hilbert space H is *graded* if it comes equipped with a direct sum decomposition $H = H^+ \oplus H^-$, or equivalently a self-adjoint unitary operator $\gamma \in \mathbb{B}(H)$ (whose ± 1-eigenspaces correspond to $H^\pm$). We say that an operator on H is *even* if it commutes with the grading operator γ and *odd* if it anti-commutes with γ.

Kasparov's definition of K-homology calls for some additional grading structure. We say that a Hilbert space is *p-multigraded* for $p \in \mathbb{N}_0$ if it comes equipped with p odd unitary operators $\varepsilon_1, \ldots, \varepsilon_p$ which anti-commute with each other and satisfy $\varepsilon_i^2 = -1$. We adopt the conventions that a 0-multigraded Hilbert space is just a graded Hilbert space with no multigraded operators and a -1-multigraded Hilbert space is simply an ungraded Hilbert space. We say that an operator is *even* (resp. *odd*) *multigraded* if it commutes (resp. anti-commutes) with the grading operator and commutes with all of the multigrading operators.

Note that the relations among the multigrading operators are precisely the relations satisfied by generators for the Clifford algebra $\mathbb{C}_p$, and indeed it is convenient to think of a p-multigraded Hilbert space as a Hilbert space equipped with the structure of a graded (right) module over $\mathbb{C}_p$.

Definition 3.4.3. *Let $p \geq -1$ be an integer. A p-multigraded Fredholm module over A is a Fredholm module (ρ, H, F) where H is a p-multigraded Hilbert space, $\rho(a)$ is even multigraded for every $a \in A$, and F is odd multigraded.*

There is a natural notion of direct sum of Fredholm modules. If (ρ, H, F) and (ρ', H', F') are two p-multigraded Fredholm modules over A with multigrading operators $\varepsilon_1, \ldots, \varepsilon_p$ and $\varepsilon'_1, \ldots, \varepsilon'_p$, respectively, then $(\rho \oplus \rho', H \oplus H', F \oplus F')$ is a p-multigraded Fredholm module over A with multigrading operators $\varepsilon_1 \oplus \varepsilon'_1, \ldots, \varepsilon_p \oplus \varepsilon'_p$. The Fredholm module for which the Hilbert space, the representation, the operator, and the multigrading operators are all 0 can be regarded as the additive identity for direct sum.

There is a natural equivalence relation on the set of all Fredholm modules over a fixed C*-algebra A which echoes the equivalence relation which defines the K-theory groups of $\mathfrak{D}^*(A)/\mathfrak{C}^*(A)$.

Definition 3.4.4. *Let (ρ, H, F) and (ρ', H', F') be p-multigraded Fredholm modules over a C*-algebra A.*

- *(ρ, H, F) and (ρ', H', F') are* unitarily equivalent *if there exists an even multigraded unitary isomorphism $U\colon H' \to H$ such that $\rho' = U^*\rho U$ and $F' = U^*FU$.*

- *(ρ, H, F) and (ρ', H', F') are* operator homotopic *if $H = H'$, $\rho = \rho'$, the multigrading operators are all the same, and there is a norm continuous path $t \mapsto F_t$, $t \in [0,1]$, such that $F = F_0$ and $F' = F_1$.*

- *(ρ, H, F) is a* compact perturbation *of (ρ', H', F') if $H = H'$, $\rho = \rho'$, the multigrading operators are all the same, and $(F - F')\rho(a)$ is compact for every $a \in A$.*

We shall say that (ρ, H, F) and (ρ', H', F') are K-equivalent *if there is a finite sequence of p-multigraded Fredholm modules over A starting with (ρ, H, F) and ending with (ρ', H', F') such that each Fredholm module in the sequence differs from the next by unitary equivalence, operator homotopy, or compact perturbation.*

We are now ready to specify Kasparov's model for K-homology:

Proposition 3.4.5. *Let $p \geq -1$ be an integer. Define $KK^{-p}(A, \mathbb{C})$ to be the abelian group generated by K-equivalence classes of p-multigraded Fredholm modules with the following relation:*

$$[(\rho, H, F)] + [(\rho', H', F')] = [(\rho, H, F) \oplus (\rho', H', F')]$$

Remark 3.4.6. *The notation $KK^{-p}(A, \mathbb{C})$ is consistent with Kasparov's bivariant KK-theory, which we will not need except to define Kasparov's model of K-homology. We shall remark soon that $KK^{-p}(A, \mathbb{C})$ is isomorphic to $K^{-p}(A)$, but until then we will stick with the notation $KK^{-p}(A, \mathbb{C})$.*

Before proceeding any further we comment on the significance of degenerate Fredholm modules:

Lemma 3.4.7. *Let (ρ, H, F) be a degenerate p-multigraded Fredholm module over A. Then $[\rho, H, F] = 0$ in $KK^{-p}(A, \mathbb{C})$.*

Proof. Let $\rho' = \bigoplus_{\mathbb{N}} \rho$, $H' = \bigoplus_{\mathbb{N}} H$, and $F' = \bigoplus F$. The triple (ρ', H', F') is a degenerate Fredholm module (note that it would not even be a Fredholm module if (ρ, H, F) weren't degenerate), and $(\rho, H, F) \oplus (\rho', H', F')$ is unitarily equivalent to (ρ', H', F'). Thus at the level of K-homology we have:

$$[\rho, H, F] + [\rho', H', F'] = [\rho', H', F']$$

Cancelling, we conclude that $[\rho, H, F] = 0$. $\square$

Calculating $KK^{-p}(A, \mathbb{C})$ directly can be somewhat involved; we shall illustrate this with an example:

Example 3.4.8. *Let us show that $KK^0(\mathbb{C}, \mathbb{C}) \cong \mathbb{Z}$. Let (ρ, H, F) be a graded Fredholm module over $\mathbb{C}$, and let $P \in \mathbb{B}(H)$ denote the projection $\rho(1)$. We have:*

$$F = \begin{pmatrix} PFP & PF(1-P) \\ (1-P)FP & (1-P)F(1-P) \end{pmatrix}$$

The off diagonal entries are compact operators since F and P commute modulo *compacts, so* (ρ, H, F) *is a compact perturbation of the Fredholm module*

$$(\rho, PH, PFP) \oplus (\rho, (1-P)H, (1-P)F(1-P))$$

The K-equivalence class of the second summand is 0 *since* ρ *acts as the trivial representation so* (ρ, H, F) *is K-equivalent to* (ρ, PH, PFP)*; write* $F_P = PFP$*.*

F_P *is an odd operator such that* $F_P - F_P^*$ *and* $F_P^2 - 1$ *are compact (since* $\rho(1)$ *is the identity operator on* PH*), so we have:*

$$F_P = \begin{pmatrix} 0 & V_P \\ U_P & 0 \end{pmatrix}$$

where $U_P V_P^* \sim U_P^* V_P \sim 1$*. In particular* U_P *is essentially unitary and therefore Fredholm, and in fact the assignment* $(\rho, H, F) \mapsto$ *Index*(U_P) *gives rise to a well-defined group homomorphism Ind:* $K^0(\mathbb{C}) \to \mathbb{Z}$*.*

According to a standard result in functional analysis, an essentially unitary Fredholm operator with index 0 *differs from a unitary by a compact operator. If* U *is a unitary operator on* PH *then* (ρ, PH, U) *is degenerate, so it follows that* (ρ, H, F) *is degenerate if and only if* U_P *has index* 0*. This shows that Ind is injective.*

To show that Ind is an isomorphism it suffices to construct a graded Fredholm module (ρ, H, F) *over* $\mathbb{C}$ *such that Ind*$[\rho, H, F] = 1$*. Let* H *denote the graded Hilbert space* $\ell^2(\mathbb{R}) \oplus \ell^2(\mathbb{R})$*, let* ρ *be the representation* $\rho(\lambda) = \lambda I$ *where* I *is the identity operator on* H*, and let*

$$F = \begin{pmatrix} 0 & S^* \\ S & 0 \end{pmatrix}$$

where S *is the left shift operator on* $\ell^2(\mathbb{R})$*. Then Ind*$[\rho, H, F] = 1$ *since* S *has Fredholm index* 1*.*

This identification $KK^0(\mathbb{C}, \mathbb{C}) \cong \mathbb{Z}$ will become important later on; specifically it will be convenient to have a label for the generator.

Definition 3.4.9. *The* unit class *in* $KK^0(\mathbb{C}, \mathbb{C})$ *is the K-equivalence class of any*

Fredholm module (ρ, H, F) *such that*

$$\rho(1)F\rho(1) = \begin{pmatrix} 0 & V \\ U & 0 \end{pmatrix}$$

where U *is an essentially unitary operator with Fredholm index* 1*.*

Example 3.4.10. *Let* M *be a smooth manifold and let* $S \to M$ *a smooth graded Hermitian vector bundle over* M*. Say that* S *is a p-multigraded vector bundle if it is equipped with odd-graded unitary bundle morphisms* $\varepsilon_1, \ldots, \varepsilon_p$ *which anti-commute and satisfy* $\varepsilon_j^2 = -1$*. Similarly, say that an essentially self-adjoint first order elliptic differential operator* D *on is p-multigraded if it is odd and it anti-commutes with the* ε_j*'s. The Fredholm module of Example 3.4.2 associated to such an operator is then p-multigraded and hence it determines a class in* $KK^{-p}(C_0(M), \mathbb{C})$*.*

With the definitions and basic examples out of the way, we now discuss the relationship between $KK^{-p}(A, \mathbb{C})$ and $K^p(A)$. It can be shown independently that each of these groups satisfies the Bott periodicity theorem: $KK^{-p}(A, \mathbb{C}) \cong KK^{-p-2}(A, \mathbb{C})$ and $K^p(A) \cong K^{p+2}(A)$. Thus all degrees can be taken mod 2. For what follows, recall that $K^p(A)$ is defined to be $K_{1-p}(\mathfrak{D}^*(A)/\mathfrak{C}^*(A))$ where $\mathfrak{D}^*(A)$ and $\mathfrak{C}^*(A)$ are defined using a fixed ample representation $\rho_A : A \to \mathbb{B}(H_A)$ on a separable Hilbert space.

- $p = 1$:

  Let $P \in M_n(\mathfrak{D}^*(A))$ be an element whose image in $M_n(\mathfrak{D}^*(A)/\mathfrak{C}^*(A))$ is a projection, so that $(P^2 - P)\rho_A(a) \sim 0$ for every $a \in A$. Then $(\rho_A, H_A^n, 2P-1)$ is an ungraded Fredholm module over A. There is a unique homomorphism

  $$\Gamma_1 \colon K^1(A) \to KK^1(A, \mathbb{C})$$

  which satisfies $\Gamma_1[P] = [\rho_A, H_A^n, 2P - 1]$

- $p = 0$:

  Let $U \in M_n(\mathfrak{D}^*(A))$ be an element whose image in $\mathfrak{D}^*(A)/\mathfrak{C}^*(A)$ is a unitary, so that $(UU^* - 1)\rho_A(a) \sim (U^*U - 1)\rho_A(a) \sim 0$ for every $a \in A$. Then

$$\left(\rho_A \oplus \rho_A, H_A^n \oplus H_A^n, \begin{pmatrix} 0 & U^* \\ U & 0 \end{pmatrix}\right)$$

is a graded Fredholm module over A. There is a unique homomorphism

$$\Gamma_0\colon K^0(A) \to KK^0(A, \mathbb{C})$$

which satisfies $\Gamma_0[U] = \left(\rho_A \oplus \rho_A, H_A^n \oplus H_A^n, \begin{pmatrix} 0 & U^* \\ U & 0 \end{pmatrix}\right)$.

Proposition 3.4.11. *The maps $\Gamma_1\colon K^1(A) \to KK^1(A, \mathbb{C})$ and $\Gamma_0\colon K^0(A) \to KK^0(A, \mathbb{C})$ defined above are isomorphisms.*

We postpone the proof until Appendix A and take the result for granted for the remainder of this chapter. With this proposition in hand we shall now retire the notation $KK^{-p}(A, \mathbb{C})$ and simply specify which model for K-homology we are using when it is not clear from context.

We conclude with a brief discussion of how relative K-homology is defined in the Kasparov model.

Definition 3.4.12. *Let A be a separable C^*-algebra and let J be an ideal in A. A* relative Fredholm module *for the pair $(A, A/J)$ is a triple (ρ, H, F) where H is a separable Hilbert space, $\rho\colon A \to \mathbb{B}(H)$ is a representation, and $F \in \mathbb{B}(H)$ satisfies*

- $(F^2 - 1)\rho(j) \sim 0$ *for every $j \in J$*
- $(F - F^*)\rho(j) \sim 0$ *for every $j \in J$*
- $[F, \rho(a)] \sim 0$ *for every $a \in A$*

We can also speak of graded and multigraded relative Fredholm modules, defined in the obvious way. The set of all K-equivalence classes of p-multigraded relative Fredholm modules generates an abelian group which is isomorphic to the relative K-homology group $K^{-p}(A, A/J)$ for the pair $(A, A/J)$. The excision theorem for Kasparov's K-homology groups, which follows from Theorem 3.3.25 and Proposition 3.4.11, takes the following form:

Theorem 3.4.13 (The Excision Theorem). *The map $K^{-p}(A, A/J) \to K^{-p}(J)$ which restricts a relative Fredholm module over $(A, A/J)$ to an ordinary Fredholm module over J is an isomorphism.*

We conclude this section by addressing functoriality. Functoriality is considerably simpler in Kasparov's model of K-homology because there is no longer any need to reconcile the representations of the two given C*-algebras, and hence Voiculescu's theorem is not required. If $\phi\colon A \to B$ is a *-homomorphism and (ρ, H, F) is a p-multigraded Fredholm module over B then $(\rho \circ \phi, H, F)$ is a p-multigraded Fredholm module over A. The assignment $(\rho, H, F) \mapsto (\rho \circ \phi, H, F)$ respects the equivalence relations which define K-homology, so it induces a group homomorphism

$$\phi^*\colon K^{-p}(B) \to K^{-p}(A)$$

3.4.2 The Dirac Class

In this section we construct a concrete Fredholm module representing the K-homology class of an important differential operator on the open interval $(-1, 1)$. This K-homology class is important for computations with boundary maps, so this section serves both as an example illustrating the theory discussed so far and preparation for future calculations. The differential operator in question is an example of a *Dirac operator*, the general theory of which we will discuss in a later chapter. We only mention here that the K-homology class of a Dirac operator on a smooth manifold (assuming it exists) plays the role of the fundamental class in K-homology.

Let $\mathbb{C}_1$ denote the complex Clifford algebra with one generator; C_1 is generated as a complex vector space by the multiplicative identity 1 and an element e such that $e^2 = -1$. $\mathbb{C}_1$ has the structure of a graded algebra, where the grading is given by $\mathbb{C}_1 = \mathbb{C}1 \oplus \mathbb{C}e$. Let $c(e)\colon \mathbb{C}_1 \to \mathbb{C}_1$ denote left multiplication by e; as an odd algebra homomorphism it has the form:

$$c(e) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

Let ε_1 denote right Clifford multiplication by e, so that ε_1 is also given by:

$$\varepsilon_1 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

Definition 3.4.14. *The* (complex) spinor bundle *over the interval* $(-1,1)$ *is the trivial* 1*-multigraded vector bundle* $S_{(-1,1)} = \mathbb{C}_1 \times (-1,1) \to (-1,1)$ *whose multigrading operator is* ε_1*. The* (complex) spinor Dirac operator *on* $(-1,1)$ *is the* 1*-multigraded differential operator acting on smooth sections of* $S_{(-1,1)}$ *given by:*

$$D_{(-1,1)} = \begin{pmatrix} 0 & -\frac{d}{dx} \\ \frac{d}{dx} & 0 \end{pmatrix}$$

The standard Hermitian inner product on $\mathbb{C}_1$ gives $S_{(-1,1)}$ the structure of a Hermitian vector bundle, and $D_{(-1,1)}$ is symmetric with respect to this Hermitian structure. The symbol of $D_{(-1,1)}$ is given by:

$$\sigma_D(x,\xi) = \begin{pmatrix} 0 & -\xi \\ \xi & 0 \end{pmatrix}$$

It follows that $D_{(-1,1)}$ is elliptic and it has finite propagation speed relative to the standard Riemannian metric on $(-1,1)$. Hence D is a Dirac-type operator.

However, $D_{(-1,1)}$ is NOT essentially self adjoint; Proposition 3.2.6 does not apply since $(-1,1)$ is not complete, and indeed $D_{(-1,1)}$ has multiple different extensions to unbounded operators on $L^2((-1,1); S_{(-1,1)})$ corresponding to different choices of boundary conditions. We will resolve this by extending $D_{(-1,1)}$ to the compact manifold S^1 and using the excision map to build a class in the K-homology of $(-1,1)$.

View S^1 as the closed interval $[-1,1]$ with the endpoints identified and embed $(-1,1)$ into S^1 in the obvious way. View $J = C_0(-1,1)$ as an ideal in $A = C(S^1)$. Extend the bundle $S \to (-1,1)$ to a trivial bundle $S' = \mathbb{C}_1 \times S^1 \to S^1$ and extend ε_1 to a multigrading operator on S' which by abuse of notation we will still call ε_1. There is a differential operator D' on S' which, in the coordinate patch with coordinate t, takes the form

$$D' = \begin{pmatrix} 0 & -\frac{d}{dt} \\ \frac{d}{dt} & 0 \end{pmatrix}$$

D' is a symmetric first order elliptic differential operator on S^1, a closed manifold; by Corollary 3.2.8, D' is essentially self adjoint.

Let χ be a normalizing function and form the Fredholm module

$$(\rho, L^2(S^1; S'), \chi(D'))$$

where ρ is the representation of $C(S^1)$ on $L^2(S^1; S'$ by multiplication operators. Observe that $\chi(D')$ is exactly self-adjoint, and $\chi(D')^2 - 1$ is compact by Proposition 3.2.10. Thus $(\rho, L^2(S^1; S'), \chi(D'))$ defines a relative Fredholm module for the pair $(C(S^1), C(S^1)/C_0(-1,1))$.

Definition 3.4.15. *The* Dirac class $\boldsymbol{d}$ *is the image in* $K^{-1}(C_0(-1,1))$ *of the relative K-homology class* $[\rho, L^2(S^1; S'), \chi(D')] \in K^{-1}(C(S^1), C(S^1)/C_0(-1,1))$ *described above.*

Let us now simplify our representative of the Dirac class a bit further. Observe that $L^2(S^1; S') = L^2(S^1, \mathbb{C}) \oplus L^2(S^1; \mathbb{C})$ as a graded Hilbert space. A straightforward calculation shows that the spectrum of $D_{(-1,1)}$ thought of as an unbounded operator on $L^2(S^1, \mathbb{C}) \oplus L^2(S^1; \mathbb{C})$ is $\{n\pi \colon n \in \mathbb{Z}\}$, and the $n\pi$-eigenspace H_n of D is spanned by the orthogonal functions

$$\begin{pmatrix} 1 \\ i \end{pmatrix} e^{in\pi x} \quad \text{and} \quad \begin{pmatrix} 1 \\ -i \end{pmatrix} e^{-in\pi x}$$

Choose a normalizing function χ such that $\chi(\lambda) = -1$ for $\lambda \leq -\pi$ and $\chi(\lambda) = 1$ for $\lambda \geq \pi$. Then $\chi(D_{(-1,1)})$ acts as the identity on H_n for $n \geq 0$ and as minus the identity for $n < 0$. Consider a vector of the form $e^{in\pi x} \oplus 0 \in L^2(S^1, \mathbb{C}) \oplus L^2(S^1; \mathbb{C})$ and write

$$\begin{pmatrix} e^{in\pi x} \\ 0 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} e^{in\pi x} + \frac{1}{2} \begin{pmatrix} 1 \\ -i \end{pmatrix} e^{in\pi x}$$

The first vector in the sum above lies in H_n while the second vector lies in H_{-n}.

Thus:

$$\chi(D_{(-1,1)})\begin{pmatrix} e^{in\pi x} \\ 0 \end{pmatrix} = \begin{cases} \begin{pmatrix} 0 \\ i \end{pmatrix} e^{in\pi x} & n \geq 0 \\ \begin{pmatrix} 0 \\ -i \end{pmatrix} e^{in\pi x} & n < 0 \end{cases}$$

Consider the Hilbert transform Y on $L^2(S^1; \mathbb{C})$ given by

$$Ye^{in\pi x} = \begin{cases} e^{in\pi x} & n \geq 0 \\ -e^{in\pi x} & n < 0 \end{cases} \tag{3.4.1}$$

We have shown that $\chi(D_{(-1,1)})(e^{in\pi x} \oplus 0) = 0 \oplus iYe^{in\pi x}$. Similarly we have $\chi(D_{(-1,1)})(0 \oplus e^{im\pi x}) = -iYe^{im\pi x} \oplus 0$. Thus we conclude that

$$\chi(D_{(-1,1)}) = \begin{pmatrix} 0 & -iY \\ iY & 0 \end{pmatrix}$$

In particular, we may represent the Dirac class as:

$$\mathbf{d} = [\rho, L^2(S^1; S'), \begin{pmatrix} 0 & -iY \\ iY & 0 \end{pmatrix}] \tag{3.4.2}$$

(equipped with the usual multigrading operator $\varepsilon_1 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$).

3.4.3 The Suspension Map

Let A be a C*-algebra and consider the short exact sequence

$$0 \to S(A) \to C(A) \to A \to 0$$

where $S(A) = C_0(-1,1) \otimes A$ is the suspension of A and $C(A) = C_0[-1,1) \otimes A$ is the cone over A. This short exact sequence gives rise to a long exact sequence in K-homology and hence there is a boundary map $K^{-p-1}(S(A)) \to K^{-p}(A)$ which we will call the *suspension map*. Our aim in this section is to exhibit the suspension map explicitly at the level of Fredholm modules and prove that the suspension map $K^{-1}(C_0(-1,1)) \to K^0(\mathbb{C})$ sends the Dirac class to the unit class (see Example 3.4.8). When combined with the Kasparov product in the next section, this will yield a powerful tool for calculating certain boundary maps in K-homology.

Let (ρ, H, F) be a $(p+1)$-multigraded relative Fredholm module for the pair $(C_b(-1,1) \otimes A, C_b(-1,1) \otimes A / C_0(-1,1) \otimes A)$ and let $X_0 \in \mathbb{B}(H)$ denote the image under ρ of the bounded continuous function $x \mapsto x$ on $(-1,1)$. Let $X \in \mathbb{B}(H)$ denote the odd self-adjoint operator $X = \gamma \varepsilon_1 X_0$ where γ is the grading operator on H and ε_1 is the first multigrading operator. Note that there is a natural representation of A on H given by $a \mapsto \rho(1 \otimes a)$.

Definition 3.4.16. *The* suspension *of a $(p+1)$-multigraded relative Fredholm module (ρ, H, F) for the pair $(C_b(-1,1) \otimes A, C_b(-1,1) \otimes A / C_0(-1,1) \otimes A)$ with multigrading operators $\varepsilon_1, \ldots, \varepsilon_{p+1}$ is the p-multigraded Fredholm module (ρ, H, V) over A where*

$$V = -X + (1 - X^2)^{1/2} F$$

and the multigrading operators are given by $\varepsilon_2, \ldots, \varepsilon_{p+1}$.

Implicit in this definition is the following lemma:

Lemma 3.4.17. *The triple (ρ, H, V) defined above is a Fredholm module.*

Proof. It is clear that V is pseudolocal. Note that $(1 - X^2)^{1/2} = (1 - X_0^2)^{1/2} \in C_0(-1,1)$ so this operator commutes with F modulo compacts and hence $(V - V^*)\rho(a) \sim 0$ for every $a \in A$. So it suffices to show that $(V^2 - 1)\rho(a) \sim 0$. Note that $X(1 - X^2)^{1/2}$ is the product of $\gamma\varepsilon_1$ with a function in $C_0(-1,1)$, so this operator anti-commutes with F modulo compact operators and hence X anti-commutes with $(1 - X^2)^{1/2} F$ modulo compact operators. Thus,

$$V^2 \rho(a) \sim (X^2 + (1 - X^2) Y^2) \rho(a) \sim \rho(a)$$

as desired. □

The assignment $(\rho, H, F) \mapsto (\rho, H, V)$ respects the relations which define K-homology, and in fact the induced map $K^{-p-1}(S(A)) \to K^{-p}(A)$ is the suspension map defined above. The proof of this fact is somewhat technical, so we defer it to Appendix A.

Proposition 3.4.18. *Let* $\boldsymbol{d} \in K^{-1}(C_0(-1,1))$ *be the Dirac class and let* $\boldsymbol{1} \in K^0(\mathbb{C})$ *be the unit class. We have:*

$$s(\boldsymbol{d}) = \boldsymbol{1}$$

Proof. We use the representative of the Dirac class given by (3.4.2). Thus,

$$V = \begin{pmatrix} 0 & X_0 - i(1-X_0^2)^{\frac{1}{2}}Y \\ X_0 + i(1-X_0^2)^{\frac{1}{2}}Y & 0 \end{pmatrix}$$

where X_0 corresponds to pointwise multiplication by the function $x \mapsto x$ on $L^2(S^1)$ and Y is the Hilbert transform defined in (3.4.1). By the definition of the unit class (together with the fact that $\rho(1)$ is the identity), we must show that V has (graded) index one. The graded index of V is the ordinary index of the lower left hand corner $V^{-+} = X_0 + i(1-X_0^2)^{\frac{1}{2}}Y$, so we must show that this operator has index 1.

Our strategy is to construct a homotopy between V^{-+} and a Toeplitz-type operator of index 1. Specifically, let P_Y denote the projection $\frac{1}{2}(1+Y)$ and consider the operator

$$W = e^{-i\pi X_0}P_Y - (1-P_Y)$$

Set $S = \sin(\frac{\pi}{2}X_0)$; using the functional calculus we observe that $e^{-i\frac{\pi}{2}X_0} = (1-S^2)^{\frac{1}{2}} + iS$. Consequently,

$$\begin{aligned} e^{i\frac{\pi}{2}X_0}W &= ((1-S^2)^{\frac{1}{2}} + iS)P_Y - ((1-S^2)^{\frac{1}{2}} - iS)(1-P_Y) \\ &= -iS + (1-S^2)^{\frac{1}{2}}Y \end{aligned}$$

Thus the path

$$t \mapsto e^{-it\frac{\pi}{2}X_0}W$$

for $t \in [0,1]$ defines a homotopy between W and the operator $iS + (1-S^2)^{\frac{1}{2}}Y$. Moreover the straight line path from S to X defines a homotopy between $-iS+(1-S^2)^{\frac{1}{2}}Y$ and $-iX_0 + (1-X_0^2)^{\frac{1}{2}}Y$. Both of these paths are paths through Fredholm operators and thus the index is preserved, so it suffices to calculate the index of W.

Now, let $\{e_n\}_{n\in\mathbb{Z}}$ be the orthogonal basis of $L^2(S^1) = L^2(-1,1)$ consisting of the exponential functions $e_n(x) = e^{\pi i n x}$. In this basis P_Y is the orthogonal projection operator onto the subspace spanned by those e_n's with $n \geq 0$ and $e^{i\pi X_0}$ is the shift operator $e_n \mapsto e_{n-1}$. Thus,

$$We_n = \begin{cases} e_{n-1} & n \geq 1 \\ 0 & n = 0 \\ e_n & n < 0 \end{cases}$$

This operator clearly has index 1, so we are done. □

3.5 The Kasparov Product

We now equip K-homology with the Kasparov product, a $\mathbb{Z}$-linear map

$$K^{-p_1}(A_1) \times K^{-p_2}(A_2) \to K^{-p_1-p_2}(A_1 \otimes A_2)$$

where A_1 and A_2 are separable C*-algebras. This product structure is compatible with both the topological features of K-homology (such as the suspension map and homotopy invariance) and the analytic features (such as elliptic operator theory). It was first introduced by Kasparov in the more general setting of his bivariant KK-theory, but as with many constructions in K-homology it was hinted at by some of the elliptic analysis developed by Atiyah and Singer in their original papers on index theory.

The goal of this section is to construct the Kasparov product with the aid of Kasparov's technical theorem and prove that it is compatible with the suspension map defined in the last section. This will allow us to calculate certain K-homology boundary maps.

3.5.1 The Construction of the Product

Given smooth vector bundles $S_1 \to M_1$ and $S_2 \to M_2$ over smooth manifolds and essentially self-adjoint first order elliptic differential operators D_1 and D_2 on S_1 and S_2, respectively, consider the operator $D = D_1 \otimes 1 + 1 \otimes D_2$ on the bundle $S_1 \otimes S_2 \to M_1 \times M_2$. This is again an essentially self-adjoint first order elliptic differential operator, so it defines a class $[D]$ in the K-homology of M. How can this class be recovered from the classes $[D_1]$ and $[D_2]$ in the K-homology of M_1 and M_2, respectively? To catch a glimpse at the nontrivial functional-analytic issues involved, note that even at the level of Fredholm modules it is not easy to relate the operator $\chi(D)$ to $\chi_1(D_1)$ and $\chi_2(D_2)$ where χ, χ_1, and χ_2 are normalizing functions. Resolving this issue leads to the Kasparov product, though a number of deep issues in functional analysis must be addressed along the way.

Suppose (ρ_i, H_i, F_i) is a p_i-multigraded Fredholm module over A_i for $i = 1, 2$. The graded Hilbert space $H_1 \hat{\oplus} H_2$ comes equipped with a natural representation of the (minimal) tensor product $A_1 \otimes A_2$ given by

$$\rho(a_1 \otimes a_2) = \rho_1(a_1) \otimes \rho_2(a_2)$$

Note that if $A_1 = C_0(X_1)$ and $A_2 = C_0(X_2)$ where X_1 and X_2 are locally compact Hausdorff spaces then the tensor product $A_1 \otimes A_2$ is naturally isomorphic to $C_0(X_1 \times X_2)$. Let $\varepsilon_1, \ldots, \varepsilon_{p_1}$ and $\varepsilon'_1, \ldots, \varepsilon p_2{}'$ denote the multigrading operators on H_1 and H_2, respectively; since $\varepsilon_i \hat{\otimes} 1$ anticommutes with $1 \hat{\otimes} \varepsilon'_j$ according to the rules of graded tensor products it follows that the operators

$$\varepsilon_1 \hat{\otimes} 1, \ldots, \varepsilon_{p_1} \hat{\otimes} 1, 1 \hat{\otimes} \varepsilon'_1, \ldots, 1 \hat{\otimes} \varepsilon'_{p_2}$$

give $H_1 \hat{\otimes} H_2$ the structure of a $(p_1 + p_2)$-multigraded Hilbert space. The representation ρ commutes with all of these multigrading operators, so all of the structure needed to define a $(p_1 + p_2)$-multigraded Fredholm module over $A_1 \otimes A_2$ is in place except for the operator.

Definition 3.5.1. *A Fredholm module (ρ, H, F) over $A_1 \otimes A_2$ is* aligned *with (ρ_1, H_1, F_1) and (ρ_2, H_2, F_2) if*

- *The operators*
$$\rho(a)(F(F_1\hat{\otimes}1) + (F_1\hat{\otimes}1)F)\rho(a^*)$$
and
$$\rho(a)(F(1\hat{\otimes}F_2) + (1\hat{\otimes}F_2)F)\rho(a^*)$$
are positive modulo compacts.

- *The operator $\rho(a)F$ derives $\mathbb{K}(H_1)\otimes\mathbb{B}(H_2)$*

Recall that the third condition means that $[\rho(a)F, K_1\otimes T_2] \in \mathbb{K}(H_1)\otimes\mathbb{B}(H_2)$ for every $K_1\otimes T_2 \in \mathbb{K}(H_1)\otimes\mathbb{B}(H_2)$. The motivation for the positivity conditions in this definition comes from the following lemma about Fredholm modules.

Lemma 3.5.2. *Let (ρ, H, F_0) and (ρ, H, F_1) be two Fredholm modules over the same C^*-algebra A and suppose that*

$$\rho(a)(F_0F_1 + F_1F_0)\rho(a^*)$$

is positive modulo compacts for every $a \in A$. Then F_0 and F_1 are operator homotopic.

Proof. Assume without loss of generality that F_0 and F_1 are self-adjoint. Since the operator $T = F_0F_1 + F_1F_0$ commutes with $\rho(a)$ modulo compacts, the hypothesis on T together with a functional calculus argument imply that there is a positive operator S with the property that $(T - S)\rho(a)$ is compact for every $a \in A$. It is easy to check that $[T, F_0]\rho(a)$ and $[T, F_1]\rho(a)$ are compact, so a straightforward calculation shows that:

$$\left(\cos\left(\frac{\pi}{2}t\right)F_0 + \sin\left(\frac{\pi}{2}t\right)F_1\right)^2\rho(a) \sim \left(1 + \cos\left(\frac{\pi}{2}t\right)\sin\left(\frac{\pi}{2}t\right)S\right)\rho(a)$$

Thus the operator

$$F_t = \left(\cos\left(\frac{\pi}{2}t\right)F_0 + \sin\left(\frac{\pi}{2}t\right)F_1\right)\left(1 + \cos\left(\frac{\pi}{2}t\right)\sin\left(\frac{\pi}{2}t\right)S\right)^{-\frac{1}{2}}$$

is a Fredholm module operator for all t and hence defines an operator homotopy between F_0 and F_1. $\square$

We shall prove that given a Fredholm module over A_1 and a Fredholm module over A_2 there is a third Fredholm module over $A_1 \otimes A_2$ which is aligned with them, and that the homotopy class of the third Fredholm module is uniquely determined by the homotopy classes of the first two. With that in mind, the definition of the Kasparov product is very straightforward:

Definition 3.5.3. *The* Kasparov product *of two K-homology classes $\boldsymbol{x} \in K^{-p_1}(A_1)$ and $\boldsymbol{y} \in K^{-p_2}(A_2)$ is the class in $K^{-p_1-p_2}(A_1 \otimes A_2)$ of any Fredholm module over $A_1 \otimes A_2$ which is aligned with a representative of $\boldsymbol{x}$ and a representative of $\boldsymbol{y}$.*

Taken at face value this definition only works for graded Fredholm modules, but it can be extended to ungraded Fredholm modules in various ways. Since we will only need the Kasparov product in the graded case we will ignore this issue.

We now turn to the problem of constructing a Fredholm module aligned with (ρ_1, H_1, F_1) and (ρ_2, H_2, F_2). This uses the following crucial construction:

Proposition 3.5.4. *Let H_1 and H_2 be separable multigraded Hilbert spaces, and let Δ be a separable subset of $\mathbb{B}(H_1 \hat{\otimes} H_2)$ which derives $\mathbb{K}(H_1) \otimes \mathbb{B}(H_2)$. Then there is a commuting pair of positive operators $N_1, N_2 \in \mathbb{B}(H_1 \hat{\otimes} H_2)$ with the following properties:*

- *N_1 and N_2 preserve the multigrading*
- $N_1^2 + N_2^2 = 1$
- $N_1(\mathbb{K}(H_1) \otimes \mathbb{B}(H_2)) \subseteq \mathbb{K}(H_1 \hat{\otimes} H_2)$
- $N_2(\mathbb{B}(H_2) \otimes \mathbb{K}(H_2)) \subseteq \mathbb{K}(H_1 \hat{\otimes} H_2)$
- *N_1 and N_2 commute modulo $\mathbb{K}(H_1 \hat{\otimes} H_2)$ with every operator in Δ*

Moreover the set of all pairs N_1, N_2 which satisfy these conditions is path connected.

The pair of operators N_1, N_2 whose existence is guaranteed by the proposition is called a *partition of unity adapted to* Δ. The proof of this proposition uses Kasparov's technical theorem and appears in Appendix B.

Proposition 3.5.5. *Let (ρ_1, H_1, F_1) and (ρ_2, H_2, F_2) be multigraded Fredholm modules over separable C^*-algebras A_1 and A_2, respectively. Let $H = H_1 \hat{\otimes} H_2$ and let $\rho \colon A_1 \otimes A_2 \to \mathbb{B}(H_1 \hat{\otimes} H_2)$ be the tensor product representation. Then there is a Fredholm module (ρ, H, F) which is aligned with F_1 and F_2, and the operator homotopy class of F is uniquely determined by those of F_1 and F_2.*

Proof. Let Δ be the subset of $\mathbb{B}(H_1) \otimes \mathbb{B}(H_2)$ which contains $F_1 \otimes 1$, $1 \otimes F_2$, and $\rho(a_1 \otimes a_2)$ for every $a_1 \in A_1$ and $a_2 \in A_2$. Let N_1, N_2 be a partition of unity for $H_1 \otimes H_2$ which is adapted to Δ and define

$$F = N_1(F_1 \hat{\otimes} 1) + N_2(1 \hat{\otimes} F_2)$$

Step 1: Prove that F is a Fredholm module operator.

- First, we show that $(F^2 - 1)\rho(a) \sim 0$ for every $a \in A_1 \otimes A_2$. We have that

  $$F^2 \sim N_1^2(F_1^2 \hat{\otimes} 1) + N_2^2(1 \hat{\otimes} F_2^2)$$

  since the cross terms $N_1 N_2((F_1 \hat{\otimes} 1)(1 \hat{\otimes} F_2) + (1 \hat{\otimes} F_2)(F_1 \hat{\otimes} 1))$ vanish according to the conventions of graded tensor products. Thus for $a = a_1 \otimes a_2 \in A_1 \otimes A_2$ we have

  $$\begin{aligned} F^2 \rho(a) &\sim N_1^2(F_1^2 \rho_1(a_1) \hat{\otimes} \rho_2(a_2)) + N_2^2(\rho_1(a_1) \hat{\otimes} F_2^2 \rho_2(a_2)) \\ &\sim (N_1^2 + N_2^2)\rho(a) = \rho(a) \end{aligned}$$

  Any element of $A_1 \otimes A_2$ is in the closure of the span of elementary tensors, so $F^2\rho(a) \sim \rho(a)$ for all a, as desired.

- Second, we show that $(F^* - F)\rho(a) \sim 0$ for every $a \in A_1 \otimes A_2$. Since

  $$F^* = (F_1^* \hat{\otimes} 1)N_1 + (1 \hat{\otimes} F_2^*)N_2$$

  and N_1 and N_2 commute modulo compacts with $F_1^* \hat{\times} 1$ and $1 \hat{\otimes} F_2^*$, this follows easily from the fact that F_1 and F_2 are Fredholm module operators.

- Finally, we show that $[F, \rho(a)] \sim 0$ for every $a \in A_1 \times A_2$. This follows easily for $a = a_1 \otimes a_2$ from the calculation

$$[F, \rho(a)] \sim N_1([F_1, \rho_1(a_1)]\hat{\otimes}\rho_2(a_2)) + N_2(\rho_1(a_1)\hat{\otimes}[F_2, \rho_2(a_2)]) \sim 0$$

The result for any $a \in A_1 \otimes A_2$ follows from the result for elementary tensors as explained above.

Step 2: Prove that F is aligned with F_1 and F_2.

We will check that $\rho(a)(F(F_1\hat{\otimes}1)+(F_1\hat{\otimes}1)F)\rho(a^*)$ is positive modulo compacts; the proof for $1\hat{\otimes}F_2$ proceeds similarly. We have that

$$F(F_1\hat{\otimes}1) + (F_1\hat{\otimes}1)F \sim 2N_1(F_1^2\hat{\otimes}1)$$

since the terms $N_2(F_1\hat{\otimes}1)(1\hat{\otimes}F_2)$ and $N_2(1\hat{\otimes}F_2)(F_1\hat{\otimes}1)$ cancel, so:

$$\rho(a)(F(F_1\hat{\otimes}1) + (F_1\hat{\otimes}1)F)\rho(a^*) \sim 2\rho(a)N_1\rho(a)^*$$

This is positive since N_1 is positive.

Step 3: Uniqueness.

We must show that the operator homotopy class of a Fredholm module aligned with F_1 and F_2 is uniquely determined by the operator homotopy classes of F_1 and F_2. If F_1 is varied by an operator homotopy $F_1(t)$ then we may form $F(t) = N_1(F_1(t)\hat{\otimes}1) + N_2(1\hat{\otimes}F_2)$ using a partition of unity adapted to the set Δ spanned by $F_1(t)\hat{\otimes}1$ for all t, $1\hat{\times}F_2$, and $\rho(a_1 \otimes a_2)$ for all $a_1 \in A_1$ and $a_2 \in A_2$. A similar argument works for F_2, so it remains only to show that any Fredholm module which is aligned with F_1 and F_2 is operator homotopic to one of the form described above.

So assume that F' is aligned with F_1 and F_2. Choose a partition of unity N_1, N_2 adapted to the set Δ spanned by $F_1\hat{\otimes}1$, $1\hat{\otimes}F_2$, $\rho(a_1 \otimes a_2)$, and $\rho(a)F'$ and use it to form $F = N_1(F_1\hat{\otimes}1) + N_2(1\hat{\otimes}F_2)$. We have:

$$\begin{aligned}
&\rho(a)(FF' + F'F)\rho(a')\\
&= \rho(a)(N_1(F_1\hat{\otimes}1)F' + F'N_1(F_1\hat{\otimes}1) + N_2(1\hat{\otimes}F_2)F' + F'N_2(1\hat{\otimes}F_2))\rho(a^*)\\
&\sim N_1\rho(a)((F_1\hat{\otimes}1)F' + F'(F_1\hat{\otimes}1))\rho(a^*) + N_2\rho(a)((1\hat{\otimes}F_2)F' + F'(1\hat{\otimes}F_2))\rho(a^*)
\end{aligned}$$

This expression is positive modulo compacts since F' is aligned with F_1 and F_2.

By Lemma 3.5.2, F and F' are operator homotopic. □

As mentioned in the introduction, the Kasparov product has a particularly nice interpretation when the Fredholm modules involved come from elliptic operators.

Proposition 3.5.6. *Let $S_1 \to M_1$ and $S_2 \to M_2$ be p_1- and p_2-multigraded Hermitian vector bundles over complete Riemannian manifolds M_1 and M_2, respectively. Let D_1 and D_2 be p_1- and p_2-multigraded Dirac-type operators on S_1 and S_2, respectively, and let D be the (p_1+p_2)-multigraded Dirac-type operator $D_1\hat{\otimes}1+1\hat{\otimes}D_2$ on $S_1\hat{\otimes}S_2 \to M_1 \times M_2$. Finally let H be the Hilbert space $L^2(M_1 \times M_2; S_1\hat{\otimes}S_2)$. For any normalizing function χ we have:*

- $\chi(D)(\chi(D_1)\hat{\otimes}1) + (\chi(D_1)\hat{\otimes}1)\chi(D)$ *is a positive operator on H*
- $\chi(D)(1\hat{\otimes}\chi(D_2) + (1\hat{\otimes}\chi(D_2))\chi(D)$ *is a positive operator on H*
- $\chi(D)$ *derives* $\mathbb{K}(L^2(M_1; S_1)\hat{\otimes}\mathbb{B}(L^2(M_2; S_2))$

Consequently the Fredholm module determined by $\chi(D)$ represents the Kasparov product of the K-homology classes determined by $\chi(D_1)$ and $\chi(D_2)$.

Proof. See Chapter 10 of [9]. □

To complete our discussion of Kasparov products, let us show that the unit class $\mathbf{1} \in K^0(\mathbb{C})$ behaves as a multiplicative identity for the Kasparov product.

Lemma 3.5.7. *For any separable C^*-algebra A and any class $\boldsymbol{x} \in K^{-p}(A)$ we have $\boldsymbol{x} \times \boldsymbol{1} = \boldsymbol{1} \times \boldsymbol{x} = \boldsymbol{x}$ under the identification $A \otimes \mathbb{C} \cong \mathbb{C} \otimes A \cong A$.*

Proof. Following Example 3.4.8, represent $\mathbf{1}$ by the Fredholm module $(id, \mathbb{C}, F_0)$ where $\mathbb{C}$ is the one dimensional graded Hilbert space with even part $\mathbb{C}$ and odd part 0, $id\colon \mathbb{C} \to \mathbb{C}$ is the identity map, and F_0 is the 0 operator. For any Fredholm module (ρ, H, F) over A we have that $(\rho \otimes id, H\hat{\otimes}\mathbb{C}, F\hat{\otimes}1)$ is a Fredholm module over $A \otimes \mathbb{C}$ which is aligned with F and F_0. Under the identifications $A \otimes \mathbb{C} \cong A$ and $H\hat{\otimes}\mathbb{C} \cong H$ we have that $\rho \otimes id$ corresponds to ρ and $F\hat{\otimes}1$ corresponds to F, so we have:

$$[\rho, H, F] \times \mathbf{1} = [\rho \otimes id, H\hat{\otimes}\mathbb{C}, F\hat{\otimes}1] = [\rho, H, F]$$

The identity for the product in the other order follows similarly. □

3.5.2 The Suspension Map and the Kasparov Product

We conclude this chapter by establishing a certain compatibility relation between the suspension map of the last section and the Kasparov product. This is a fairly technical result which makes heavy use of partitions of unity and the Kasparov technical theorem, but the payoff is an important geometric calculation which will be of critical importance later on.

Proposition 3.5.8. *Let A and B be C^*-algebras and let $\boldsymbol{y} \in K^{-q}(B)$ be any K-homology class. Then the following diagram commutes:*

$$\begin{tikzcd}
K^{-p-1}(C_0(-1,1)\otimes A) \arrow[r, "s"] \arrow[d, "\times \boldsymbol{y}"] & K^{-p}(A) \arrow[d, "\times \boldsymbol{y}"] \\
K^{-p-q-1}(C_0(-1,1)\otimes A\otimes B) \arrow[r, "s"] & K^{-p-q}(A\otimes B)
\end{tikzcd}$$

Proof. Pick a class $\mathbf{x} \in K^{-p-q}(C_0(-1,1)\otimes A)$ and represent it by a nondegenerate relative Fredholm module (ρ_1, H_1, F_1) for the pair $(C[-1,1]\otimes A, C\{-1,1\}\otimes A)$. Represent $\mathbf{y}$ by an ordinary Fredholm module (ρ_2, H_2, F_2). Our strategy is to build explicit Fredholm modules which represent the classes $s(\mathbf{x}\times\mathbf{y})$ and $s(\mathbf{x})\times\mathbf{y}$ and prove that these Fredholm modules are K-equivalent.

Let $X_0 \in \mathbb{B}(H_1)$ denote the image under ρ_1 of the map $x \mapsto x$ on $(-1,1)$ and let $X = \gamma\varepsilon_1 X_0$ where γ is the grading operator on H_1 and ε_1 is the first multigrading operator for $\mathbf{x}$. We need the operator X to define the suspension map.

Recall that the product $\mathbf{x}\times\mathbf{y}$ is represented by a Fredholm module (ρ, H, F) where $H = H_1\hat{\otimes}H_2$, $\rho\colon A_1\otimes A_2 \to \mathbb{B}(H)$ is the product representation, and $F = N_1(F_1\hat{\otimes}1) + N_2(1\hat{\otimes}F_2)$ where N_1, N_2 is a partition of unity adapted to F_1, F_2, ρ, and the multigrading operators. Choose such N_1 and N_2 which satisfy the additional requirement that they are adapted to X, and form the operator

$$V = X + (1-X^2)^{1/2}F$$

on H. By definition the Fredholm module (ρ, H, V) represents the K-homology class $s(\mathbf{x}\times\mathbf{y})$.

According to the Kasparov technical theorem there is a partition of unity M_1, M_2 on H which is adapted to F_1, F_2, X, N_1, N_2, and ρ, and which satisfies the additional requirements that $M_1\rho(C_0(-1,1)\otimes A\otimes B)\sim 0$ and $M_2(1-F^2)\sim 0$. Set $\widetilde{V}=M_1X+M_2F$.

Claim: The operator $\rho(a\otimes b)(V\widetilde{V}+\widetilde{V}V)\rho(a\otimes b)^*$ is positive modulo compacts. Indeed, we have:

$$\begin{aligned}
&V\widetilde{V}+\widetilde{V}V\\
&\sim 2M_1X^2+M_2(XF+FX)+M_1(1-X^2)^{1/2}(FX+XF)\\
&+M_2(F^2(1-X^2)^{1/2}+F(1-X^2)^{1/2}F)
\end{aligned}$$

Note that F commutes with the operator X_0 modulo locally compact operators and hence anticommutes with X modulo locally compact operators. It follows that:

$$\rho(a\otimes b)(V\widetilde{V}+\widetilde{V}V)\rho(a\otimes b)^*\sim\rho(a\otimes b)(2M_1X^2+2M_2(1-X^2)^{1/2})\rho(a\otimes b)^*$$

This is positive, so the claim has been proved.

By Lemma 3.5.2, the claim implies that V and $\widetilde{V}$ represent the same class in $K^{-p-q}(A\otimes B)$. We will show that the Fredholm module operator $\widetilde{V}$ gives a representative of $s(\mathbf{x})\times\mathbf{y}$.

First we construct a representative of $s(\mathbf{x})\times\mathbf{y}$ directly. By definition $s(\mathbf{x})$ is represented by the Fredholm module operator $V_1=X+(1-X^2)^{1/2}F_1$, and thus the Kasparov product of $s(\mathbf{x})$ with $\mathbf{y}$ is represented by the Fredholm module operator:

$$\widetilde{F}=P_1(V_1\hat{\otimes}1)+P_2(1\hat{\otimes}F_2)$$

where P_1,P_2 is a partition of unity adapted to V_1, F_2, and ρ. Impose the additional requirement that P_1,P_2 is adapted to N_1, N_2, M_1, and M_2.

Claim: The operator $\rho(a\otimes b)(\widetilde{F}\widetilde{V}+\widetilde{V}\widetilde{F})\rho(a\otimes b)^*$ is positive modulo compacts. We have:

$$\begin{aligned}
\widetilde{F}\widetilde{V}+\widetilde{V}\widetilde{F}&=P_1M_1((V_1\hat{\otimes}1)X+X(V_1\hat{\otimes}1))+P_1M_2((V_1\hat{\otimes}1)F+F(V_1\hat{\otimes}1))\\
&+P_2M_1((1\hat{\otimes}F_2)X+X(1\hat{\otimes}F_2))+P_2M_2((1\hat{\otimes}F_2)F+F(1\hat{\otimes}F_2))
\end{aligned}$$

We must check that each of these terms are either locally compact or positive modulo locally compact operators. The first term is given by:

$$(V_1\hat{\otimes}1)X + X(V_1\hat{\otimes}1) \sim X^2 + (1-X^2)^{1/2}(F_1X + XF_1)$$

This is positive modulo locally compact operators because X^2 is positive and $(F_1X + XF_1)$ is locally compact. The second term takes the form:

$$(V_1\hat{\otimes}1)F + F(V_1\hat{\otimes}1) \sim XF + FX + 2N_1(1-X^2)^{1/2}$$

This is positive modulo locally compact operators because $XF + FX$ is locally compact and $2N_1(1-X^2)^{1/2}$ is positive. The third term $(1\hat{\otimes}F_2)X + X(1\hat{\otimes}F_1)$ is zero because X is an odd operator on H_1 and F_2 is an odd operator on H_2. Finally the fourth term is given by:

$$F(1\hat{\otimes}F_2) + (1\hat{\otimes}F_2)F \sim 2N_2(1\hat{\otimes}F_2^2)$$

This is equal to the positive operator $2N_2$ modulo locally compact operators.

This completes the proof. □

We are now ready to prove one of the main results of this chapter. Let M be the interior of a complete Riemannian manifold $\overline{M}$ with boundary ∂M. Let $S \to \overline{M}$ be a smooth p-multigraded vector bundle and let D be a $(p+1)$-multigraded Dirac-type operator on S. Finally let $\overline{U} \cong [-1,1) \times \partial M$ be a collaring neighborhood of ∂M and let $U = (-1,1) \times \partial M$ be its interior. Assume that the bundle S splits over U as $S = S_{\partial M}\hat{\otimes}S_{(-1,1)}$ where $S_{\partial M}$ is a p-multigraded vector bundle over ∂M and $S(-1,1)$ is the spinor bundle over $(-1,1)$. Assume further that D splits as $D_{\partial M}\hat{\otimes}1 + 1\hat{\otimes}D_{(-1,1)}$ where $D_{\partial M}$ is a p-multigraded Dirac-type operator on $S_{\partial M}$ and $D_{(-1,1)}$ is the spinor Dirac operator. In this setting we have the following result:

Theorem 3.5.9. *Let $\partial\colon K^{-p-1}(M) \to K^{-p}(\partial M)$ denote the boundary map in K-homology associated to the short exact sequence of C^*-algebras*

$$0 \to C_0(M) \to C_0(\overline{M}) \to C_0(\partial M) \to 0$$

Then $\partial[D] = [D_{\partial M}]$ *in* $K^{-p}(\partial M)$.

Proof. By the naturality of the boundary map for $*$-homomorphisms (Lemma 3.3.26) we can reduce to the case where $\overline{M} = \overline{U}$ and $M = U$. Over U the K-homology class of D satisfies $[D] = [D_{(-1,1)}] \times [D_{\partial M}]$ by Proposition 3.5.6 and the boundary map ∂ is precisely the suspension map. Since $[D_{(-1,1)}]$ is the Dirac class, we have:

$$\partial[D] = s([D_{(-1,1)}] \times [D_{\partial M}]) = s([D_{(-1,1)}]) \times [D_{\partial M}] = [D_{\partial M}]$$

as desired. $\square$

Chapter 4

Index Theory and Large-Scale Geometry

4.1 Introduction

In the last chapter we used index theory for elliptic differential operators as motivation for the definition of the K-homology groups of a locally compact Hausdorff space. If P is a one-point space then $K_0(P) = K^0(\mathbb{C}) \cong \mathbb{Z}$, and according to a computation in the previous chapter this isomorphism sends the K-homology class of a nondegenerate Fredholm module (ρ, H, F) over P to the Fredholm index of F. It follows that if D is a graded elliptic operator over a smooth compact manifold M then the map $K_0(M) \to \mathbb{Z}$ induced by the obvious surjection $M \to P$ sends the K-homology class of D to the Fredholm index of D. The original Atiyah-Singer index theorem can be conveniently expressed as a topological statement about this map.

However, there are a variety of reasons why this is inadequate for index problems on non-compact manifolds. Elliptic operators on non-compact manifolds often fail to be Fredholm, so it is not immediately obvious how to even pose an interesting index problem. This is a reflection of the fact that K-homology is not functorial for arbitrary continuous maps between non-compact spaces, and in particular the crushing map $M \to P$ described above does not naturally induce a map $K_0(M) \to \mathbb{Z}$. Roe realized that this difficulty can be resolved by introducing a new notion of index which takes values in a group which reflects the large-scale geometry of M instead of the integers. Since any compact manifold has the large-scale geometry of a point, Roe's construction recovers the classical index map above if M happens to be compact.

Our main motivation for developing index theory for non-compact manifolds is based on the observation that non-compact manifolds appear in the statement of the partitioned manifold index theorem. But there is good reason to investigate non-compact spaces even if one is only interested in invariants of compact manifolds. Suppose M is a smooth compact manifold, G is a discrete group, and $\widetilde{M} \to M$ is a G-cover of M. Any elliptic operator D on M lifts to a G-invariant elliptic operator $\widetilde{D}$ on $\widetilde{M}$, and one can contemplate an "equivariant index" of $\widetilde{D}$ which accounts for the fact that the kernel and cokernel of $\widetilde{D}$ carry representations of G. There is a considerable amount of literuature surrounding these equivariant indices; they are crucially involved in Gromov and Lawson's proof that the n-torus admits no metric of positive scalar curvature, and they are at the heart of analytic approaches to the Novikov conjecture. But to even define them one must come to terms with the fact that the universal cover of a compact manifold need not be compact.

In this chapter we will begin by developing the language of coarse geometry, invented by Roe to capture the large-scale geometry of non-compact spaces. We will then develop coarse counterparts of the C*-algebras used to define Paschke's model of K-homology and fit their K-theory groups into a long exact sequence with K-homology. Along the way we will develop a relative counterpart of the theory in anticipation of the Mayer-Vietoris sequences that we will build in the next chapter, and throughout we will work equivariantly with respect to a free and proper group action so that our results will apply to the equivariant indices described above.

Many of the ideas in this chapter have antecedents in [19], but the language of coarse geometry and coarse C*-algebras was not developed until later. Much of our treatment of the subject is adapted from chapter 6 of [9] and the monograph [18].

4.2 Large-Scale Geometry and Analysis

We begin by introducing Roe's notion of a *coarse structure*. The theory of coarse structures axiomatizes geometric properties which are only detectable on very large scales. Roe's axioms are flexible enough to allow a variety of different notions of "large-scale", but for our purposes it suffices to restrict our attention to the *metric coarse structure* associated to a metric space and use the metric to distinguish between large and small scales. For instance, we will see that any bounded metric space has the same metric coarse structure as a point.

If X is a metric space and $C_0(X)$ is represented on a Hilbert space H then the metric coarse structure on X determines a special class of bounded operators on H called the *controlled* operators. The controlled condition is a functional-analytic abstraction of the finite propagation speed property for differential operators described in the previous chapter; informally, an operator is controlled if it doesn't enlarge the support of a function in $C_0(X)$ too much. The main result of this section is that the set of all controlled operators forms a $*$-subalgebra of $\mathbb{B}(H)$, and the remainder of the chapter will involve using this $*$-subalgebra to build C*-algebras which carry information about the metric coarse structure of X.

4.2.1 The Metric Coarse Structure

The guiding principle in coarse geometry is that two spaces have the same coarse structure if they can only be distinguished below some finite scale. In metric spaces, this can be captured with the aid of the following definition:

Definition 4.2.1. *Let X be a metric space and let S be any set. Declare that two maps $\alpha_1, \alpha_2\colon S \to X$ are* close *if $\{d(\alpha_1(s), \alpha_2(s))\colon\ s \in S\}$ is a bounded set of real numbers.*

Note that if X is bounded then any pair of maps are close, so this definition is only interesting for unbounded (and hence non-compact) X. Closeness is an equivalence relation on the set of all maps $S \to X$, and it has the following basic properties:

Lemma 4.2.2. *Let X be a metric space.*

- *If $\varphi\colon S' \to S$ is a map between sets and $\alpha_1, \alpha_2\colon S \to X$ are close then $\alpha_1 \circ \varphi$ and $\alpha_2 \circ \varphi$ are close.*

- *If S and S' are sets and $\alpha_1, \alpha_2\colon S \cup S' \to X$ restrict to close maps on S and S' then α_1 and α_2 are close.*

- *Any two constant maps from any set to X are close.*

Roe defines an abstract coarse structure on X to be a choice of equivalence relation on the set of all maps $S \to X$ (for each set S) which is compatible with compositions and unions and for which all constant maps are equivalent. Thus Lemma 4.2.2, whose proof is straightforward, shows that the closeness relation defines a coarse structure on a metric space. The coarse structure determined by the closeness relation is called the *metric coarse structure*; since it is the only coarse structure which we will use, we will not comment any further on the abstract theory of coarse spaces.

Another way to capture the metric coarse structure is to specify the subsets of $X \times X$ on which the metric restricts to a bounded function. This will be particularly important when we begin to do functional analysis on coarse spaces.

Definition 4.2.3. *Let X be a metric space. A subset $E \subseteq X \times X$ is* controlled *if the projection maps $\pi_1, \pi_2\colon X \times X \to X$ restrict to close maps on E.*

Thus the controlled sets are precisely the sets which lie in a uniformly bounded neighborhood of the diagonal in $X \times X$. Note that a set $B \subseteq X$ is bounded if and only if $B \times B$ is controlled, and a family of subsets $\{B_\alpha\}$ is uniformly bounded if and only if $\bigcup_\alpha B_\alpha \times B_\alpha$ is controlled. Thus boundedness and uniform boundedness are “coarse notions”.

It is useful to know that the controlled sets determine the coarse structure of X:

Lemma 4.2.4. *Two maps $\alpha_1, \alpha_2\colon S \to X$ are close if and only if the image of $(\alpha_1, \alpha_2)\colon S \to X \times X$ is controlled.*

Proof. Let E denote the image of (α_1, α_2). If E is controlled then α_1 and α_2 are close by the first property in Lemma 4.2.2. To prove the converse, choose a set-theoretic splitting $r\colon E \to S$ of the map $(\alpha_1, \alpha_2)\colon S \to E$. We have that $\alpha_1 \circ r$ and

$\alpha_2 \circ r$ are close by the first property in Lemma 4.2.2 again, and thus π_1 is close to π_2 on E since $\pi_i = \pi_i \circ (\alpha_1, \alpha_2) \circ r = \alpha_i \circ r$. □

There is a general notion of morphism between coarse spaces. Coarse morphisms are equivalence classes of maps which respect the coarse structure of a space in the following sense.

Definition 4.2.5. *Let X and Y be metric spaces. A map $f\colon X \to Y$ is* coarse *if:*

- *Whenever $\alpha_1, \alpha_2\colon S \to X$ are close, $f \circ \alpha_1$ and $f \circ \alpha_2$ are close.*
- *Whenever B is a bounded subset of Y, $f^{-1}(B)$ is a bounded subset of X.*

A coarse morphism *from X to Y is a closeness class of a coarse map.*

The following two examples capture the spirit of this definition:

Example 4.2.6. *The map $f\colon \mathbb{R} \to \mathbb{Z}$ which sends a real number to its integer part is a coarse map.*

Example 4.2.7. *If X is an unbounded metric space then the projection maps $\pi_1, \pi_2\colon X \times X \to X$ are not coarse.*

Definition 4.2.8. *The* metric coarse category *is the category whose objects are metric spaces and whose morphisms are coarse morphisms.*

We call attention to the natural notion of equivalence in the coarse category:

Definition 4.2.9. *A coarse map $f\colon X \to Y$ between metric spaces is a* coarse equivalence *if there is a map $g\colon Y \to X$ such that $g \circ f$ and $f \circ g$ are close to the identity maps on X and Y, respectively.*

Example 4.2.10. *The closeness class of the map $f\colon \mathbb{R} \to \mathbb{Z}$ of Example 4.2.6 is a coarse equivalence whose inverse morphism is the class of the inclusion map $\mathbb{Z} \to \mathbb{R}$.*

This example shows that the coarse structure of a space remembers nothing about its local structure. Still, coarse geometry can be used to capture global topological invariants, at least for a suitable class of metric spaces:

Definition 4.2.11. *A metric space is* proper *if its compact subspaces are precisely is closed and bounded subsets.*

Every proper metric space X is second countable and locally compact; in particular $C_0(X)$ is separable. By the Hopf-Ronow theorem every complete Riemannian manifold is a proper metric space, so for the applications to index theory that we have in mind little is lost by restricting our attention to proper metric spaces. We shall do so from now on.

4.2.2 Controlled Operators

We will now tie analysis and coarse geometry together by defining a preferred class of Hilbert space operators associated to the coarse structure of a proper metric space X. From now on all Hilbert space operators are assumed to be bounded unless otherwise stated.

Definition 4.2.12. *Let X and Y be proper metric spaces equipped with representations $\rho_X \colon C_0(X) \to \mathbb{B}(H_X)$ and $\rho_Y \colon C_0(Y) \to \mathbb{B}(H_Y)$ on separable Hilbert spaces.*

- *The* support *of a vector $\boldsymbol{v} \in H_X$ is the set $Supp(\boldsymbol{v})$ of all $x \in X$ such that for every neighborhood $\mathcal{U}$ of x there exists $f \in C_0(\mathcal{U})$ such that $\rho_X(f)v \neq 0$.*
- *The* support *of an operator $T \colon H_X \to H_Y$ is the set $Supp(T)$ of all points $(y, x) \in Y \times X$ with the following property: for every open neighborhood $\mathcal{V} \times \mathcal{U} \subseteq Y \times X$ of (y, x) there exist functions $f_1 \in C_0(\mathcal{U})$ and $f_2 \in C_0(\mathcal{V})$ such that $\rho_Y(f_2) T \rho_X(f_1) \neq 0$.*
- *An operator $T \colon H_X \to H_Y$ is* properly supported *if the slices*

$$S_x = \{y \in Y \colon (y, x) \in Supp(T)\}$$
$$S^y = \{x \in X \colon (y, x) \in Supp(T)\}$$

  are closed sets.

Definition 4.2.13. *Let X be as in the previous definition. An operator $T \in \mathbb{B}(H_X)$ is controlled if its support is a controlled subset of $X \times X$.*

Thus an operator is controlled if it is supported in a uniformly bounded neighborhood of the diagonal in $X \times X$. Crucial examples of such operators arise are provided by the functional calculus for Dirac-type operators on complete Riemannian manifolds. The key result about controlled operators is as follows:

Theorem 4.2.14. *The set of all controlled operators for X is a unital $*$-subalgebra of $\mathbb{B}(H_X)$.*

We need to set up some machinery before proving this theorem. The main challenge is proving that the product of two controlled operators is controlled, and to do this we need some tools for manipulating supports.

Definition 4.2.15. *Let X, Y, and Z be sets and let $A \subseteq Z \times Y$, $B \subseteq Y \times X$ be subsets. Then AB is defined to be*

$$\{(z,x) \in Z \times X\colon\ (z,y) \in A \text{ and } (y,x) \in B \text{ for some } y \in Y\}$$

Remark 4.2.16. *The "product" of a subset of $Y \times X$ with a subset of X is defined similarly.*

We can characterize the support of an operator $T \in \mathbb{B}(H_X)$ according to the property $\mathrm{Supp}(T\mathbf{v}) \subseteq \mathrm{Supp}(T)\mathrm{Supp}(\mathbf{v})$. To achieve this we need a lemma:

Lemma 4.2.17. *Let $\boldsymbol{v} \in H_X$ and let $f \in C_0(X)$. If $f|_{Supp(\boldsymbol{v})} = 0$ then $\rho_X(f)\boldsymbol{v} = 0$.*

Proof. Suppose f has compact support. The complement of $\mathrm{Supp}(\mathbf{v})$ is the union of all open sets $\mathcal{U}$ with the property that $\rho_X(g)\mathbf{v} = 0$ for all $g \in C_0(\mathcal{U})$, so this collection of open sets covers $\mathrm{Supp}(f)$ since f vanishes on $\mathrm{Supp}(\mathbf{v})$. By compactness there is a finite subcover $\mathcal{U}_1, \ldots, \mathcal{U}_n$ and a partition of unity $h_1, \ldots, h_n$ subordinate to the subcover. But $\rho_X(f)\mathbf{v} = \sum_{i=1}^n \rho_X(fh_i)\mathbf{v}$ and $\rho_X(fh_i)\mathbf{v} = 0$ (since $fh_i \in C_0(\mathcal{U}_i)$ where $\mathcal{U}_i$ is in the complement of $\mathrm{Supp}(\mathbf{v})$), so $\rho_X(f)\mathbf{v} = 0$. Thus we have proven the lemma for all f in $C_c(X - \mathrm{Supp}(\mathbf{v}))$. But $C_c(X - \mathrm{Supp}(\mathbf{v}))$ is dense in $C_0(X - \mathrm{Supp}(\mathbf{v}))$, so we are done. $\square$

Proposition 4.2.18. *Let $T\colon H_X \to H_Y$ be a properly supported operator. Then Supp($T\boldsymbol{v}$) $\subseteq$ Supp(T)Supp($\boldsymbol{v}$) for any compactly supported vector $\boldsymbol{v} \in H_X$, and Supp(T) is the smallest closed subset of $Y \times X$ with this property.*

Proof. Suppose y is not in $\mathrm{Supp}(T)\mathrm{Supp}(\mathbf{v})$. Then the closed sets $A = \mathrm{Supp}(\mathbf{v})$ and $B = \{x \in X\colon\ (y,x) \in \mathrm{Supp}(T)\}$ are disjoint, so there exists $h \in C_0(X)$ such that $h|_A = 1$ and $h|_B = 0$. We shall construct an open set $\mathcal{V} \subseteq Y$ such that $y \in \mathcal{V}$ and $\rho_Y(f)T\rho_X(h) = 0$ for every $f \in C_0(\mathcal{V})$.

For any $x \in A$ there is an open set $\mathcal{V}_x \times \mathcal{U}_x \subseteq Y \times X$ such that $(y,x) \in \mathcal{V}_x \times \mathcal{U}_x$ and $\rho_Y(f)T\rho_X(g) = 0$ for every $f \in C_0(\mathcal{V}_x)$, $g \in C_0(\mathcal{U}_x)$. Clearly $A \subseteq \bigcup_{x\in A} \mathcal{U}_x$, so since A is compact there is a finite subcover $\mathcal{U}_{x_1}, \ldots, \mathcal{U}_{x_n}$. Let $\mathcal{V}_{x_1}, \ldots, \mathcal{V}_{x_n}$ be the corresponding open sets in Y and let $\mathcal{V}$ denote $\bigcup_{i=1}^n \mathcal{V}_{x_i}$. Clearly $y \in \mathcal{V}$ and $\rho_Y(f)T\rho_X(g_i) = 0$ for every $f \in C_0(\mathcal{V})$ and every $g_i \in C_0(\mathcal{U}_{x_i})$. We can assume that h was chosen so that it is supported in $\mathcal{U}_{x_1} \cup \ldots \cup \mathcal{U}_{x_n}$, so by a straightforward partition of unity argument we have $\rho_Y(f)T\rho_X(h) = 0$ for every $f \in C_0(\mathcal{V})$, as desired.

Now write $\rho_Y(f)T\mathbf{v} = \rho_Y(f)T\rho_X(h)\mathbf{v} + \rho_Y(f)T(1-\rho_X(h))\mathbf{v}$ where $f \in C_0(\mathcal{V})$. The first term vanishes by our choice of $\mathcal{V}$, and the second term vanishes by Lemma 4.2.17. Thus $\rho_Y(f)T\mathbf{v} = 0$ for every $f \in C_0(\mathcal{V})$, and hence y is not in $\mathrm{Supp}(T\mathbf{v})$.

It remains only to show that if A is a closed subset of $Y \times X$ such that $\mathrm{Supp}(T\mathbf{v}) \subseteq A \cdot \mathrm{Supp}(\mathbf{v})$ for every compactly supported $\mathbf{v}$ then $\mathrm{Supp}(T) \subseteq A$. Let A be such a set and let (y,x) be a point in the complement of A. Since A is closed there is a neighborhood $\mathcal{U}$ of x such that (y,x') is in the complement of A for every $x' \in \mathcal{U}$. Take any vector $\mathbf{v} \in H$ and let f be a continuous compactly supported function on $\mathcal{U}$; we have that $\rho_X(f)\mathbf{v}$ has compact support in $\mathcal{U}$. Thus (y,x') is in the complement of A for every $x \in \mathrm{Supp}(\rho_X(f)\mathbf{v})$, and since the complement of $A \cdot \mathrm{Supp}(\mathbf{v})$ is contained in the complement of $\mathrm{Supp}(T\mathbf{v})$ there exists a neighborhood $\mathcal{V}$ of y such that $\rho_Y(g)T\rho_X(f)\mathbf{v} = 0$ for every $g \in C_0(\mathcal{V})$. Since $\mathbf{v}$ was arbitrary we have that $\rho_Y(g)T\rho_X(f) = 0$ for every $g \in C_0(\mathcal{V})$ and every $f \in C_0(\mathcal{U})$. Thus (y,x) is in the complement of $\mathrm{Supp}(T)$ from which it follows that $\mathrm{Supp}(T) \subseteq A$, as desired. $\square$

We are now ready to prove Theorem 4.2.14.

Proof of Theorem 4.2.14. Most of the proof is very straightforward. The identity is a controlled operator because its support is precisely the diagonal in $X \times X$. The support of T^* is the image of $\mathrm{Supp}(T)$ under the map $(x_1,x_2) \mapsto (x_2,x_1)$, so T^* is controlled if T is. We have $\mathrm{Supp}(S+T) \subseteq \mathrm{Supp}(S) \cup \mathrm{Supp}(T)$, so $S+T$ is controlled if S and T are.

Let us prove that ST is controlled if S and T are. Any controlled operator is properly supported, so by the previous proposition $\mathrm{Supp}(ST)$ is the smallest closed set with the property that

$$\mathrm{Supp}(ST\mathbf{v}) \subseteq \mathrm{Supp}(ST)\mathrm{Supp}(\mathbf{v})$$

for every compactly supported vector $\mathbf{v}$. Applying the lemma again,

$$\mathrm{Supp}(T\mathbf{v}) \subseteq \mathrm{Supp}(T)\mathrm{Supp}(\mathbf{v}) = \pi_1(\mathrm{Supp}(T) \cap \pi_2^{-1}(\mathrm{Supp}(\mathbf{v})))$$

where $\pi_1, \pi_2 \colon X \times X \to X$ are the projection maps, and thus $\mathrm{Supp}(T\mathbf{v})$ is within a bounded distance of the set $\pi_2(\mathrm{Supp}(T)) \cap \mathrm{Supp}(\mathbf{v})$ since T is controlled. Thus $T\mathbf{v}$ is compactly supported and hence

$$\mathrm{Supp}(ST\mathbf{v}) \subseteq \mathrm{Supp}(S)\mathrm{Supp}(T)\mathrm{Supp}(\mathbf{v})$$

So $\mathrm{Supp}(ST) \subseteq \mathrm{Supp}(S)\mathrm{Supp}(T)$, and this latter set is controlled by the triangle inequality.

This shows that the space of controlled operators is a $*$-subalgebra of $\mathbb{B}(H_X)$, as desired. □

4.3 C*-algebras and Coarse Geometry

Now that we have organized the set of all controlled operators associated to a space X into a $*$-subalgebra of $\mathbb{B}(H_X)$, it is natural to try to introduce C*-algebras generated by controlled operators. Indeed, we shall build a coarse analogue of the short exact sequence

$$0 \to \mathfrak{C}^*(X) \to \mathfrak{D}^*(X) \to \mathfrak{D}^*(X)/\mathfrak{C}^*(X) \to 0$$

involving the dual algebra (Definition 3.3.1) and the locally compact algebra (Definition 3.3.2). As described in the introduction, it is desirable to account for a free and proper group action on X. It turns out that this does not have a substantial effect on the proofs of many of the main results, and the reader is invited to focus on the case where G is trivial for much of what follows.

So let X be a proper metric space equipped with a free and proper action of a countable discrete group G of isometries of X. We will need to consider representations of $C_0(X)$ which are compatible with the G-action.

Definition 4.3.1. *Let H be a Hilbert space equipped with a representation*

$$\rho\colon C_0(X) \to \mathbb{B}(H)$$

and a unitary representation

$$U\colon G \to \mathbb{B}(H)$$

*We say that the the triple (H, U, ρ) is a G-*equivariant *X-*module *or simply a* (X, G)-module *if $U_\gamma \rho(f) = \rho(\gamma^* f) U_\gamma$ for every $\gamma \in G$, $f \in C_0(X)$.*

Equivariant modules can be obtained by setting $H = L^2(X, \mu)$ where μ is a G-invariant Borel measure on X. $C_0(X)$ has a representation on $\rho\colon C_0(X) \to \mathbb{B}(H)$ by multiplication operators, and G has a representation $U\colon G \to \mathbb{B}(H)$ by translation:

$$U_\gamma \varphi = \varphi \circ \gamma^*$$

In fact, every equivariant module is the direct sum of equivariant modules of this form by the spectral theorem.

As usual if H carries a unitary representation U of G then we say that an operator $T \in \mathbb{B}(H)$ is G-equivariant if $U_\gamma T U_\gamma^* = T$ for every $\gamma \in G$. The space of G-equivariant controlled operators in $\mathbb{B}(H)$ is a $*$-subalgebra, and it is this $*$-subalgebra which we will use in the discussion ahead.

For the remainder of this section G is a fixed countable discrete group, all spaces are proper metric spaces on which G acts freely and properly by isometries, and all representations are G-equivariant.

4.3.1 The Coarse Algebra

We begin with the coarse counterpart of the C*-algebra $\mathfrak{C}^*(X)$. Let $\rho\colon C_0(X) \to \mathbb{B}(H_X)$ be a representation, and recall that $\mathfrak{C}^*_\rho(X)$ is the C*-algebra of locally compact operators, i.e. operators $T \in \mathbb{B}(H_X)$ such that $T\rho(f) \sim \rho(f)T \sim 0$ for every $f \in C_0(X)$. As in our discussion of K-homology it will often be important to use ample representations.

Definition 4.3.2. *The (equivariant)* coarse C* algebra *of a space X is the closure in $\mathbb{B}(H_X)$ of the $*$-algebra of all G-equivariant locally compact controlled operators for X. It is denoted by $C^*_G(X)$.*

If G is the trivial group then we will suppress it from the notation and write $C^*(X)$. As with the locally compact algebra $C^*_G(X)$ is generally not independent of the representation used to define it, but as we shall see shortly the ambiguity disappears at the level of K-theory.

Definition 4.3.3. *An isometry $V\colon H_X \to H_Y$ coarsely covers a coarse map*

$$\phi\colon X \to Y$$

if π_1 and $\phi \circ \pi_2$ are close as maps Supp$(V) \subseteq Y \times X \to Y$. A coarse covering isometry V is G-equivariant if in addition $VU^X_g = U^Y_g V$.

As with our earlier notion of covering isometry we will use equivariant coarse covering isometries to implement functoriality of the assignment $X \mapsto K_p(C^*_G(X))$. Note that an isometry which coarsely covers ϕ also coarsely covers any map which is close to ϕ, so it is reasonable to hope that $K_p(C^*_G(X))$ is functorial for coarse morphisms. This is indeed the case, but first we must construct coarse covering isometries.

Lemma 4.3.4. *Any proper metric space Y has a countable uniformly bounded cover by disjoint Borel sets with non-empty interiors.*

Proof. Choose any uniformly bounded open cover $\{\mathcal{U}_n\}$ of Y; since Y is a proper metric space, $\{\overline{\mathcal{U}}_n\}$ is also uniformly bounded. Let $\mathcal{V}_1 = \mathcal{U}_1$ and recursively define

$$\mathcal{V}_n = \mathcal{U}_n - (\mathcal{U}_1 \cup \ldots \cup \mathcal{U}_{n-1})$$

The set $\{\mathcal{V}_n\}$ is a uniformly bounded cover of Y by disjoint Borel sets. Let $\{\mathcal{V}_{n_k}\}$ be the subcollection of $\{\mathcal{V}_n\}$ consisting of precisely those sets with nonempty interior.

Note that if $\mathcal{V}_n$ is a set with empty interior and $y \in \mathcal{V}_n$ then y is a limit point of $\mathcal{U}_m$ for some $m < n$ and hence y is in the closure of $\mathcal{V}_m$; by induction we have $y \in \overline{\mathcal{V}}_{n_k}$ for some k and hence $\{\overline{\mathcal{V}_{n_k}}\}$ covers Y. Now let $\mathcal{W}_1 = \overline{\mathcal{V}}_{n_1}$ and recursively define

$$\mathcal{W}_k = \overline{\mathcal{V}}_{n_k} - (\overline{\mathcal{V}}_{n_1} \cup \ldots \cup \overline{\mathcal{V}}_{n_{k-1}})$$

The collection $\{\mathcal{W}_k\}$ consists of disjoint Borel sets which cover Y, it is uniformly bounded since $\mathcal{W}_k \subseteq \overline{\mathcal{U}}_{n_k}$, and each $\mathcal{W}_k$ has nonempty interior since the V_{n_k}'s are disjoint. □

Proposition 4.3.5. *Let $\rho_X \colon C_0(X) \to \mathbb{B}(H_X)$ be a nondegenerate G-equivariant representation and let $\rho_Y \colon C_0(Y) \to \mathbb{B}(H_Y)$ be an ample G-equivariant representation. Then every G-equivariant coarse map $\phi \colon X \to Y$ is coarsely covered by a G-equivariant isometry $H_X \to H_Y$.*

Proof. Let $\{Y_n\}$ be a countable cover of Y by disjoint uniformly bounded G-invariant Borel sets with non-empty interior as in the previous lemma. The representation $\rho_Y \colon C_0(Y) \to \mathbb{B}(H_Y)$ extends uniquely to the algebra of Borel functions on Y; let Q_n denote the operator corresponding to the characteristic function of Y_n. Q_n is a projection, and the image of Q_n is orthogonal to the image of Q_m whenever $n \neq m$ since Y_n is disjoint from Y_m. Since Y_n has nonempty interior there exists a nonzero function $f \in C_0(Y)$ supported in Y_n, and for such f we have $\rho_Y(f)Q_n = \rho_Y(f)$. Since ρ_Y is ample, this shows that Q_n has infinite dimensional range.

Let P_n denote the operator on H_X corresponding to the characteristic function of $\phi^{-1}(Y_n)$ via the Borel functional calculus, as above. As before the P_n's are projections with pairwise orthogonal ranges and thus H_X decomposes as the orthogonal direct sum of the subspaces $P_n H_X$ (since ρ_X is nondegenerate).

Now choose a isometries $V_n \colon P_n H_X \to Q_n H_Y$ and lift them to the partial isometries $V_n P_n \colon H_X \to H_Y$. We shall show that the series $\sum_{n=1}^{\infty} V_n P_n$ converges strongly to an isometry $V \colon H_X \to H_Y$ which coarsely covers ϕ. Relative to the orthogonal decomposition $H_X = \bigoplus P_n H_X$, the series above is just $\bigoplus V_n \colon H_X \to H_Y$ and this series converges strongly since $\|V_n\| = 1$ for each n. To see that the limiting operator V is an isometry, observe that $V^*V = \bigoplus V_n^* V_n = I$.

We must show that V coarsely covers ϕ. Note that $\mathrm{Supp}(V)$ is the union of the

$\mathrm{Supp}(V_n)$'s, so we begin by characterizing the support of V_n. Since $V_n = Q_n V_n P_n$ we have that $(y, x) \in \mathrm{Supp}(V_n)$ if and only if $y \in Y_n$ and $x \in \phi^{-1}(Y_n)$, and thus $d(y, \phi(x)) \leq C$ where C is an upper bound for the diameter of each Y_n. This shows that the restrictions of π_1 and $\phi \circ \pi_2$ to $\mathrm{Supp}(V)$ have pointwise distance bounded by C, and by definition this says that V coarsely covers ϕ.

Finally, we must check that V is G-equivariant. Note that Q_n commutes with U_γ^Y for each $\gamma \in G$ since Y_n is G-invariant, and similarly U_γ^X commutes with P_n since ϕ is G-equivariant. It follows that each V_n, and hence V, is G-equivariant. $\square$

We now show that coarse covering isometries are compatible with coarse C* algebras and their K-theory.

Lemma 4.3.6. *If $V\colon H_X \to H_Y$ is an equivariant coarse covering isometry for an equivariant coarse map $\phi\colon X \to Y$ then Ad_V maps $C^*(X)$ into $C^*(Y)$.*

Proof. Let T be a locally compact G-invariant controlled operator; we will show that $\mathrm{Ad}_V(T)$ is also G-invariant, locally compact, and controlled.

G-invariance is a straightforward calculation:

$$\begin{aligned} U_\gamma^Y \mathrm{Ad}_V(T)(U_\gamma^Y)^* &= U_\gamma^Y V T V^* (U_\gamma^Y)^* \\ &= U_\gamma^Y V T (U_\gamma^Y V)^* \quad = V U_\gamma^X T (U_\gamma^X)^* V^* = V T V^* = \mathrm{Ad}_V(T) \end{aligned}$$

To see that $\mathrm{Ad}_V(T)$ is controlled, let $S \subseteq Y \times X \times X \times Y$ denote the set of all 4-tuples (y, x, x', y') such that $(y, x) \in \mathrm{Supp}(V)$, $(x, x') \in \mathrm{Supp}(T)$, and $(x', y') \in \mathrm{Supp}(V^*)$. We have $\mathrm{Supp}(\mathrm{Ad}_V(T)) \subseteq (\pi_1 \times \pi_4)(S)$, so it suffices to show that π_1 is close to π_4. Observe that π_1 is close to $\phi \circ \pi_2$ since V covers ϕ, and $\phi \circ \pi_2$ is close to $\phi \circ \pi_3$ since π_2 is close to π_3 (T is controlled) and ϕ is coarse. Using the fact that V covers ϕ again we see that $\phi \circ \pi_3$ is close to π_4, so $\mathrm{Ad}_V(T)$ is controlled.

To show that $\mathrm{Ad}_V(T)$ is locally compact, take $f \in C_c(Y)$ and note that $\rho_Y(f)V$ covers ϕ since V does. The set $\pi_2(\mathrm{Supp}(\rho_Y(f)V))$ is bounded because $\pi_1(\mathrm{Supp}(\rho_Y(f)V))$ is bounded, π_1 is close to $\phi \circ \pi_2$, and ϕ is coarse. Thus there exists $g \in C_c(X)$ such that $\rho_Y(f)V = \rho_Y(f)V\rho_X(g)$, yielding $\rho_Y(f)\mathrm{Ad}_V(T) = (\rho_Y(f)V)(\rho_X(g)T)V^*$. This is compact since $\rho_X(g)T$ is compact. A similar argument shows that $\mathrm{Ad}_V(T)g$ is compact for $g \in C_c(X)$, so we are done since $C_c(\cdot)$ is dense in $C_0(\cdot)$. $\square$

As with covering isometries for dual algebras, the map on K-theory induced by Ad_V is independent of the coarse covering isometry chosen:

Lemma 4.3.7. *If V_1 and V_2 are two equivariant coarse covering isometries for the same equivariant coarse map ϕ then Ad_{V_1} and Ad_{V_2} induce the same map $K_p(C^*_G(X)) \to K_p(C^*_G(Y))$.*

Proof. According to Lemma 2.3.16 it suffices to show that $V_i V_j^* \in C^*_G(Y)$ for each $i, j \in \{0, 1\}$. $V_i V_j^*$ is clearly G-equivariant, so it suffices to show that it is controlled. Let $E \subseteq Y \times X \times Y$ denote the set of all triples (y, x, y') such that $(y, x) \in \mathrm{Supp}(V_i)$ and $(y', x) \in \mathrm{Supp}(V_j)$ and let π_1, π_2, π_3 denote the restrictions to E of the three projection maps on $Y \times X \times Y$. The support of $V_i V_J^*$ is precisely the image of $\pi_1 \times \pi_3 \colon E \to Y \times Y$, so it suffices to show that π_1 and π_3 are close. But both π_1 and π_3 are close to $\phi \circ \pi_2 \colon E \to Y$ since V_1 and V_2 both cover ϕ, so in particular they are close to each other. □

Thus if $\phi \colon X \to Y$ is an equivariant coarse map between proper G-spaces we may unambiguously define $\phi_* \colon K_p(C^*_G(X)) \to K_p(C^*_G(Y))$ to be $(\mathrm{Ad}_V)_*$ where V is any equivariant isometry which coarsely covers ϕ. Note that any such V also coarsely covers any equivariant coarse map which is close to ϕ, so $(\mathrm{Ad}_V)_*$ depends only on the closeness class of ϕ. Thus we define a *coarse G-morphism* to be a closeness class of G-equivariant coarse maps. Given two ample (X, G)-modules H_X and H'_X, the previous results applied to the identity map $X \to X$ verify our earlier claim that $K_p(C^*_G(X))$ is independent of the ample representation used to define it.

In summary:

Proposition 4.3.8. *The assignment*

$$\{\phi \colon X \to Y\} \mapsto \{\phi_* \colon K_p(C^*_G(X)) \to K_p(C^*_G(Y))\}$$

is a covariant functor from the category of proper G-spaces with equivariant coarse morphisms to the category of abelian groups.

There is also a notion of a relative coarse C*-algebra. Given a subspace $Y \subseteq X$, we build an ideal $C^*(Y \subseteq X)$ in $C^*(X)$ by considering only those operators whose support is close to Y in an appropriate sense.

Definition 4.3.9. *An operator $T \in \mathbb{B}(H_X)$ is said to be* supported near Y *if there exists R such that Supp$(T) \subseteq B_R(Y) \times B_R(Y)$ where $B_R(Y)$ is the set of all points of X which lie within distance R of some point of Y.*

The next result shows that the set of all controlled operators which are supported near Y is a $*$-ideal in the set of all controlled operators.

Lemma 4.3.10. *Let $T \in \mathbb{B}(H_X)$ be a controlled operator and let $S \in \mathbb{B}(H_X)$ be a controlled operator which is supported near Y. Then TS and ST are controlled operators which are supported near Y.*

Proof. We have already shown that TS and ST are controlled operators, so we just need to check that they are supported near Y. Let $E \subseteq X \times X \times X$ denote the set of all triples (x_1, x_2, x_3) such that $(x_1, x_2) \in \mathrm{Supp}(S)$ and $(x_2, x_3) \in \mathrm{Supp}(T)$. We have that $\mathrm{Supp}(ST) \subseteq \mathrm{Supp}(S)\mathrm{Supp}(T) = (\pi_1 \times \pi_3)(E)$ where $\pi_1, \pi_2, \pi_3 \colon X \times X \times X \to X$ are the projection maps. π_1 is close to π_3 since both maps are close to π_2 (S and T are controlled), so $d(\pi_1(x), \pi_3(x)) \leq C$ for some constant C. The image of π_1 is contained in $B_R(Y)$ for some R since S is supported near Y, so by the triangle inequality the image of π_3 is contained in $B_{R+C}(Y)$. Thus $\mathrm{Supp}(ST)$ is contained in $B_{R+C}(Y) \times B_{R+C}(Y)$ from which it follows that ST is supported near Y. The proof that TS is supported near Y is similar. □

Definition 4.3.11. *Given a G-invariant subspace $Y \subseteq X$ the* relative ideal *of Y in X, denoted by $C^*_G(Y \subseteq X)$, is the norm closure of the set of all G-invariant locally compact controlled operators which are supported near Y. It is an ideal in $C^*_G(X)$.*

As with dual algebras, $C^*_G(Y \subseteq X)$ has the same K-theory as $C^*_G(Y)$. The key to the proof is the observation that the inclusion $Y \hookrightarrow B_R(Y)$ is a coarse equivalence.

Proposition 4.3.12. *The inclusion map $Y \hookrightarrow X$ induces an isomorphism*

$$K_p(C^*_G(Y)) \to K_p(C^*_G(Y \subseteq X))$$

Proof. Let H_Y and H_X be ample (Y, G)- and (X, G)-modules, respectively, and let $V \colon H_Y \to H_X$ be an equivariant isometry which coarsely covers the inclusion map

$Y \hookrightarrow X$. We begin by checking that the image of Ad_V lies in $C^*_G(Y \subseteq X)$; we must show that if $T \in \mathbb{B}(H_Y)$ is a controlled operator then $\mathrm{Ad}_V(T)$ is supported near Y. By definition $\mathrm{Ad}_V(T) = VTV^*$, so $\mathrm{Supp}(\mathrm{Ad}_V(T)) \subseteq \mathrm{Supp}(V)\mathrm{Supp}(T)\mathrm{Supp}(V^*)$. For every point (x_2, x_1) in this set there exist y_1 and y_2 such that $(x_2, y_2) \in \mathrm{Supp}(V)$, $(y_2, y_1) \in \mathrm{Supp}(T)$, and $(y_1, x_1) \in \mathrm{Supp}(V^*)$ (equivalently, $(x_1, y_1) \in \mathrm{Supp}(V)$). Since V covers the inclusion map, there is a constant R depending only on V such that $d(x_2, y_2) < R$ and $d(x_1, y_1) < R$. Thus $(x_2, x_1) \in B_R(Y)$ for every $(x_2, x_1) \in \mathrm{Supp}(\mathrm{Ad}_V(T))$ which shows that $Ad_V(T)$ is supported near Y.

Second, we show that the induced map on K-theory is an isomorphism. For each $n \in \mathbb{N}$, compress H_X to the subspace $H_n = P_n H_X P_n$ where P_n is the projection obtained from the characteristic function of $B_n(Y)$ by the Borel functional calculus. Note that H_n has the structure of a $(B_n(Y), G)$-module since $B_n(Y)$ is G-invariant; since every operator supported near Y is supported in some $B_n(Y)$ we have:

$$C^*_G(Y \subseteq X) = \overline{\bigcup_n C^*_G(B_n(Y))}$$

Hence $K_p(C^*(Y \subseteq X)) \cong \lim_{\rightarrow} K_p(C^*(B_n(Y)))$. Since the inclusion $Y \hookrightarrow B_n(Y)$ is a G-equivariant coarse equivalence for each n and K-theory commutes with direct limits of C*-algebras, Ad_V induces an isomorphism

$$K_p(C^*(Y)) \cong K_p(C^*(B_n(Y)))$$

The result follows. $\square$

We conclude this section by calculating the K-theory of the coarse C*-algebra in a few important special cases.

Let X be a proper metric space and assume a discrete group G acts freely and properly on X by isometries. Fix a basepoint $x_0 \in X$ and define a map $\phi\colon G \to X$ by $\phi(\gamma) = \gamma(x_0)$. ϕ is clearly injective, so it induces a metric on G: define $d_G(\gamma_1, \gamma_2) = d_X(\phi(\gamma_1), \phi(\gamma_2))$. This metric is independent of the basepoint chosen.

Lemma 4.3.13. *G acts on X cocompactly then ϕ is a coarse equivalence.*

Proof. Let D be a fundamental domain for the action of G on X, so that D is connected and its G-translates form a disjoint cover of X. Let $\psi\colon X \to G$ be the map which sends γD to γ. We shall prove that ψ and ϕ are inverses in the coarse category, i.e. $\psi\phi$ and $\phi\psi$ are each close to the identity. In fact $\psi\phi = 1_G$, so it suffices to show that $\phi\psi$ is close to 1_X.

By construction $d_X(\phi\psi(x), 1_X(x)) \leq d_X(\gamma(x_0), x)$ for every $\gamma \in G$. In particular $d_X(\phi\psi(x), 1_X(x)) \leq d_{X/G}([x_0], [x])$ where $[x_0]$ and $[x]$ are the equivalence classes in X/G of x_0 and x, respectively. Thus if C is the diameter of X/G we have that $d_X(\phi\psi(x), 1_X(x)) \leq C$, as desired. □

The map ϕ and its coarse inverse ψ are clearly G-equivariant, so by the functoriality of $K_p(C^*_G(\cdot))$ for G-equivariant coarse maps we have an isomorphism $K_p(C^*_G(X)) \cong K_p(C^*_G(G))$ where G is equipped with the metric d_G described above. While the metric d_G depends on the space X on which G acts, it turns out that $K_p(C^*_G(G))$ depends only on G; we shall identify $C^*_G(G)$ with a particular C*-algebra familiar to operator algebraists and representation theorists. Let $L\colon G \to \mathbb{B}(\ell^2(G))$ denote the left regular representation of G, so that $L_\gamma(e_\delta) = e_{\gamma\delta}$ where $\{e_\delta\}$ is the standard orthonormal basis for $\ell^2(G)$. L extends to a representation of the group algebra $\mathbb{C}G$ on $\ell^2(G)$ in the obvious way, so we may make the following definition:

Definition 4.3.14. *The* reduced group C*-algebra *of G, denoted $C^*_r(G)$, is the closure of $L(\mathbb{C}G) \subseteq \mathbb{B}(\ell^2(G))$.*

Proposition 4.3.15. *Let d_G be a metric on G induced by a free and cocompact action of G on a proper metric space X. Then $K_p(C^*_G(G)) \cong K_p(C^*_r(G))$.*

Proof. Observe that the usual counting measure μ is a G-invariant Borel measure on G and that $L^2(G, \mu) = \ell^2(G)$. This yields a nondegenerate representation $\rho\colon C_0(G) \to \mathbb{B}(\ell^2(G))$, and the left regular representation is the composition $G \hookrightarrow C_0(G) \to \mathbb{B}(\ell^2(G))$ where the embedding $G \hookrightarrow C_0(G)$ sends γ to the characteristic function of γ. This gives $\ell^2(G)$ the structure of a (G, G)-module, but it is not ample because each finitely supported function in $C_0(G)$ acts as a compact operator. So instead let H be an auxiliary separable infinite dimensional Hilbert space and represent $C_0(G)$ amply on $\ell^2(G) \otimes H$ by allowing $C_0(G)$ to act trivially on H.

For each $\delta \in G$ let $P_\delta \in \mathbb{B}(\ell^2(G) \otimes H)$ denote the orthogonal projection operator onto the subspace $\mathbb{C}e_\delta \otimes H$ where $e_\delta \in \ell^2(G)$ is the indicator function for δ.

An operator $T \in \mathbb{B}(\ell^2(G) \otimes H)$ is controlled if and only if $P_\delta T P_\gamma = 0$ for all but finitely many pairs (γ, δ) since the bounded subsets of G are precisely the finite sets. Such an operator is G-equivariant if and only if it is the finite linear combination of operators of the form $L_\gamma \otimes S\colon \ell^2(G) \otimes H$ (where L denotes the left regular representation as above) and S is any bounded operator on H. Finally an operator of the form $L_\gamma \otimes S$ is locally compact if and only if S is a compact operator. Thus the $*$-algebra of G-equivariant locally compact controlled operators is precisely $L(\mathbb{C}G) \otimes \mathbb{K}(H)$, and its norm closure is $C^*_G(G) = C^*_r(G) \otimes \mathbb{K}(H)$. The result now follows from the fact that K-theory is invariant under tensoring with the C*-algebra of compact operators. □

We conclude this section with one more example which will help us compute the K-theory of the coarse C*-algebra of $\mathbb{R}$ and various other spaces later on.

Proposition 4.3.16. *Let Y be a proper metric space and let $X = Y \times \mathbb{R}^+$ equipped with the product metric where $\mathbb{R}^+ = [0, \infty)$ with the standard metric. Then $C^*(X)$ has trivial K-theory.*

Proof. Reference □

4.3.2 The Structure Algebra

We now repeat the discussion above for the coarse analogue of the dual algebra $\mathfrak{D}^*(X)$. We will obtain (in the equivariant case) a C*-algebra $D^*_G(X)$ which is functorial for G-invariant maps which are both continuous and coarse; such maps are called uniform maps. In our discussion of K-homology we were able to pass freely between $\mathfrak{D}^*(X)$ and $\mathfrak{D}^*(X)/\mathfrak{C}^*(X)$ at the level of K-theory because we proved that $K_p(\mathfrak{C}^*(X)) = 0$, but $K_p(C^*_G(X)) \cong K_p(C^*_r(G))$ can have an extremely rich and complicated structure and thus some care must be taken when handling $D^*_G(X)$. One of our main results in this section is that the K-theory of the quotient $D^*_G(X)/C^*_G(X)$ recovers the K-homology of X, so $K_p(D^*_G(X))$ can be seen as a mediator between the small-scale and large-scale structure of X.

As in the last section, all spaces are proper metric spaces on which the countable discrete group G acts freely and properly by isometries.

Definition 4.3.17. *The (equivariant)* structure algebra $D^*_G(X)$ *of* X *is the norm closure in* $\mathbb{B}(H_X)$ *of the the* $*$*-subalgebra of all pseudolocal controlled operators where* H_X *has the structure of a very ample* (X, G)*-module.*

As with the coarse C*-algebra we will often use the notation $D^*(X)$ when the group G is trivial. Our first order of business is to investigate the functorial properties of $K_p(D^*_G(X))$; as usual this involves covering isometries.

Definition 4.3.18. *Let* H_X *and* H_Y *be* X*- and* Y*-modules, respectively.*

- *A map* $\phi\colon X \to Y$ *is* uniform *if it is continuous and coarse.*
- *An isometry* $V\colon H_X \to H_Y$ uniformly covers *a uniform map* ϕ *if it topologically covers* ϕ *in the sense of Definition 3.3.10 and coarsely covers* ϕ *in the sense of Definition 4.3.3.*

Our results on coarse and topological covering isometries imply several basic facts about uniform covering isometries. If V is an equivariant isometry which uniformly covers an equivariant map ϕ then Ad_V maps G-equivariant pseudolocal operators to G-equivariant pseudolocal operators (since it topologically covers ϕ) and it maps G-equivariant controlled operators to G-equivariant controlled operators (since it topologically covers ϕ). Taking closures, Ad_V maps $D^*_G(X)$ into $D^*_G(Y)$. Additionally, the induced map on K-theory $(\mathrm{Ad}_V)_*$ is independent of the choice of uniform covering isometry:

Lemma 4.3.19. *Let* V_1 *and* V_2 *be two equivariant isometries which uniformly cover an equivariant map* $\phi : X \to Y$ *between proper* G*-spaces. Then*

$$(Ad_{V_1})_* = (Ad_{V_2})_*\colon K_p(D^*_G(X)) \to K_p(D^*_G(Y))$$

Proof. This follows by combining the proofs of Lemma 3.3.14 and Lemma 4.3.7 and applying Lemma 2.3.16. $\square$

Thus we will have shown that $K_p(D^*_G(X))$ is functorial for equivariant uniform maps if we can prove that equivariant uniform covering isometries always exist. Our strategy for constructing such an isometry is to assemble it from locally defined topological covering isometries associated to open covers of X and Y with uniformly bounded diameters. So far we have not discussed equivariant topological covering isometries, so we take a moment to fill in this gap.

Lemma 4.3.20. *Let $X = \mathcal{U} \times G \to \mathcal{U}$ and $Y = \mathcal{V} \times G \to \mathcal{V}$ be trivial G-covers of proper metric spaces $\mathcal{U}$ and $\mathcal{V}$ and let H_X and H_Y be ample (X, G)- and (Y, G)-modules, respectively. Then any continuous G-equivariant map $\phi\colon X \to Y$ is topologically covered by a G-equivariant isometry $H_X \to H_Y$.*

Proof. Let $H_{\mathcal{U}}$ denote the closure of $C_0(\mathcal{U} \times \{e\})H_X$ where $e \in G$ is the identity and define $H_{\mathcal{V}}$ similarly. $H_{\mathcal{U}}$ and $H_{\mathcal{V}}$ carry ample representations of $C_0(\mathcal{U})$ and $C_0(\mathcal{V})$, respectively, and hence Voiculescu's theorem guarantees the existence of an isometry $V\colon H_{\mathcal{U}} \to H_{\mathcal{V}}$ which topologically covers the restriction of ϕ to $U \times \{e\}$. We have that $H_X \cong \ell^2(G) \otimes H_{\mathcal{U}}$ and $H_Y \cong \ell^2(G) \otimes H_{\mathcal{V}}$, so $1 \otimes V\colon H_X \to H_Y$ is an equivariant topological covering isometry for ϕ. □

We obtain G-equivariant uniform isometries by gluing together isometries constructed in this lemma. However, this requires an even stronger sort of representation than what we have considered so far.

Definition 4.3.21. *Let A be a C^*-algebra and H a Hilbert space. Say that a representation $\rho\colon A \to \mathbb{B}(H)$ is* very ample *if it is the countable direct sum of a fixed ample representation.*

Proposition 4.3.22. *Let H_X be a (X, G)-module and let H_Y be a very ample (Y, G)-module. Then every equivariant uniform map $\phi\colon X \to Y$ is uniformly covered by an equivariant isometry $V\colon H_X \to H_Y$.*

Proof. Assume $H_Y = \bigoplus_{\mathbb{N}} H$ where H is a Hilbert space carrying a fixed ample representation $\rho_Y\colon C_0(Y) \to \mathbb{B}(H)$. Choose countable open covers $\{\mathcal{U}_m\}$ of X and $\{\mathcal{V}_n\}$ of Y with the following properties:

- $\{\mathcal{U}_m\}$ and $\{\mathcal{V}_n\}$ are G-invariant (meaning G permutes that U_m's and V_n's, respectively) and locally finite
- Each $\mathcal{U}_m$ evenly covers its image in X/G and similarly for $\mathcal{V}_n$
- $\{\mathcal{U}_m\}$ and $\{\mathcal{V}_n\}$ have uniformly bounded diameters

- For every m there exists $n(m)$ such that $\phi(\mathcal{U}_m) \subseteq \mathcal{V}_{n(m)}$

Let H_m^X denote the closure of $\rho_X(C_0(\mathcal{U}_m))H_X$ and let H_n^Y denote the closure of $\rho_Y(C_0(\mathcal{V}_n))H$. Thus H_m^X carries an ample G-equivariant representation of $C_0(\mathcal{U}_m)$ and H_n^Y carries an ample G-equivariant representation of $C_0(\mathcal{V}_n)$. Since ϕ restricts to a continuous map $\mathcal{U}_m \to \mathcal{V}_{n(m)}$, there is a G-equivariant isometry $V_m \colon H_m^X \to H_{n(m)}^Y$ which topologically covers $\phi|_{\mathcal{U}_m}$ by Lemma 4.3.20.

Let $\{h_m\}$ be a G-invariant partition of unity subordinate to the open cover $\{\mathcal{U}_m\}$ and define $V \colon H_X \to H_Y = \bigoplus_{\mathbb{N}} H$ to be the strong limit

$$V = \bigoplus_m V_m \rho_X(h_m^{1/2})$$

It is clear that V is a G-equivariant isometry. To show that V topologically covers ϕ, we must show that $\rho_X(g \circ \phi) \sim V^* \rho_Y(g) V$ for every $g \in C_0(Y)$. We have

$$\begin{aligned} V^* \rho_Y(g) V &= \bigoplus_m \rho_X(h_m^{1/2}) V_m^* \rho_Y(g) V_m \rho_X(h_m^{1/2}) \\ &\sim \sum_m \rho_X(h_m^{1/2}) \rho_X(g \circ \phi|_{\mathcal{U}_m}) \rho_X(h_m^{1/2}) \\ &= g \circ \phi \end{aligned}$$

since V_m topologically covers $\phi|_{\mathcal{U}_m}$.

Finally, to show that V coarsely covers ϕ note that $\mathrm{Supp}(V) = \bigcup_m \mathrm{Supp}(V_m)$ and $\mathrm{Supp}(V_m) \subseteq \mathcal{V}_{n(m)} \times \mathcal{U}_m$. Since the diameters of the sets $\mathcal{V}_n$ and $\mathcal{U}_m$ are uniformly bounded and ϕ maps $\mathcal{U}_m$ into $\mathcal{V}_{n(m)}$, the restrictions of π_1 and $\phi \circ \pi_2$ to $\mathrm{Supp}(V)$ are close. This completes the proof. □

By this proposition and the preceeding discussion, we have proven the following:

Proposition 4.3.23. *The assignment*

$$\{\phi \colon X \to Y\} \mapsto \{\phi_* \colon K_p(D_G^*(X)) \to K_p(D_G^*(Y))\}$$

is a contravariant functor from the category of proper metric spaces with uniform maps to the category of abelian groups with group homomorphisms.

There is a relative counterpart of the structure algebra obtained by synthesizing relative K-homology with the relative theory for coarse C*-algebras as follows:

Definition 4.3.24. *Given a closed G-invariant subspace $Y \subseteq X$, the ideal $D^*_G(Y \subseteq X)$ is defined to be the norm closure of the set of all G-invariant operators $T \in \mathbb{B}(H)$ with the following properties:*

- *T is pseudolocal, meaning T commutes with $\rho_X(f)$ modulo compact operators for every $f \in C_0(X)$.*
- *T is locally compact for the complement of Y, meaning Tg and gT are compact for every $g \in C_0(X - Y)$.*
- *T is supported near Y, meaning its support lies within a bounded distance of $Y \times Y$.*

Following our discussion of relative dual algebras and relative coarse algebras, we will prove that the inclusion $Y \hookrightarrow X$ induces an isomorphism $K_p(D^*(Y)) \cong K_p(D^*(Y \subseteq X))$. Our approach will rely on an important general principle which uses the structure algebra to link coarse geometry and K-homology, so we will explore this principle in detail before concluding our discussion of relative structure algebras. At the heart of the matter is the following result:

Theorem 4.3.25. *Let $Y \subseteq X$ be a closed G-invariant subspace and let $Y_G \subseteq X_G$ be the quotient of Y by G viewed as a subspace of the quotient of X by G. Then there is an isomorphism*

$$D^*_G(Y \subseteq X)/C^*_G(Y \subseteq X) \cong \mathfrak{D}^*(Y_G \subseteq X_G)/\mathfrak{C}^*(X_G)$$

The main idea of the proof is that operators in $\mathfrak{D}^*(X_G)$ can be "truncated" into operators with small supports up to an error which lies in $\mathfrak{C}^*(X_G)$. This truncation procedure allows us to handle two issues at once: first, we can insist that the truncated operators are supported in evenly covered neighborhoods in X_G which allows them to be lifted to G-invariant operators on X; and second, we can constrain the supports of the truncated operators and thereby manufacture a controlled operator.

Let us turn to the precise definitions.

Definition 4.3.26. *Let X be a proper metric space and let $\rho\colon C_0(X) \to \mathbb{B}(H_X)$ be a representation. Let $\{U_n\}$ be a countable locally finite collection of open subsets of X and let $\{h_n\}$ be a subordinate partition of unity. Given any $T \in \mathbb{B}(H_X)$, define $\mathfrak{T}(T)$ to be the strong limit of the series $\sum_n \rho(h_n^{1/2})T\rho(h_n^{1/2})$. $\mathfrak{T}$ is called the* truncation *of T relative to the open cover $\{\mathcal{U}_n\}$.*

There is a considerable amount of flexibility in Definition 4.3.26 in that we can adapt the truncation construction to a variety of different geometric contexts by carefully selecting the open sets used to define it. But regardless of these choices $\mathfrak{T}$ is always a continuous linear operator on $\mathbb{B}(H_X)$:

Lemma 4.3.27. *The truncation operator defines a continuous linear map*

$$\mathfrak{T}\colon \mathbb{B}(H_X) \to \mathbb{B}(H_X)$$

Proof. Assume at first that T is positive and let $T_N = \sum_{n=1}^N \rho(h_n^{1/2})T\rho(h_n^{1/2})$ be the Nth partial sum. For $v \in H$ we have the estimate

$$\langle T_N v, v\rangle = \sum_{n=1}^{N} \left\langle \rho_X(h_n^{1/2})T\rho_X(h_n^{1/2})v, v\right\rangle \leq \|T\| \sum_{n=1}^{N} \left\|\rho_{h_n}^{1/2} v\right\|^2 \leq \|T\| \, \|v\|^2$$

So $\|T_N\| \leq \|T\|$ for every N. An infinite series of positive operators whose partial sums are all bounded by a constant C converges strongly to an operator whose norm is bounded by C, so the strong limit $\mathfrak{T}(T) = \lim_N T_N$ exists and satisfies $\|\mathfrak{T}(T)\| \leq \|T\|$. Every self adjoint operator is the difference of two positive operators, and every bounded operator T can be expressed in terms of self-adjoint operators via the functional calculus: $T = \mathrm{Re}(T) + i\mathrm{Im}(T)$. Since the positive and negative parts of $\mathrm{Re}(T)$ and $\mathrm{Im}(T)$ have norm at most $\|T\|$, we have a naive estimate $\|\mathfrak{T}(T)\| \leq 4\,\|T\|$ for any bounded operator T (in fact the factor of 4 can be removed, but this is not required for our purposes). We conclude that $\mathfrak{T}T$ is a bounded and therefore continuous operator on the Banach space $\mathbb{B}(H_X)$. □

The next result estimates the support of $\mathfrak{T}(T)$ in terms of the open sets used to define it.

Lemma 4.3.28. *Let T be an operator in $\mathbb{B}(H_X)$ and assume that $\mathfrak{T}(T)$ is defined using a partition of unity $\{h_n\}$ subordinate to the open cover $\{\mathcal{U}_n\}$ of X. Then*

$$Supp(\mathfrak{T}(T)) \subseteq \bigcup_n \overline{\mathcal{U}_n \times \mathcal{U}_n}$$

Proof. Let $\mathcal{V}_1 \times \mathcal{V}_2$ be an open set in the complement of $\bigcup_n \overline{\mathcal{U}_n \times \mathcal{U}_n}$. This means that

$$\rho(f)\rho(h_n^{1/2})T\rho(h_n^{1/2})\rho(g) = 0$$

for every $f \in C_0(V_1)$ and every $g \in C_0(V_2)$, so

$$\rho(f)\left(\sum_{n=1}^{N} \rho(h_n^{1/2})T\rho(h_n^{1/2})\right)\rho(g) = 0$$

Passing to the strong limit we conclude that $\mathcal{V}_1 \times V_2$ is in the complement of $\mathrm{Supp}(\mathfrak{T}(T))$. $\square$

Corollary 4.3.29. *In the situation of Lemma 4.3.28, we have for any $T \in \mathbb{B}(H_X)$ that:*

- *$\mathfrak{T}(T)$ is controlled if $\{\mathcal{U}_n\}$ has uniformly bounded diameter.*
- *$\mathfrak{T}(T)$ is supported near a subset $Y \subseteq X$ if there exists $R > 0$ such that $\mathcal{U}_n \subseteq B_R(Y)$ for every n.*

The truncation construction is also compatible with the representation of $C_0(X)$ in the following sense:

Lemma 4.3.30. *If $T \in \mathbb{B}(H_X)$ is pseudolocal then $\mathfrak{T}(T)$ is pseudolocal and $T - \mathfrak{T}(T)$ is locally compact.*

Proof. Let f be any continuous compactly supported function on X. The support of f meets only finitely many of the $\mathcal{U}_n$'s, so

$$[\mathfrak{T}(T), \rho(f)] = \sum_n \rho(h_n^{1/2})[T, \rho(f)]\rho(h_n^{1/2})$$

is a finite sum of compact operators and hence is compact. Since $C_c(X)$ is dense in $C_0(X)$ it follows that $\mathfrak{T}(T)$ is pseudolocal.

To see that $T - \mathfrak{T}(T)$ is locally compact, write:

$$T - \mathfrak{T}(T) = \sum_n \rho(h_n^{1/2})[T, \rho(h_n^{1/2})] = -\sum_n [T, \rho(h_n^{1/2})]\rho(h_n^{1/2})$$

Each term of this sum is compact since T is pseudolocal, so again $\rho(f)(T - \mathfrak{T}(T))$ and $(T - \mathfrak{T}(T))\rho(f)$ can each be written as the finite sum of compact operators whenever $f \in C_c(X)$. As before it follows that $T - \mathfrak{T}(T)$ is locally compact since $C_c(X)$ is dense in $C_0(X)$. □

We are now return to Theorem 4.3.25, which reduces to two statements that we will handle separately. The first statement is purely about coarse geometry and has nothing to do with the group action: informally, it says that every pseudolocal operator agrees with a pseudolocal controlled operator up to a locally compact perturbation. The precise claim, formulated for relative C*-algebras, is as follows:

$$\mathfrak{D}^*(Y_G \subseteq X_G)/\mathfrak{C}^*(X_G) \cong D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G)$$

The second statement focuses on the group action: it asserts that a G-equivariant pseudolocal controlled operator on X is the same thing as a pseudolocal controlled operator on the quotient space X_G, at least up to a locally compact error. The precise formulation, again in the relative setting, is as follows:

$$D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G) \cong D^*_G(Y \subseteq X)/C^*_G(Y \subseteq X)$$

Of course the proofs are more complicated than their naive interpretations in part because elements of the structure algebra are actually limits of pseudolocal controlled operators. Both statements will be proved using the truncation operator described above.

Proposition 4.3.31. *The inclusion $D^*(Y_G \subseteq X_G) \hookrightarrow \mathfrak{D}^*(Y_G \subseteq X_G)$ induces an isomorphism*

$$\mathfrak{D}^*(Y_G \subseteq X_G)/\mathfrak{C}^*(X_G) \cong D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G)$$

Proof. Injectivity of the induced map follows from the identity $C^*(Y_G \subseteq X_G) =$

$D^*(Y_G \subseteq X_G) \cap \mathfrak{C}^*(X_G)$. It is immediate from the definitions that every locally compact controlled operator on X_G which is supported near Y_G is pseudolocal and locally compact for the complement of Y_G, so the containment $C^*(Y_G \subseteq X_G) \subseteq D^*(Y_G \subseteq X_G) \cap \mathfrak{C}^*(X_G)$ follows by passing to norm limits. To prove the other containment, form the truncation operator $\mathfrak{T}$ on $\mathbb{B}(H_{X_G})$ using a countable, locally finite, uniformly bounded collection of open sets $\{\mathcal{U}_n\}$ such that $Y \subseteq \bigcup_n \mathcal{U}_n \subseteq B_R(Y_G)$ for some $R > 0$. Given any operator $T \in D^*(Y_G \subseteq X_G) \cap \mathfrak{C}^*(X_G)$ it is clear that $\mathfrak{T}(T) \in C^*(Y_G \subseteq X_G)$. Writing T as the norm limit of pseudolocal controlled operators T_n which are supported near Y we have that $T_n - \mathfrak{T}(T_n) \in C^*(Y_G \subseteq X_G)$ by Lemma 4.3.28 and Lemma 4.3.30, so the same is true of $T - \mathfrak{T}(T)$. It follows that $T \in C^*(Y_G \subseteq X_G)$, as desired.

Surjectivity of the induced map follows from the statement that $\mathfrak{D}^*(Y_G \subseteq X_G) = D^*(Y_G \subseteq X_G) + \mathfrak{C}^*(X_G)$. The inclusion "$\supseteq$" is obvious since $D^*(Y_G \subseteq X_G)$ and $\mathfrak{C}^*(X_G)$ are both subalgebras of $\mathfrak{D}^*(Y_G \subseteq X_G)$, so it suffices to show that $\mathfrak{D}^*(Y_G \subseteq X_G) \subseteq D^*(Y_G \subseteq X_G) + \mathfrak{C}^*(X_G)$. Form the truncation operator $\mathfrak{T}$ using the same collection of open sets as above and let T be any operator in $\mathfrak{D}^*(Y_G \subseteq X_G)$. Lemma 4.3.28 and Lemma 4.3.30 imply that $\mathfrak{T}(T)$ is pseudolocal, controlled and supported near Y_G; $\mathfrak{T}(T)$ is also locally compact for the complement of Y_G since $\rho(f)\mathfrak{T}(T)$ and $\mathfrak{T}(T)\rho(f)$ are each finite sums of compact operators for any $f \in C_c(X_G - Y_G)$. Thus $\mathfrak{T}(T) \in D^*(Y_G \subseteq X_G)$, and $T - \mathfrak{T}(T) \in \mathfrak{C}^*(X_G)$ by Lemma 4.3.30. $\square$

The isomorphism

$$D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G) \cong D^*_G(Y \subseteq X)/C^*_G(Y \subseteq X)$$

is a little more subtle but still uses the truncation operator in an essential way. The main idea is that equivalence classes in both quotient algebras can be localized into operators with small support. The quotient map $\pi \colon X \to X_G$ is a local isometry since G acts on X by isometries, so given any fixed $\varepsilon > 0$ there is a countable open cover $\{\mathcal{U}_n\}$ of X_G by sets of diameter smaller than ε such that each $\mathcal{U}_n$ is evenly covered by π. Let $\widetilde{\mathcal{U}}_n = \pi^{-1}(\mathcal{U}_n) = \mathcal{U}_n \times G$ be the lift of $\mathcal{U}_n$ to X.

Lemma 4.3.32. *Every equivalence class in $D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G)$ has a representative whose support lies in $\bigcup_n \mathcal{U}_n \times \mathcal{U}_n$ and every equivalence class in*

$D_G^*(Y \subseteq X)/C_G^*(Y \subseteq X)$ *has a representative whose support lies in* $\bigcup_n \widetilde{\mathcal{U}}_n \times \widetilde{\mathcal{U}}_n$.

Proof. Let $T \in D^*(Y_G \subseteq X_G)$ be an operator which is pseudolocal, controlled, locally compact for the complement of Y_G, and supported near Y_G. Define a truncation operator $\mathfrak{T}$ on $\mathbb{B}(H_{X_G})$ using a partition of unity subordinate to $\{\mathcal{U}_n\}$ and note that $\mathfrak{T}(T)$ is an operator in $D^*(Y_G \subseteq X_G)$ whose support lies in $\bigcup_n(\mathcal{U}_n \times \mathcal{U}_n)$ by Lemma 4.3.28. Moreover $T - \mathfrak{T}(T) \in C^*(Y_G \subseteq X_G)$ by Lemma 4.3.30, so T and $\mathfrak{T}(T)$ represent the same equivalence class in the quotient.

The same argument using a G-invariant partition of unity subordinate to $\{\widetilde{\mathcal{U}}_n\}$ proves the corresponding statement about operators in the equivariant quotient algebra. □

The isomorphism between the two quotient algebras is now a simple matter of lifting operators on X_G with small support to G-invariant operators on X with small support. The key observation is that if we define H_n to be the closure of $C_0(\mathcal{U}_n)H_{X_G}$ and $\widetilde{H}_n$ to be the closure of $C_0(\widetilde{\mathcal{U}}_n)H_X$ then $\widetilde{H}_n \cong H_n \otimes \ell^2(G)$. Thus any operator T_n on H_n lifts to the G-invariant operator $T_n \otimes 1$ on $\widetilde{H_n}$.

Proposition 4.3.33. *There is an isomorphism*

$$D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G) \cong D_G^*(Y \subseteq X)/C_G^*(Y \subseteq X)$$

Proof. We define a map $\mathfrak{L}\colon D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G) \to D_G^*(Y \subseteq X)/C_G^*(Y \subseteq X)$ as follows. Given a class $x \in D^*(Y_G \subseteq X_G)/C^*(Y_G \subseteq X_G)$, choose a representative T for x which is supported in $\bigcup_n \mathcal{U}_n \times \mathcal{U}_n$ using a partition of unit $\{h_n\}$. Set $T_n = \rho(h_n^{1/2})T\rho(h_n^{1/2}$; note that T_n is supported in $\mathcal{U}_n \times \mathcal{U}_n$ and T is the strong limit of the sum $\sum_n T_n$. T_n can be regarded as an operator on $H_n\colon = \overline{C_0(\mathcal{U}_n)H_{X_G}}$, so T_n lifts to the G-invariant operator $\widetilde{T}_n = T_n \otimes 1$ on $\widetilde{H}_n\colon = \overline{C_0(\widetilde{\mathcal{U}}_n)H_X} \cong H_n \otimes \ell^2(G)$.

So we define $\mathfrak{L}(x)$ to be the strong limit of the sum $\sum_n \widetilde{T}_n$; by the previous lemma, $\mathfrak{L}(x)$ defines a class in $D_G^*(Y \subseteq X)/C_G^*(Y \subseteq X)$. Note that if $\mathfrak{L}'$ is another lift operator defined using a different open cover and partition of unity then by passing to a common refinement of the two partitions of unity we would find that $\rho(f)(\mathfrak{L}(x) - \mathfrak{L}'(x))$ is a finite sum of compact operators for any compactly supported function f and hence $\mathfrak{L}(x)$ and $\mathfrak{L}'(x)$ determine the same class in $D_G^*(Y \subseteq X)/C_G^*(Y \subseteq X)$. Thus $\mathfrak{L}(x)$ is a well-defined algebra homomorphism.

If $\mathfrak{L}([T]) = 0$ then we have that $\rho(f)T\rho(g) = 0$ for every $f, g \in C_0(U_n)$, and this implies that $T = 0$ since the representation ρ is nondegenerate. Thus $\mathfrak{L}$ is injective. To see that $\mathfrak{L}$ is surjective, choose a representative S of a class $y \in D^*_G(Y \subseteq X)/C^*_G(Y \subseteq X)$ and write S as the strong limit of the sum $\sum_n S_n$ using the partition of unity obtained by lifting the one used to define $\mathcal{L}$ to X. Thus each S_n is a G-invariant operator on $\widetilde{H}_n \cong H_n \otimes \ell^2(G)$. Every such operator has the form $S_n = T_n \otimes 1$ for some operator T_n on H_n, and we have that $\mathfrak{L}([T]) = y$ where T is the strong limit of the sum $\sum_n T_n$. $\square$

The previous two propositions combined yield the proof of Theorem 4.3.25. This result is important for a variety of reasons, but in particular it implies that $D^*_G(Y)$ and $D^*_G(Y \subseteq X)$ have isomorphic K-theory groups.

Proposition 4.3.34. *The inclusion $Y \hookrightarrow X$ induces an isomorphism*

$$K_p(D^*_G(Y)) \cong K_p(D^*_G(Y \subseteq X))$$

Proof. Note that the inclusion $Y \hookrightarrow X$ is an equivariant uniform map, so it is uniformly covered by an equivariant isometry $V \colon H_Y \to H_X$. V is in particular an equivariant coarse covering isometry, and following the proof of Proposition 4.3.22 it is the lift of a topological covering isometry V_G for the inclusion $Y_G \hookrightarrow X_G$. Thus there is a commutative diagram:

$$\begin{tikzcd}
0 \arrow[r] & C^*_G(Y) \arrow[r] \arrow[ddd, "\mathrm{Ad}(V)"] & D^*_G(Y) \arrow[r] \arrow[ddd, "\mathrm{Ad}(V)"] & D^*_G(Y)/C^*_G(Y) \arrow[r] \arrow[d, "\cong"] & 0 \\
 & & & \mathfrak{D}^*(Y_G)/\mathfrak{C}^*(Y_G) \arrow[d, "\mathrm{Ad}(V_G)"] & \\
 & & & \mathfrak{D}^*(Y_G \subseteq X_G)/\mathfrak{C}^*(Y_G \subseteq X_G) \arrow[d, "\cong"] & \\
0 \arrow[r] & C^*_G(Y \subseteq X) \arrow[r] & D^*_G(Y \subseteq X) \arrow[r] & D^*_G(Y \subseteq X)/C^*_G(Y \subseteq x) \arrow[r] & 0
\end{tikzcd}$$

This gives rise to a commutative diagram in K-theory:

$$\begin{tikzcd}
K_{p+1}(Y_G) \arrow[r] \arrow[d] & K_{p+1}(C_G^*(Y)) \arrow[r] \arrow[d] & K_{p+1}(D_G^*(Y)) \arrow[r] \arrow[d] & K_p(Y_G) \arrow[r] \arrow[d] & K_p(C_G^*(Y)) \arrow[d] \\
K_{p+1}(Y_G \subseteq X_G) \arrow[r] & K_{p+1}(C_G^*(Y \subseteq X)) \arrow[r] & K_{p+1}(D_G^*(Y \subseteq X)) \arrow[r] & K_p(Y_G \subseteq X_G) \arrow[r] & K_p(C_G^*(Y \subseteq X))
\end{tikzcd}$$

We have proven that all of the vertical maps except the middle one are isomorphisms, so the middle map is an isomorphism by the five lemma. □

4.4 The Analytic Surgery Exact Sequence

We conclude this chapter by summarizing the main results and commenting on their significance. Given a complete Riemannian manifold M, there is a short exact sequence

$$0 \to C^*(M) \to D^*(M) \to Q^*(M) \to 0$$

where $Q^*(M)$ denotes the quotient $D^*(M)/C^*(M)$. The K-theory boundary map associated to this short exact sequence is a map $K_p(M) \to K_p(C^*(M))$; if M happens to be compact then it is coarsely equivalent to a one-point space P and hence we recover in degree 0 Atiyah and Singer's analytic index map $K_0(M) \to \mathbb{Z}$. Thus we can consider the map $K_p(M) \to K_p(C^*(M))$ to be a generalized index map, and it is the basis of Roe's coarse index theory. We will use it in the next chapter to formulate the partitioned manifold index theorem.

Now let M be a compact manifold, let G be a quotient of the fundamental group of M, and let $\widetilde{M}$ be a G-cover of M. The short exact sequence

$$0 \to C_G^*(\widetilde{M}) \to D_G^*(\widetilde{M}) \to Q_G^*(\widetilde{M}) \to 0$$

gives rise to a long exact sequence in K-theory:

$$\ldots \to K_{p+1}(C_r^*(G)) \to K_{p+1}(D_G^*(\widetilde{M})) \to K_p(M) \to K_p(C_r^*(G)) \to \ldots$$

If G is trivial then $K_p(C_r^*(G)) = K_p(\mathbb{K})$ and the boundary map $K_0(M) \to K_0(\mathbb{K}) = \mathbb{Z}$ is again Atiyah and Singer's analytic index map. Once again, we regard $K_p(M) \to K_p(C_r^*(G))$ as a sort of generalized index map. In the case where M is a classifying space for G (so that M has fundamental group G and all of the higher homotopy groups of M vanish), this map is called the *Baum-Connes assembly map* based on the conjecture of Baum and Connes ([16]) that it is an isomorphism for every group G. The Baum-Connes conjecture has many applications to geometry and topology, and it is known to be true for a very wide variety of discrete groups.

The group $K_p(M)$ is reasonably well understood because it is accessible to the techniques of algebraic topology. $K_p(C_r^*G)$ is often very difficult to calculate directly (without the aid of the Baum-Connes conjecture), but in some cases it can be approached using tools in geometric group theory and representation theory. The intermediate group, $K_p(D_G^*(\widetilde{M}))$, is in some ways the most mysterious because it has no obvious algebraic or geometric interpretation. There is, however, an immediately available analytic interpretation: a class in $K_p(M)$ lifts to $K_p(D_G^*(\widetilde{M}))$ if and only if its generalized index in $C_r^*(M)$ vanishes, so we can view $K_p(D_G^*(\widetilde{M}))$ as a group of *secondary index invariants* of elliptic operators on M. There are a number of non-trivial vanishing theorems in index theory, including Lichnerowicz's result that the spinor Dirac operator on a Riemannian spin manifold with positive scalar curvature has vanishing index and Atiyah and Hirzebruch's result ([2]) that the spinor Dirac operator on a spin manifold with a nontrivial circle action has vanishing index. Every such vanishing theorem determines a potentially interesting element of $K_p(D_G^*(\widetilde{M}))$.

Recent results of Higson and Roe ([10], [11], and [12]) provide a further interpretation of the group $K_p(D_G^*(\widetilde{M}))$. They showed that the long exact sequence

$$\ldots \to K_{p+1}(C_r^*(G)) \to K_{p+1}(D_G^*(\widetilde{M})) \to K_p(M) \to K_p(C_r^*(G)) \to \ldots$$

fits into a commutative diagram (modulo 2-torsion) with the *surgery exact sequence* from algebraic topology. For this reason we refer to this long exact sequence as the *analytic surgery exact sequence*. Given a manifold M of dimension n with fundamental group G the topologists' surgery exact sequence has the form:

$$\to L_{n+1}(G) \to \mathcal{S}(M) \to \mathcal{N}(M) \to L_n(G) \to$$

The details need not concern us here; we only note that $\mathcal{N}(M) \to L_n(G)$ is an assembly map analagous to the Baum-Connes assembly map described above and the *structure set* $\mathcal{S}(M)$ (not in general a group in this setting) measures the failure of this map to be an isomorphism. The structure set is of great interest to algebraic topologists because it classifies the different possible manifold structures available on M, and thus its counterpart $K_p(D^*_G(\widetilde{M}))$ in the analytic surgery exact sequence may carry interesting topological information. In honor of this connection we introduce the following notation:

Definition 4.4.1. *Let X be a proper metric space and let G be a discrete group which acts freely and properly on X by isometries. The pth* analytic structure group *of the pair (X, G) is defined to be*

$$S_p(X, G) = K_{1-p}(D^*_G(X))$$

In the case where G is trivial we simply write $S_p(X)$.

In the next chapter we will use the analytic surgery exact sequence in tandem with the relative theory developed in this chapter and the last one to prove our generalization of the partitioned manifold index theorem. In the final chapter we will comment on possible analogues of this result in the realm of secondary index invariants. We will not discuss how our results relate to surgery theory, but it would be interesting to revisit such questions in future work.

Chapter 5

The Partitioned Manifold Index Theorem

In this chapter we prove the main result of this thesis: a generalized partitioned manifold index theorem. As explained in the introduction, this result solves an index problem on a non-compact manifold partitioned by a compact hypersurface. The theorem was first proved by Roe and subsequently simplified by Higson, who also observed that it implies a cobordism invariance property for indices of elliptic operators. Our approach lends itself to generalizations of Roe's theorem which do not appear in the literature, such as a counterpart of the result for equivariant indices. It also provides an appealing geometric approach to index theory on partitioned manifolds and suggests possible generalizations to secondary index invariants, a topic which we will discuss in the next chapter.

In the problem at hand we are given a manifold M which is the union of two submanifolds $M = M^+ \cup M^-$ with common boundary N. These data provide the input for the classical Mayer-Vietoris sequence in algebraic topology which relates topological invariants of M to those of M^+, M^-, and N. In the previous two chapters we showed that index theory on any proper metric space X can be organized around a long exact sequence in K-theory:

$$\ldots \to K_{p+1}(C^*(X)) \to K_p(D^*(X)) \to K_p(X) \to K_p(C^*(X)) \to \ldots$$

In this chapter we will build Mayer-Vietoris sequences for each of the groups appearing in this long exact sequence - coarse K-theory, the structure group, and K-homology - and use them to relate index theory on M to index theory on N. Our main result will follow from an explicit calculation of Mayer-Vietoris boundary maps made possible by our computations with Kasparov products in the first chapter.

5.1 The Mayer-Vietoris Sequence in K-Theory for C*-Algebras

In order to build the required Mayer-Vietoris sequences we will begin by reviewing a construction of the Mayer-Vietoris sequence in the abstract setting of K-theory for C*-algebras. This Mayer-Vietoris sequence is apparently a folklore construction in C*-algebra theory; it appears in [15], but it is probably older. The construction takes as input a C*-algebra A presented as the sum of two ideals I_1 and I_2 and produces a long exact sequence in K-theory relating A, I_1, I_2, and $I_1 \cap I_2$. Upon constructing this long exact sequence we will discuss conditions under which a decomposition $X = Y_1 \cup Y_2$ of a proper metric space into subspaces corresponds to an appropriate decomposition of the coarse C*-algebra, the structure algebra, or the dual algebra. These *excision conditions* will be sufficiently general to accommodate partitioned manifolds.

5.1.1 The Abstract Mayer-Vietoris Sequence

Given a C*-algebra A with two ideals I_1 and I_2 with the property that $I_1 + I_2$ is dense in A, our aim in this section is to construct a long exact sequence

$$\ldots \to K_p(I_1 \cap I_2) \to K_p(I_1) \oplus K_p(I_2) \to K_p(A) \to K_{p-1}(I_1 \cap I_2) \to \ldots$$

Most of the construction is based on naturality and homotopy invariance properties of K-theory; the only real C*-algebraic input is provided by the following lemma:

Lemma 5.1.1. *Let A be a C*-algebra and let I_1 and I_2 be ideals such that $I_1 + I_2$*

is dense in A. Then the map

$$I_1/(I_1 \cap I_2) \oplus I_2/(I_1 \cap I_2) \to A/(I_1 \cap I_2)$$

defined by $(a_1 + I_1 \cap I_2, a_2 + I_1 \cap I_2) \mapsto a_1 + a_2 + I_1 \cap I_2$ is a $*$-isomorphism. In particular, $A = I_1 + I_2$.

Proof. First note that the map described is a $*$-homomorphism because $I_1 I_2 \subseteq I_1 \cap I_2$. To see that it is injective, assume that $a_1 \in I_1$ and $a_2 \in I_2$ are chosen so that $a_1 + a_2 \in I_1 \cap I_2$. Then $a_1 = (a_1 + a_2) - a_2 \in I_2$ and $a_2 = (a_1 + a_2) - a_1$, which shows that $a_1, a_2 \in I_1 \cap I_2$ as desired.

By basic C*-algebra theory an injective $*$-homomorphism between C* algebras automatically has closed range, so the map is an isomorphism since its range was assumed to be dense. □

Let $\Omega(A, I_1, I_2)$ denote the C*-algebra of continuous paths $f\colon [0,1] \to A$ such that $f(0) \in I_1$ and $f(1) \in I_2$. There is a map $\Omega(A, I_1, I_2) \to I_1 \oplus I_2$ given by $f \mapsto (f(0), f(1))$ which fits into a short exact sequence:

$$0 \to SA \to \Omega(A, I_1, I_2) \to I_1 \oplus I_2 \to 0 \tag{5.1.1}$$

We will use the shorthand $\boldsymbol{\Omega}(A, I_1, I_2)$ for the short exact sequence (5.1.1). The Mayer-Vietoris sequence will be obtained as the long exact sequence in K-theory associated to this short exact sequence; this is achieved using the following lemma:

Lemma 5.1.2. *The $*$-homomorphism $I_1 \cap I_2 \to \Omega(A, I_1, I_2)$ defined by sending $a \in I_1 \cap I_2$ to the constant path in A based at a induces an isomorphism in K-theory.*

Proof. First, we show that $I_1 \cap I_2$ is homotopy equivalent to the ideal P in $\Omega(A, I_1, I_2)$ consisting of paths $f\colon [0,1] \to I_1 \cap I_2$. Let $\varphi\colon I_1 \cap I_2 \to P$ denote inclusion-by-constant-paths and let $\psi\colon P \to I_1 \cap I_2$ denote evaluation at 0. Obviously $\psi \circ \varphi$ is the identity, and the maps $h_s : P \to P$ given by $h_s(f)(t) = f(st)$ define a homotopy between $\varphi \circ \psi$ and the identity.

Thus it suffices to show that the inclusion $P \hookrightarrow \Omega(A, I_1, I_2)$ induces an isomorphism in K-theory. The quotient $\Omega(A, I_1, I_2)/P$ is the C*-algebra of paths

$g\colon [0,1] \to A/(I_1 \cap I_2)$ such that $g(0) \in I_1/(I_1 \cap I_2)$ and $g(1) \in I_2/(I_2 \cap I_2)$. For any such g, the isomorphism

$$A/(I_1 \cap I_2) \cong I_1/(I_1 \cap I_2) \oplus I_2/(I_1 \cap I_2)$$

of Lemma 5.1.1 yields a decomposition $g(t) = (g_1(t), g_2(t))$ where g_1 and g_2 are continuous paths in $I_1/(I_1 \cap I_2)$ and $I_2/(I_1 \cap I_2)$, respectively. The condition $g(0) \in I_1/(I_1 \cap I_2)$ implies that $g_2(0) = 0$ and the condition $g(1) \in I_2/(I_1 \cap I_2)$ implies that $g_1(1) = 0$. Thus

$$\begin{aligned} g_1 &\in C_0[0,1) \otimes I_1/(I_1 \cap I_2) \\ g_2 &\in C_0(0,1] \otimes I_2/(I_1 \cap I_2) \end{aligned}$$

Indeed, the map $g \mapsto (g_1, g_2)$ defines an isomorphism

$$\Omega(A, I_1, I_2)/P \cong (C_0[0,1) \otimes I_1/(I_1 \cap I_2)) \oplus (C_0(0,1] \otimes I_2/(I_1 \cap I_2))$$

and thus $\Omega(A, I_1, I_2)/P$ has trivial K-theory. It follows from this together with the long exact sequence in K-theory that the inclusion of P into $\Omega(A, I_1, I_2)$ induces an isomorphism on K-theory. □

Proposition 5.1.3. *Let A be a C^* algebra and let I_1, I_2 be ideals in A such that $A = I_1 + I_2$. Then there is a long exact sequence in K-theory*

$$\ldots \to K_p(I_1 \cap I_2) \to K_p(I_1) \oplus K_p(I_2) \to K_p(A) \to K_{p-1}(I_1 \cap I_2) \to \ldots \qquad (5.1.2)$$

where the map $K_p(I_1 \cap I_2) \to K_p(I_1) \oplus K_p(I_2)$ is induced by inclusion and the map $K_p(I_1) \oplus K_p(I_2) \to K_p(A)$ is induced by the map $(a_1, a_2) \mapsto a_1 - a_2$.

Proof. Consider the short exact sequence $\mathbf{\Omega}(A, I_1, I_2)$ given by:

$$0 \to SA \to \Omega(A, I_1, I_2) \to I_1 \oplus I_2 \to 0$$

Using the identifications $K_p(SA) \cong K_{p+1}(A)$, $K_p(\Omega(A, I_1, I_2)) \cong K_p(I_1 \cap I_2)$, and $K_p(I_1 \oplus I_2) \cong K_p(I_1) \oplus K_p(I_2)$ we obtain a long exact sequence in K-theory of the form (5.1.2). Thus we need only verify that the maps appearing in this long exact sequence agree with the maps specified in the statement of the proposition.

The map $K_p(I_1 \cap I_2) \to K_p(I_1) \oplus K_p(I_2)$ is a priori the composition of the map $K_p(I_1 \cap I_2) \cong K_p(\Omega(A, I_1, I_2))$ induced by inclusion and the map $K_p(\Omega(A, I_1, I_2)) \to K_p(I_1) \oplus K_p(I_2)$ induced by evaluation. Since the diagram

$$\begin{tikzcd} I_1 \cap I_2 \arrow[d] \arrow[dr] & \\ \Omega(A, I_1, I_2) \arrow[r] & I_1 \oplus I_2 \end{tikzcd}$$

commutes, it follows that the map $K_p(I_1 \cap I_2) \to K_p(I_1) \oplus K_p(I_2)$ is induced by inclusion.

The map $K_p(I_1) \oplus K_p(I_2) \to K_p(A)$ is a priori the boundary map in K-theory associated to the short exact sequence $\mathbf{\Omega}(A, I_1, I_2)$. This fits into a commutative diagram with the short exact sequence for the suspension of I_1 whose boundary map is the isomorphism $K_{p+1}(I_1) \cong K_p(S(I_1))$:

$$\begin{tikzcd} 0 \arrow[r] & C_0(0,1) \otimes I_1 \arrow[r] \arrow[d] & C_0[0,1) \otimes I_1 \arrow[r] \arrow[d] & I_1 \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & S(A) \arrow[r] & \Omega(A, I_1, I_2) \arrow[r] & I_1 \oplus I_2 \arrow[r] & 0 \end{tikzcd}$$

An element of $C_0[0,1) \otimes I_1$ can be viewed as a path $\gamma\colon [0,1] \to I_1$ such that $\gamma(1) = 0$, and such a path can be viewed as an element of $\Omega(A, I_1, I_2)$. Thus the middle map above is given by inclusion. It follows that the map $K_p(I_1) \oplus K_p(I_2) \to K_p(A)$ sends K-theory classes of the form $(x, 0)$ to the image of x under the natural map $K_p(I_1) \to K_p(A)$.

There is a similar diagram involving I_2, but this time there is a natural inclusion $C_0(0,1] \otimes I_2 \to \Omega(A, I_1, I_2)$:

$$\begin{tikzcd} 0 \arrow[r] & C_0(0,1) \otimes I_2 \arrow[r] \arrow[d] & C_0[0,1) \otimes I_2 \arrow[r] \arrow[d] & I_2 \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & C_0(0,1) \otimes I_2 \arrow[r] \arrow[d] & C_0(0,1] \otimes I_2 \arrow[r] \arrow[d] & I_2 \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & S(A) \arrow[r] & \Omega(A, I_1, I_2) \arrow[r] & I_1 \oplus I_2 \arrow[r] & 0 \end{tikzcd}$$

Here the vertical maps $C_0(0,1)\otimes I_2 \to C_0(0,1)\otimes I_2$ and $C_0[0,1)\otimes I_2 \to C_0(0,1]\otimes I_2$ are given by the orientation reversing homeomorphism $t \mapsto 1-t$. This induces the isomorphism $K_p(C_0(0,1)\otimes I_2) \cong K_p(C_0(0,1)\otimes I_2)$ given by $y \mapsto -y$, so the map $K_p(I_1)\oplus K_p(I_2) \to K_p(A)$ sends K-theory classes of the form $(0,y)$ to the image of $-y$ under the natural map $K_p(I_2) \to K_p(A)$. $\square$

5.1.2 Geometric Examples

In this section we discuss examples of C*-algebras and ideals coming from geometry which fit into the Mayer-Vietoris sequence of the previous section. Recall that if X is a second countable locally compact Hausdorff space and $Y \subseteq X$ is a closed subspace then its dual algebra $\mathfrak{D}^*(X)$ has an ideal $\mathfrak{D}^*(Y \subseteq X)$ whose K-theory is isomorphic to the K-theory of $\mathfrak{D}^*(Y)$. If X is the union of two suitable subspaces Y_1 and Y_2, we can use the dual algebra of X and the ideals associated to Y_1 and Y_2 to build a Mayer-Vietoris sequence relating the K-homology groups of X, Y_1, Y_2, and $Y_1 \cap Y_2$. If X has the additional structure of a proper metric space then a similar construction also yields Mayer-Vietoris sequences for the K-theory of the coarse algebra and the K-theory of the structure algebra.

However, not any decomposition $X = Y_1 \cup Y_2$ is suitable for these Mayer-Vietoris sequences. Let $\mathcal{F}(X)$ denote either the dual algebra, coarse algebra, or structure algebra associated to a space X and let $\mathcal{F}(Y \subseteq X)$ denote the ideal associated to a subspace Y.

Definition 5.1.4. *Let a decomposition $X = Y_1 \cup Y_2$ of a proper metric space into subspaces is $\mathcal{F}$-*excisive *if the following conditions hold:*

- $\mathcal{F}(Y_1 \subseteq X) + \mathcal{F}(Y_2 \subseteq X) = \mathcal{F}(X)$
- $\mathcal{F}(Y_1 \cap Y_2 \subseteq X) = \mathcal{F}(Y_1 \subseteq X) \cap \mathcal{F}(Y_2 \subseteq X)$

We will use the phrase *topologically excisive* when $\mathcal{F}$ is the dual algebra, *coarsely excisive* when $\mathcal{F}$ is the coarse algebra, and *uniformly excisive* when $\mathcal{F}$ is the structure algebra. If additionally a group G acts on X in such a way that Y_1 and Y_2 are G-invariant subspaces then we will say that the decomposition $X = Y_1 \cup Y_2$ is *G-equivariantly $\mathcal{F}$-excisive.*

Note that the first of the three conditions in Definition 5.1.4 provides the input for the abstract Mayer-Vietoris sequenc developed above. By Lemma 5.1.1 it suffices to show that $\mathcal{F}(Y_1 \subseteq X)+\mathcal{F}(Y_2 \subseteq X)$ is dense in $\mathcal{F}(X)$. The second condition ensures that there is a straightforward correspondence between geometry and algebra: the intersection of subspaces (geometry) translates into the intersection of ideals (algebra) and vice-versa.

We are now ready to identify explicit geometric conditions which guarantee excisiveness of various decompositions. We begin with the topological excision conditions which require that both subspaces be closed but which impose no constraints on their intersection.

Proposition 5.1.5. *Any decomposition $X = Y_1 \cup Y_2$ of a proper metric space as the union of closed subspaces is topologically excisive.*

Proof. First, we show that $\mathfrak{D}^*(X) = \mathfrak{D}^*(Y_1 \subseteq X) + \mathfrak{D}^*(Y_2 \subseteq X)$.

The containment "$\supseteq$" is vacuous. For the other direction, let f_1 and f_2 be Borel measureable functions on X taking values in $[0,1]$ with the property that $f_1+f_2 = 1$, $f_1|_{X-Y_1} = 0$, and $f_2|_{X-Y_2} = 0$. Such functions exist since the complements of Y_1 and Y_2 are disjoint open sets. Let $P_1 = \rho_X(f_1)$ and $P_2 = \rho_X(f_2)$; note that $g_1 f_1 = 0$ for any $g_1 \in C_0(X - Y_1)$ and $g_2 f_2 = 0$ for any $g_2 \in C_0(X - Y_2)$, so $P_1 \in \mathfrak{D}^*(Y_1 \subseteq X)$ and $P_2 \in \mathfrak{D}^*(Y_2 \subseteq X)$. Since $T = P_1 T + P_2 T$, the proof is complete.

Second, we show that $\mathfrak{D}^*(Y_1 \cap Y_2 \subseteq X) = \mathfrak{D}^*(Y_1 \subseteq X) \cap \mathfrak{D}^*(Y_2 \subseteq X)$.

The containment "$\subseteq$" is trivial since $C_0(X - Y_i) \subseteq C_0(X - Y_1 \cap Y_2)$, so that an operator which is locally compact for $X - Y_1 \cap Y_2$ is automatically locally compact for $X - Y_1$ and $X - Y_2$. To prove "$\supseteq$", begin by observing that $X - Y_1 \cap Y_2$ is the disjoint union of $X - Y_1$ and $X - Y_2$. So given any $f \in C_0(X - Y_1 \cap Y_2)$ we have $f = f_1 + f_2$ where $f_i = f|_{Y_i}$. For any $T \in \mathfrak{D}^*(Y_1 \subseteq X) \cap \mathfrak{D}^*(Y_2 \subseteq X)$ we have that $T f_1 \sim f_1 T \sim T f_2 \sim f_2 T \sim 0$, so it follows that $T(f_1 + f_2) \sim (f_1 + f_2)T \sim 0$. $\square$

We now turn to the coarse excision conditions. Here the topological structure of the two subspaces is irrelevant, but the geometry of their intersection plays a crucial role.

Proposition 5.1.6. *A decomposition $X = Y_1 \cup Y_2$ is coarsely excisive if for every $R > 0$ there exists $S > 0$ such that*

$$B_R(Y_1) \cap B_R(Y_2) \subseteq B_S(Y_1 \cap Y_2)$$

Proof. Let T be a locally compact controlled operator in $C^*(X)$. As in the proof of the previous proposition, choose Borel measureable functions f_1 and f_2 taking values in $[0, 1]$ such that $f_1 + f_2 = 1$, $f_1|_{X-Y_1} = 0$, and $f_2|_{X-Y_2} = 0$. Setting $P_1 = \rho_X(f_1)$ and $P_2 = \rho_X(f_2)$, we have that P_1 and P_2 are locally compact controlled operators which are supported near Y_1 and Y_2, respectively. Thus $P_1 \in C^*(Y_1 \subseteq X)$ and $P_2 \in C^*(Y_2 \subseteq X)$, and we have $T = P_1T + P_2T$. Hence $T \in C^*(Y_1 \subseteq X) + C^*(Y_2 \subseteq X)$, which shows that $C^*(Y_1 \subseteq X) + C^*(Y_2 \subseteq X)$ is dense in $C^*(X)$.

To prove that $C^*(Y_1 \cap Y_2 \subseteq X) = C^*(Y_1 \subseteq X) \cap C^*(Y_2 \subseteq X)$, begin by noting that the containment "$\subseteq$" follows immediately from the definition: any operator supported near $Y_1 \cap Y_2$ is supported near both Y_1 and Y_2. The intersection of two closed ideals in a C* algebra is necessarily equal to their product, so to verify the other containment we must show that if $T_1 \in C^*(Y_1 \subseteq X)$ and $T_2 \in C^*(Y_2 \subseteq X)$ then $T_1T_2 \in C^*(Y_1 \cap Y_2)$. Assume that

$$\mathrm{Supp}(T_1) \subseteq B_{R_1}(Y_1) \times B_{R_1}(Y_1)$$

and

$$\mathrm{Supp}(T_2) \subseteq B_{R_2}(Y_2) \times B_{R_2}(Y_2)$$

We have that

$$\begin{aligned} &\mathrm{Supp}(T_1T_2) \subseteq \mathrm{Supp}(T_1)\mathrm{Supp}(T_2) \\ &\ = \{(p, r) \in X \times X : \ (p, q) \in \mathrm{Supp}(T_1), \ (q, r) \in \mathrm{Supp}(T_2) \text{ for some } q \in X\} \end{aligned}$$

Thus if $(p, r) \in \mathrm{Supp}(T_1T_2)$ then there exists $q \in B_{R_1}(Y_1) \cap B_{R_2}(Y_2)$ such that $(p, q) \in \mathrm{Supp}(T_1)$ and $(q, r) \in \mathrm{Supp}(T_2)$. We have that

$$B_{R_1}(Y_1) \cap B_{R_2}(Y_2) \subseteq B_{R_1+R_2}(Y_1) \cap B_{R_1+R_2}(Y_1)$$

and since the decomposition $X = Y_1 \cup Y_2$ is C^*-excisive it follows that this set is contained in $B_S(Y_1 \cap Y_2)$ for some $S > 0$ (indepednent of p and r). Since T_1 and T_2 are controlled, there is a constant R independent of p, q, and r such that $d(p,q) < R$ and $d(q,r) < R$ and thus p and r both lie in $B_{S+R}(Y_1 \cap Y_2)$ by the triangle inequality. We conclude that $\mathrm{Supp}(T_1T_2) \subseteq B_{S+R}(Y_1 \cap Y_2)$. Passing to the norm closure, this completes the proof. □

Higson, Roe, and Yu define a decomposition which satisfies the hypothesis in Proposition 5.1.6 to be *ω-excisive* in [15]; with this terminology Proposition 5.1.6 asserts that every ω-excisive decomposition is coarsely excisive. The proof is adapted from [15]. Our final result is that a decomposition which satisfies the hypotheses of both Proposition 5.1.5 and Proposition 5.1.6 is uniformly excisive.

Proposition 5.1.7. *Let X be a proper metric space and let $X = Y_1 \cup Y_2$ be a decomposition of X into closed sets Y_1, Y_2 with the property that for every $R > 0$ there exists $S > 0$ such that*

$$B_R(Y_1) \cap B_R(Y_2) \subseteq B_S(Y_1 \cap Y_2)$$

Then the decomposition is uniformly excisive.

Proof. First we show that $D^*(Y_1 \subseteq X) + D^*(Y_2 \subseteq X)$ is dense in $D^*(X)$. As above, choose Borel measureable functions f_1 and f_2 taking values in $[0,1]$ such that $f_1 + f_2 = 1$, $f_1|_{X-Y_1} = 0$, and $f_2|_{X-Y_2} = 0$, and set $P_1\ \rho_X(f_1)$, $P_2 = \rho_X(f_2)$. Then P_1 is a pseudolocal controlled operator which is locally compact for the complement of Y_1 and which is supported in a bounded neighborhood of Y_1 by the condition (5.1.7). Thus $P_1 \in D^*(Y_1 \subseteq X)$, and similarly $P_2 \in D^*(Y_2 \subseteq X)$. Since $T = P_1T + P_2T$ and pseudolocal controlled operators are dense in $D^*(X)$, we have proved the first condition in Definition 5.1.4.

Second, we show that $D^*(Y_1 \cap Y_2 \subseteq X) = D^*(Y_1 \subseteq X) \cap D^*(Y_2 \subseteq X)$. Since both sides of this equation are closed subsets of $D^*(X)$, it suffices to show that they share a common dense set. But we showed that a pseudolocal operator is locally compact for $X - Y_1 \cap Y_2$ if and only if it is locally compact for both $X - Y_1$ and $X - Y_2$ in the proof of Proposition 5.1.5 and we showed that a controlled operator is supported near $Y_1 \cap Y_2$ if and only if it is supported near Y_1 and Y_2 in the proof of Proposition 5.1.6 (the assumption (5.1.7) was used here). Thus a pseudolocal controlled operator is locally compact for $X - Y_1 \cap Y_2$ and supported near $Y_1 \cap Y_2$ if and only if it is locally compact for both $X - Y_1$ and $X - Y_2$ and it is supported near Y_1 and Y_2. □

We conclude this section by remarking that there are straightforward equivariant counterparts of the three results above in the case where G is a group acting freely and properly on X and the subspaces Y_1 and Y_2 are G-invariant. For instance, if $X = Y_1 \cup Y_2$ is a G-invariant decomposition which satisfies the hypotheses of Proposition 5.1.7 then we have

$$D^*_G(X) = D^*_G(Y_1 \subseteq X) + D^*_G(Y_2 \subseteq X)$$

and

$$D^*_G(Y_1 \cap Y_2 \subseteq X) = D^*_G(Y_1 \subseteq X) \cap D^*_G(Y_2 \subseteq X)$$

5.2 The Main Diagram

Let X be a proper metric space equipped with a proper action of a discrete group G of isometries and assume $X = Y_1 \cup Y_2$ is a uniformly excisive G-invariant decomposition (so that Y_1 and Y_2 are closed G-invariant subspaces which satisfy the coarse excision condition explained in the last section). According to Theorem 4.3.25 there is an isomorphism

$$D^*_G(Y \subseteq X)/C^*_G(Y \subseteq X) \cong \mathfrak{D}^*(Y_G \subseteq X_G)/\mathfrak{C}^*(X_G)$$

The subspaces $(Y_1)_G$ and $(Y_2)_G$ are closed and their union is X, so it follows that $X_G = (Y_1)_G \cup (Y_2)_G$ is a topologically excisive decomposition. Hence the decomposition $X = Y_1 \cup Y_2$ is Q^*_G-excisive, where we define $Q^*_G(\cdot) = D^*_G(\cdot)/C^*_G(\cdot)$. In other words,

$$Q^*_G(X) = Q^*_G(Y_1 \subseteq X) + Q^*_G(Y_2 \subseteq X)$$
$$Q^*_G(Y_1 \cap Y_2 \subseteq X) = Q^*_G(Y_1 \subseteq X) \cap Q^*_G(Y_2 \subseteq X)$$

Using the notation $\mathbf{\Omega}(\mathcal{F}, X, Y_1, Y_2)$ for the short exact sequence of C*-algebras

(5.1.1) whose long exact sequence in K-theory is the Mayer-Vietoris sequence for $\mathcal{F}(X) = \mathcal{F}(Y_1 \subseteq X) + \mathcal{F}(Y_2 \subseteq X)$, we have a complex

$$0 \to \mathbf{\Omega}(C^*_G, X, Y_1, Y_2) \to \mathbf{\Omega}(D^*_G, X, Y_1, Y_2) \to \mathbf{\Omega}(Q^*_G, X, Y_1, Y_2) \to 0$$

Passing to K-theory the columns induce Mayer-Vietoris sequences and the rows induce analytic surgery exact sequences, so by the naturality properties of K-theory we obtain:

$$\begin{tikzcd}
K_p(Y_1) \oplus K_p(Y_2) \arrow[r, bend left] \arrow[dr] & K_p(C^*_G(Y_1)) \oplus K_p(C^*_G(Y_2)) \arrow[r, bend left] \arrow[dr] & K_p(C^*_G(X)) \arrow[r, bend left] \arrow[dr] & S_{p-1}(\widetilde{X}, G) \\
K_p(C^*_G(Y_1 \cap Y_2)) \arrow[ur] \arrow[dr] & K_p(X) \arrow[ur] \arrow[dr] & S_{p-1}(\widetilde{Y}_1, G) \oplus S_{p-1}(\widetilde{Y_2}, G) \arrow[ur] \arrow[dr] & K_{p-1}(C^*_G(Y_1 \cap Y_2)) \\
S_p(\widetilde{X}, G) \arrow[ur] \arrow[r, bend right] & S_{p-1}(\widetilde{Y}_1 \cap \widetilde{Y}_2, G) \arrow[ur] \arrow[r, bend right] & K_{p-1}(Y_1 \cap Y_2) \arrow[ur] \arrow[r, bend right] & K_{p-1}(Y_1) \oplus K_{p-1}(Y_2)
\end{tikzcd}$$

Let us call this the *Mayer-Vietoris and analytic surgery diagram.* Consider in particular the following commuting square:

$$\begin{tikzcd}
K_p(X) \arrow[r] \arrow[d] & K_p(C^*_G(X)) \arrow[d] \\
K_{p-1}(Y_1 \cap Y_2) \arrow[r] & K_{p-1}(C^*_G(Y_1 \cap Y_2))
\end{tikzcd}$$

The horizontal maps are generalized index maps and the vertical maps are Mayer-Vietoris boundary maps so the Mayer-Vietoris and analytic surgery diagram related index theory on X to index theory on $Y_1 \cap Y_2$. This is the strategy of our proof of the partitioned manifold index theorem, but in order to implement the strategy we must be able to calculate Mayer-Vietoris boundary maps.

5.3 The Mayer-Vietoris Boundary Map

In this section we will derive an explicit formula which calculates the boundary map in the abstract Mayer-Vietoris sequence. Our strategy is to use the decomposition $A = I_1 + I_2$ to associate to a class $\xi \in K_p(A)$ a related class $\xi' \in K_p(A/(I_1 \cap I_2))$ and argue that $\partial_{MV}(\xi) = \partial(\xi')$ where $\partial : K_p(A/(I_1 \cap I_2)) \to K_{p-1}(I_1 \cap I_2)$ is the boundary map associated to the short exact sequence

$$0 \to I_1 \cap I_2 \to A \to A/(I_1 \cap I_2) \to 0$$

To construct ξ' we use a *partition of unity* for the decomposition $A = I_1 + I_2$ defined as follows:

Definition 5.3.1. *Let A be a unital C^*-algebra and assume $A = I_1 + I_2$ where I_1 and I_2 are closed C^*-ideals in A. A* partition of unity *for this decomposition is a pair (a_1, a_2) where $a_j \in I_j$ are positive elements of A satisfying $a_1^2 + a_2^2 = 1$.*

Let us briefly discuss existence and uniqueness questions related to partitions of unity.

Lemma 5.3.2. *Any decomposition $A = I_1 + I_2$ of a unital C^*-algebra has a partition of unity, and any two partitions of unity are homotopic through partitions of unity.*

Proof. First we show that partitions of unity exist. Choose any two elements $x_1 \in I_1$ and $x_2 \in I_2$ such that $x_1 + x_2 = 1$; replacing x_1 with $\frac{1}{2}(x_1 + x_1^*)$ and x_2 with $\frac{1}{2}(x_2 + x_2^*)$, we can assume that x_1 and x_2 are self-adjoint. Choose any continuous function $f_1 \colon \mathbb{R} \to [0, 1]$ such that $f_1(0) = 0$ and $f_1(1) = 1$, and set $f_2 = 1 - f_1$. Since $x_1 \in I_1$ and $f_1(0) = 0$ we have that $f_1(x_1) \in I_1$. Similarly since $x_2 = 1 - x_1 \in I_2$ and $f_2(1) = 0$ we have $f_2(x_2) \in I_2$. Finally $f_1(x_1) + f_2(x_2) = f_1(x_1) + (1 - f_1)(1 - x_1) = 1$, so the pair (a_1, a_2) given by $a_1 = f_1(x_1)^{1/2}$ and $a_2 = f_2(x_2)^{1/2}$ is a partition of unity.

Now suppose (a_1, a_2) and (b_1, b_2) are two different partitions of unity. For $t \in [0, 1]$ we have that $ta_1 + (1 - t)b_1$ is positive (and similarly for a_2 and b_2, so we set $c_1(t) = (ta_1 + (1 - t)b_1)^{\frac{1}{2}}$ and $c_2(t) = (ta_2 + (1 - t)b_2)^{\frac{1}{2}}$. Then $(c_1(t), c_2(t))$ is a continuous path through partitions of unity which join (a_1, a_2) to (b_1, b_2). $\square$

This lemma allows us to use partitions of unity in computations and it shows that the results of such computations are independent of the choice of partition of unity up to homotopy (and therefore at the level of K-theory). As we shall see, the partitions of unity relevant to decompositions of the dual algebra, coarse algebra, and structure algebra of a space will come from multiplication operators by Borel measurable functions. For now we will simply assume that our decomposition

$A = I_1 + I_2$ has a partition of unity (a_1, a_2) and use it to explicitly calculate $\partial_{MV} : K_p(A) \to K_p(I_1 \cap I_2)$. We begin with the case $p = 1$.

Lemma 5.3.3. *Let $\boldsymbol{u} \in M_n(A)$ be a unitary representing a class in $K_1(A)$ and define $\boldsymbol{v} = a_1\boldsymbol{u} + a_2\boldsymbol{1}$. We have:*

- $\boldsymbol{v} \sim \boldsymbol{1}$ *modulo* $M_n(I_1)$
- $\boldsymbol{v} \sim \boldsymbol{u}$ *modulo* $M_n(I_2)$
- $\boldsymbol{v}$ *is a unitary modulo* $M_n(I_1 \cap I_2)$

Proof. Since $a_1 \in I_1$ and $a_2 \in I_2$ are positive elements satisfying ${a_1}^2 + {a_2}^2 = 1$, it follows that $a_2 = \sqrt{1 - {a_1}^2}$ and hence $a_2 \sim 1$ modulo I_1. Thus $a_1\mathbf{u} + a_2\mathbf{1} \sim \mathbf{1}$ modulo $M_n(I_1)$, and $\mathbf{v} \sim \mathbf{u}$ modulo $M_n(I_2)$ follows similarly. To see that $\mathbf{v}$ is a unitary modulo $M_n(I_1 \cap I_2)$, we use the fact that a_1 and a_2 commute with any element of A modulo $I_1 \cap I_2$ since $a_1 \sim a_2 \sim 1$. Thus:

$$\begin{aligned}
\mathbf{v}\mathbf{v}^* &= (a_1\mathbf{u} + a_2\mathbf{1})(\mathbf{u}^*a_1 + a_2\mathbf{1}) \\
&= ({a_1}^2 + {a_2}^2)\mathbf{1} + a_1\mathbf{u}a_2 + a_2\mathbf{u}^*a_1 \\
&\sim \mathbf{1} + a_1a_2\mathbf{u} + a_1a_2\mathbf{u}^* \\
&\sim \mathbf{1} \text{ modulo } M_n(I_1 \cap I_2)
\end{aligned}$$

Similarly:

$$\begin{aligned}
\mathbf{v}^*\mathbf{v} &= (\mathbf{u}^*a_1 + a_2\mathbf{1})(a_1\mathbf{u} + a_2\mathbf{1}) \\
&\sim ({a_1}^2 + {a_2}^2)\mathbf{1} + \mathbf{u}^*a_1a_2 + a_2a_1\mathbf{u} \\
&\sim \mathbf{1} + a_1a_2\mathbf{u}^* + a_1a_2\mathbf{u} \\
&\sim \mathbf{1} \text{ modulo } M_n(I_1 \cap I_2)
\end{aligned}$$

□

We are now ready to calculate the Mayer-Vietoris boundary map.

Proposition 5.3.4. *Let $\boldsymbol{u}$ and $\boldsymbol{v}$ be as above, and define $\boldsymbol{w} \in M_{2n}(A)$ by*

$$\boldsymbol{w} = \begin{pmatrix} \boldsymbol{v} & -(1 - \boldsymbol{v}\boldsymbol{v}^*)^{\frac{1}{2}} \\ (1 - \boldsymbol{v}^*\boldsymbol{v})^{\frac{1}{2}} & \boldsymbol{v}^* \end{pmatrix}$$

Then we have

$$\partial_{MV}[\boldsymbol{u}] = \left[\boldsymbol{w}^* \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \boldsymbol{w}\right] - \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

In particular, $\partial_{MV}[\boldsymbol{u}] = \partial[\boldsymbol{v}]$ *where* $[\boldsymbol{v}]$ *is the class of* $\boldsymbol{v}$ *in* $K_1(A/(I_1 \cap I_2))$ *and* ∂ *is the boundary map for the short exact sequence*

$$0 \to I_1 \cap I_2 \to A \to A/(I_1 \cap I_2) \to 0$$

Proof. By definition ∂_{MV} is the composition $K_1(A) \cong K_0(S(A)) \to K_0(B) \cong K_0(I_1 \cap I_2)$ where B is the C*-algebra of continuous paths $f : [0,1] \to A$ with the property that $f(0) \in I_1$ and $f(1) \in I_2$ and $K_0(S(A)) \to K_0(B)$ is the inclusion. Recall that under the identification $K_1(A) \cong K_0(S(A))$, $\mathbf{u}$ corresponds to the formal difference $[\mathbf{p}(t)] - [\mathbf{q}(t)]$ of normalized loops of projections in A given by

$$\mathbf{p}(t) = \mathbf{u}(t) \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \mathbf{u}(t)^* \qquad \mathbf{q}(t) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

where $\mathbf{u}(t)$ is any path of unitaries in $M_{2n}(A)$ satisfying

$$\mathbf{u}(0) = \begin{pmatrix} \mathbf{u} & 0 \\ 0 & \mathbf{u}^* \end{pmatrix} \qquad \mathbf{u}(1) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

To prove the proposition we shall make an explicit choice of $\mathbf{u}(t)$ and deform the corresponding normalized loop of projections $\mathbf{p}(t) = \mathbf{u}(t)(1 \oplus 0)\mathbf{u}(t)^*$ through paths of projections over B into the constant path $\mathbf{w}(1 \oplus 0)\mathbf{w}^*$. Here is our choice of $\mathbf{u}(t)$:

- On $[0, \frac{1}{2}]$, let $\mathbf{a}(t)$ be the linear path which satisfies $\mathbf{a}(0) = \mathbf{u}$ and $\mathbf{a}(\frac{1}{2}) = \mathbf{v}$. Define:

  $$\mathbf{u}(t) = \begin{pmatrix} \mathbf{a}(t) & -(1 - \mathbf{a}(t)\mathbf{a}(t)^*)^{\frac{1}{2}} \\ (1 - \mathbf{a}(t)^*\mathbf{a}(t))^{\frac{1}{2}} & \mathbf{a}(t)^* \end{pmatrix}$$

  We have that $\mathbf{a}(t) \sim \mathbf{u}$ modulo I_2 since $\mathbf{v} \sim \mathbf{u}$ and thus $\mathbf{u}(t) \sim \mathbf{u} \oplus \mathbf{u}^*$ for $t \in [0, \frac{1}{2}]$. Additionally, we have $\mathbf{u}(0) = \mathbf{u} \oplus \mathbf{u}^*$ and $\mathbf{u}(\frac{1}{2}) = \mathbf{w}$.

- On $[\frac{1}{2}, 1]$ let $\mathbf{b}(t)$ be the linear path which satisfies $\mathbf{b}(\frac{1}{2}) = \mathbf{v}$ and $\mathbf{b}(1) = 1$.

Define:

$$\mathbf{u}(t) = \begin{pmatrix} \mathbf{b}(t) & -(1-\mathbf{b}(t)\mathbf{b}(t)^*)^{\frac{1}{2}} \\ (1-\mathbf{b}(t)^*\mathbf{b}(t))^{\frac{1}{2}} & \mathbf{b}(t)^* \end{pmatrix}$$

We have that $\mathbf{a}(t) \sim 1$ modulo I_1 since $\mathbf{v} \sim 1$, so $\mathbf{u}(t) \sim 1 \oplus 1$ for $t \in [\frac{1}{2}, 1]$. Additionally, we have $\mathbf{u}(\frac{1}{2}) = \mathbf{w}$ and $\mathbf{u}(1) = 1 \oplus 1$.

Thus we have constructed a path of unitaries $\mathbf{u} : [0,1] \to M_{2n}(A)$ which satisfies:

$$\mathbf{u}(0) = \begin{pmatrix} \mathbf{u} & 0 \\ 0 & \mathbf{u}^* \end{pmatrix}, \qquad \mathbf{u}\left(\frac{1}{2}\right) = \mathbf{w}, \qquad \mathbf{u}(1) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Set $\mathbf{p}(t) = \mathbf{u}(t)(1 \oplus 0)\mathbf{u}(t)^*$ as above. Our next step is to construct a homotopy of paths $\mathbf{p}_s(t)$, $s \in [0,1]$, with the property that $\mathbf{p}_s$ is a projection in $M_{2n}(\widetilde{B})$ for every s, $\mathbf{p}_0(t) = \mathbf{p}(t)$, and $\mathbf{p}_1(t) = \mathbf{w}(1 \oplus 0)\mathbf{w}^*$. We begin by defining a homotopy of unitary paths $\mathbf{u}_s(t)$ which deforms $\mathbf{u}(t)$ into the constant path $\mathbf{w}$. Define:

$$\mathbf{u}_s(t) = \begin{cases} \mathbf{u}(t + \frac{s}{2}) & 0 \le t \le \frac{1-s}{2} \\ \mathbf{w} & \frac{1-s}{2} \le t \le \frac{1+s}{2} \\ \mathbf{u}(t - \frac{s}{2}) & \frac{1+s}{2} \le t \le 1 \end{cases}$$

Finally, define $\mathbf{p}_s(t) = \mathbf{u}_s(t)(1 \oplus 0)\mathbf{u}_s(t)^*$. Since $0 \le \frac{s}{2} \le \frac{1}{2}$ we have $\mathbf{u}_s(0) = \mathbf{u}(\frac{s}{2}) \sim \mathbf{u} \oplus \mathbf{u}^*$ modulo I_2 and thus $\mathbf{p}_s(0) = \mathbf{u}_s(0)(1 \oplus 0)\mathbf{u}_s(0)^* \sim 1 \oplus 0$ modulo I_2. Similarly $\mathbf{p}_s(1) \sim 1 \oplus 0$ modulo I_1. Thus $\mathbf{p}_s(0) \in M_{2n}(\widetilde{I_2})$ and $\mathbf{p}_s(1) \in M_{2n}(\widetilde{I_1})$ from which it follows that $\mathbf{p}_s \in M_{2n}(\widetilde{B})$. Clearly $\mathbf{p}_0(t) = \mathbf{u}(t)(1 \oplus 0)\mathbf{u}(t)^* = \mathbf{p}(t)$ and $\mathbf{p}_1(t) = \mathbf{w}(1 \oplus 0)\mathbf{w}^*$, so the proof is complete. □

Remark 5.3.5. *This formula can easily be adapted to the nonunital case so long as we are more flexible with the notion of a partition of unity. Let $A = I_1 + I_2$ be a non-unital C*-algebra which is the sum of two ideals and let A' be a unital C*-algebra which contains A as an ideal. In this setting we will take a partition of unity to mean a pair (a_1, a_2) of positive elements of A' such that $a_1 A \subseteq I_1$, $a_2 A \subseteq I_2$, and $a_1^2 + a_2^2 = 1$. The calculation above carries through verbatim with this notion of partition of unity.*

The specific example of this which will be most important to us is the case where $A = I_1 + I_2$ is a unital C-algebra equipped with a partition of unity (a_1, a_2) and we* *consider $SA = SI_1 + SI_2$. Viewing $SA = C_0(0,1) \otimes A$ as an ideal in $C(S^1) \otimes A$, we see that $(1 \otimes a_1, 1 \otimes a_2)$ form a partition of unity.*

We can produce a similar formula for the Mayer-Vietoris boundary map in the other degree. However, we must assume that $A = I_1 + I_2$ admits a partition of unity (a_1, a_2) such that a_1 and a_2 are projections; this is required to ensure that if $\mathbf{p}$ is a projection over A then $a_1\mathbf{p} + a_2$ is a projection over $A/I_1 \cap I_2$.

Corollary 5.3.6. *In the situation above, the Mayer-Vietoris boundary map $\partial_{MV} : K_0(A) \to K_1(I_1 \cap I_2)$ satisfies $\partial_{MV}[\boldsymbol{p}] = \partial[a_1\boldsymbol{p} + a_2]$ where ∂ is the boundary map associated to the short exact sequence*

$$0 \to I_1 \cap I_2 \to A \to A/I_1 \cap I_2 \to 0$$

Proof. Our strategy is to reduce to Proposition 5.3.4 using Bott periodicity. Consider the following diagram:

$$\begin{tikzcd} K_0(A) \arrow[r, "\mathrm{Bott}"] \arrow[d] & K_1(SA) \arrow[r, "\partial_{MV}"] \arrow[d] & K_1(I_1 \cap I_2) \arrow[d, equal] \\ K_0(A/I_1 \cap I_2) \arrow[r, "\mathrm{Bott}"] & K_1(SA/S(I_1 \cap I_2)) \arrow[r, "\partial"] & K_1(I_1 \cap I_2) \end{tikzcd}$$

Note that $\partial_{MV} : K_0(A) \to K_1(I_1 \cap I_2)$ is the composition of the two maps in the top row and $\partial : K_0(A/I_1 \cap I_2) \to K_1(I_1 \cap I_2)$ is the composition of the two maps in the bottom row. Let the vertical map $K_0(A) \to K_0(A/I_1 \cap I_2)$ be given by $[\mathbf{p}] \mapsto [a_1\mathbf{p} + a_2]$, and similarly let $K_1(SA) \to K_1(SA/S(I_1 \cap I_2))$ be the map $[\mathbf{u}_t] \mapsto [a_1\mathbf{u}_t + a_2]$ where $\mathbf{u}_t$ is a normalized loop of unitaries over A. To prove the corollary it suffices to show that the diagram commutes.

The assumption that a_1 and a_2 are projections guarantees that the two vertical maps are induced by $*$-homomorphisms $A \to A/I_1 \cap I_2$ and $SA \to SA/S(I_1 \cap I_2)$, respectively, and thus it follows from the naturality of the Bott map that the left square commutes. The right square commutes by Proposition 5.3.4, so the proof is complete. $\square$

5.4 The Partitioned Manifold Index Theorem

We are now ready to prove the main result of this thesis: the partitioned manifold index theorem. In order to state the theorem we need to settle some terminology and carry out some preliminary calculations.

Definition 5.4.1. *Let M be a smooth manifold and let N be a submanifold of codimension* 1*. M is* partitioned *by N if M is the union of two submanifolds M^+ and M^- with common boundary N.*

Assume now that M is a complete Riemannian manifold (in particular a proper metric space) and that N is compact. Let G be a countable discrete group and let $\widetilde{M} \to M$ be a locally isometric G-cover of M. Define a map $M \to \mathbb{R}$ by the formula $x \mapsto \operatorname{dist}(x, N)$; this is a coarse map since M is a length metric space and it lifts to a G-invariant coarse map $\widetilde{M} \to \mathbb{R} \times G$. Thus it induces a homomorphism $K_p(C^*_G(M)) \to K_p(C^*_G(\mathbb{R} \times G))$. Let us use the Mayer-Vietoris sequence to calculate $K_p(C^*_G(\mathbb{R} \times G))$.

Lemma 5.4.2. *For any proper metric space Y, the Mayer-Vietoris boundary map*

$$\partial_{MV} \colon K_p(C^*(\mathbb{R}) \times Y) \to K_{p-1}(C^*(\{0\} \times Y))$$

associated to the coarsely excisive decomposition $\mathbb{R} \times Y = \mathbb{R}^{\geq 0} \times Y \cup \mathbb{R}^{\leq 0} \times Y$ is an isomorphism.

Proof. The Mayer-Vietoris sequence takes the form:

$$\begin{aligned} &K_p(C^*(\mathbb{R}^{\geq 0} \times Y)) \oplus K_p(C^*(\mathbb{R}^{\leq 0} \times Y)) \\ &\to K_p(C^*(\mathbb{R} \times Y)) \xrightarrow{\partial_{MV}} K_{p-1}(C^*(\{0\} \times Y)) \\ &\to K_{p-1}(C^*(\mathbb{R}^{\geq 0} \times Y)) \oplus K_{p-1}(C^*(\mathbb{R}^{\leq 0} \times Y)) \end{aligned}$$

so ∂_{MV} is an isomorphism since $K_p(C^*(\mathbb{R}^{\leq 0} \times Y)) = K_p(C^*(\mathbb{R}^{\geq 0} \times Y)) = 0$. □

If Y carries a free and proper G action and we allow G to act trivially on $\mathbb{R}$ then the same argument shows that $\partial_{MV} : K_p(C^*_G(\mathbb{R} \times Y)) \to K_{p-1}(C^*_G(\{0\} \times Y))$ is an isomorphism. Specializing to the case where $Y = G$ (equipped with a G-invariant metric) we conclude that $K_p(C^*_G(\mathbb{R}\times G)) \cong K_{p-1}(C^*_r(G))$, where the isomorphism is given by the Mayer-Vietoris boundary map.

Definition 5.4.3. *Let M be a complete Riemannian manifold partitioned by a compact hypersurface N and let $\widetilde{M}$ be a G-cover of M where G is a countable discrete group. The* partitioned index map *is the composition*

$$Ind^G_{M,N}\colon K_p(M) \to K_p(C^*_G(\mathbb{R}\times \widetilde{M})) \cong K_{p-1}(C^*_r(G))$$

Recall that the boundary map in the analytic surgery exact sequence associated to N defines a generalized index map:

$$\mathrm{Ind}^G_N\colon K_{p-1}(N) \to K_{p-1}(C^*_r(G))$$

The partitioned manifold index theorem relates the partitioned index of an appropriate operator on M to the index of a corresponding operator on N. The following definition specifies conditions on an operator on M which make this relationship possible.

Definition 5.4.4. *Let M be a smooth manifold partitioned by a hypersurface N, let $S_M \to M$ be a smooth p-multigraded vector bundle over M, and let $S_N \to N$ be a smooth $(p-1)$-multigraded vector bundle over N. Let D_M and D_N be p- and $(p-1)$-multigraded differential operators acting on smooth sections of S_M and S_n, respectively. Say that D_M is* partitioned by D_N *if there is a collaring neighborhood $U \cong (-1,1)\times N$ of N in M with the following properties:*

- *$S_M|_U \cong S_{(-1,1)}\hat{\otimes}S_N$ where $S_{(-1,1)}$ is the standard 1-multigraded spinor bundle over $(-1,1)$ (see chapter 2).*
- *$D_M = D_{(-1,1)}\hat{\otimes}1 + 1\hat{\otimes}D_N$ where $D_{(-1,1)}$ is the spinor Dirac operator on $S_{(-1,1)}$.*

If D_M is partitioned by some differential operator on some smooth vector bundle over N then we will say that D_M is partitioned near N

We are now ready to state the main theorem:

Theorem 5.4.5 (The Partitioned Manifold Index Theorem)**.** *Let M be a complete Riemannian manifold and let $\widetilde{M}$ be a locally isometric G-cover of M where G is a countable discrete group. Suppose M is partitioned by a hypersurface N and $\widetilde{M}$ is partitioned by the lift $\widetilde{N}$ of N. If D_M is a p-multigraded Dirac-type operator on M which is partitioned by a $(p-1)$-multigraded Dirac-type operator D_N on N then*

$$Ind^G_{M,N}[D_M] = Ind^G_N[D_N]$$

in $K_{p-1}(C^*_r(G))$.

Earlier proofs of this theorem, such as in [20] or [8], deal only with the case where G is the trivial group $\{1\}$. In this setting it is possible to give a more concrete description of the map $\mathrm{Ind}_{M,N}\colon K_p(M) \to K_{p-1}(C^*_r(\{1\}))$. There is a natural isomorphism $C^*_r(\{1\}) \cong \mathbb{K}$ where $\mathbb{K}$ is the C*-algebra of compact operators on a separable Hilbert space, so the case $p = 0$ is vacuous since $K_1(\mathbb{K}) = 0$. In the case $p = 1$ the K-homology class of D_M can be represented by a certain unitary operator U on a certain Hilbert space which carries an ample representation of $C_0(M)$ (see Chapter 2), and it can be shown that the operator $\varphi^+U + \varphi^-$ is Fredholm where $\varphi^\pm$ is the multiplication operator by the characteristic function of $M^\pm$. Under the standard isomorphism $K_0(\mathbb{K}) \cong \mathbb{Z}$ we have that $\mathrm{Ind}_{M,N}[D_M] = \mathrm{Index}(\varphi^+U + \varphi^-)$. This is the definition of the partitioned index in [8], and the proof that it is equivalent to our definition will be verified during the course of our proof of Theorem 5.4.5.

The strategy of the proof is to fit the relevant index maps into a commutative diagram with the boundary map in the Mayer-Vietoris sequence for K-homology:

$$\begin{tikzcd}
K_p(M) \arrow[dr, "\mathrm{Ind}^G_{M,N}"] \arrow[dd, "\partial_{MV}"'] & \\
& K_{p-1}(C^*_r(G)) \\
K_{p-1}(N) \arrow[ur, "\mathrm{Ind}^G_N"'] &
\end{tikzcd}$$

The result would then follow if we showed that ∂_{MV} sends the K-homology class of D_M to the K-homology class of D_N. It is here that the technical assumptions on D_M in the statement of Theorem 5.4.5 play a crucial role. This is also where our calculations with suspension maps in Chapter 2 and our calculations with Mayer-Vietoris boundary maps in the last section of this chapter enter into the proof.

Lemma 5.4.6. *In the setting of Theorem 5.4.5 we have*

$$\partial_{MV}[D_M] = [D_N]$$

where $[D_M]$ *is the class in* $K_p(M)$ *determined by* D_M*,* $[D_N]$ *is the class in* $K_{p-1}(N)$ *determined by* D_N*, and*

$$\partial_{MV} \colon K_p(M) \to K_{p-1}(M^+ \cap M^-) = K_{p-1}(N)$$

is the boundary map in the Mayer-Vietoris sequence in K-homology associated to the topologically excisive decomposition $M = M^+ \cup M^-$ *determined by the partitioning of* M*.*

Proof. Let φ^+ and φ^- denote the multiplication operators by the characteristic functions of M^+ and M^-, respectively. Observe that (φ^+, φ^-) defines a partition of unity for the decomposition $\mathfrak{D}^*(M) = \mathfrak{D}^*(M^+ \subseteq M) + \mathfrak{D}^*(M^- \subseteq M)$. Thus according to our calculation of the Mayer-Vietoris boundary map, ∂_{MV} is the composition of the map $K_p(M) \to K_p(M, N)$ which sends an operator F representing a class in $K_p(M)$ to the class of the operator $\varphi^+ F + \varphi^-$ in $K_p(M, N)$ and the K-homology boundary map $K_p(M, N) \to K_p(N)$. The excision map $K_p(M, N) \to K_p(M^- - N) \oplus K_p(M^+ - N)$ given by $[F] \mapsto ([\varphi^- F], [\varphi^+ F])$ and the boundary map $K_p(M^- - N) \oplus K_p(M^+ - N) \to K_{p-1}(N)$ given by $([F^-], [F^+]) \mapsto \partial[F^+]$ fit into a commutative diagram with ∂_{MV}:

$$\begin{tikzcd}
K_p(M) \arrow[r] \arrow[rr, bend left] & K_p(M, N) \arrow[r] \arrow[d] & K_{p-1}(N) \\
& K_p(M^- - N) \oplus K_p(M^+ - N) \arrow[ur] &
\end{tikzcd}$$

The image of $[D_M]$ in $K_p(M^- - N) \oplus K_p(M^+ - N)$ is $(0, [\varphi^+ D_M])$, and $\varphi^+ D_M$ satisfies the hypotheses of Theorem 3.5.9 by our assumptions on D_M. Thus the image of $(0, [\varphi^+ D_M])$ is $[D_N]$, as desired. $\square$

We are now ready to prove the main theorem:

Proof of Theorem 5.4.5. Using the Mayer-Vietoris and analytic surgery diagram together with Lemma 5.4.2, there is a commuting diagram:

$$\begin{tikzcd} K_p(M) \arrow[r] \arrow[d, "\partial_{MV}"] \arrow[dr, "\mathrm{Ind}^G_{M,N}"] & K_p(C^*_G(\widetilde{M} \times G)) \arrow[d, "\partial_{MV}"] \\ K_{p-1}(N) \arrow[r, "\mathrm{Ind}^G_N"] & K_{p-1}(C^*_r(G)) \end{tikzcd}$$

Commutativity of this diagram together with Lemma 5.4.6 completes the proof. $\square$

We conclude this chapter with a higher dimensional generalization of Theorem 5.4.5. The basic idea of Theorem 5.4.5 is to compute the index of an operator on a non-compact manifold M by localizing the calculation to a compact partitioning hypersurface. This is possible because the partitioning structure provides a mechanism for relating the uniform geometry of M to the uniform geometry of $\mathbb{R}$. This in turn suggests that a similar result can be proved for a manifold whose uniform geometry is related to that of $\mathbb{R}^k$.

Definition 5.4.7. *Let M be a smooth manifold and let N be a submanifold of codimension k. A k-*partitioning map *for the pair (M, N) is a coarse submersion $F\colon M \to \mathbb{R}^k$ such that $N = F^{-1}(0)$. Say that M is k-*partitioned *by N if there exists a k-partitioning map for (M, N).*

Note that if M is k-partitioned by N then the pre-image of a sufficiently small neighborhood of $0 \in \mathbb{R}^k$ yields a collaring neighborhood $\mathcal{U} \cong (-1,1)^k \times N$. Just as the index theorem for partitioned manifolds made crucial use of the spinor Dirac operator on $(-1, 1)$, the index theorem for k-partitioned manifolds will require a canonical operator on $(-1,1)^k$. For the purpose of the next definition, if $S_M \to M$ and $S_N \to N$ are vector bundles equipped with differential operators D_M and D_N, respectively, then the product of D_M and D_N is the differential operator $D_M \otimes 1 + 1 \otimes D_N$ on $S_M \otimes S_N \to M \times N$.

Definition 5.4.8. *The* (complex) spinor bundle *over the open cube* $(-1,1)^k \subseteq \mathbb{R}^k$ *is the trivial k-multigraded vector bundle* $S_{(-1,1)^k} = S_{(-1,1)} \hat{\otimes} \ldots \hat{\otimes} S_{(-1,1)}$ *(k factors) where $S_{(-1,1)}$ is the spinor bundle over $(-1,1)$. The* (complex) spinor Dirac operator *on $(-1,1)^k$ is the k-multigraded differential operator acting on smooth sections of $S_{(-1,1)^k}$ given by forming the k-fold product of the spinor Dirac operator $D_{(-1,1)}$ on $(-1,1)$.*

There are certainly more concrete ways to define the spinor bundle and the spinor Dirac operator, but however they are defined $D_{(-1,1)^k}$ is a Dirac-type operator. Observe that $(-1,1)^{k-1}$ partitions $(-1,1)^k$ in a natural way, and relative to that decomposition $D_{(-1,1)^{k-1}}$ partitions $D_{(-1,1)^k}$ in the sense of Definition 5.4.4.

Definition 5.4.9. *Let M be a smooth manifold and suppose N is a submanifold of M of codimension k. Let D_M be a p-multigraded differential operator acting on a smooth p-multigraded vector bundle $S_M \to M$, $p \geq k$, and let D_N be a $(p-k)$-multigraded differential operator acting on a smooth $(p-k)$-multigraded vector bundle $S_N \to N$. Say that D_M is* k-partitioned *by D_N if there is a collaring neighborhood $U \cong (-1,1)^k \times N$ of N in M with the following properties:*

- $S_M|_U \cong S_{(-1,1)^k} \hat{\otimes} S_N$
- $D_M = D_{(-1,1)^k} \hat{\otimes} 1 + 1 \hat{\otimes} D_N$

Suppose M is k-partitioned by a compact hypersurface N and $\widetilde{M}$ is a locally isometric G-cover of M where G is a countable discrete group. Suppose further that the k-partitioning map $F\colon M \to \mathbb{R}^k$ lifts to a k-partitioning map $\widetilde{F}\colon \widetilde{M} \to \mathbb{R}^k$ for the pair $(\widetilde{M}, \widetilde{N})$. Then $\widetilde{F}$ induces a map

$$\widetilde{F}_*\colon K_p(C_G^*(\widetilde{M})) \to K_p(C_G^*(\mathbb{R}^k \times G))$$

By Lemma 5.4.2 and induction on k, there is an isomorphism $K_p(C_G^*(\mathbb{R}^k \times G)) \cong K_{p-k}(C_r^*(G))$ given by iterated Mayer-Vietoris boundary maps. Thus we have a k-partitioned index map:

$$\mathrm{Ind}_{M,N}^G\colon K_p(M) \to K_p(C_G^*(\widetilde{M})) \cong K_{p-k}(C_r^*(G))$$

analogous to the partitioned index map defined above. Our main result about k-partitioned manifolds computes this index map.

Proposition 5.4.10. *Let M be a complete Riemannian manifold and let $\widetilde{M}$ be a locally isometric G-cover of M where G is a countable discrete group. Let N be a submanifold of M which lifts to a submanifold $\widetilde{N}$ of $\widetilde{M}$, and suppose there is a k-partitioning map $F\colon M \to \mathbb{R}^k$ for the pair (M, N) which lifts to a k-partitioning map $\widetilde{F}\colon \widetilde{M} \to \mathbb{R}^k$ for the pair $(\widetilde{M}, \widetilde{N})$. If D_M is a p-multigraded Dirac-type operator on M, $p \geq k$, which is k-partitioned by a $(p-k)$-multigraded Dirac-type operator D_N on N then*

$$Ind^G_{M,N}[D_M] = Ind^G_N[D_N]$$

in $K_{p-k}(C^*_r(G))$.

Proof. We use induction on k; the base case is simply Theorem 5.4.5, so assume $k \geq 2$. Since F is a submersion the sets $M^+ = F^{-1}(\mathbb{R}^{k-1} \times \mathbb{R}^{\geq 0})$ and $M^- = F^{-1}(\mathbb{R}^{k-1} \times \mathbb{R}^{\leq 0})$ are submanifolds with boundary which partition M, and the partitioning hypersurface $N' = M^+ \cap M^-$ is $(k-1)$-partitioned by N. Moreover D_M is partitioned by the $(k-1)$-multigraded operator $D_{N'} = D_{(-1,1)^{k-1}} \hat{\otimes} 1 + 1 \hat{\otimes} D_N$ and $D_{N'}$ is $(k-1)$-partitioned by D_N. By the induction hypothesis it suffices to show that $\mathrm{Ind}^G_{M,N}[D_M] = \mathrm{Ind}_{N',N}[D_{N'}]$ in $K_{p-k}(C^*_r(G))$. As in the proof of Theorem 5.4.5 there is a commutative diagram

$$\begin{tikzcd}
K_p(M) \arrow[d, "\partial_{MV}"'] \arrow[dr, "\mathrm{Ind}^G_{M,N}"] & \\
& K_{p-k}(C^*_r(G)) \\
K_{p-1}(N') \arrow[ur, "\mathrm{Ind}^G_{N',N}"'] &
\end{tikzcd}$$

so the result follows from Lemma 5.4.6. ∎

Chapter 6

Positive Scalar Curvature Invariants

In this final chapter we discuss applications of the partitioned manifold index theorem and of the machinery that we used to prove it to problems in Riemannian geometry concerning manifolds with positive scalar curvature metrics. If M is a Riemannian manifold and $p \in M$ then the scalar curvature $\kappa(p)$ of M at p is a number which measures the difference between the volume of a small metric ball centered at p and the volume of a Euclidean ball with the same radius. The scalar curvature function κ of a surface S is precisely twice the Gaussian curvature, so it determines the metric on S up to isometry. In higher dimensions, however, the scalar curvature of a Riemannian manifold is a very weak invariant:

Theorem 6.0.11. *Let M be a smooth compact manifold of dimension at least 3. Then any smooth function on M which takes negative values somewhere is the scalar curvature function for some Riemannian metric on M.*

Proof. [13] $\square$

This theorem leaves open the possibility that there are obstructions to the existence of metrics whose scalar curvature function is everywhere positive. The first such obstructions were identified by Lichnerowicz using the Atiyah-Singer index theorem:

Theorem 6.0.12. *Let M be a compact Riemannian spin manifold whose scalar curvature function is everywhere positive. Then the $\hat{A}$-genus of M vanishes.*

A *spin structure* on a manifold is a generalized orientation; the important observation for our present purposes is that a choice of spin structure on a Riemannian manifold M determines a canonical choice of differential operator on M called the *spinor Dirac operator*. The spinor Dirac operator on M is a Dirac-type operator in the sense we have defined and hence it is Freholm if M is compact. Atiyah and Singer calculated using their index theorem that its Fredholm index is the $\hat{A}$-genus of M, and Lichnerowicz proved that if the scalar curvature function on M is everywhere positive then the index of the spinor Dirac operator is zero. Thus Theorem 6.0.12 follows; for more detail see [14].

As we shall see, the index theorem for k-partitioned manifolds proved in the last chapter provides obstructions to the existence of metrics on appropriate non-compact manifolds whose positive scalar curvature function is bounded from below by a positive constant. Additionally, it yields a new proof of an important theorem of Gromov and Lawson which asserts that a manifold which admits a metric of non-positive sectional curvature cannot admit a metric of positive scalar curvature.

The Mayer-Vietoris and analytic surgery diagram introduced in the previous chapter may also have interesting applications to the theory of positive scalar curvature obstructions. Recall that the analytic surgery exact sequence takes the form:

$$\ldots \to S_p(\widetilde{M}, G) \to K_p(M) \to K_p(C_r^*(G)) \to \ldots$$

where M is a compact manifold, $\widetilde{M}$ is the universal cover of M, and G is the fundamental group of M. If D is a Dirac-type operator on M whose index in $K_p(C_r^*(G))$ vanishes then by exactness the K-homology class of D lifts to a class in the analytic structure group $S_p(\widetilde{M}, G)$. This applies in particular to spinor Dirac operators on compact spin manifolds with positive scalar curvature metrics, and indeed we shall see that Lichnerowicz's argument shows how to construct an explicit element of the analytic structure group corresponding to a positive scalar curvature metric. Our Mayer-Vietoris sequence for the analytic structure group may help manipulate these positive scalar curvature invariants and prove results analogous to the partitioned manifold index theorem.

6.1 The Lichnerowicz Vanishing Theorem

In this section we take a deeper look into the statement and proof of Theorem 6.0.12. A thorough discussion would require a detour into algebraic topology and differential geometry, so we will give references where appropriate. Our main goal is to explain how a positive scalar curvature metric on a compact spin manifold M determines a class in the analytic structure group associated to M.

6.1.1 The Spinor Dirac Operator

The theory of positive scalar curvature invariants is based on a specific Dirac-type operator which is determined by a choice of Riemannian metric and spin structure on a smooth manifold M. In order to construct this operator and the bundle on which it acts we need to make a brief detour into the theory of Clifford algebras.

Recall that if V is a real vector space equipped with a Euclidean inner product then the (real) Clifford algebra $Cl(V)$ associated to V is the universal $\mathbb{R}$-algebra which contains V as a linear subspace and which satisfies the relation $v \cdot v = -\langle v, v\rangle 1$ for every $v \in V$. $Cl(V)$ is generated as an algebra by any orthonormal basis for V (together with 1) and it is isomorphic as a vector space (but not as an algebra) to the exterior algebra $\bigwedge^* V$. The Clifford algebra associated to $\mathbb{R}^n$ with the standard inner product is denoted $\mathbb{R}_n$; note that the Clifford algebra associated to any n-dimensional Euclidean space is isomorphic to $\mathbb{R}_n$.

$Cl(V)$ has the structure of a $\mathbb{Z}/2\mathbb{Z}$-graded algebra: using an orthonormal basis for V as a generating set, $Cl(V)^+$ consists of those elements which can be written as the product of an even number of generators and $Cl(V)^-$ consists of those elements which can be written as the product of an odd number of generators. It is straightforward to check that $Cl(V) = Cl(V)^+ \oplus Cl(V)^-$ and that the decomposition is independent of the choice of orthonormal basis.

Definition 6.1.1. *Let M be a Riemannian manifold. A* Dirac bundle *over M is a graded Euclidean vector bundle $S \to M$ equipped with an $\mathbb{R}$-linear bundle map $c\colon T^*M \to End(S)$ which for each $p \in M$ extends to a graded representation $c_p\colon Cl(T_p^*M) \to End(S_p)$.*

There is a natural way to construct Dirac bundles locally. Let $\mathcal{U}$ be a coordinate neighborhood for M and let $\xi_1, \ldots, \xi_n$ be an orthonormal frame for $T^*M|_{\mathcal{U}}$. Consider the trivial bundle $\mathbb{R}_n \times \mathcal{U}$ over $\mathcal{U}$; there is a unique bundle map $T^*M|_{\mathcal{U}} \to \mathrm{End}(\mathbb{R}_n \times \mathcal{U})$ which sends ξ_i to the left multiplication map by the ith generator for $\mathbb{R}_n$ (using the standard basis of $\mathbb{R}^n$ as the generating set), and this map clearly gives $\mathbb{R}_n \times \mathcal{U}$ the structure of a Dirac bundle over $\mathcal{U}$. Moreover $\mathbb{R}_n \times \mathcal{U}$ has n multigrading operators $\varepsilon_1 \ldots \varepsilon_n$ corresponding to right multiplication by the generators for $\mathbb{R}_n$.

Definition 6.1.2. *Let M be a Riemannian manifold of dimension n. A* spinor bundle *over M is a n-multigraded Dirac bundle over M which is locally isomorphic to the trivial Dirac bundle $\mathbb{R}_n \times \mathcal{U}$ defined using a local orthonormal frame for T^*M.*

There are global topological obstructions to gluing the trivial Dirac bundles $\mathbb{R}_n \times U$ into a global spinor bundle; it turns out that this is possible if and only if the first two Stiefel-Whitney classes of M vanish (in particular a necessary condition for M to admit a spinor bundle is that M is orientable). If M is a smooth manifold and S, S' are two spinor bundles over M (defined using possibly distinct Riemannian metrics), one says that S and S' are *concordant* if there is a spinor bundle over $M \times \mathbb{R}$ which restricts to $S \hat{\otimes} S_{\mathbb{R}}$ over some open interval and $S' \hat{\otimes} S_{\mathbb{R}}$ over some other open interval, where $S_{\mathbb{R}}$ is the trivial spinor bundle over $\mathbb{R}$.

Definition 6.1.3. *A* spin structure *on a smooth manifold M is a concordance class of spinor bundles over M, and a* spin manifold *is a smooth manifold equipped with a spin structure.*

There is also a counterpart of these notions over the complex numbers: a complex spinor bundle is a bundle which is locally modelled on the complex Clifford algebra $\mathbb{C}_n$ instead of $\mathbb{R}_n$, and a *spin^c structure* on a manifold is a concordance class of complex spinor bundles. We already encountered the complex spinor bundle on Euclidean space in previous chapters.

There is another definition of real and complex spinor bundle involving the representation theory of the Lie groups $Spin(n)$ and $Spin^c(n)$. These definitions are not equivalent to the ones given here, but there is an algebraic procedure which passes back and forth between our definition and the representation theoretic definition. This procedure begins with a p-multigraded vector bundle and reduces it to a $(p-2)$-multigraded (complex) vector bundle with a smaller dimension; iterating the construction, it reduces any p-multigraded bundle to either a graded or ungraded bundle (depending on the parity of p).

So let S be a p-multigraded vector bundle. Then $X = i\varepsilon_{p-1}\varepsilon_p$ is an even self-adjoint bundle morphism which squares to 1, and the $+1$-eigenbundle for X is a $(p-2)$-multigraded Dirac bundle S'. Thus we can reduce from a p-multigraded bundle to a $(p-2)$-multigraded bundle. Conversely, if S' is any $(p-2)$-multigraded bundle then dilate to a p-multigraded bundle $S'' = S' \oplus (S')^{opp}$ (where $(S')^{opp}$ is S' with the opposite grading) by defining new multigrading operators

$$\varepsilon'_{n-1} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \qquad \varepsilon'_n = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$$

One can easily check that the procedures of reduction and dilation are inverses of one another (up to isomorphism of multigraded bundles).

It follows that the spinor bundle S over M can be reduced to a bundle S_{red} whose dimension is half that of S and that this bundle is graded if M is odd dimensional and graded if M is even dimensional. Note that while the spinor bundle is a real vector bundle, the corresponding reduced bundle is a complex vector bundle.

We are now ready to introduce the spinor Dirac operator, a differential operator naturally associated to a spinor bundle. The construction uses the *spin connection*, a connection on a spinor bundle S over a Riemannian manifold M naturally associated to the Levi-Civita connection which is compatible with both the Dirac structure and the Euclidean structure on S.

Definition 6.1.4. *Let M be a spin manifold and let $S \to M$ be its spinor bundle. The* (real) spinor Dirac operator *on $S \to M$ is the n-multigraded first order differential operator*

$$Ds = \sum_{i=1}^{n} c(\xi_i)\nabla_{X_i} s$$

where $X_1, \ldots, X_n$ is a local orthonormal frame for TM, $\xi_1, \ldots, \xi_n$ is the dual frame for T^*M, and $c\colon T^*M \to End(S)$ is given by the Dirac structure on S.

Note that the symbol of D is given by

$$\sigma_D(p,\xi) = c(\xi)$$

so since $c(\xi)^2 = -\|\xi\|^2$ it follows that D is elliptic and has finite propagation speed. Therefore D is a Dirac-type operator and hence it determines a class in $K_n(M)$ if M is complete.

In fact D determines a class in the *real K-homology* of M, that is in the homology theory that naturally corresponds to real K-theory. Indeed, the real K-homology class of D plays the role of the fundamental class of M. For our purposes, the importance of the spinor Dirac operator is based on its connection with scalar curvature.

6.1.2 The Spinor Dirac Operator and Scalar Curvature

The main result which relates Dirac operators and scalar curvature is the following calculation:

Theorem 6.1.5 (Lichnerowicz)**.** *Let M be a Riemannian spin manifold and let D be the spinor Dirac operator on M. Then*

$$D^2 = \nabla^*\nabla + \frac{1}{4}\kappa$$

where ∇ is the spin connection and κ is the scalar curvature function on M.

Proof. See the textbook [14]. □

The vanishing result for spin manifolds with positive scalar curvature stated above follows from this formula together with the following special case of the Atiyah-Singer index theorem:

Theorem 6.1.6. *Let M be a compact Riemannian spin manifold and let D be the spinor Dirac operator on M. Then*

$$\mathit{Index}(D) = \hat{A}(M)$$

Proof. Reference □

Proof of Theorem 6.0.12. Suppose M is a compact Riemannian spin manifold with positive scalar curvature. Let S be the spinor bundle over M, ∇ the spin connection on S, and D the spinor Dirac operator. The operator $\nabla^*\nabla$ is a positive essentially self-adjoint unbounded operator on $L^2(M;S)$ so its spectrum consists of nonnegative real numbers. Since κ is positive on M it attains a minimum value $\varepsilon > 0$ somewhere on M, and thus the spectrum of D^2 is a subset of the interval $[\frac{1}{4}\varepsilon, \infty)$. Consequently there is a gap around 0 in the spectrum of D which imples that the graded kernel and cokernel of D are trivial. Thus the index of D is 0 and so is $\hat{A}(M)$. □

We will now adapt this argument to construct an explicit lift of the K-homology class of the spinor Dirac operator D on M to the analytic structure group which "explains" why the index is 0. This requires the following result based on finite propagation speed arguments:

Lemma 6.1.7. *Let M be a connected and complete Riemannian manifold and let D be a Dirac-type operator on M. If ϕ is any bounded Borel function on $\mathbb{R}$ with compactly supported Fourier transform then $\phi(D)$ is controlled.*

Proof. See Chapter 10 of [9]. □

Any normalizing function is the uniform limit of normalizing functions with compactly supported Fourier transform, so on a connected and complete Riemannian manifold we have that $\chi(D) \in D^*(M)/C^*(M)$ for any normalizing function χ. If $\widetilde{M}$ is a locally isometric G-cover of M and $\widetilde{D}$ is the lift of D to an operator on $\widetilde{M}$ then $\chi(D)$ is certainly G-equivariant, so $\chi(D) \in D^*_G(\widetilde{M})$.

Now assume that M is a connected and complete Riemannian spin manifold whose scalar curvature function is bounded from below by a positive constant and let $\widetilde{M}$ be a G-cover as above. The spin structure, spinor bundle, and spinor Dirac operator on M lift to a spin structure, spinor bundle, and spinor Dirac operator on $\widetilde{M}$; let $\widetilde{S}$ denote the spinor bundle and let $\widetilde{D}$ denote the spinor Dirac operator. The scalar curvature function on $\widetilde{M}$ is bounded below by the same positive constant as the scalar curvature function on M, so by the Lichnerowicz formula there is an interval $(-\varepsilon, \varepsilon)$ which does not meet the spectrum of $\widetilde{D}$. Choose a normalizing function which takes the value 1 on the interval $[\varepsilon, \infty)$ and the value

-1 on the interval $(-\infty, -\varepsilon]$, and observe that $\chi(D)^2 = 1$ exactly since $\chi^2 = 1$ on the spectrum of D.

View $\widetilde{D}$ as an operator on the reduced spinor bundle $\widetilde{S}_{red}$, so that $\widetilde{D}$ is ungraded if the dimension of M is odd and graded if the dimension of M is even. In the ungraded case, form the operator $\widetilde{P} = \frac{1}{2}(\chi(D) + 1) \in D^*_G(\widetilde{M})$; when we defined the K-homology class associated to $\widetilde{D}$ we observed that P is a projection modulo locally compact operators, but now P is is itself a projection and hence defines a class in $S_1(\widetilde{M}, G) = K_0(D^*_G(\widetilde{M}))$ which maps to the class of D in $K_1(M) \cong K_0(D^*_G(\widetilde{M})/C^*_G(\widetilde{M}))$. In the graded case, form the operator $\widetilde{U}$ determined by

$$\chi(\widetilde{D}) = \begin{pmatrix} 0 & \widetilde{U}^* \\ \widetilde{U} & 0 \end{pmatrix}$$

$\widetilde{U}$ determines a unitary operator on the infinite direct sum of $L^2(\widetilde{M}; \widetilde{S}_{red})$ and therefore a class in $S_0(\widetilde{M}, G) = K_1(D^*_G(\widetilde{M}))$, and this class maps to the class of D in $K_0(M)$.

Note that in either case the class in $S_p(\widetilde{M}, G)$ is independent of the normalizing function used to define it so long as that normalizing function is locally constant outside of the spectral gap for $\widetilde{D}$.

Definition 6.1.8. *Let M be a smooth connected spin manifold of dimension n and let $\widetilde{M}$ be a locally isometric G-cover of M. Given a Riemannian metric g on M which makes M a complete Riemannian manifold, the* structure invariant *associated to g is the class $\rho_g \in S_n(\widetilde{M}, G)$ determined by the operator $\chi(D)$, where D is the spinor Dirac operator for (M, g), as in the previous paragraph.*

Since the natural map $S_p(\widetilde{M}, G) \to K_p(M)$ sends ρ_g to the K-homology class of D and since it fits into the long exact sequence

$$\ldots \to S_p(\widetilde{M}, G) \to K_p(M) \to K_p(C^*_G(\widetilde{M})) \to \ldots$$

it follows that the equivariant indices of D in $K_p(C^*_G(\widetilde{M}))$ are zero. Combining this observation with the index theorem for k-partitioned manifolds, we obtain the following:

Corollary 6.1.9. *Let M be a complete Riemannian spin manifold of dimension n whose positive scalar curvature function is bounded below by a positive constant and let $\widetilde{M}$ be a locally isometric G-cover of M. Let N be a compact submanifold of M of codimension k which lifts to a submanifold $\widetilde{N}$ of $\widetilde{M}$ and pull back the spin structure on M to a spin structure on N via the inclusion map $N \hookrightarrow M$. If there is a k-partitioning map for the pair (M, N) which lifts to a k-partitioning map for the pair $(\widetilde{M}, \widetilde{N})$ then the equivariant index of the spinor Dirac operator on N in $K_{n-k}(C_r^*(G))$ vanishes.*

Proof. Let S_M and S_N denote the complexifications of the spinor bundles on M and N, respectively, and extend the spinor Dirac operators D_M and D_N to S_M and S_N. Unwinding the definitions it suffices to show that the operator D_M is k-partitioned by D_N. By the hypotheses on M and N there is a neighborhood $\mathcal{U}$ of N of the form $U \cong (-1,1)^k \times N$. Choose an open set $\mathcal{V} \subseteq N$ which trivializes S_N and such that $\mathcal{U}' = (-1,1)^k \times \mathcal{V}$ trivializes S_M ($\mathcal{V}$ exists since the spin structure on N is pulled back from the spin structure on M). Thus we have that $S_N|_{\mathcal{V}} \cong \mathbb{C}_{n-k} \times \mathcal{V}$ and $S_M|_{\mathcal{U}'} \cong \mathbb{C}_n \times \mathcal{U}'$. For any p and q there is an isomorphism

$$\mathbb{C}_p \hat{\otimes} \mathbb{C}_q \cong \mathbb{C}_{p+q}$$

defined as follows. Let $\{e_i\}$ be an orthonormal basis for $\mathbb{C}^p$ and let $\{e'_j\}$ be an orthonormal basis for $\mathbb{C}^q$; then the set

$$\{e_1 \hat{\otimes} 1, \ldots, e_p \hat{\otimes} 1, 1 \hat{\otimes} e'_1, \ldots, 1 \hat{\otimes} e'_q\}$$

is a generating set for $\mathbb{C}_p \hat{\otimes} \mathbb{C}_q$. Using an orthonormal basis $\{e''_1, \ldots, e''_{p+q}\}$ as a generating set for $\mathbb{C}_{p+q}$, the map $\mathbb{C}_p \hat{\otimes} \mathbb{C}_q \to \mathbb{C}_{p+q}$ determined by $e_i \hat{\otimes} 1 \mapsto e''_i$ and $1 \hat{\otimes} e'_j \mapsto e''_{p+j}$ is an isomorphism. Thus $S_M|_{\mathcal{U}'} \cong S_{(-1,1)^k} \hat{\otimes} S_N|_{\mathcal{V}}$; since N can be covered by open sets $\mathcal{V}$ which trivialize both S_N and S_M we conclude that $S_M|_{\mathcal{U}} \cong S_{(-1,1)^k} \hat{\otimes} S_N$. Relative to this decomposition it is clear from the definition of the spinor Dirac operator that $D_M = D_{(-1,1)^k} \hat{\otimes} 1 + 1 \hat{\otimes} D_N$ over $\mathcal{U}$, so the proof is complete. □

As another corollary, we give a new proof of a theorem due to Gromov and Lawson ([7]):

Corollary 6.1.10 (Gromov and Lawson). *Let M be a compact manifold of dimension n which admits a Riemannian metric of non-positive sectional curvature. Then M has no metric of positive scalar curvature.*

Proof. Let $\widetilde{M}$ be the universal cover of M equipped with a Riemannian metric of non-positive sectional curvature (lifted from a metric g on M). $\widetilde{M}$ is simply connected, so according to the Cartan-Hadamard theorem the exponential map $\exp_p \colon T_pM \to M$ is a diffeomorphism for any p. In fact we can say more: the exponential map is expansive in the sense that $d(\exp_p(v_1), \exp_p(v_2)) \geq \|v_1 - v_2\|$ for any $v_1, v_2 \in T_pM$, so the inverse map $\log\colon \widetilde{M} \to T_pM$ is a coarse diffeomorphism.

Now, let g' be any Riemannian metric on M. Any two norms on a finite dimensional vector space are equivalent, so there is a constant C such that $\frac{1}{C}\|v\|' \leq \|v\| \leq C$ for every $v \in T_pM$. Integrating, it follows that $\frac{1}{C}d'(x,y) \leq d(x,y) \leq Cd'(x,y)$ for every $x, y \in M$, so it follows that the map $\log\colon \widetilde{M} \to T_pM$ defined using the metric g is still a coarse map with respect to the metric g' (lifted to $\widetilde{M}$). Thus the 0-dimensional submanifold $\{p\}$ of $\widetilde{M}$ is n-partitioned by the coarse diffeomorphism $\log\colon \widetilde{M} \to T_pM \cong \mathbb{R}^n$.

Clearly $\widetilde{M}$ is contractible, so let $S_{\widetilde{M}}$ denote the trivial spinor bundle $\mathbb{R}_n \times \widetilde{M}$ and let $D_{\widetilde{M}}$ denote the spinor Dirac operator on $S_{\widetilde{M}}$ (defined using the Riemannian metric g'). Consider the trivial bundle $S_p = \mathbb{C} \times \{p\} \to \{p\}$ and let D_p be the zero map on S_p; $D_{\widetilde{M}}$ is k-partitioned by D_p because $S_{\widetilde{M}} \cong \mathbb{R}_n \hat{\otimes} S_p$ and $D_{\widetilde{M}} = D_{\mathbb{R}^n} \hat{\otimes} 1 + 1 \hat{\otimes} D_p$. By the index theorem for k-partitioned manifolds, we have

$$\mathrm{Ind}_{\widetilde{M},\{p\}}(D_{\widetilde{M}}) = \mathrm{Ind}_{\{p\}}(D_p) = 1$$

in $\mathbb{Z}$. But $\mathrm{Ind}_{\widetilde{M},\{p\}}$ factors through the coarse index map

$$K_p(\widetilde{M}) \to K_p(C^*(\widetilde{M}))$$

and the image under this map of $D_{\widetilde{M}}$ would be 0 if the scalar curvature function on M associated to g' were positive. Thus M can have no metric of positive scalar curvature. $\square$

This result in particular implies that the n-torus has no metric of positive scalar curvature; this was a long-standing open problem before Gromov and Lawson solved it, though it was known in low dimensions due to results of Schoen and Yau using variational techniques instead of index theory.

6.2 Partitioned Manifolds and Structure Invariants

Suppose $M = M^+ \cup M^-$ is a complete Riemannian spin manifold of dimension $n+1$ partitioned by a compact hypersurface N, and restrict the Riemannian and spin structures on M to corresponding structures on N. There is a structure invariant $\rho_M \in S_{n+1}(M)$ associated to the positive scalar curvature metric on M, and there is a map

$$\partial_{MV} \colon S_{n+1}(M) \to S_n(N)$$

given by the boundary map in the Mayer-Vietoris sequence for the structure group associated to the partitioning of M. It is natural to try to calculate $\partial_{MV}(\rho_M)$. In the most elementary case where N is a compact Riemannian spin manifold with positive scalar curvature and $M = \mathbb{R} \times N$, one has invariants $\rho_N \in S_n(N)$ and $\rho_M \in S_{n+1}(M)$, and one expects that $\partial_{MV}(\rho_M) = \rho_N$.

The goal of this section is to report our progress on this problem, though a complete solution eludes us at this time. Our strategy is to imitate the proof given in the previous chapter that the Mayer-Vietoris boundary map in K-homology $K_{n+1}(M) \to K_n(N)$ sends the K-homology class of the product of an operator D_N on N with the complex spinor Dirac operator on $\mathbb{R}$ to the K-homology class of D_N itself. This argument relied heavily on the Kasparov product and its compatibility with the suspension map in K-homology explored in Chapter 3, and we have tried to imitate the main features of this argument.

To begin, we will introduce a new model of the analytic structure group $S_p(X)$, where X is a proper metric space, in the spirit of Kasparov's model for K-homology. Using this model we will construct a product

$$K_p(X) \times S_q(Y) \to S_{p+q}(X \times Y)$$

which is compatible with the Kasparov product in the sense that the following diagram commutes:

$$\begin{array}{ccc} K_p(X) \times S_q(Y) & \longrightarrow & S_{p+q}(X \times Y) \\ \downarrow & & \downarrow \\ K_p(X) \times K_q(Y) & \longrightarrow & K_{p+q}(X \times Y) \end{array}$$

We conjecture that this product is compatible with the Mayer-Vietoris boundary map in the sense that if $X = X_1 \cup X_2$ is a decomposition of X into closed subspaces then following diagram commutes:

$$\begin{array}{ccc} K_p(X) \times S_q(Y) & \longrightarrow & S_{p+q}(X \times Y) \\ \downarrow & & \downarrow \\ K_p(X_1 \cap X_2) \times S_q(Y) & \longrightarrow & S_{p+q}((X_1 \cap X_2) \times Y) \end{array}$$

(where the vertical maps are Mayer-Vietoris boundary maps). The proof of this fact has eluded us so far, even in the case where $X = \mathbb{R}$, $X_1 = \mathbb{R}^{\geq 0}$, and $X_2 = \mathbb{R}^{\leq 0}$. This case alone would be sufficient for our purposes: we will show that if N is a compact Riemannian spin manifold with positive scalar curvature then the structure invariant for $\mathbb{R} \times N$ is the product of the Dirac class and the structure invariant for N and we have already shown that the Mayer-Vietoris boundary map $K_1(\mathbb{R}) \to K_0(\{0\})$ sends the Dirac class to the unit class, so it would follow from the compatibility statement above that $\partial_{MV}(\rho_{\mathbb{R}\times N}) = \rho_N$.

Until this point we have stated and proved most of our results equivariantly with respect to a free and proper group action, but for simplicity (and since we do not yet know how a group action will affect our results) we will not incorporate a group action in this chapter.

6.2.1 Analytic Structure Cycles

The goal of this section is to develop a model of the analytic structure group in the spirit of Kasparov's model of K-homology. Our new model for $S_p(X)$ is based on the following:

Definition 6.2.1. *Let X be a proper separable metric space. A p-multigraded* analytic structure cycle *for X consists of the following data:*

- *A self-adjoint p-multigraded Fredholm module (ρ, H, F) over $C_0(X)$ such that F is the norm limit of p-multigraded pseudolocal controlled operators.*

- *A norm continuous path F_t such that $F_0 = F$, $F_1^2 \geq \varepsilon > 0$ for some constant ε, and each F_t is the norm limit of p-multigraded, controlled, self-adjoint Fredholm operators*

Informally, an analytic structure cycle consists of a Fredholm module (ρ, H, F) with $F \in D^*(X)$ (note that any K-homology class has a representative of this sort) together with a homotopy F_t which "explains" why the index of F is 0. Note that the operators F_t are assumed to be Fredholm but not pseudolocal except at $t = 0$; indeed, the operators F_t for $t \neq 0$ may have no relationship to X. This ensures that F_t causes the index of F to to be 0 without necessarily causing the K-homology class of (ρ, H, F) to be 0.

Definition 6.2.2. *Let (ρ, H, F^0, F_t^0) and (ρ, H, F^1, F_t^1) be two p-multigraded structure cycles with the same Hilbert space and representation. Say that (F^0, F_t^0) and (F^1, F_t^1) are* homotopic *if they are joined by a path of p-multigraded structure cycles (ρ, H, F^s, F_t^s)*

There is a natural notion of direct sum of structure cycles which is evidently compatible with the homotopy equivalence relation, and the structure cycle for which the Hilbert space, representation, and operators are all zero is an additive identity. Denote by $\hat{S}_p(X)$ the Grothendieck group of the set of all homotopy classes of p-multigraded structure cycles with this additive structure. The assignment $(\rho, H, F, F_t) \mapsto (\rho, H, F)$ respects homotopies and direct sums, so it determines a group homomorphism $\hat{S}_p(X) \to K_p(X)$. We shall prove that $\hat{S}_p(X)$ is isomorphic to the analytic structure set $S_p(X)$ and that the map $\hat{S}_p(X) \to K_p(X)$ is compatible with the map $S_p(X) \to K_p(X)$ defined using the isomorphism $K_p(X) \cong K_{1-p}(D^*(X)/C^*(X))$:

$$
\begin{tikzcd}
\hat{S}_p(X) \arrow[r] \arrow[d, "\cong"] & K_p(X) \arrow[d, "\cong"] \\
K_{1-p}(D^*(X)) \arrow[r] & K_{1-p}(D^*(X)/C^*(X))
\end{tikzcd}
$$

This requires some preliminary calculations with C*-algebra K-theory. Consider the *multiplier algebra* $\mathcal{M}(J)$, the C*-algebra $\mathbb{B}_J(J)$ of bounded adjointable operators on J regarded as a Hilbert J-module in the usual way. Note that if X is a proper separable metric space represented on a separable Hilbert space H then $M(C^*(X))$ is the norm closure of the set of all controlled operators in $\mathbb{B}(H)$. The main fact about multiplier algebras that we will need is the following analogue of Kuiper's theorem due to Cuntz and Higson:

Proposition 6.2.3. *If J is stable then the unitary group of $\mathcal{M}(J)$ is contractible and hence $\mathcal{M}(J)$ has trivial K-theory.*

Proof. See [5]. □

Assume that J is an ideal in another C*-algebra A. There is a natural map $A \to \mathcal{M}(J)$, defined by allowing $a \in A$ to act on J as the bounded adjointable Hilbert module operator $j \mapsto aj$. Thus there is a commuting diagram:

$$\begin{tikzcd} 0 \arrow[r] & J \arrow[r] \arrow[d, equal] & A \arrow[r, "\pi"] \arrow[d] & A/J \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & J \arrow[r] & \mathcal{M}(J) \arrow[r, "\pi"] & \mathcal{M}(J)/J \arrow[r] & 0 \end{tikzcd} \tag{6.2.1}$$

The six-term exact sequence in K-theory associated to the first row of this diagram indicates that the K-theory of A is home to secondary invariants for A/J in the sense that for any class $x \in K_p(A/J)$ we have that $\partial x = 0$ if and only if x lifts to a class in $K_p(A)$. The six term exact sequence of the second row is more degenerate by Proposition 6.2.3; it simply boils down to the assertion that $K_p(\mathcal{M}(J)/J) \cong K_{1-p}(J)$. This suggests an alternative model for the K-theory of A which is compatible with our proposed model of the analytic structure group. We will develop this model first for K_0 and then for K_1.

Definition 6.2.4. *Define $\hat{K}_0(A)$ to be the Grothendieck group of homotopy classes of pairs (p, p_t) where p is a projection over A/J and p_t is a path of projections over $\mathcal{M}(J)/J$ such that $p_0 = p$ and $p_1 = 1 \oplus 0$.*

Remark 6.2.5. *In principle J should be part of the notation for $\hat{K}_0(A)$, but Proposition 6.2.7 below will imply that $\hat{K}_0(A)$ is independent of J.*

By a homotopy of pairs we mean a map $s \mapsto (p^s, p^s_t)$ (norm continuous in each variable) where $s \mapsto p^s$ is a path of projections over A/J and p^s_t is a path of projections over $\mathcal{M}(J)/J$ for each s such that $p^s_0 = p^s$ and $p^s_1 = 1 \oplus 0$.

There is a map $\psi : K_0(A) \to \hat{K}_0(A)$ defined as follows. Given a class $[q] \in K_0(A)$, its image in $K_0(\mathcal{M}(J))$ is trivial by Proposition 6.2.3 and thus there is a homotopy q_t of projections over $\mathcal{M}(J)$ such that $q_0 = q$ and $q_1 = 1 \oplus 0$. The pair $(\pi(q), \pi(q_t))$ determines a class in $\hat{K}_0(A)$.

Lemma 6.2.6. *The equivalence class of $(\pi(q), \pi(q_t))$ is independent of the path q_t used to define it.*

Proof. Suppose q'_t is another path of projections over $\mathcal{M}(J)$ such that $q'_0 = q$ and $q'_1 = 1 \oplus 0$. Let u_t and u'_t be paths of unitaries such that $q_t = u^*_t q u_t$ and $q'_t = u'^*_t q u'_t$; by Proposition 6.2.3, u_t is homotopic to u'_t and hence q_t is homotopic to q'_t. $\square$

Thus ψ is well-defined; we now prove that it is an isomorphism.

Proposition 6.2.7. $\psi : K_0(A) \to \hat{K}_0(A)$ *is an isomorphism.*

Proof. Our strategy is to fit ψ into the commuting diagram

$$\begin{tikzcd}
K_1(A/J) \arrow[r] \arrow[d, "\cong"] & K_0(J) \arrow[r] \arrow[d, "\cong"] & K_0(A) \arrow[r] \arrow[d, "\psi"] & K_0(A/J) \arrow[r, "\partial"] \arrow[d, equal] & K_1(J) \arrow[d, equal] \\
K_0(S(A/J)) \arrow[r] & K_0(S(\mathcal{M}(J)/J)) \arrow[r] & \hat{K}_0(A) \arrow[r] & K_0(A/J) \arrow[r, "\partial"] & K_1(J)
\end{tikzcd}$$

with exact rows and apply the five lemma. The top row is just the long exact sequence in K-theory associated to $0 \to J \to A \to A/J \to 0$ while the bottom row is defined as follows:

- The map $K_0(S(A/J)) \to K_0(S(\mathcal{M}(J)/J))$ is induced by the map $A/J \to \mathcal{M}(J)/J$.
- The map $K_0(S(\mathcal{M}(J)/J)) \to \hat{K}_0(A)$ sends the K-theory class of a normalized loop of projections p_t over $\mathcal{M}(J)/J$ with $p_0 = p_1 = 1 \oplus 0$ to the class $[1 \oplus 0, p_t] \in \hat{K}_0(A)$.
- The map $\hat{K}_0(A) \to K_0(A/J)$ sends a class $[p, p_t] \in \hat{K}_0(A)$ to $[p] \in K_0(A/J)$.

- The map $K_0(A/J) \to K_1(J)$ is the usual boundary map in K-theory.

To define the vertical map $K_0(J) \to K_0(S(\mathcal{M}(J)/J))$, note that the homotopy theoretic boundary map $\partial : K_0(S(\mathcal{M}(J)/J)) \to K_0(J)$ associated to the short exact sequence $0 \to J \to \mathcal{M}(J) \to \mathcal{M}(J)/J \to 0$ is an isomorphism by Proposition 6.2.3. So the vertical map is simply defined to be its inverse.

Commutativity of the diagram with these definitions is clear except at the square

$$\begin{tikzcd}
K_0(J) \arrow[r] \arrow[d, "\cong"] & K_0(A) \arrow[d, "\psi"] \\
K_0(S(\mathcal{M}(J)/J)) \arrow[r] & \hat{K}_0(A)
\end{tikzcd}$$

A class in $K_0(S(\mathcal{M}(J)/J))$ is represented by a normalized loop of projections p_t over $\mathcal{M}(J)/J$, and we can assume without loss of generality that $p_0 = p_1 = 1 \oplus 0$. The projection $1 \oplus 0$ over $\mathcal{M}(J)/J$ lifts to $1 \oplus 0$ regarded as a projection over $\mathcal{M}(J)$, and by the path lifting property for projections the loop p_t lifts to a path q_t of projections over $\mathcal{M}(J)$ with $q_1 = 1 \oplus 0$. Since q_0 lifts $1 \oplus 0$ it is a projection over $\widetilde{J}$, and the boundary map $K_0(S(\mathcal{M}(J)/J)) \to K_0(J)$ by definition sends $[p_t]$ to $[q_0]$.

The map $K_0(J) \to K_0(A)$ sends the class of q_0 to itself (regarded as a projection over A), and $\psi[q_0] = [\pi(q_0), \pi(q_t)]$ where q_t is as above. But $\pi(q_0) = 1 \oplus 0$ and $\pi(q_t) = p_t$, so $\psi[q_0] = [1 \oplus 0, p_t]$. This proves that the diagram commutes.

Let us now show that the bottom row is exact.

1. Exactness at $K_0(S(\mathcal{M}(J)/J))$

   Let p_t be a normalized loop of projections over $\mathcal{M}(J)/J$ and assume its K-theory class $[p_t]$ maps to 0 in $\hat{K}_0(A)$. The image of $[p_t]$ in $\hat{K}_0(A)$ is represented by the pair $(1\oplus 0, p_t)$ and the zero class is represented by the pair $(1\oplus 0, 1\oplus 0)$ (a constant path in the second coordinate), so the fact that they represent the same class means (up to stabilization) that there is a pair (p^s, p^s_t) with the property that:

   - p^s is a projection over A/J for each s and p^s_t is a projection over $\mathcal{M}(J)/J$ for each s, t.

- $p^s_0 = p^s$ and $p^s_1 = 1 \oplus 0$ for each s
- $p^0_t = p_t$ and $p^1_t = 1 \oplus 0$ for each t

Notice that $p^0 = p^0_0 = p_0 = 1 \oplus 0$ and $p^1 = p^1_0 = 1 \oplus 0$, so p^s is a normalized loop of projections over A/J and hence determines a class in $K_0(S(A/J))$. Moreover the image of $s \mapsto p^s$ is homotopic through normalized loops of projections over $\mathcal{M}(J)/J$ to $t \mapsto p_t$ since both are homotopic to $u \mapsto p^u_u$: the homotopy $h_1(u,v) = p^u_{uv}$ joins $h_1(u,0) = p^u$ to $h_1(u,1) = p^u_u$ while $h_2(u,v) = p^{uv}_u$ joins $h_2(u,0) = p_u$ to $h_2(u,1) = p^u_u$. Thus $[p_t]$ is the image of the class $[p^s] \in K_0(S(A/J))$.

2. Exactness at $\hat{K}_0(A)$

   Let $[p, p_t] \in \hat{K}_0(A)$ be a class which maps to 0 in $K_0(A/J)$, so that $[p] = 0$. Up to stabilization and homotopy we can assume that $p = 1 \oplus 0$ and thus p_t is a normalized loop of projections over $\mathcal{M}(J)/J$. Such a loop defines a class $[p_t] \in K_0(S(\mathcal{M}(J)/J))$, and $[p_t]$ maps to $[1 \oplus 0, p_t]$ in $\hat{K}_0(A)$.

3. Exactness at $K_0(A/J)$

   Let $[q]$ be a class in $K_0(A/J)$ such that $\partial[q] = 0$ in $K_1(J)$. By the usual long exact sequence in K-theory this means $[q]$ lifts to a class $[q']$ in $K_0(A)$, and it is immediate from the definitions that $\psi[q']$ maps to $[q]$.

The five lemma completes the proof. □

We now indicate the modifications in this construction required to build a model for $K_1(A)$.

Definition 6.2.8. *Define $\hat{K}_1(A)$ to be the Grothendieck group of homotopy classes of pairs (u, u_t) where u is a unitary over A/J and u_t is a path of unitaries over $\mathcal{M}(J)/J$ such that $u_0 = u$ and $u_1 = 1 \oplus 0$.*

As before we define a map $\psi : K_1(A) \to \hat{K}_1(A)$ by $\psi[v] = [\pi(v), \pi(v_t)]$ where v is a unitary over A and v_t is a path of unitaries over $\mathcal{M}(J)$ such that $v_0 = v$ and $v_1 = 1 \oplus 0$. As before, Proposition 6.2.3 implies that ψ is well-defined.

Proposition 6.2.9. $\psi : K_1(A) \to \hat{K}_1(A)$ *is an isomorphism.*

Proof. Following our previous approach, we fit ψ into the commuting diagram

$$\begin{tikzcd}
K_0(A/J) \arrow[r] \arrow[d, "\cong"] & K_1(J) \arrow[r] \arrow[d, "\cong"] & K_1(A) \arrow[r] \arrow[d, "\psi"] & K_1(A/J) \arrow[r, "\partial"] \arrow[d, equal] & K_0(J) \arrow[d, equal] \\
K_1(S(A/J)) \arrow[r] & K_1(S(\mathcal{M}(J)/J)) \arrow[r] & \hat{K}_1(A) \arrow[r] & K_1(A/J) \arrow[r, "\partial"] & K_0(J)
\end{tikzcd}$$

with exact rows and apply the five lemma. All maps are defined in the same way as above except for the two leftmost vertical maps which are the Bott periodicity isomorphisms. Commutativity of the diagram follows from our earlier arguments except for commutativity of the leftmost square, and this follows from the naturality of the Bott map. The proof that the bottom row is exact follows from exactly the same arguments as before. □

These results allow us to show that our two models of the analytic structure group are the same.

Proposition 6.2.10. *For any proper metric space X, there is an isomorphism $K_{1-p}(D^*(X)) \to \hat{S}_p(X)$ which makes the following diagram commute:*

$$\begin{tikzcd}
K_{1-p}(D^*(X)) \arrow[r] \arrow[d, "\cong"] & K_{1-p}(D^*(X)/C^*(X)) \arrow[d, "\cong"] \\
\hat{S}_p(X) \arrow[r] & K_p(X)
\end{tikzcd}$$

Proof. Define $D^*(X)$ using an ample representation $\rho\colon C_0(X) \to \mathbb{B}(H)$. According to Proposition 6.2.7 and Proposition 6.2.9, we have that $K_p(D^*(X))\hat{K}_p(D^*(X))$ where the latter is defined using the ideal $C^*(X)$. A class in $\hat{K}_0(D^*(X))$ is represented by a pair $(P, \pi(P_t))$ where P is a projection over $D^*(X)/C^*(X)$ and P_t is a path in $M(C^*(X))$ whose image $\pi(P_t)$ in $M(C^*(X))/C^*(X)$ is a path of projections joining P to $1 \oplus 0$.

Define a map $\hat{K}_0(D^*(X)) \to \hat{S}_1(X)$ by sending the class of (P, P_t) to the class of $(\rho, H, 2P-1, 2P_t-1)$; this is well-defined since a homotopy of pairs (P, P_t) maps to a homotopy of structure cycles. Note that it restricts to the isomorphism $K_{1-p}(D^*(X)/C^*(X)) \cong K_p(X)$ defined using Proposition 3.4.11 together with Theorem 4.3.25, and it is injective by Proposition 3.4.11. To see that it is surjective, it suffices to show by Proposition 3.4.11 that any controlled self-adjoint Fredholm operator $F \in M_n(\mathbb{B}(H))$ which satisfies $F^2 \geq \varepsilon > 0$ is connected through a path of controlled self-adjoint Fredholm operators to a trivial Fredholm operator.

Consider the functions $f(t) = t$ and $g(t) = \operatorname{sign}(t)$ on $\mathbb{R}$; the restrictions of f and g to the spectrum of F are homotopic through continuous nonvanishing functions since the spectrum of F is a compact subset of $\mathbb{R}$ which misses the interval $(-\sqrt{\varepsilon}, \sqrt{\varepsilon})$, so $f(F) = F$ is homotopic through a path of controlled self-adjoint Fredholm operators to $g(F)$, a trivial Fredholm operator.

This completes the proof for $p = 1$; the proof for $p = 0$ follows from the same argument using the other isomorphism appearing in Proposition 3.4.11. □

From now on we will drop the notation $\hat{S}_p(X)$ and specify which model for the analytic structure group we are using when ambiguity arises. We conclude this section by specifying an analytic structure cycle which represents the structure invariant of a positive scalar curvature invariant. Let M be a complete Riemannian spin manifold whose scalar curvature function is bounded below by a positive constant, let D be the spinor Dirac operator on M, and let χ be a normalizing function which is locally constant on $\mathbb{R} - (-\varepsilon, \varepsilon)$ where ε is a number such that $D^2 > \varepsilon^2$. Finally let $\rho\colon C_0(M) \to \mathbb{B}(L^2(M; S_M))$ be the representation by multiplication operators where S_M is the complexified spinor bundle on M. Then $(\rho, L^(M; S_M), \chi(D), \chi(D))$ is an n-multigraded structure cycle over M (we use the constant path $\chi(D)$) and hence determines a class in $S_n(M)$. Note that the reduction procedure which passes from S_M to $(S_M)_{red}$ identifies this structure invariant with the structure invariant in $S_0(M)$ or $S_1(M)$ defined earlier.

6.2.2 The Analytic Structure Group and Products

We are now ready to define a product

$$K_p(X) \times S_q(Y) \to S_{p+q}(X \times Y)$$

which is compatible with the Kasparov product, where X and Y are proper metric spaces.

Let (ρ^X, H^X, F^X) be a graded Fredholm module over X such that F^X is self-adjoint and controlled and let $(\rho^Y, H^Y, F^Y, F^Y_t)$ be a graded analytic structure cycle over Y. Let $(\rho^X \otimes \rho^Y, H^X \hat{\otimes} H^Y, F)$ be a Fredholm module which is aligned with the pair (F^X, F^Y), so that its K-homology class represents the product of (ρ^X, H^X, F^X) and (ρ^Y, H^Y, F^Y). Form the operator

$$F' = \frac{1}{\sqrt{2}}(F^X \hat{\otimes} 1 + 1 \hat{\otimes} F^Y) \tag{6.2.2}$$

According to the rules for graded tensor products we have $F'^2 - 1$ is a compact operator, so F' is a self-adjoint controlled Fredholm operator. The condition that F is aligned with the pair (F^X, F^Y) implies that the path

$$t \mapsto \cos(\frac{\pi}{2}t)F + \sin(\frac{\pi}{2}t)F' \tag{6.2.3}$$

joins F and F' through a path of self-adjoint controlled Fredholm operators. Finally, the path

$$t \mapsto \frac{1}{\sqrt{2}}(F^X \hat{\otimes} 1 + 1 \hat{\otimes} F^Y_t)$$

joins F' (again through a path of self-adjoint controlled Fredholm operators) to the operator $\frac{1}{\sqrt{2}}(F^X \hat{\otimes} 1 + 1 \hat{\otimes} F^Y_1)$ whose square is bounded below by a positive constant.

Definition 6.2.11. *Let X and Y be proper metric spaces, let (ρ^X, H^X, F^X) be a graded Fredholm module over X such that F^X is a self-adjoint controlled operator, and let $(\rho^Y, H^Y, F^Y, F^Y_t)$ be a graded analytic structure cycle over Y. The* structure product *of (ρ^X, H^X, F^X) and $(\rho^Y, H^Y, F^Y, F^Y_t)$ is defined to be the analytic structure cycle*

$$(\rho^X \otimes \rho^Y, H^X \hat{\otimes} H^Y, F, F_t)$$

where F is an operator aligned with the pair (F^X, F^Y) and F_t is the concatenation of the paths (6.2.3) *and* (6.2.2).

This construction is plainly compatible with multigrading structure and homotopies in both factors, so it determines the desired product $K_p(X) \times S_q(Y) \to K_{p+q}(X \times Y)$. We conclude by examining how it behaves when applied to structure invariants for positive scalar curvature metrics.

Lemma 6.2.12. *Let X and Y be proper metric spaces, let (ρ^X, H^X, F^X) be a p-multigraded Fredholm module over X such that F^X is a self-adjoint controlled* *operator, and let (ρ^Y, H^Y, F^Y, F^Y) be a q-multigraded analytic structure cycle over Y such that $(F^Y)^2 > \varepsilon > 0$. Suppose that $(\rho^X \otimes \rho^Y, H^X \hat{\otimes} H^Y, F)$ is a $p+q$-multigraded Fredholm module over $X \times Y$ which is aligned with the pair (F^X, F^Y) and which satisfies:*

- $F^2 \geq \varepsilon > 0$
- $[F, F^X \hat{\otimes} 1] \geq 0$
- $[F, 1 \hat{\otimes} F^Y] \geq 0$

Then the analytic structure cycle $(\rho^X \otimes \rho^Y, H^X \hat{\otimes} H^Y, F, F)$ with a constant path of Fredholm operators represents the structure product

$$[\rho^X, H^X, F^X] \times [\rho^Y, H^Y, F^Y, F_t^Y]$$

Proof. For clarity let us suppress the representation and Hilbert space from the notation. It suffices to show that there is a homotopy (F^s, F_t^s) of analytic structure cycles joining (F, F) the analytic structure cycle (F, F_t) which appears in Definition 6.2.11. The hypotheses on F imply that each of the operators appearing in (6.2.3) and (6.2.2) square to operators bounded below by a positive constant, so the contraction homotopy given by $F^s = F$ and $F_t^s = F_{ts}$ is a homotopy through analytic structure cycles. $\square$

Note that the commutators between F and $F^X \hat{\otimes} 1$ and $1 \hat{\otimes} F^Y$ must actually be positive, not just positive modulo compact operators. Fortunately, this condition is satisfied for Fredholm modules coming from differential operators.

Proposition 6.2.13. *Let M and N be complete Riemannian spin manifolds and assume that N and $M \times N$ (equipped with the product Riemannian metric ans spin structure) have positive scalar curvature. Then*

$$[D_M] \times \rho_N = \rho_{M \times N}$$

where $[D_M]$ is the K-homology class of the spinor Dirac operator on M, ρ_N is the structure invariant of N, and $\rho_{M \times N}$ is the structure invariant of $M \times N$.

Proof. Let D_M, D_N, and $D_{M\times N}$ denote the spinor Dirac operators for M, N, and $M \times N$, respectively. By a variation on an argument in Corollary 6.1.9, there is an isomorphism $\mathbb{R}_p \hat{\otimes} \mathbb{R}_q \cong \mathbb{R}_{p+q}$ which yields an identification

$$D_{M\times N} = D_M \hat{\otimes} 1 + 1 \hat{\otimes} D_N$$

Let χ be a normalizing function which is locally constant on $\mathbb{R} - (-\varepsilon, \varepsilon)$ where $(-\varepsilon, \varepsilon)$ is an interval contained in the spectral gap of D_N and $D_{M\times N}$. By the assumption that $M \times N$ has positive scalar curvature and Proposition 3.5.6, the Fredholm modules determined by the operators $\chi(D_M)$, $\chi(D_N)$ and $\chi(D_{M\times N})$ satisfy the hypotheses of Lemma 6.2.12, so the result follows. $\square$

Appendix A

Dual Algebras and Fredholm Modules

We defined the K-homology groups of a C*-algebra A in two different ways: first using Paschke's theory of dual algebras and second using Kasparov's theory of Fredholm modules. Paschke's model, denoted by $K^p(A)$, is by definition $K_{1-p}(\mathfrak{D}^*(A)/\mathfrak{C}^*(A))$ where $\mathfrak{D}^*(A)$ and $\mathfrak{C}^*(A)$ are defined using a fixed ample representation of A. Kasparov's model, denoted by $KK^{-p}(A, \mathbb{C})$, is by definition the Grothendieck group of p-multigraded Fredholm modules over A.

Our first task in this appendix is to prove Proposition 3.4.11 which asserts that $K^1(A) \cong KK^1(A, \mathbb{C})$ and $K^0(A) \cong KK^0(A, \mathbb{C})$. Since both groups satisfy the Bott periodocity theorem, this shows that the two models of K-homology are the same. Our second aim is to relate two different definitions of the suspension map in K-homology. Using Paschke's model of K-homology we defined the suspension map to be the boundary map $s\colon K^{-p-1}(S(A)) \to K^{-p}(A)$ in K-homology associated to the short exact sequence

$$0 \to S(A) \to C(A) \to A \to 0$$

We asserted that this map agrees with an explicit map defined at the level of Fredholm modules, and we will give a proof of this fact, adapted from chapter 8 of [9].

A.1 Equivalence of Models

Let A be a separable C*-algebra and let $\rho_A\colon A \to H_A$ be an ample representation. The precise statement of Proposition 3.4.11 was that the map

$$\Gamma_1\colon K^1(A) \to KK^1(A, \mathbb{C})$$

which sends the K-homology class of a projection P in $M_n(\mathfrak{D}^*(A)/\mathfrak{C}^*(A))$ to the ungraded Fredholm module $(\rho_A, H_A^n, 2P-1)$, and the map

$$\Gamma_0\colon K^0(A) \to KK^0(A, \mathbb{C})$$

which sends the K-homology class of a unitary U in $M_n(\mathfrak{D}^*(A)/\mathfrak{C}^*(A))$ to the graded Fredholm module

$$\left(\rho_A \oplus \rho_A, H_A^n \oplus H_A^n, \begin{pmatrix} 0 & U^* \\ U & 0 \end{pmatrix}\right)$$

are both isomorphisms.

Proof of Proposition 3.4.11. Let us first prove that Γ_1 is an isomorphism. To show that Γ_1 is surjective we must show that every Fredholm module over A is K-equivalent to one whose representation is ample. So let (ρ, H, F) be a Fredholm module over A; setting $P = \rho(A)H$ we have that (ρ, PH, PFP) is a nondegenerate Fredholm module (i.e. the representation ρ is nondegenerate) which is a compact perturbation of (ρ, H, F). So let us simply assume that (ρ, H, F) is nondegenerate. Let $(\rho_A, H_A, 1_A)$ denote the degenerate Fredholm module whose representation ρ_A is the ample representation used to define $\mathfrak{D}^*(A)$ and $\mathfrak{C}^*(A)$ and whose operator is the identity map $1_A \in \mathbb{B}(H_A)$. Then $(\rho \oplus \rho_A, H \oplus H_A, F \oplus 1_A)$ is a Fredholm module representing the same K-homology class as (ρ, H, F), and Voiculescu's theorem implies that there is a unitary $U\colon H_A\colon H \oplus h_A$ such that $U\rho_A(a)U^* \sim \rho(a) \oplus \rho_A(a)$. Thus the Fredholm module $(\rho_A, H_A, U^*(F \oplus F_A)U)$ represents the same K-homology class as (ρ, H, F) and it is the image under Γ_1 of the projection $\frac{1}{2}(U^*(F \oplus F_A)U + 1)$ in $\mathfrak{D}^*(A)/\mathfrak{C}^*(A)$.

The proof that Γ_1 is injective is simply a "relative" version of the argument above. Let P be a projection over $\mathfrak{D}^*(A)/\mathfrak{C}^*(A)$ such that $\Gamma_1[P] = 0$. The K-homology class of the Fredholm module $(\rho_A, H_A, 2P - 1)$ is by definition $\Gamma_1[P]$, so it is K-equivalent to 0; thus there is another Fredholm module (ρ', H', F') over A and chain of unitary equivalences, operator homotopies, and compact perturbations connecting $(\rho_A \oplus \rho', H_A \oplus H', (2P-1) \oplus F')$ to the direct sum of a degenerate Fredholm module with (ρ', H', F'). By the surjectivity argument above the K-homology class of (ρ', H', F') is the image under Γ_1 of some projection Q over $\mathfrak{D}^*(A)/\mathfrak{C}^*(A)$, and similarly the chain of unitary equivalences, operator homotopies, and compact perturbations lifts to a chain of equivalences joining $P \oplus Q$ to $(0 \oplus 1) \oplus Q$. Hence $[P] = 0$, as desired.

The proof that Γ_0 is an isomorphism is almost identical, except we must keep track of the graded structure. Let (ρ, H, F) be any graded Fredholm module, so that $H = H^+ \oplus H^-$ and

$$F = \begin{pmatrix} 0 & V \\ U & 0 \end{pmatrix}$$

Consider a sequence of Hilbert spaces H_n, $n \in \mathbb{Z}$, where $H_n = H^-$ for $n \leq -1$ and $H_n = H^+$ for $n \geq 0$. Define $\hat{H} = \bigoplus_n H_n$, and define an operator $\hat{U} \in \mathbb{B}(\hat{H})$ which is the identity $H_n \to H_{n-1}$ for $n \neq 0$ and which is $U \colon H_0 \to H_{-1}$ at $n = 0$. Define $\hat{V}$ similarly, and form a new Fredholm module (ρ, H', F') where $H' = \hat{H} \oplus \hat{H}$ and

$$F' = \begin{pmatrix} 0 & \hat{V} \\ \hat{U} & 0 \end{pmatrix}$$

Then (ρ', H', F') is a graded Fredholm module which is unitarily equivalent to the direct sum of (ρ, H, F) and two degenerate Fredholm modules, so that (ρ', H', F') has the same K-homology class as (ρ, H, F). The effect has been to replace (ρ, H, F) with a graded Fredholm module whose graded Hilbert space is the direct sum of two copies of the same fixed Hilbert space $\hat{H}$. The proof that Γ_1 is an isomorphism now applies almost word-for-word to Γ_0, except the role of the ungraded degenerate Fredholm module $(\rho_A, H_A, 1_A)$ must now be played by the graded degenerate Fredholm module

$$\left(\rho_A \oplus \rho_A, H_A \oplus H_A, \begin{pmatrix} 0 & 1_A \\ 1_A & 0 \end{pmatrix}\right)$$

$\square$

A.2 The Suspension Map

Let A be a separable unital C*-algebra and consider the short exact sequence

$$0 \to S(A) \to C(A) \to A \to 0$$

where $S(A) = C_0(0,1) \otimes A$ and $C(A) = C_0(0,1] \otimes A$. There is a corresponding short exact sequence of dual algebras

$$0 \to \mathfrak{D}^*(C(A)//S(A)) \to \mathfrak{D}^*(C(A)) \to \mathfrak{D}^*(C(A))/\mathfrak{D}^*(C(A)//S(A)) \to 0$$

Exploiting the excision isomorphism

$$K_*(\mathfrak{D}^*(C(A))/\mathfrak{D}^*(C(A)//S(A))) \cong K_*(\mathfrak{D}^*(S(A)))$$

and the isomorphism

$$K_*(\mathfrak{D}^*(C(A)//S(A))) \cong K_*(\mathfrak{D}^*(A))$$

there is a boundary map $K_p(\mathfrak{D}^*(S(A))) \to K_{p-1}(\mathfrak{D}^*(A))$. At the level of K-homology this is by definition the suspension map

$$s\colon K^{-p-1}(S(A)) \to K^{-p}(A)$$

On the other hand, we introduced a notion of suspension for Fredholm modules. This is the assignment $(\rho, H, F) \mapsto (\rho, H, V)$ where (ρ, H, F) is a $(p+1)$-multigraded relative Fredholm module for the pair

$$(C[0,1] \otimes A, C[0,1] \otimes A/C_0(0,1) \otimes A)$$

and V is the operator

$$V = -X + (1 - X^2)^{1/2} F$$

Here $X = \gamma\varepsilon_1 X_0$ where γ is the grading operator on H, ε_1 is the first multigrading operator for (ρ, H, F), and X_0 is the image under ρ of the function $t \mapsto t$ in $C[0,1] \otimes A$. Note that at various points in this discussion we passed between suspensions over $(-1,1)$ and suspensions over $(0,1)$; let us identify them using the orientation preserving diffeomorphism $(0,1) \to (-1,1)$ given by $t \mapsto 2t-1$.

Our goal in this section is to prove the following:

Proposition A.2.1. *With the notation above, we have:*

$$s[\rho, H, F] = [\rho, H, V]$$

in $K^{-p}(A)$.

The proof involves explicit calculations with boundary maps in K-theory, and consequently it is sensible to separate the cases where p is even and odd.

A.3 The Proof of Proposition A.2.1 for p even

According to the Bott periodicity theorem $K^{-p-1}(S(A)) \cong K^1(S(A))$, so let us assume (ρ, H, F) is an ungraded relative Fredholm module. Note that (ρ, H, F) is a compact perturbation of $(\rho, H, \frac{1}{2}(F + F^*))$, so assume that F is self-adjoint. By the results of the last section, we can assume that the representation ρ is ample, and the K-homology class of (ρ, H, F) corresponds to the K-theory class of the projection $\frac{1}{2}(F+1)$ in $\mathfrak{D}^*(C(A))/\mathfrak{D}^*(C(A)//S(A))$. We need the following formula:

Proposition A.3.1. *Let J be an ideal in a unital C^*-algebra B. Let $\boldsymbol{p} \in M_n(B/J)$ be a projection and let $\boldsymbol{x}$ be a self-adjoint lift of $\boldsymbol{p}$ to $M_n(B)$. Then the boundary map in K-theory $\partial\colon K_0(B/J) \to K_1(J)$ sends the K-theory class of p to the K-theory class of the unitary* $\exp(2\pi i x)$.

Proof. Reference. ☐

Since we assumed that F is self-adjoint, we have that the image of

$$[\frac{1}{2}(F+1)] \in K_0(\mathfrak{D}^*(C(A))/\mathfrak{D}^*(C(A)//S(A)))$$

in $K_1(\mathfrak{D}^*(C(A)//S(A)))$ is given by the unitary $\exp(\pi i(F+1))$. Let us review the identification

$$K_1(\mathfrak{D}^*(C(A))//S(A)) \cong K_1(A)$$

Let $\sigma \in C_0(0,1]$ be the function $\sigma(t) = t$ and consider the completely positive map $A \to C_0(0,1] \otimes A$ given by $a \mapsto \sigma \otimes a$. This map has an explicit Stinespring dilation

$$\psi\colon A \to \mathbb{B}(H \oplus H)$$

given by

$$\psi(a) = U \begin{pmatrix} \rho(1 \otimes a) & 0 \\ 0 & 0 \end{pmatrix} U$$

where U is the self-adjoint unitary

$$U = \begin{pmatrix} \rho(\sigma^{1/2} \otimes 1) & \rho((1-\sigma)^{1/2} \otimes 1) \\ \rho((1-\sigma^{1/2}) \otimes 1) & -\rho(\sigma^{1/2} \otimes 1) \end{pmatrix}$$

According to the proof of Theorem 3.3.22, the isomorphism

$$K_1(\mathfrak{D}^*(C(A)//S(A))) \cong K_1(A)$$

is induced by the $*$-homomorphism $\mathfrak{D}^*_\rho(C(A)//S(A)) \to \mathfrak{D}^*_\psi(A)$ given by

$$T \mapsto \begin{pmatrix} T & 0 \\ 0 & 0 \end{pmatrix}$$

The induced map sends the unitary $\exp(\pi i(F+1))$ in $\mathfrak{D}^*_\rho(C(A)//S(A))$ to the unitary

$$\begin{pmatrix} \exp(\pi i F) & 0 \\ 0 & -1 \end{pmatrix}$$

Hence we obtain the following formula for the suspension map:

$$s[\rho, H, F] = \left[\psi, H \oplus H, \begin{pmatrix} \exp(\pi i F) & 0 \\ 0 & -1 \end{pmatrix}\right]$$

Conjugating the Fredholm module on the right-hand side with the unitary U and using the fact that F is a *relative* Fredholm module, we obtain:

$$\left[\begin{pmatrix} \rho(1\otimes a) & 0 \\ 0 & 0 \end{pmatrix}, H\oplus H, \begin{pmatrix} \exp(\pi i F)X_0 - (1-X_0) & 0 \\ 0 & -1 \end{pmatrix}\right]$$

where $X_0 = \rho(\sigma\otimes 1)$. Finally, the same homotopy calculation appearing in the proof of Proposition 3.4.18 shows that $\exp(\pi i F)X_0 - (1-X_0)$ is operator homotopic to $V = iX_0 + (1-X_0^2)^{1/2}F$. This completes the proof.

Appendix B

The Kasparov Technical Theorem

Our construction of the Kasparov product in Chapter 3 was contingent on the existence of partitions of unity (Proposition 3.5.4). In this appendix we construct partitions of unity using an intricate result in functional analysis called the *Kasparov Technical Theorem.* The theorem is originally due to Kasparov, but it was subsequently simplified by Higson; the argument here is adapted from chapter 3 of [9]. We will need a tool in C*-algebra theory called a quasi-central approximate unit, which we now review.

Definition B.0.2. *Let A be a separable C*-algebra and let $J \subseteq A$ be an ideal.*

- *An* approximate unit *for A is a sequence $\{u_n\}$ of self-adjoint elements of A such that $0 \leq u_n \leq 1$, $u_m \leq u_n$ if $m \leq n$, and $\|au_n - a\| \to 0$ for every $a \in A$.*

- *An approximate unit $\{u_n\}$ for J is* quasicentral *for A if $\|au_n - u_n a\| \to 0$ for every $a \in A$.*

Every separable C*-algebra has an approximate unit (reference), and every ideal in a separable C*-algebra has an approximate unit which is quasicentral for the whole C*-algebra. We will prove the latter assuming the former.

Proposition B.0.3. *Let J be an ideal in a separable C*-algebra B. Given any approximate unit $\{u_n\}$ for J there is a quasicentral approximate unit $\{v_n\}$ such that each v_n is a finite convex combination of the elements $\{u_n, u_{n+1}, \ldots\}$.*

We will need a strengthening of this result:

Corollary B.0.4. *Let H be a separable Hilbert space, let E be a separable C*-subalgebra of $\mathbb{B}(H)$, and let Δ be a separable linear subspace of $\mathbb{B}(H)$ which derives E. Then there is an approximate unit $\{u_n\}$ for E which is quasicentral for Δ, meaning $\|u_n x - x u_n\| \to 0$ for every $x \in \Delta$.*

Proof. Let A denote the C*-subalgebra of $\mathbb{B}(H)$ generated by Δ and E and let J be the closure of the linear subspace $AE + E \subseteq A$. Note that A derives E since Δ derives E, so J is an ideal in A.

Choose an approximate unit $\{v_n\}$ for E, and note that $\{v_n\}$ is also an approximate unit for J. By Proposition B.0.3 there is a quasicentral approximate unit $\{u_n\}$ for J in A obtained as convex combinations of the $\{v_n\}$. Certainly $\{u_n\}$ is an approximate unit for E, and it is quasicentral for Δ since it is quasicentral for A. $\square$

We need one last lemma before we are ready to formulate and prove the Kasparov technical theorem.

Lemma B.0.5. *Let A be a unital C*-algebra and let $\varphi : [0,1] \to \mathbb{C}$ be a continuous function. Then for every $\varepsilon > 0$ there exists δ such that $\|[a, a']\| < \delta$ implies $\|[\varphi(a), a']\| < \varepsilon$ for any a, a' in the unit ball of A with $a \geq 0$.*

Proof. Let Φ denote the set of all continuous functions φ for which the lemma holds. It is clear that Φ contains the constant functions and the functions $t \mapsto t$ and $t \mapsto it$. It is also clear that Φ is closed under addition, complex conjugation, and uniform limits. Thus by the Stone-Weierstrass theorem it suffices to show that Φ is closed under pointwise multiplication.

Let $\varphi_1, \varphi_2 \in \Phi$ and let $\varepsilon > 0$ be given. Choose δ_1 so that $\|[a, a']\| < \delta_1$ implies

$$\|[\varphi_1(a), a']\| < \frac{\varepsilon}{2\,\|\varphi_2\|_{\sup}}$$

and choose δ_2 similarly. If $\delta = \min\{\delta_1, \delta_2\}$ and $\|[a, a]\| < \delta$ then by the triangle inequality:

$$\|[\varphi_1(a)\varphi_2(a), a']\| \leq \|\varphi_1(a)\|\,\|[\varphi_2(a), a']\| + \|\varphi_2(a)\|\,\|[\varphi_1(a), a']\| < \varepsilon$$

□

We are now ready for our main result:

Theorem B.0.6 (Kasparov Technical Theorem)**.** *Let H be a separable Hilbert space and let E_1, E_2 be separable C^*-subalgebras of $\mathbb{B}(H)$ such that $E_1E_2 \subseteq \mathbb{K}(H)$. Given any separable linear subspace $\Delta \subseteq \mathbb{B}(H)$ which derives E_1 there is a self-adjoint operator $T \in \mathbb{B}(H)$ such that:*

- $0 \leq T \leq 1$
- $(1 - T)E_1 \subseteq \mathbb{K}(H)$
- $TE_2 \subseteq \mathbb{K}(H)$
- $[T, \Delta] \subseteq \mathbb{K}(H)$.

Moreover the set of all self-adjoint operators T which satisfy these three conditions is convex.

Proof. Let $\{x_m\}$, $\{y_m\}$, and $\{z_m\}$ be dense sequences in the unit balls of E_1, E_2, and Δ, respectively.

By Corollary B.0.4 there is an approximate unit $\{u_n\}$ for E_1 which is quasi-central for Δ. In particular $\|u_nx_m - x_m\| \to 0$ as $n \to \infty$ and $\|[u_n, z_m]\| \to 0$ as $n \to \infty$ for each m. Passing to a subsequence of $\{u_n\}$ if necessary, assume that

$$\|u_nx_m - x_m\| < 2^{-n} \qquad \text{and} \qquad \|[u_n, z_m]\| < 2^{-n}$$

whenever $m \leq n$.

Now let $\{w_n\}$ be an approximate unit for $\mathbb{K}(H)$ and set $d_n = (w_n - w_{n-1})^{1/2}$. Note that $\|d_nu_{m_1}y_{m_2}\| \to 0$ as $n \to \infty$ since $u_{m_1}y_{m_2} \in \mathbb{K}(H)$ and w_n is an approximate unit for $\mathbb{K}(H)$. Also note that for any operator $S \in \mathbb{B}(H)$ we have that $\|[d_n, S]\| \to 0$ as $n \to \infty$ by Lemma B.0.5 applied to $A = \mathbb{B}(H)$, $a = w_n - w_{n-1}$, $a' = S$, and $\varphi(t) = t^{1/2}$. Thus at the possible cost of passing to a subsequence we can assume that:

- $\|d_nu_{m_1}y_{m_2}\| < 2^{-n}$
- $\|[d_n, x_m]\| < 2^{-n}$

- $\|[d_n, y_m]\| < 2^{-n}$
- $\|[d_n, z_m]\| < 2^{-n}$

whenever m_1, m_2, and m are no larger than n. The operator T that we want is given by the strong limit of the series

$$T = \sum_n d_n u_n d_n$$

Note that this series really does converge in the strong topology since

$$\sum_{n=1}^{N} d_n u_n d_n \leq \sum_{n=1}^{N} d_n^2 = w_N \leq 1$$

We now verify that T has all of the required propeties. It is clear from the calculation above that $0 \leq T \leq 1$, so we just need to check the three compactness conditions.

- $(1 - T)E_1 \subseteq \mathbb{K}(H)$

  $(1 - T)x_m$ is the strong limit of the series $\sum_n d_n(1 - u_n)d_n x_m$. Each term in the series is compact, so it suffices to show that the series converges in norm. Indeed,

  $$d_n(1 - u_n)d_n x_m = d_n(x_m - u_n x_m)d_n + d_n(1 - u_n)[d_n, x_m]$$

  Since $\|x_m - u_n x_m\| < 2^{-n}$ and $\|[d_n, x_m]\| < 2^{-n}$ whenever $n \geq m$, $(1 - T)x_m$ is bounded in norm by a convergent geometric series for each m. This is enough since x_m is dense in the unit ball of E_1.

- $TE_2 \subseteq \mathbb{K}(H)$

  By a similar argument as above, it suffices to show that the strongly convergent series $Xy_m = \sum_n d_n u_n d_n y_m$ is in fact norm convergent. Here we use the calculation

  $$d_n u_n d_n y_m = d_n u_n y_m d_n + d_n u_n [d_n, y_m]$$

  and argue as before.

- $[T, \Delta] \subseteq \mathbb{K}(H)$

  Again, it suffices to show that $[T, z_m] = \sum_n [d_n u_n d_n, z_m]$ converges in norm. We use:

$$[d_n u_n d_n, z_m] = [d_n, z_m] u_n d_n + d_n [u_n, z_m] d_n + d_n u_n [d_n, z_m]$$

□

We can deduce the existence of the partitions of unity that we needed in order to construct the Kasparov product as a corollary of the Kasparov technical theorem.

Proof of Proposition 3.5.4. Let E_1 be the smallest C*-subalgebra of $\mathbb{K}(H_1) \otimes \mathbb{B}(H_2)$ which contains all elementary tensors $K \otimes 1$ and is derived by Δ. Note that E_1 is separable since $\mathbb{K}(H_1)$ is separable, Δ is separable, and every element of E_1 is the limit of finitely iterated commutators of elementary tensors and elements of Δ. Let E_2 simply be the C*-subalgebra of $\mathbb{B}(H_1) \otimes \mathbb{K}(H_2)$ generated by the elementary tensors $1 \otimes K$. Since Δ derives $\mathbb{K}(H_1) \otimes \mathbb{B}(H_2)$ we have that $E_1 E_2 \subseteq \mathbb{K}(H_1 \otimes H_2)$, so the hypotheses of the Kasparov technical theorem are satisfied; let T be the self-adjoint operator that it guarantees. The desired partition of unity is given by $N_1 = (1 - T)^{1/2}$ and $N_2 = T^{1/2}$. To ensure compatibility with the multigrading structure, include the multigrading operators in Δ and average T over the finite group of automorphisms of $\mathbb{B}(H)$ that they generate to obtain a compact perturbation of T which commutes with them on the nose. □

Bibliography

[1] M.F. Atiyah. Global theory of elliptic operators. *Proc. Conf. Functional Analysis*, 1969.

[2] M.F. Atiyah and Friedrich Hirzebruch. *Essays on Topology and Related Topics (Memoires dedies a Georges de Rham)*. Springer, 1970. Spin-Manifolds and Group Actions.

[3] M.F. Atiyah and I.M Singer. The index of elliptic operators: I. *The Annals of Mathematics*, 1968.

[4] B. Blackardar. *K-theory for Operator Algebras*. Springer-Verlag, 1986.

[5] Joachim Cuntz and Nigel Higson. Kuiper's theorem for hilbert modules. *Contemp. Math.*, (62), 1987.

[6] R.G. Douglas. *Banach Algebra Techniques in Operator Theory*. Boston, 1972.

[7] Mikhael Gromov and Blaine Lawson. Positive scalar curvature and the dirac operator on complete riemannian manifolds. *Inst. Hautes Etudes Sci. Publ. Math.*, (58):83–196, 1984.

[8] Nigel Higson. A note on the cobordism invariance of the index. *Topology*, 1991.

[9] Nigel Higson and John Roe. *Analytic K-Homology*. Oxford University Press, 2000.

[10] Nigel Higson and John Roe. Mapping surgery to analysis I: Analytic signatures. *K-Theory*, 2005.

[11] Nigel Higson and John Roe. Mapping surgery to analysis II: Geometric signatures. *K-Theory*, 2005.

[12] Nigel Higson and John Roe. Mapping surgery to analysis III: Exact sequences. *K-Theory*, 2005.

[13] J.L. Kazdan and F.W. Warner. Scalar curvature and conforrmal deformation of riemannian structure. *Journal of Differential Geometry*, (10):113–134, 1975.

[14] H.B. Lawson and M.L. Michelsohn. *Spin Geometry.* Princeton University Press, 1995.

[15] John Roe Nigel Higson and Guo Liang Yu. A coarse mayer-vietoris principle. *Math. Proc. Cambridge Philos. Soc.*, 1993.

[16] A. Connes P. Baum and N. Higson. Classifying space for proper g-actions and k-theory of group c*-algebras. *Contemporary Mathematics*, 167:241–291, 1994.

[17] W.L. Paschke. K-theory for commutants in the calkin algebra. *Pacific Journal of Mathematics*, (95):427–437, 1981.

[18] J. Roe. *Lectures on coarse geometry.*

[19] J. Roe. An index theorem on open manifolds i and ii. *Journal of Differential Geometry*, (27):87–113, 1988.

[20] John Roe. Partitioning noncompact manifolds and the dual toeplitz problem. In *Operator Algebras and Applications, Vol 1.* Cambridge University Press, 1988.

[21] W. Rudin. *Functional Analysis.* McGraw-Hill, 1973.

Vita

Paul Siegel

Paul Siegel was born in Lansing, Michigan in 1984. His interest in mathematics first emerged in middle school when he received a copy of a book on algebra and calculus for engineers as a gift. He attended East Lansing High School from 1999 to 2003, and his interest in mathematics was rivalled only by his interest in philosophy and political theory nurtured by his participation in academic debate. Mathematics ultimately came to dominate: he took several math classes at Michigan State University while still in high school and participated in a mathematical biology research project at Michigan State during the summer before his senior year.

He attended the University of Michigan from 2003-2007 with a partial scholarship to study mathematics. He participated in two undergraduate research projects in mathematics, the first of which led to an article co-authored with Prof. Igor Kriz in the July 2008 issue of Scientific American. Aside from his academic interests, he developed strong interests in the board game Go and the card game Bridge; he competed regularly in tournaments for both games. His mathematical interests gravitated toward geometry and topology, and these interests led him to attend graduate school in mathematics at Penn State in 2007.

His interest in and understanding of geometry and topology deepened at Penn State, particularly under the supervision of his thesis adviser, John Roe. He attended several conferences during his time at Penn State, including Michael Atiyah's 80th birthday conference in Edinburgh. He won the math department's "ZZRQ" award for contributing to a sense of community and a teaching award for his role as a teaching assistant during the 2011 MASS program. His non-academic interests at Penn State included swing dancing and volunteer work. He is now a Ritt Assistant Professor at Columbia University.