A Noncommutative Proof of Bott Periodicity

June 12, 2004

1 Introduction

Bott periodicity in K-theory is a rather mysterious object. The classical proofs typically consist of showing that the unitary groups form an Ω-spectrum from which to get a cohomology theory; then showing that that theory is K-theory; and most formidably showing that $U(n)$ is homotopic to $U(n+2)$ for all n.

However, Cuntz showed that Bott periodicity can be derived in a much simpler way if one resorts to using not only traditional topological spaces, but also "non-commutative topological spaces," i.e. C^*-algebras. That in fact, Bott periodicity is not only reflected in the collective topology of the unitary groups, but also encoded in the structure of the Toplitz C^*-algebra.

In next two sections we will overview the essential ingredients of non-communitive topology that are needed to follow Cuntz's argument. Namely the analogues of topological spaces, continuous functions, suspensions, and vector bundles. These two sections are rather sketchy, without any actual proofs given. In the third section we will use the functorial properties of K_0 to compute the K-theory of the Toplitz C^*-algebra, and with this we will be able to give Cuntz's quick proof of Bott periodicity.

It should be noted that Cuntz's proof does not actually even use the definition of K_0, just a couple of functorial properties. So the curious reader can skip to the proof of Bott periodicity without any extra difficulties.

Finally, it should also be noted that tensor products are rather ubiquetous in the following. It's a rather unfortunate fact that in general the tensor product of two C^*-algebras is not uniquely defined. One would like to simply that the algebraic tensor product and complete with respect to a C^*-norm, but in general there is more than one C^*-norm on the algebraic tensor product. However, the tensor product is uniquely defined if at least one of the two factors happens to be nuclear. Examples of nuclear C^*-algebras include the commutative ones; the algebra $\mathbb{K}$ of compact operators on an infinite dimensional, seperable Hilbert space; and the Toplitz operators. Luckily, in all of the tensor products below at least one of the factors is nuclear. Also, there is good side to C^*-algebra tensor products. As long as at least one of the factors in each tensor product is nuclear, it's a fact that taking the tensor product of a short exact sequence with a fixed C^*-algebra yields another short exact sequence (i.e., there are no Tor terms). This will come in handy at one point below. The interested reader will find a very nice appendix on tensor products in Wegge-Olsen's $K-Theory\ and\ C^*-Algebras$.

All material for this paper can be found in Arveson's *A Short Course in Spectral Theory* or Wegge-Olsen's $K-Theory\ and\ C^*-Algebras$.

2 Basic ideas of Noncommutative Topology

Let X be a compact Hausdorff space. Then it is useful to consider $C(X)$, the set of continuous functions $f : X \to \mathbb{C}$. $C(X)$ forms an associative algebra with identity under pointwise sum and product. It has an involution, namely $f^*(x) = \overline{f(x)}$, and it is a complete metric space with the uniform norm. Moreover, the norm and the involution are related by the formula $\|f^*f\| = \|f\|^2$. There is one more property of $C(X)$ that we would like to point out, namely, it is commutative. A noncommutative topological space, or C^*-algebra, is defined to be an algebra with the above properties except the commutivity. Namely, a unital C^*-algebra A is a unital associative algebra with an involution and a submultiplicative norm $\| \centerdot \|$ such that A is complete with respect to $\| \centerdot \|$ and $\|a^*a\| = \|a\|^2$ for all $a \in A$.

A continuous map $\phi : X \to Y$ induces a corresponding map $\bar{\phi} : C(Y) \to C(X)$ by $\bar{\phi}(f) = f \circ \phi$. $\bar{\phi}$ is a unital *-homomorphism from $C(Y)$ to $C(X)$. Motivated by this, one declares a morphism from A to B in the category of unital C^*-algebras to be a unital *-homomorphism $\phi : A \to B$. One checks that any morphism of unital C^*-algebras is automatically norm continuous. Most importantly, Gelfand theory implies our associations $X \longrightarrow C(X)$ and $(\phi : X \to Y) \longrightarrow (\bar{\phi} : C(Y) \to C(X))$ gives an equivalence of the category of compact Hausdorff spaces with the subcategory of commutative unital C^*-algebras. This justifies the philosophy that arbitrary unital C^*-algebras correspond to noncommutative compact Hausdorff spaces.

We will also have need for locally compact, but noncompact Hausdorff spaces. If X is locally compact but not compact, then we can form $C_0(X)$, the set of continuous $f : X \to \mathbb{C}$ that vanish at infinity; that is, for any ϵ there is a compact $K \subset X$ such that $\|f(x)\| < \epsilon$ for $x \in X - K$. $C_0(X)$ has all the properties above except that it is nonunital. A proper map $\phi : X \to Y$ induces a *-homomorphism $\bar{\phi} : C_0(Y) \to C_0(X)$. Continuous functions can also be treated, but it is more complicated. Fortunately we will not have need of the more general morphisms motivated by arbitrary continuous functions, so we will say a morphism from a nonunital C^*-algebra to a C^*-algebra is simply a *-homomorphism.

3 K-theory and Suspension

First we need to understand what K_0 for an arbitrary C^*-algebra should be. We restrict the motivating discussion to the case of unital C^*-algebras, just like in topology where K-theory is first described for compact spaces. In the topological case, the K-theory for a locally compact space is then defined to the K-theory for its one point compactification, quotiented out by a copy of $\mathbb{Z}$. In the same way, for nonunital C^*-algebras one first looks at the K-theory of the unitalization, and then quotients out by a copy of $\mathbb{Z}$.

Recall that K^0 is a contravariant functor for topological spaces, but since the equivalence of topological spaces with C^*-algebras reverses arrows, our corresponding functor K_0 for C^*-algebras will be covariant. In the topological space case, for a compact X we have that $Vect(X)$ is the cancelative semigroup generated by stable equivalence classes of vector bundles over X with addition induced from the Whitney sum. Then $K^0(X)$ is the Grothendick group of $Vect(X)$; that is, the group of formal differences of elements of $Vect(X)$. By Swan's Theorem, stable equivalence classes of vector bundles over X correspond to stable equivalence classes of finitely generated projective modules over $C(X)$, with addition given by the direct sum. (Two projective modules over a ring are called stably equivalent if they can be made isomorphic by direct summing them with finitely generated free modules of large enough size). So we define for a unital C^*-

algebra A the K_0 group to be the Grothendick group of the cancelative semigroup generated by the stable equivalence classes of finitely generated projective right A-modules with addition of classes induced from the direct sum.

We also need to think about what K_0 does to morphisms of C^*-algebras. If $\phi : A \to B$ is a *-homomorphism, then we can B as a left A-module by $a \cdot b = \phi(a)b$. Thus given a right A-module M we can form a corresponding right B-module by $M \mapsto M \otimes_A B$. This association clearly takes stably equivalent A-modules to stably equivalent B-modules, so we get an induced map $\phi_* : K_0(A) \to K_0(B)$. This enables us (among other things) to define K_0 for non-unital algebras; namely, for a nonunital C^*-algebra A we define $K_0(A) = ker(\phi_* : K_0(\mathbb{C}) \to K_0(\tilde{A}))$, where $\tilde{A}$ is the algebra obtained by adjoining a unit to A, and $\phi : \mathbb{C} \to \tilde{A}$ is the natural inclusion $\lambda \mapsto \lambda 1_{\tilde{A}}$.

There are three critical properties of K_0, its stability, homotopy invariance, and half exactness. The stability property of K_0 says that $K_0(A) = K_0(A \otimes \mathbb{K})$ for any C^*-algebra A, where $\mathbb{K}$ is the C^*-algebra of compact operators on an infinite dimensional, seperable Hilbert space. Indeed, we see that $K_0(A) = K_0(A \otimes M_n(\mathbb{C}))$ for any $n \in \mathbb{N}$ because A and $A \otimes M_n(\mathbb{C})$ are Morita equivalent rings, as in just plain old algebra. (Two rings are called Morita equivalent if they have equivalent categories of right modules. The standard example is to take a ring R and $M_n(R)$, for the map $N \mapsto N \otimes_R M_n(R)$ induces an equivalence. Since projectiveness of a module has a categorical interpretation– N is a projective A-module iff the functor $Hom_A(N, -)$ is exact– it follows that $K_0(A)$ depends only on the Morita equivalence class of A.) A typical direct limit argument shows that K_0 preserves direct limits (as does $\otimes$), so we indeed get $K_0(A) = K_0(A \otimes \mathbb{K})$.

The homotopy invariance of K_0 says that an induced map $\phi_* : K_0(A) \to K_0(B)$ only depends on the homotopy class of $\phi : A \to B$. To see this, it is helpful to look at another approach to constructing K_0. First note that a finitely generated projective module over A corresponds to a direct summand of A^n, which in turn corresponds to a idempotent in $M_n(A)$. While it's a little bit subtle to determine exactly when two such idempotents correspond to stably equivalent modules, it's certainly true that it is sufficient that the idempotents be similar. Any *-homomorphism $\phi : A \to B$ will send idempotents to idempotents, and a little chasing shows that this sending idempotents to idempotents corresponds to the $\phi_* : K_0(A) \to K_0(B)$ constructed above from the module viewpoint. We want to see that K_0 is actually homotopy invariant, i.e. a stongly continuous family of *-homomorphisms $\{\phi_t\}_{t \in [0,1]}$ all induce the same map on the K_0 groups. (A family of morphisms is called strongly continuous if $(t \mapsto \phi_t(a))$ is continuous for any fixed $a \in A$.) By above, it suffices to show that if p and q are homotopic idempotents, then p and q are similar. By dividing the path from p to q into pieces if necessary, we can assume that $\|p - q\| < 1$. Let $z = (2p - 1)(2q - 1) + 1$. Then direct computation shows that $\|z - 2\| < 2$, so z is invertible, and that $qz = zp$. Thus p and q are similar and so K_0 is homotopy invariant.

The final important property of K_0 is its half exactness. A (covariant) functor F is called half exact if given a short exact sequence $0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$, then the induced sequence $F(A) \longrightarrow F(B) \longrightarrow F(C)$ is exact. The proof that K_0 is half exact is again to just take the corresponding result for topological spaces, and translate the proof into algebraic language. Since the goal of this paper is to show how Bott periodicity becomes simplier when looked at from a noncommutative viewpoint, in the spirit of the paper I'm going to omit this translation.

Recall that in classical topology we have that given any half exact, homotopy invariant covariant (contravariant) functor F, we can get a (co)homology theory (with a long exact sequence, split exact additive functors, etc) by defining $H_{-n}(X) = F(S^n X)$, where the right hand side denotes the iterated suspension. (It's also easy to see that every homotopy invariant homology theory comes about in this way, using the short exact sequence with X, it's suspension, and it's cone.)

So to define the higher K-groups, we need to understand suspension. The classical definition of suspension is to multiply the space X by $[0,1]$ and then collapse the ends. The corresponding operation to Cartesian product of the spaces is tensor product of the function algebras, for if X and Y are locally compact Hausdorf then we have a map from the algebraic tensor product $C(X)\otimes_{alg}C(Y) \to C(X\times Y)$ by $(f\otimes g)(x,y) = f(x)g(y)$. Stone Weierstrauss gives that this map has dense range, so that the C^*-algebraic tensor product identifies $C(X)\otimes C(Y) = C(X\times Y)$. In our case we have $Y = [0,1]$. The collapsing of the ends means that we need each of the functions $f\otimes g$ to be constant in the x variable at the ends of $[0,1]$. This forces the function g to vanish at the endpoints. Such a function g is easily identified with an element of $C_0(\mathbb{R})$, so we define the suspension of a C^*-algebra A to be $SA \equiv C_0(\mathbb{R})\otimes A$. Now that we have suspension, we can define the higher K-groups by $K_{-n}(A) = K_0(S^nA)$. Pretty much the same (long) proof as in the topoligical case shows that the homotopy invariance of K_0 implies that the K_n's form a homology theory.

4 The Proof of Bott Periodicity

Cuntz's proof of Bott periodicity only requires two facts about K_0:

•)K_0 is a half exact, homotopy invariant functor

•)$K_0(A) = K_0(A\otimes\mathbb{K})$ for any A, where $\mathbb{K}$ denotes the compact operators on the seperable, infinite dimensional Hilbert space

As we mentioned in the last section, the first fact implies that defining $K_{-n}(A) = K_0(S^nA)$ where $SA = C_0(\mathbb{R})\otimes A$ gives us a homotopy invariant homology theory (which has a long exact sequence and split exact, additive functors). What we will show in this section, is that this together with the second fact gives us Bott periodicity.

The main step is to calculate the K-theory of the Toplitz C^*-algebra $\mathcal{T}$. $\mathcal{T}$ is defined as the C^*-algebra generated by the unilateral shift S, defined on an orthonormal basis $\{e_n\}_{n\in\mathbb{N}}$ by $Se_n = e_{n+1}$. Note that $S^*S = I$, so $\mathcal{T}$ is unital, and $I - SS^*$ is the projection onto the linear subspace spanned by e_1. Knowing exactly what this algebra looks like is not important, the only other thing that is needed is that there is a short exact sequence: $0 \longrightarrow \mathbb{K} \longrightarrow \mathcal{T} \longrightarrow C(\mathbb{T}) \longrightarrow 0$, where $C(\mathbb{T})$ denotes the set of continuous complex valued functions on the circle group $\mathbb{T}$. (See Arveson's book for details). We have:

Theorem 4.1 *There is a canonical surjection $q : \mathcal{T} \to \mathbb{C}$ such that $q_* : K_0(\mathcal{T}) \to K_0(\mathbb{C})$ is an isomorphism.*

Notice that $K_0(\mathbb{C}) = \mathbb{Z}$, for a finitely generated right $\mathbb{C}$ module is characterized up to stable equivalence by its dimension, so the Grothendick group you get is $\mathbb{Z}$.

Proof: Let $q : \mathcal{T} \to \mathbb{C}$ be obtained by composing the quotient map $\mathcal{T} \to \mathcal{T}/\mathbb{K} \simeq C(\mathbb{T})$ with evaluation at $1 \in \mathbb{T}$. Let $j : \mathbb{C} \to \mathcal{T}$ be the map sending 1 to the identity of $\mathcal{T}$. It is clear that $q_*\circ j_* = id_{K_0(\mathbb{C})}$. To obtain $j_*\circ q_* = id_{K_0(\mathcal{T})}$, we will need to use the above facts about K_0.

Consider the rank 1 projection $I - SS^*$ in $\mathcal{T}$ and the standard embedding $\sigma : \mathcal{T} \to \mathbb{K}\otimes\mathcal{T}$ given by $T \mapsto (I - SS^*)\otimes T$. That is, view T as being a 1 by 1 matrix in $\mathbb{K}\otimes\mathcal{T}$. Since $\sigma_* : K_0(\mathcal{T}) \to K_0(\mathbb{K}\otimes\mathcal{T})$ is assumed to be an isomorphism, it is sufficient to prove that $\sigma_*\circ j_*\circ q_* = \sigma_*$. In order to do this, we construct another C^*-algebra $\overline{\mathcal{T}}$ and a morphism $\iota : \mathbb{K}\otimes\mathcal{T} \to \overline{\mathcal{T}}$ such that ι_* is injective and $\iota_*\circ\sigma_*\circ j_*\circ q_* = \iota_*\circ\sigma_*$.

To start, let $\widehat{\mathcal{T}} = C^*(\mathbb{K} \otimes \mathcal{T}, \mathcal{T} \otimes 1) \subset \mathcal{T} \otimes \mathcal{T}$. Note that $\mathbb{K} \otimes \mathcal{T}$ is an ideal in $\widehat{\mathcal{T}}$ and that $\widehat{\mathcal{T}}/(\mathbb{K} \otimes \mathcal{T}) \simeq \mathcal{T}/\mathbb{K} \simeq C(\mathbb{T})$. Thus we can define $\overline{\mathcal{T}}$ as the fibre product of the quotient maps $\pi : \mathcal{T} \to C(\mathbb{T})$ and $\hat{\pi} : \widehat{\mathcal{T}} \to C(\mathbb{T})$ That is: $\overline{\mathcal{T}} \equiv \{x \oplus y \in \widehat{\mathcal{T}} \oplus \mathcal{T} | \hat{\pi}x = \pi y\}$.

Putting $\iota(x) \equiv x \oplus 0, \bar{\pi}(x \oplus y) \equiv y$, and $\gamma(y) = (I - SS^*) \otimes y \oplus y$ gives us a split exact sequence $0 \longrightarrow \mathbb{K} \otimes \mathcal{T} \longrightarrow \overline{\mathcal{T}} \longleftrightarrow \mathcal{T} \longrightarrow 0$, so $\iota_* : K_0(\mathbb{K} \otimes \mathcal{T}) \to K_0(\overline{\mathcal{T}})$ is injective by split exactness of K_0, as desired.

Now let's actually compute the compositions $\iota \circ \sigma$ and $\iota \circ \sigma \circ j \circ q$. (Both map $\mathcal{T} \to \overline{\mathcal{T}}$). For the first one we have $\iota \circ \sigma(T) = \iota((I - SS^*) \otimes T) = (I - SS^*) \otimes T \oplus 0$. For the second we have $\iota \circ \sigma \circ j \circ q(T) = (I - SS^*) \otimes I \oplus 0$. Let $\beta(T) = S^2 S^* \otimes I \oplus T$. Then the range of β is orthogonal to that of $\iota \circ \sigma$ in $\overline{\mathcal{T}}$, so the desired equality will follow from the additivity and the homotopy invariance of K_0 once we can show that $\beta_0 \equiv \beta + \iota \circ \sigma$ and $\beta_1 \equiv \beta + \iota \circ \sigma \circ j \circ q$ are homotopic.

To simplify notation, define $v \equiv S \otimes I, w \equiv (I - SS^*) \otimes S$, and $e \equiv (I - SS^*) \otimes (I - SS^*)$. Consider the two elements of $\widehat{\mathbb{T}}$ given by $u_0 = v^2 v^* + w v^* + v w^* + e$ and $u_1 = v^2 v^{*2} + (1 - vv^*) v^* + v(1 - vv^*)$. Then a simple calculation gives that $\beta_i(T) = u_i(T \otimes I) \oplus T$. The reason for picking these particular u_i rather than one of the simplier expressions that would give the same relation with the β_i is that these two u_i's are actually symmetries. (An operator is called a symmetry if $u = u^*$ and $u^2 = I$.) Indeed, the u_i are manifestly selfadjoint. Checking that $u_i^2 = I$ for each i is a little annoying, but it just comes down to playing around with the unilateral shift. The importance of this fact is that since the u_i are symmetries, we have that the spectrum of u_i is contained in $\{-1, 1\}$ for each i. In particular, we can define a branch of the logrithm on the spectrum of each u_i, and so get a continuous path from each u_i to the identity through unitaries in $\widehat{\mathcal{T}}$. In particular, the u_i are homotopic to each other through unitaries in $\widehat{\mathcal{T}}$, which gives that the two β_i's are homotopic aswell.

□

With the K-theory of the Toplitz C^*-algebra in hand, we can now prove Bott Periodicity rather quickly.

Theorem 4.2 *Bott Periodicity: For any C^*-algebra A we have a natural isomorphism $K_0(A) = K_0(S^2 A)$. In other words, $K_0(A) = K_{-2}(A)$.*

Proof: Let q, j be as in the previous theorem and let $\mathcal{T}_0 = ker(q)$. Then we have a short exact sequence $0 \longrightarrow \mathcal{T}_0 \longrightarrow \mathcal{T} \longleftrightarrow \mathbb{C} \longrightarrow 0$.

As mentioned before, K_0 is a split exact functor, as is tensoring with a fixed C^*-algebra A, so we have another split exact sequence:

$0 \longrightarrow K_0(A \otimes \mathcal{T}_0) \longrightarrow K_0(A \otimes \mathcal{T}) \longleftrightarrow K_0(A \otimes \mathbb{C}) \longrightarrow 0.$

By the previous theorem, the map $(id_A \otimes q)_*$ is an isomorphism, so exactness implies that $K_0(A \otimes \mathcal{T}_0) = 0$ for any C^*-algebra A.

By definition of $\mathcal{T}_0$, it also fits into another short exact sequence $0 \longrightarrow \mathbb{K} \longrightarrow \mathcal{T}_0 \longrightarrow C_0(\mathbb{R}) \longrightarrow 0$. Since for there is no Tor terms for nice enough C^*-algebras that the tensor product is unambiguously defined, we can tensor this sequence with A to obtain another short exact sequence:

$0 \longrightarrow A \otimes \mathbb{K} \longrightarrow A \otimes \mathcal{T}_0 \longrightarrow SA \longrightarrow 0,$

where we recalled the definition $SA \equiv A \otimes C_0(\mathbb{R})$. Now the long exact sequence for K-theory along with the fact that $K_{-1}(A) = K_0(SA)$ gives us an exact sequence

$K_0 S(A \otimes \mathbb{K}) \longrightarrow K_0 S(A \otimes \mathcal{T}_0) \longrightarrow K_0 S^2 A \longrightarrow K_0(A \otimes \mathbb{K}) \longrightarrow K_0(A \otimes \mathcal{T}_0) \longrightarrow K_0 SA.$

We observed above that $K_0(A \otimes \mathcal{T}_0) = 0$. Also by associativity (up to isomorphism) of tensor products we have $K_0 S(A \otimes \mathcal{T}_0) = K_0(SA \otimes \mathcal{T}_0) = 0$. Thus exactness gives us that

$K_0 S^2 A \simeq K_0(A \otimes \mathbb{K}) \simeq K_0(A)$, and the isomorphisms are natural by the naturality of the long exact sequence and the naturality of the stability isomorphism.

$\square$