Noncommutative vector bundles

The Greatest Hits™ of the K-theory of C^*-algebras

Keeley Hoek

June 11, 2018

Contents

0 Abstract

In this report we present famous and fundamental selections (the "Greatest Hits™") from the K-theory of C^*-algebras, motivated by comparison to the "commutative special case" of topological K-theory. We aim to survey by example the ways in which techniques in analysis, homological algebra, and topology, are woven together to obtain the proofs of the C^*-algebraic results. Familiarity with the definition of topological K-theory and its basic properties is assumed.

Not only is C^*-algebraic K-theory a method for transporting tools from topological K-theory to the noncommutative setting, but it also finds applications in the study of honestly topological objects. For instance, in (commutative!) differential geometry one can easily stumble upon orbifolds or other so-called "singular" spaces, which are no longer manifestly topological. These spaces can often be represented on equal footing with the original topological objects by passing to the more general noncommutative world, and this can have great practical utility. Moreover, the K-theory of C^*-algebras and *its* generalisations—such as to Kasparov's KK-theory—are the natural regime in which to attempt to generalise index theorems (such as that of Atiyah and Singer), and these results have even found concrete applications in theoretical physics.

We first survey the (familiar) fundamental properties of the functor K_0, before going on to establish one of the historically first great achievements of the theory; the classification of Almost Finite (AF) C^*-algebras. We then briefly highlight the higher K-groups, in order to arrive at the Bott Periodicity theorem, and its generalisations in the noncommutative world. We conclude by connecting the C^*-theory, via algebraic K-theory, back to the topological one—concretely establishing the rigorous sense in which the latter is a special case of the former.

Our exposition is structured with the goal of rapidly introducing the functors and constructions under consideration, while largely evading complex calculations in matrices or technical results in C^*-algebras by appealing to citations. We are effectively forced to do this in order to have access to some of the most powerful and fundamental tools of the theory—which we will then be able to apply in following sections to prove interesting and critical theorems in K-theory that are built on these results.

1 Why study the K-theory of C^*-algebras?

It would be dishonest to introduce C^*-algebras and their K-theory without first providing some insight into the reasons why they have historically been of interest to mathematicians. We provide a principal motivating example here, deferring precise mathematical descriptions until the next section. The story of C^*-algebras begins with the following theorem of Gelfand and Naimark.

Theorem 1.1 (Gelfand–Naimark)**.** *Every C^*-algebra A is isometrically $*$-isomorphic to a C^*-algebra of bounded operators on a Hilbert space H.*

Moreover, if the C^*-algebra A is separable, then we can ensure that the Hilbert space H is as well. Given that there is only one infinite-dimensional separable Hilbert space, this might seem particularly disheartening to one sceptical of the potential for possible C^*-algebra structures. Fortunately, this theorem does not remotely confine the C^*-algebra structures which occur in practice! (And from now on, all of our Hilbert spaces will be assumed separable.)

The real reason for our interest in C^*-algebra lies in another theorem which is also ascribed to Gelfand, and on which we will often rely;

Theorem 1.2 (Gelfand duality)**.** *There is a contravariant functor*

$$C_0 : \mathsf{CptHaus} \to \mathsf{AbC}^*_I$$

from the category of compact spaces to the category of commutative unital C^*-algebras. Furthermore, this functor is part of an equivalence (in the sense that it is as a covariant functor $\mathsf{CptHaus} \to {\mathsf{AbC}^*_I}^{\mathrm{OP}}$*).*

For each compact Hausdorff space X, this functor simply yields[1] the set $C_0(X)$ of continuous complex valued functions on X. From the definitions to appear in the next section, this object is easily seen

[1] We may occasionally deal with only locally compact Hausdorff spaces, in which case C_0 denotes the set of complex valued continuous functions which *vanish at infinity*.

to be a commutative unital C^*-algebra and is called the *spectrum of* X.

This is a principal example of the algebra–geometry duality in mathematics. Indeed, the correspondence is much more fundamental than a bijection of objects; one finds that almost every topological notion associated to a space X can be transported to a dual algebraic notion for $C_0(X)$. The fact that $C_0(X)$ is always commutative typically has little bearing on the algebraic statement of these properties, and so in fact we can ask whether arbitrary noncommutative C^*-algebras posses them as well!

We've simply dropped the commutativity hypothesis in our domain of discourse, and so an immediate geometric interpretation of the subject at hand is no longer available. Nonetheless, we can now say that we are studying *noncommutative topology*! This duality (in a more general sense) also extends to a much wider class of objects. For example, the study of *noncommutative probability* amounts to the study of *von Neumann algebras* (a particular restricted class of C^*-algebras) and maps from them, while one can also study *noncommutative differentiable manifolds* and *noncommutative affine and projective schemes*. Despite their apparent "abstractness", these subjects can still find extremely concrete applications; for example, the aforementioned case of noncommutative probability actually realises a formal foundation of the theory of quantum mechanics in physics.

In the presence of a K-theory for compact topological spaces, the Gelfand duality suggests that we should search for an extension to the noncommutative world. This is the primary goal of the coming sections.

2 A primer on C^*-algebras

Before we begin our development of the main theory, we must be aware of its basic definitions and constructions. Moreover, in order to avoid laborious and/or irrelevant detail, we rapidly survey these requirements providing external references where necessary. This section really must be much more of a directory than a primer!

Throughout the entirety of this document we will assume that all of our spaces are Hausdorff. Of course, we must first work to define a C^*-algebra.

Definition 2.1. A *Banach algebra* A is a non-necessarily-unital associative algebra over the field $\mathbf{C}$ (or sometimes $\mathbf{R}$) which is also a Banach space, and such that the algebra and norm structures are compatible in that $\|xy\| \leq \|x\|\|y\|$ for every $x, y \in A$. A map $f : A \to B$ of Banach algebras is then just a simultaneous morphism of the Banach space and algebra structures.

An archetypal example is that of the set $\mathscr{B}(X)$ of bounded linear operators from a Banach space X to itself. If X is in addition a Hilbert space, then $\mathscr{B}(X)$ becomes a C^*-algebra;

Definition 2.2. A *C^*-algebra* A is a complex Banach algebra equipped with an involution $^* : A \to A$ satisfying the following properties:

- linearity and antimultiplicativity,
$$(x+y)^* = x^* + y^* \text{ and } (xy)^* = y^* x^* \text{ for all } x, y \in A,$$
- and scalar antimultiplicativity,
$$(\lambda x)^* = \overline{\lambda} x^* \text{ for all } \lambda \in \mathbf{C} \text{ and } x \in A,$$
which together make A into a $*$-algebra, and
- for which the C^*-identity
$$\|x^* x\| = \|x^*\|\|x\| \text{ for all } x \in A$$
holds.

  An element $x \in A$ such that $x = x^*$ is called *self-adjoint*, and if $x^* = x^{-1}$ then x is called *unitary*.

A morphism $f : A \to B$ of C^*-algebras is then just a morphism of algebras which also respects the ambient involution, in that $f(x^*) = f(x)^*$ for every $x \in A$.[2]

It can be the case that a C^*-algebra A lacks a unit. In this case, it will be convenient for our purposes to formally adjoin one by forming the direct sum of vector spaces $A^+ = A \oplus \mathbf{C}$; we equip pairs $(x,\lambda),(y,\mu) \in A \times \mathbf{C}$ with the multiplication

$$(x,\lambda)(y,\mu) = (xy + \mu x + \lambda y, \lambda\mu).$$

If A was nonunital, then from this multiplication A^+ has the new unit $(0,1)$. Of course, we must technically now define a norm on A^+ and explain how it is once again a C^*-algebra. This is a somewhat awkward (though elementary) task, and so we refer the reader to [7] for the details. Most importantly for this construction, we have a short exact sequence $0 \to A \to A^+ \to A^+/A \to 0$ which always splits when A is already unital—therefore, for A unital we have $A^+ \cong A \oplus \mathbf{C}$ with this the full direct sum of C^*-algebras (with multiplication is defined componentwise and the norm is the sum of the norms of each summand). Actually, when A is nonunital A^+ is in a rigorous sense the minimal unital C^*-algebra containing A, but this is already beyond the scope of this introduction. Viewed under the Gelfand duality, one finds that A^+ is exactly one-point compactification; if X is compact and we let $A = C_0(X)$ then by forming A^+ we have simply disjoint unioned-in a point, while for X non-compact the A^+ construction gives rise to a compact space (with only one extra point compared with X). Moreover, just as in the case[3] of the one-point compactification, adjoining a unit with a superscript "+" is a functor; for every map $f : A \to B$ there is a(n obvious) canonical map $A^+ \to B^+$.

We also have quotients and tensor products of C^*-algebras as one would normally expect, with norms defined as they would be for Banach spaces. In fact, it will prove very important to be aware that the category of C^*-algebras has direct limits—though they are actually slightly subtle to define, see [19]. Because our C^*-algebras will be in general noncommutative, each can have both left and right ideals—we will use "ideal" to mean a left ideal, but the distinction will be largely immaterial. Of particular importance is the ideal of A^+ denoted by (abuse of notation) A, which is obtained from the natural inclusion of A into A^+.

Prototypical examples of C^*-algebras are $\mathbf{C}$ and the $n \times n$ matrices with coefficients in $\mathbf{C}$—the latter of which we will denote by $\mathbf{M}_n(\mathbf{C})$. Indeed, for arbitrary C^*-algebras A the matrix construction $\mathbf{M}_n(A)$ continues to make sense, and will be heavily utilised (we take any one of the typical and equivalent norms on $n \times n$ matrices). We will also require $n \times m$ matrices with entries in A, which we will denote by $\mathbf{M}_{n\times m}(A)$. Throughout the rest of this document we will use H to refer to a fixed Hilbert space, for which both the sets of bounded linear operators $\mathscr{B}(H)$ on H and the compact operators $\mathscr{K}(H)$ make sense (both of these sets are also C^*-algebras). Finally, the Calkin algebra $\mathscr{B}(H)/\mathscr{K}(H)$ (being the quotient of two C^*-algebras) is also very easily within reach. For H separable—which we will henceforth assume—this set famously consists of two elements, $0+\mathscr{K}(H)$ and $I+\mathscr{K}(H)$. Taking the tensor product of a C^*-algebra A with $\mathscr{K}(H)$ is called *stabilisation*, for reasons which will become apparent as we develop the theory. We say two C^*-algebras are *stably isomorphic* if their stabilisations are isomorphic.

Of great importance to in the definition of the higher K-groups will be the algebras $\mathrm{GL}_n(A)$, defined to the the subalgebra of $\mathbf{M}_n(A)$ consisting of invertible elements. We will also need in particular the algebra $\mathrm{GL}_n^+(A)$, defined by

$$\mathrm{GL}_n^+(A) = \{x \in \mathrm{GL}_n(A^+) : \pi_A(x) = I_n\}$$

for I_n the $n \times n$ diagonal matrix in $\mathrm{GL}_n(A^+)$ with nonzero entries the unit adjoined to A. Here, the "+" is to be thought of not as referring to positivity, but instead the fact that we have adjoined a unit to A and then normalised to obtained a kind of *normalised general linear group*.

3 The functor K_0

We are now prepared enough to set forth into the world of the K_0-groups of C^*-algebras. We will endeavour to avoid throwing ourselves headlong into the possibly enormous number of technicalities

[2]Note that such a map is automatically bounded and even of norm at most one; this is an interesting situation where the algebra actually completely controls the topology.

[3]The astute reader will note that in the case of only *locally compact* spaces, one-point compactification is only a functor in the category of locally compact spaces with *proper* morphisms. In the context of Gelfand duality, the analogy with adjoining a unit it still complete because the Gelfand duality functor takes every map of C^*-algebras to a proper map of spaces.

which can arise, instead trying to focus on the bigger picture. Also, in order to develop the machinery as rapidly as possible, we will defer concretely establishing the connection to topological K-theory for now.

3.1 The idea of C^*-algebraic K-theory

The core idea used to construct the K-theory of a C^*-algebra A is to study the *projections* of A. In analogy with the concrete case of $\mathscr{B}(H)$, an element $p \in A$ is called a projection if it is idempotent and self-adjoint. Because we would like $K_0(A)$ to ultimately be a group, we need an additive structure on these elements. Given two projections p and q, their sum is still self adjoint, but we have

$$(p+q)^2 = p^2 + pq + qp + q^2 = p + q + pq + qp \neq p + q,$$

which is unfortunate. Certainly if p and q are *orthogonal* projections in that $pq = qp = 0$, we have a notion of an additive operation, but this is not even the case for $p + p$. Thus in order to make this work we must at least find some way to make a projection p orthogonal to itself (or some version thereof). There is an immediate a cop-out solution; formally "adjoin" more dimensions to A so that we have two copies of p which completely "miss" one another. We would then have access to the trickery of performing computations such as

$$\begin{bmatrix} p & 0 \\ 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 0 & 0 \\ 0 & p \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix},$$

which then solves our problem, but replaces it with a new one; how do we add a projection matrix to itself? Clearly, we need even more dimensions. In fact, because we might need access to arbitrarily many dimensions, we need to bundle together *all* of the matrix algebras on A. The most common (and convenient) way to do this is to form an algebraic direct limit (colimit)

$$\varinjlim_n \mathbf{M}_{n\times n}(A) = \mathbf{M}_\infty(A)$$

by sending a matrix in $\mathbf{M}_k(A)$ to one in $\mathbf{M}_{k+1}(A)$ by appending a zero row and column. We shall say that $p \in \mathbf{M}_\infty(A)$ has *dimension n* if n is the least integer such that p exists in $\mathbf{M}_n(A)$. Note that $\mathbf{M}_\infty(A)$ is not a C^*-algebra because it is not complete (we only take the algebraic direct limit and not the C^*-algebraic one in order to construct it)—but the completion is the stabilisation $A \otimes \mathscr{K}(H)$.

Given any two projections $p, q \in \mathbf{M}_\infty(A)$, we now have strategy for adding them; we form the matrix

$$p + q = \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix}.$$

However, there are now more problems. Indeed,

$$p + q = \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix} \neq \begin{bmatrix} q & 0 \\ 0 & p \end{bmatrix} = q + p \text{ and } \begin{bmatrix} p & 0 \\ 0 & 0 \end{bmatrix} \neq \begin{bmatrix} 0 & 0 \\ p & 0 \end{bmatrix}$$

which of course is terrible—but we are persistent! We must implement some equivalence relation which ensures this is always the case, and so we are intuitively motivated to specify the natural algebraic condition for the equivalence of projections; namely that

$$p \sim q \text{ if there exists unitary } U \in \mathbf{M}_\infty(\mathbf{C}) \text{ such that } p = UqU^{-1}.$$

By conjugating, we alleviate problems to do with potential noncommutativity. Taking equivalence classes with square brackets, elementary linear algebra shows that we have just formed a monoid $V(A) = \{[p] : p \in \mathbf{M}_\infty(A) \text{ is a projection}\}$ associated to A (the identity is the zero projection). We will call the classes in $V(A)$ *generalised projections on A*. The name V for this monoid is a homage to the *Vector Bundles* of topological K-theory, which we will eventually come to.

Because it will prove useful in the coming sections, we also state without proof the following lemma from C^*-algebra theory, which permits the use of alternative definitions of the equivalence relation.

Lemma 3.1. *Let $p, q \in \mathbf{M}_\infty(A)$ be projections. Then $p \sim q$ if and only if*

- *there exists $V \in \mathbf{M}_\infty(\mathbf{C})$ such that $p = VV^*$ and $q = V^*V$, if and only if*
- *there is a continuous map (path) $f : [0,1] \to \mathbf{M}_\infty(A)$ such that $f(0) = p$ and $f(1) = q$.*

Some examples are now in order; first we consider the simplest possible case of $V(A)$ for $A = \mathbf{C}$, and thus really the object $\mathbf{M}_\infty(\mathbf{C})$. Because there is a unitary operator $U : \mathbf{C}^n \to \mathbf{C}^n$ taking any k-dimensional subspace of $\mathbf{C}^n$ to any other, generalised projections in $\mathbf{M}_\infty(\mathbf{C})$ are exactly classified by their rank. Thus $V(\mathbf{C}) \cong \mathbf{N} = \{0,1,2,\ldots\}$. In the case of $A = \mathbf{M}_n(\mathbf{C})$ there is an obvious isomorphism $\mathbf{M}_\infty(\mathbf{M}_n(\mathbf{C})) \cong \mathbf{M}_\infty(\mathbf{C})$, and therefore $V(\mathbf{M}_n(\mathbf{C})) \cong V(\mathbf{C}) \cong \mathbf{N}$ as well (we will soon see an analogous result for the infinite case as well). Indeed, we have $\mathbf{M}_\infty(\mathbf{M}_n(A)) \cong \mathbf{M}_\infty(A)$ and this result for any C^*-algebra A.

The other canonical example of a C^*-algebra is $\mathscr{B}(H)$ for H a Hilbert space. Here, the theory is essentially the same for that of $\mathbf{C}$; there is a unitary operator taking one subspace S of H to another subspace S' exactly when S and S' are of the same dimension. Thus we have $V(\mathscr{B}(H)) \cong \mathbf{N} \cup \{\infty\}$ with the monoid addition $n + \infty = \infty + n = \infty$. Meanwhile, the compact projections on H are exactly those which are of finite rank, and so $V(\mathscr{K}(H)) \cong \mathbf{N}$. Finally, for the Calkin algebra we have $V(\mathscr{B}(H)/\mathscr{K}(K)) \cong V(\{0 + \mathscr{K}(H), I + \mathscr{K}(H)\}) = \{0, \infty\}$ with the same monoid addition as that for $\mathscr{B}(H)$.

At this point, as in the topological case, we will take the Grothendieck completion $V(A)$ to yield a group $GV(A)$ consisting of formal differences of generalised projections on A. For unital C^*-algebras A, this will be exactly the zeroth K-group of A! In the general case it turns out that this is not exactly the right definition, in that the presence of a unit in A has apparently very important implications (as we will see in Section 7, this is completely analogous to the distinction between the topological K-theory of compact and locally compact spaces).

We will instead have to add a unit to A via the A^+ construction, but this is slightly dishonest—clearly A^+ isn't the C^*-algebra we started with! In fact, the only information we generically know about A^+ is that it has the (canonically included) ideal A. It will be of great utility of us to study the quotient of A^+ by A:

Proposition 3.2. *Every map $f : A \to B$ of C^*-algebras gives rise to a map $f_* : V(A) \to V(B)$, and therefore $f_* : GV(A) \to GV(B)$ (of the same name), in a functorial way. In particular, the quotient map $\rho_B : A \to A/B$ gives rise to map $\rho_B : GV(A) \to GV(A/B)$.*

Proof. Given $f : A \to B$, the map just sends $[p] \in \mathbf{M}_\infty(A)$ to $[f(p)] \in M_\infty(B)$ (where $f(p)$ applies f to every entry of p). The fact that M_∞ is a functor then gives that this map is functorial and respects equivalence classes. □

3.2 General definition

The definition of K_0 can now finally be given.

Definition 3.3. For A an arbitrary C^*-algebra, we define

$$K_0(A) = \ker(\rho_* : GV(A^+) \to GV(A^+/A)).$$

The latter quotient A^+/A is always $\mathbf{C}$, and so ρ_* is a map $GV(A^+) \to G\mathbf{N} = \mathbf{Z}$. It is immediate that K_0 is a group, and one easily sees that because V and G are functors then K_0 is as well (the small matter of whether elements of $K_0(A)$ remain in the kernel of ρ_* under an induced map $f_* : K_0(A) \to K_0(B)$ is easy to check). As we have already begun to do, we will use a subscript asterisk to denote maps of the various algebraic groups at our disposal which are induced by maps of C^*-algebras.

The presence of a kernel in the definition may seem at first unusual, but it is not much work to define the relative K_0 groups $K_0(\cdot,\cdot)$ in such a way that $K_0(A) \cong K_0(A^+, A)$ (see [2], for example, where it is explained that this equivalence of definitions is a consequence of the Strong Excision Theorem for C^*-algebras). Our definition also has the advantage of still being relatively concrete—it is easy to see that G, V, and K_0 respect direct sums, for example. The following proposition shows that the definition of K_0 is particularly easy in the unital case.

Proposition 3.4. *The sequence*

$$0 \longrightarrow K_0(A) \xrightarrow{\subset} GV(A^+) \xrightarrow{\rho_*} \mathbf{Z} \longrightarrow 0$$

is exact for every C^*-algebra A, with the first nonzero map the inclusion of sets, and ρ_* the map used to define $K_0(A)$. Furthermore, it always splits, and when A is unital $K_0(A) \cong GV(A)$.

Proof. The sequence is clearly exact, and the inclusion $\iota : \mathbf{C} \to A^+$ of the subalgebra generated by the adjoined unit induces a map $\iota_* : GV(\mathbf{C}) \cong \mathbf{Z} \to GV(A^+)$ which is a right-inverse of ρ_*. Therefore we have a splitting $GV(A^+) \cong K_0(A) \oplus \mathbf{Z}$. Finally, when A is unital we also have $A \oplus \mathbf{C} \cong A^+$, and the induced map $GV(A) \oplus \mathbf{Z} \cong GV(A^+)$ composed with the inclusion $0 \oplus \mathbf{Z} \to GV(A) \oplus \mathbf{Z}$ exactly recovers ι_* (by the definition of the former map). Thus, the isomorphism $GV(A) \oplus \mathbf{Z} \cong K_0(A) \oplus \mathbf{Z}$ itself splits as a direct sum, and in particular we get an isomorphism $GV(A) \cong K_0(A)$. □

Our definition strongly resembles the necessary modification which must be made to topological K-theory in order to deal with only locally compact spaces. However, there are also already some points of divergence; for example, the is no natural multiplicative structure which we can impose on $K_0(A)$ for general A. Once we connect K_0 to algebraic K-theory, we will see that this is the consequence of the lack of a general "noncommutative tensor product" of modules over noncommutative rings. We will also see a different attempt to impose additional structure on K_0 in Section 4, which works particularly well for the class of Almost Finite (AF) C^*-algebras. Note as well that up until this point we have not used any C^*-algebra structure, and that the entire construction thus far goes through for arbitrary rings—this will change shortly.

We calculated $V(A)$ above for common C^*-algebras A, and now given the definition of $K_0(A)$ and Proposition 3.4 we can do the same for the K_0 group. Indeed, we immediately find that $K_0(\mathbf{C}) \cong K_0(\mathbf{M}_n(\mathbf{C})) \cong K_0(\mathscr{K}(H)) \cong \mathbf{Z}$, while $K_0(\mathscr{B}(H)) \cong K_0(\mathscr{B}(H)/\mathscr{K}(H)) \cong 0$ due to the presence of an "∞" element in the corresponding monoid obtained from V. We are yet to stumble upon a C^*-algebra A which has torsion in its K-theory, but tame examples[4] certainly do exist. We take this opportunity to mention the particular example of K-theory torsion provided by the Cuntz algebras $\mathscr{O}_n$. They are notable for being the first separable infinite simple C^*-algebras to be explicitly constructed. Indeed, every simple infinite C^*-algebra has every Cuntz algebra $\mathscr{O}_n$ as a quotient! The details can be found in [4] or [19] (and require some of the general theory which we are yet to develop), but the idea of the construction is very intuitive. The point is to select $n > 1$ isometries $(U_j)_{j=1}^n$ on a Hilbert space H which together have independent and spanning image projections, i.e in that

- each U_j is an isometry; for each $1 \le j \le n$ we have $U_j^{-1} U_j = I$ for I the identity on H, and
- the U_j's have independent and spanning projections, in that $I = \sum_{j=1} U_j U_j^{-1}$ (this implies that the projections $U_j U_j^{-1}$ are orthogonal).

The Cuntz algebra $\mathscr{O}_n$ is then defined to be the smallest sub-C^*-algebra of $\mathscr{B}(H)$ containing all of the operators U_j—and it turns out that this construction is unique up to isomorphism. (The algebra $\mathscr{O}_n$ must be defined concretely in this way, without simply imposing these relations on a free C^*-algebra generated by n symbols, because free C^*-algebras do not exist in general.) The reason why this should "obviously" give rise to a K_0 group with torsion is because when taking equivalence classes of the projections $U_j U_j^{-1}$ orthogonality allows us to calculate that

$$[I] = \left[\sum_{j=1} U_j U_j^{-1} \right] = \sum_{j=1} [U_j U_j^{-1}] = n[I].$$

Thus, notwithstanding unexpected degeneracy in $K_0(\mathscr{O}_n)$, we expect an element of order n. It is the content[5] of [4] that there is no such degeneracy, and in fact $K_0(\mathscr{O}_n) \cong \mathbf{Z}/(n-1)\mathbf{Z}$.

[4] As a consequence of the results of Section 7, we see for example that any compact space X with torsion in its topological K_0-group gives rise to a C^*-algebra $C_0(X)$ which has torsion in its C^*-algebraic K_0-group as well. We also see in Subsection 6.2 a way to construct C^*-algebras with arbitrary K-theories (by essentially taking very many copies of the compact operators $\mathscr{K}(H)$ and employing a sly exact sequence).

[5] Once we reach Section 5, we will be able to make sense of one of Cuntz's other results that $K_1(\mathscr{O}_n) \cong 0$ for all $n > 1$.

The proofs of the fundamental theorems on the functor K_0 to come could be said to be *more than lightly seasoned* with matrix calculations; and they are necessitated by the fact that $\mathbf{M}_\infty$ lies at the heart of the definition of K_0. In order to carefully negotiate the road ahead without unnecessarily triggering any matrix explosions, we conclude this section by stating two useful properties of elements of $K_0(A)$ which follow from elementary calculations in linear algebra, by rearranging the elements of representative matrices given the definition of the addition in K_0 and the equivalence relation on projections defining $V(A)$ (and using Lemma 3.1 where necessary). Complete proofs can be found in [16] and [19].

Proposition 3.5. *Let A be a C^*-algebra, and $[p]-[q] \in K_0(A)$ for projections $p, q \in \mathbf{M}_n(A^+)$. Then:*

(i) There exists $m, k \in \mathbf{N}$ with $m \geq k$ and another projection $p' \in \mathbf{M}_m(A^+)$ such that we have the refinement $[p]-[q] = [p'] - [I_k]$, with I_k the $k \times k$ identity matrix.

(ii) We have $[p] = [q]$ if and only if there exists $m, k \in \mathbf{N}$ such that there is a path in $\mathbf{M}_{m+k}(A^+)$ (the additional k dimensions provide "extra room" for the path) between the block matrices

$$\begin{bmatrix} p & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \textit{and} \quad \begin{bmatrix} q & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

3.3 Fundamental properties of K_0

In this section we will establish some fundamental properties of the K_0 functor. We first verify the half-exactness of K_0, which is an absolute necessity if we are to construct the crown jewel of K-theory, the canonical six-term exact sequence. The following elementary lemma from abstract algebra will be required.

Lemma 3.6. *Let*

$$0 \longrightarrow A \xrightarrow{\ f\ } B \xrightarrow{\ g\ } C \longrightarrow 0$$

be an exact sequence of C^*-algebras. Then there exists a unique (isomorphism) $h : C \to B/\operatorname{im} f$ such that the diagram

$$\begin{tikzcd}
& 0 \arrow[d] & & 0 \arrow[d] & \\
0 \arrow[r] & A \arrow[dr, "f"] \arrow[dd, "f"'] & & C \arrow[r] \arrow[dd, dotted, "h"] & 0 \\
& & B \arrow[ur, "g"] \arrow[dr, "\rho_{\operatorname{im} f}"] & & \\
0 \arrow[r] & \operatorname{im} f \arrow[ur, hook, "\subset"] \arrow[d] & & B/\operatorname{im} f \arrow[r] \arrow[d] & 0 \\
& 0 & & 0 &
\end{tikzcd}$$

commutes and is exact.

Proof. This is essentially the first isomorphism theorem for C^*-algebras. □

Proposition 3.7 (Half-exactness)**.** *Let*

$$0 \longrightarrow A \xrightarrow{\ f\ } B \xrightarrow{\ g\ } C \longrightarrow 0$$

be an exact sequence of C^*-algebras. Then the sequence

$$K_0(A) \xrightarrow{\ f_*\ } K_0(B) \xrightarrow{\ g_*\ } K_0(C)$$

is also exact.

Proof. It is an immediate consequence of Lemma 3.6 that we may assume without loss of generality that A is an ideal of B, that we have $C = B/A$, that f is the inclusion, and that g is the projection $\rho_A : B \to B/A$. It thus remains to show that the sequence

$$K_0(A) \xrightarrow{\ \subset\ } K_0(B) \xrightarrow{\ \rho_{A*}\ } K_0(B/A)$$

is exact.

We now follow the argument of [19]. First, by Proposition 3.5 each element of $K_0(A)$ is of the form $[p]-[I_n]$ with p a projection such that $p - I_n \in \mathbf{M}_k(A)$ for some $k \in \mathbf{N}$. Thus $0 = \rho_A(p - I_n) = \rho_A(p) - \rho_A(I_n)$, and so $\rho_A(p) = \rho_A(I_n)$. Hence $\rho_{A*}([p]-[I_n]) = [\rho_A(p)] - [\rho_A(I_n)] = 0$, which shows that we at least have a complex.

Now suppose that $[p]-[I_n]$ is an element of $K_0(B)$ such that $\rho_{A*}([p]-[I_n]) = [\rho_{A*}(p)] - [I_n] = 0$. Then Proposition 3.5 together with Lemma 3.1 imply that there are integers $m, k \in \mathbf{N}$ and a unitary matrix $U \in \mathbf{M}_{m+k}((B/A)^+)$ such that in $\mathbf{M}_{m+k}((B/A)^+)$ we have

$$\begin{bmatrix} I_n & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} = U \begin{bmatrix} \rho_{A*}(p) & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} U^{-1}.$$

Because multiplication by elementary matrices may be performed via continuous paths (i.e. there is a path with endpoints before and after performing the multiplication), once again Proposition 3.5 and Lemma 3.1 give that there is a unitary matrix $V \in \mathbf{M}_{2m+2k}(A^+)$ such that

$$\rho_A(V) = \begin{bmatrix} U & 0 \\ 0 & U^{-1} \end{bmatrix}.$$

We then form the projection q by (we correct the dimensions of the multiplication by adding sufficiently many zero rows and columns)

$$q = V \begin{bmatrix} p & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} V^{-1}.$$

In particular, we have slightly abused block matrix notation above, with zeros representing block zero matrices if necessary. We will continue in this manner for the remainder of the proof. The point is that under the image of ρ_A (because ρ_A is a group homomorphism) we have that

$$\begin{aligned} \rho_A(q) = \rho_A(V) &\begin{bmatrix} \rho_A(p) & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} \rho_A(V)^{-1} \\ = U &\begin{bmatrix} \rho_A(p) & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} U^{-1} \\ = &\begin{bmatrix} I_n & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & 0 \end{bmatrix} \sim I_{n+m}. \end{aligned}$$

Because this last matrix has entries just zero and the multiplicative identity, it is in $\mathbf{M}_{2m+2k}(A^+)$ (and not just $\mathbf{M}_{2m+2k}(B^+)$). Now by construction the matrix V exhibits the fact that

$$q \sim \begin{bmatrix} p & 0 \\ 0 & I_m \end{bmatrix}, \quad \text{and moreover we have} \quad p \sim \begin{bmatrix} p & 0 \\ 0 & I_m \end{bmatrix}.$$

Therefore (outer square brackets on matrices denote taking an equivalence class)

$$[p] - [I_n] = \left[\begin{bmatrix} p & 0 \\ 0 & I_m \end{bmatrix} \right] - \left[\begin{bmatrix} I_n & 0 \\ 0 & I_m \end{bmatrix} \right] = [q] - [I_{n+m}].$$

Because q is a matrix with entries in A^+ only, this completes the proof. $\square$

As we mentioned in the introduction to this subsection, half-exactness of K_0 is absolutely necessary for the (extremely useful) six-term exact sequence. Thus Proposition 3.7 provides justification for the perhaps at-first unintuitive definition of the K_0 functor; the functor GV is not half-exact! Indeed, consider the C^*-algebra $A = C_0(\mathbf{R}^2)$, consisting of the continuous functions on $\mathbf{R}^2$ vanishing at infinity. Then A has no unit, and by the vanishing condition cannot have any projections which are not zero. The same is true for any matrix algebra $\mathbf{M}_n(A)$, and therefore $V(A) \cong 0 \cong GV(A)$. However, the unitisation A^+ corresponds to the one-point compactification of $\mathbf{R}^2$ under the Gelfand duality. This is just the sphere S^2, for which it can be shown (e.g. in [2]) that $K_0(C_0(S^2)) \cong \mathbf{Z}^2$ using Bott Periodicity. The additional projection corresponding simply to the always-unity function on S^2 has nontrivial image in $K_0(C_0(S^2)/C_0(\mathbf{R}^2)) \cong K_0(\mathbf{C})$, and therefore we actually have $K_0(A) \cong \mathbf{Z}$, which is nontrivial unlike the group for A obtained from GV. Also of note is the fact that we have no hope of extending this result to an exact sequence on its own by prepending and appending zero groups—for example there are typically more unitary matrices in $\mathbf{M}_\infty(A)$ compared to $\mathbf{M}_\infty(B)$, which increases the coarseness of the equivalence relation in $V(B)$ compared to $V(A)$.

The proof of Proposition 3.7 is indicative of the large amount of detailed matrix manipulations which are typical in the proofs of fundamental theorems in the K-theory of C^*-algebras (even despite the fact that we have evaded most of these details using the extremely handy Proposition 3.5). Indeed, *it gets much, much, worse*—the "extra room" always available in $\mathbf{M}_\infty$ is frequently required, and the proof of significant theorems such as Bott Periodicity lead to the construction of absolutely enormous matrices. For this reason, in order to keep the size of this account under control and to avoid quoting an excessive number of properties of projections on C^*-algebras, we only give the main ideas the proofs of the next two properties which we list.

Proposition 3.8 (Continuity)**.** *The functor K_0 is continuous, in that it preserves direct limits of C^*-algebras.*

Proof sketch. The proof of this property is essentially the repeated application the fact that projections in an inductive limit A of C^*-algebras $\{A_j\}$ can be approximated in norm arbitrarily closely by projections in some A_j for a suitably chosen index j (that is free to vary). One then executes an ε-δ argument in order to ensure that the induced map $\varinjlim_j K_0(A_j) \to K_0(A)$ is both injective and surjective (the former requires the greater amount of work). □

Proposition 3.9 (Stability)**.** *If A is a C^*-algebra, then $K_0(A \otimes \mathcal{K}(H)) \cong K_0(A)$, so that K_0 is invariant under stabilisation (tensoring with $\mathcal{K}(H)$).*

Proof sketch. The proof consists of checking that the map $A \to A \otimes \mathcal{K}(H)$ defined by $x \mapsto x \otimes I_1$ for I_1 a rank 1 projection on H actually induces an isomorphism $K_0(A) \to K_0(A \otimes \mathcal{K}(H))$. The reason that this possibly works is that the map $\mathbf{M}_n(A) \to \mathbf{M}_m(A)$ for $n \le m$ defined by appending zero rows and columns actually induces isomorphisms of K-theory, and this in turn follows because in the definition of K_0 we immediately take the matrix algebra $\mathbf{M}_\infty(A)$. The remainder of the proof inserts these isomorphisms in the definition of the inductive limit in order to use universality to guarantee that we actually get an isomorphism $K_0(A) \to K_0(A \otimes \mathcal{K}(H))$ out. □

Finally, we will deal with "homotopy invariance" of K_0. We first introduce some definitions which the reader should find to be completely familiar from the topological setting.

Definition 3.10. We call morphisms $f, f' : A \to B$ *homotopic* if they are homotopic as maps in the traditional sense. If g is a map $B \to A$, then f and g give a *homotopy equivalence* if $f \circ g$ and $g \circ f$ are both homotopic to the identity. By analogy we also introduce the terminology *deformation retraction*, and in particular a map $f : A \to B$ is *contractible* if it is homotopic to a constant map (we can always ensure that this constant is zero). A C^*-algebra A is *contractible* if the identity on A is contractible.

Proposition 3.11 (Homotopy invariance)**.** *Homotopic maps $f, g : A \to B$ induce the identical maps $f_*, g_* : K_0(A) \to K_0(B)$.*

Proof. A homotopy $h_t : [0,1] \times A \to B$ from f to g gives rise to a family of induced maps $h_{t*} : [0,1] \times \mathbf{M}_\infty(A^+) \to \mathbf{M}_\infty(B^+)$. For each $[p]-[q] \in K_0(A)$ we therefore have that $f_*(p)$ is homotopic to $g_*(p)$ (and likewise for q), and thus the proof is completed by Lemma 3.1. □

Corollary 3.11.1. *If A is a contractible C^*-algebra, then $K_0(A) \cong 0$.*

Proof. The zero C^*-algebra 0 has $V(0) \cong 0$, and so Proposition 3.11 gives that $K_0(A) \cong K_0(0) \cong 0$. □

Continuing the complete analogy with topology, we now make the following definitions (especially consider the case where $A = C_0(X)$ for a compact space X under Gelfand duality!). The notation "$C([0,1] \to A)$" just means the set of continuous functions from $[0,1]$ to A, which is always a C^*-algebra under pointwise operations and the supremum norm.

Definition 3.12. Let A be a C^*-algebra. Then the *cone* on A is the C^*-algebra

$$CA = \{s \in C([0,1] \to A) : s(0) = 0\}.$$

The *suspension* of A is the C^*-algebra

$$\Sigma A = \{s \in C([0,1] \to A) : s(0) = s(1) = 0\},$$

which is just

$$\Sigma A = \{s \in CA : s(1) = 0\}.$$

Both of these constructions are easily seen to be functors, with a map $f : A \to B$ of C^*-algebras inducing a map $f_* : CA \to CB$ by post-composition with f (and similarly for the suspension Σ).

Indeed, these constructions possess the topologically expected homotopy properties;

Lemma 3.13. *Let A be a C^*-algebra. Then*

- *the cone CA is contractible, and*
- *if A is contractible, the suspension ΣA is as well.*

Proof. For the first point, the trick is simply to squeeze everything into the "point" of the cone; there is a homotopy $f_t : [0,1] \times CA \to CA$ defined by

$$f_t(s) = (\lambda \mapsto s(t\lambda))$$

which explicitly exhibits the fact that CA is contractible.

Now suppose that we have a homotopy $f_t : [0,1] \times A \to A$ from $\mathrm{id}_A : A \to A$ to the constant zero map on A. Then we have another homotopy $\hat{f}_t : [0,1] \times \Sigma A \to \Sigma A$ defined by

$$\hat{f}_t(s) = (\lambda \mapsto \hat{f}_t(s(\lambda)))$$

is a homotopy between the identity on ΣA and the map constantly equal to $\lambda \mapsto 0$ (i.e. the zero element in ΣA). □

We finally note the following two obvious but useful facts for cones and suspensions.

Proposition 3.14 (Exact sequence for cones)**.** *Let A be a C^*-algebra. Then there is an exact sequence*

$$0 \longrightarrow \Sigma A \xrightarrow{\subset} CA \xrightarrow{\mathrm{ev}_1} A \longrightarrow 0$$

with $\mathrm{ev}_1 : CA \to A$ *the map which evaluates at* 1.

Proof. The claim and the next follow by inspecting definitions (e.g. for $s \in CA$, we have $s \in \Sigma A$ exactly when $s(1) = \mathrm{ev}_1(s) = 0$). □

Lemma 3.15. *Every exact sequence*

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$$

of C^*-algebras gives rise to an exact sequence

$$0 \longrightarrow \Sigma A \longrightarrow \Sigma B \longrightarrow \Sigma C \longrightarrow 0 \ .$$

In this section, we have established (among other things) four major properties of K_0; continuity, stability, half-exactness, and homotopy invariance. We mention in the concluding remarks (Section 8) the result of Cuntz that modulo some initial conditions, these more than determine the K_0 functor on a very large class of C^*-algebras!

4 A first application: classifying AF-algebras

As an application of some the properties of the functor K_0 that we have just developed, we will devote some time to establishing one of the historically first achievements of the K-theory of C^*-algebras in its own right. We will get the opportunity to explore unique properties of the C^*-algebra theory (as opposed to the topological one) as we try to compensate for the lack of a multiplicative structure in K_0 as was pointed out above.

4.1 Addressing the lack of multiplicative structure

That the noncommutativity of general C^*-algebras A prevents the implementation of a multiplicative structure in $K_0(A)$ is certainly quite unfortunate. Indeed, as we will see in the sequel, the additive structure of $K_0(A)$ cannot even distinguish between finite dimensional C^*-algebras of the form

$$A = \mathbf{M}_{n_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{n_N}(\mathbf{C})$$

and

$$B = \mathbf{M}_{m_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{m_N}(\mathbf{C})$$

for $(n_j)_{j=1}^N$ and $(m_j)_{j=1}^N$ sequences of integers which are not equal up to permutation. Proposition 3.9 also shows that we cannot distinguish between stabilisations of A and itself, for instance. By extracting as much additional structure from A (which is related to $K_0(A)$) as we can, we will hope to be able to salvage this information.

Throughout this section we will consider the case of a *unital* C^*-algebra A, with the extensions in the nonunital case to be left until the end. Regardless however, in the case of a possibly non-unital C^*-algebra A, observe that there is a canonical monoid map $\iota : V(A) \to GV(A)$ defined by $\iota([p]) = [p] - [0]$ (for $[0]$ the zero projection). Furthermore, because the sequence

$$V(A) \longrightarrow V(A^+) \xrightarrow{\ \rho\ } V(A^+/A) \cong \mathbf{N}$$

is a complex, we actually get a well-defined map $\iota : V(A) \to \ker(\rho : GV(A^+) \to G\mathbf{N}) = K_0(A)$ (we call several maps ι and ρ). Due to the obvious analogy with the inclusion $\mathbf{N} \to G\mathbf{N}$, the idea is to is to single out the subset $K_0^+(A) = \iota(V(A)) \subset K_0(A)$. Because $K_0^+(A)$ is closed under the operations of the monoid structure conferred by $K_0(A)$, we would like to think of $K_0^+(A)$ as the "positive" elements of $K_0(A)$. Indeed, for arbitrary groups we can make the following definition.

Definition 4.1. A *positive cone* K of an arbitrary group G is a subset such that for $e \in G$ the identity we have

- $e \in K$,
- $G = \{g - h : g, h \in K\}$, and
- if $k^{-1} \in K$ for some $k \in K$ then $k = e$.

If G is abelian, a cone $K \subset G$ immediately induces a G-invariant partial order by

$$g \geq h \text{ if } g - h \in K, \quad g, h \in G.$$

Of course, the subset $K_0^+(A) \subset K_0(A)$ need not always be a positive cone—for example, in Subsection 3.2 we saw the Cuntz algebras $\mathcal{O}_n$ for which $K_0(\mathcal{O}_n) \cong \mathbf{Z}/(n-1)\mathbf{Z}$ and so cannot possibly have $K_0^+(\mathcal{O}_n)$ a positive cone—but when it does we will call the combined data $(K_0(A), K_0^+(A))$ the *ordered K_0 group of A*. A morphism of ordered groups is a morphism of the underlying groups which is also order-preserving.

There is a second piece of additional data which is readily available to us, and will prove invaluable in completing our classification. For each C^*-algebra A, we define the *scale*

$$\Xi(A) = \{[p] : p \in A \text{ is a projection, i.e. } p \in V(A) \text{ which is one-dimensional}\}.$$

For example, because projections on $\mathbf{C}^n$ are exactly classified up to equivalence by their rank, we have $\Xi(\mathbf{M}_k(\mathbf{C})) = \{0,1,\dots,k\}$. Thus we can distinguish matrix algebras over $\mathbf{C}$ of different dimensions! In general, the scale data will permit us to make distinction between C^*-algebras which are stably equal (isomorphic after some stabilisation), but are nonetheless distinct. Depending on our purposes, we will think of $\Xi(A)$ as a subset of $V(A)$, or of $K_0(A)$ via ι. Putting all of this data together, we obtain the *scaled and ordered*, or simply *extended*, K_0 group of A: $(K_0(A), K_0^+(A), \Xi(A))$ (under the right circumstances).

4.2 Studying finite dimensions

The purpose of this subsection is to determine the data stored in the scaled and ordered K_0 groups of all of the finite-dimensional C^*-algebras. We will first record the following (easy) result on the structure of sums of matrix algebras.

Proposition 4.2. *Let $A = \mathbf{M}_{n_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{n_N}(\mathbf{C})$ be a sum of matrix algebras. Then*

- $K_0(A) = \mathbf{Z}^N$,
- $K_0^+(A) = \mathbf{N}^N$, *and*
- $\Xi(A) = \{(x_1,\dots,x_N) \in \mathbf{N}^N : x_j \le n_j\}$.

Proof. We noted above that K_0 respects direct sums, and that $K_0(\mathbf{M}_n(\mathbf{C})) \cong \mathbf{Z}$, and so therefore $K_0(A) \cong \mathbf{Z}^N$. It is also easily seen that $K_0^+(A) \cong \mathbf{N}^N$, given that $K_0^+(\mathbf{C}) \cong \mathbf{N}$. Finally, by definition Ξ also respects direct sums, and this gives

$$\Xi(A) \cong \mathbf{N}^{n_1} \times \mathbf{N}^{n_2} \times \cdots \times \mathbf{N}^{n_N}.$$

□

It is the content of the following theorem that every finite-dimensional C^*-algebra is isomorphic to a sum of matrix algebras, and so in fact we have computed the extended K-theory of all finite dimensional C^*-algebras. In doing so, we have found that the extended K_0 groups of such algebras completely classify them up to isomorphism. The classification theorem admits an unreasonably long proof via algebraic techniques and the theory of the spectrum of Hilbert spaces (frankly, it is awful, e.g. see [14]). However because of the positive implications of the result on our classification, we give the idea[6] of an elegant proof in [16] which at its core exploits the extension of Gelfand's theorem for locally compact Hausdorff spaces (we skip the details when we return to the realm of linear algebra).

Theorem 4.3 (Classification of finite-dimensional C^*-algebras)**.** *Let A be a finite-dimensional C^*-algebra. Then A is unital, and there exists a finite sequence of nonnegative integers $(n_j)_{j=1}^N$ such that*

$$A \cong \mathbf{M}_{n_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{n_N}(\mathbf{C}).$$

Proof. Let $B \subset A$ be a maximal abelian sub-C^*-algebra of A (such an object exists by Zorn's lemma). Then by (an extension of) Gelfand's theorem, B is isomorphic to $C_0(X)$ for X an only *locally compact* space. But B is finite dimensional, so X must at most be the disjoint union of finitely many points. Therefore, in particular, X is compact and therefore B is unital. It is an easy computation that this unit must also be a unit of A.

The remainder of the proof rests on the consideration of objects which are by now dear to our heart—projections on A and D. We define a finite family of projections on X by collecting together the characteristic functions on each of X's finitely many points, and then use Gelfand duality to obtain projections on D. It can then be shown that each pair of such projections have image either equal or disjoint excluding 0, from which we obtain a decomposition of A in terms of the direct sum of individual matrix algebras over $\mathbf{C}$ (we check that we have formed a *system of matrix units*, to which general theory then applies and provides the desired decomposition). □

[6] We skip the proof of several results from the theory of maximal abelian subalgebras of C^*-algebras, and of systems of matrix units, but they are all easy calculations in C^*-algebras and linear algebra. They can be found in [7].

We now turn to the problem of studying the morphisms between these finite-dimensional sums. The notion of when a morphism $\phi : K_0(A) \to K_0(B)$ is a *morphism of scaled groups* will be useful, and we simply define it to mean that $\phi(\Xi(A)) \subset \Xi(B)$. The extension to the case of almost-finite algebras presented in the next section then essentially hangs on the following proposition (we roughly follow the argument of [19]).

Proposition 4.4. *Suppose that* $A = \mathbf{M}_{n_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{n_N}(\mathbf{C})$ *and* $B = \mathbf{M}_{m_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{m_M}(\mathbf{C})$.

(i) Let $f : A \to B$ *be a morphism inducing* $f_* : K_0(A) \to K_0(B)$*. Then:*

- *The map* f_**, thought of as a map* $\mathbf{Z}^N \to \mathbf{Z}^M$*, acts as multiplication by a matrix in* $\mathbf{M}_{N\times M}(\mathbf{Z})$.
- *The* i*th row and* j*th column of this matrix is the number of times* $T_{i,j}$ *the (unique) irreducible representation* $\mathbf{M}_{n_i} \to \mathbf{M}_{m_j}$ *appears as a decomposition of* f *as a direct sum of irreducible representations.*

(ii) If $\Phi : K_0(A) \to K_0(B)$ *is a morphism of scaled groups then there exists a map* $f : A \to B$ *of* C^**-algebras such that* $\Phi = f_*$.

Proof. The first part of Point (i) is trivial, and the second part follows simply because the rank of a projection on the $\mathbf{M}_{n_i}$ summand of A has its rank multiplied by $T_{i,j}$ once restricted to $\mathbf{M}_{m_j}$ and viewed under the action of f_*. To show Point (ii), suppose that $\Phi : K_0(A) \to K_0(B)$ is a morphism of scaled groups and is represented by a matrix $T \in \mathbf{M}_{N\times M}$. The idea is to define maps

$$f_j : \mathbf{M}_{n_1}(\mathbf{C}) \oplus \cdots \oplus \mathbf{M}_{n_N}(\mathbf{C}) \to \mathbf{M}_{m_j}(\mathbf{C})$$

by

$$f_j(S_1 \oplus \cdots \oplus S_N) = \begin{bmatrix} Q_{T_{1,j}}(S_1) & 0 & \cdots & 0 \\ 0 & Q_{T_{2,j}}(S_2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & Q_{T_{2,N}}(S_N) \end{bmatrix},$$

with $Q_j(S)$ itself a j-fold repetition of the matrix S along the diagonal;

$$Q_j(S) = j \text{ rows} \left\{ \begin{bmatrix} S & 0 & \cdots & 0 \\ 0 & S & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & S \end{bmatrix} \right. .$$

If each f_j is well-defined, we would then be able to form the map $f : A \to B$ by summing over the f_j's, and then by Point (i) we would exactly have $K_0(f) = \Phi$ by construction.

Of course, this matrix could (and in general, would) be too large to be an element of $\mathbf{M}_{m_j}(\mathbf{C})$. But because Φ is a morphism of scaled groups for $v = (n_1, \ldots, n_N) \in \Xi(A)$ we have $\Phi(v_j) \in \Xi(A)$, and therefore in particular

$$\sum_{k=1} T_{j,k} n_k \le m_j$$

for each $1 \le j \le M$. This is precisely the desired size constraint, and completes the proof. □

Even without the classification of Theorem 4.3 and calculation of Proposition 4.2, this result shows that the ordered, scaled K_0 groups are a complete invariant of the C^*-algebras which are finite sums of matrix algebras over $\mathbf{C}$. In preparation for the coming generalisation, we note without proof a linear-algebraic lemma on the relationship between morphisms of finite-dimensional C^*-algebras which induce the same maps on K_0-groups.

Lemma 4.5. *Suppose that* A *and* B *are both finite direct sums of matrix algebras over* $\mathbf{C}$*, for which two morphisms* $\phi, \psi : A \to B$ *induce the same map* $\phi_* = \psi_* : K_0(A) \to K_0(B)$*. Then* ϕ *and* ψ *differ by conjugation by a unitary element of* B^+.

4.3 The case of almost-finite dimensions

Definition 4.6. A C^*-algebra A is *almost finite*, or simply *AF*, if it is the inductive limit of a sequence $\{A_j, \alpha_j : A_j \to A_{j+1}\}$ of finite dimensional C^*-algebras. As we will only be dealing with unital C^*-algebras in this section, we require that the morphisms α_j be unit-preserving.

Viewed under the Gelfand duality, unital AF-algebras are actually exactly the compact totally disconnected spaces.

In the finite-dimensional case, we were not able glean any distinguishing information from the positive cone $K_0^+(A)$ without further refining it into the scale $\Xi(A)$. The utility (really, distinguishing power) of the ordered K_0 groups in isolation in the almost-finite case is made apparent by the following result, which shows that the extended K_0 group is a complete invariant for AF-algebra algebras. The weaker version of the theorem, also stated, establishes that ordered K_0 groups really form a middle-ground between K_0 on its own and the full extended K_0 groups.

Theorem 4.7 (Elliott [5], Fundamental theorem of AF-algebras)**.** *Let A and B be unital AF-algebras. Every isomorphism $\Phi : K_0(A) \to K_0(B)$ of ordered and scaled groups gives rise to an isomorphism $f : A \to B$ of C^*-algebras. If Φ is only an isomorphism of ordered groups, then A and B are at least stably isomorphic.*

Proof. If A and B are AF-algebras, then there exactly exist direct limit systems $\{A_j, \alpha_j\}$ and $\{B_j, \beta_j\}$ so that

$$\varinjlim_j A_j = A, \text{ and } \varinjlim_j B_j = B.$$

We also have canonical morphisms $\hat{\alpha}_j : A_j \to A$ and $\hat{\beta}_j : B_j \to B$, which we will frequently exploit. By virtue of the continuity of K_0 (and technically V and G as well), we are guaranteed that the maps α_j, β_j, and their "hatted" versions are all scale and order preserving. The proof then hinges on the fact that, by suitably intertwining these direct limit systems, we can ensure the existence of an isomorphism $A \cong B$ effectively by the universality of direct limits.

Concretely, we will exploit the following elementary lemma from the theory of direct limits.

Lemma 4.8. *Suppose that $\{A_j, \alpha_j : A_j \to A_{j+1}\}$ and $\{B_j, \beta_j : B_j \to B_{j+1}\}$ are direct limit systems, for which there exist intertwining morphisms $g_j : A_{2j-1} \to B_{2j}$ and $f_j : B_{2j} \to A_{2j+1}$ which are compatible with the direct limit systems, in that we have an infinite commutative diagram*

$$\begin{tikzcd}
\cdots \arrow[r, "\alpha_{2j-2}"] & A_{2j-1} \arrow[r, "\alpha_{2j-1}"] \arrow[dr, "g_j"] & A_{2j} \arrow[r, "\alpha_{2j}"] & A_{2j+1} \arrow[r, "\alpha_{2j+1}"] \arrow[dr, "g_{j+1}"] & \cdots \\
\cdots \arrow[r, "\beta_{2j-2}"] \arrow[ur, "f_{j-1}"] & B_{2j-1} \arrow[r, "\beta_{2j-1}"] & B_{2j} \arrow[r, "\beta_{2j}"] \arrow[ur, "f_j"] & B_{2j+1} \arrow[r, "\beta_{2j+1}"] & \cdots
\end{tikzcd}.$$

Then there exists a pair of mutually inverse C^*-algebra morphisms $g_\infty : A \to B$ and $f_\infty : B \to A$.

The only problem is to actually construct the morphisms g_j and f_j! The idea for this is to use Point (ii) of Proposition 4.4 to bootstrap from morphisms of the K_0 groups of the finite-dimensional algebras in the direct limit systems. To begin, the "finite-end" of the relevant diagram is

$$\begin{tikzcd}
A_1 \arrow[r] \arrow[dr, dotted, "g_1"] & A_2 \arrow[r] & A_2 \arrow[r] \arrow[dr, dotted, "g_2"] & \cdots \\
B_1 \arrow[r] & B_2 \arrow[r] \arrow[ur, dotted, "f_1"] & B_3 \arrow[r] & \cdots
\end{tikzcd},$$

for which we have a corresponding diagram

$$\begin{tikzcd}
K_0(A_1) \arrow[r] \arrow[dr, dotted, "\phi_1"] & K_0(A_2) \arrow[r] & K_0(A_3) \arrow[r] \arrow[dr, dotted, "\phi_2"] & \cdots \\
K_0(B_1) \arrow[r] & K_0(B_2) \arrow[r] \arrow[ur, dotted, "\psi_1"] & K_0(B_3) \arrow[r] & \cdots
\end{tikzcd}$$

with all of the dotted morphisms (and in particular ϕ_j and ψ_j) yet to be constructed. Dealing with the general case requires additional notational clutter, and the extension will be clear, so we restrict ourselves to finding the indicated first pieces of the intertwiners.

First, by Proposition 4.2 there is an isomorphism $K_0(A_1) \cong \mathbf{Z}^{N_1}$ which we will endeavour to avoid writing, instead simply identifying the two groups. Then we have generators

$$e_1 = (1,0,\ldots,0),\ldots,e_{N_1} = (0,\ldots,0,1)$$

of $K_0(A_1)$ which we can map via $\hat{\alpha}_1 : K_0(A_1) \to K_0(A)$ and then via $\Phi : K_0(A) \to K_0(B)$ into $K_0(B)$. Because K_0 is continuous, we can find some $k \in \mathbf{N}$ so that for every generator $e_j \in K_0(A_1)$ there is some generator $e'_j \in K_0(B_k)$ such that $\Phi(\hat{\alpha}_1(e_j)) = \hat{\beta}_k(e'_j)$ (this is just the statement that elements of $K_0(B)$ factor through the subobjects, then using the fact that Φ is an isomorphism). Without loss of generality (by forgetting pieces of the direct limit system), we can assume that $k = 2$.

The plan, of course, is to define the morphism $\phi_1 : K_0(A_1) \to K_0(B_2)$ by $\phi_1(e_j) = e'_j$. We will then have $\Phi(\hat{\alpha}_{1*}(e_j)) = \hat{\beta}_{2*}(\phi_1(e_j))$ for every generator of $K_0(A_1)$, and hence ϕ_1 will exactly satisfy the required commutativity condition $\Phi \circ \hat{\alpha}_{1*} = \hat{\beta}_{2*} \circ \phi_1$. However, this morphism will not lift by Proposition 4.4 unless we can ensure that ϕ_1 is scale-preserving. Fortunately, by composing with the maps β_j each e'_j (and every multiple thereof which is also in the scale) must eventually be in the scale of some $\Xi(B_j)$, and because there are finitely many such generators and distinct multiples, we can delete additional intermediary objects from the direct limit system $\{B_j, \beta_j : B_j \to B_{j+1}\}$ in order to ensure that ϕ_1 is actually scale preserving as well.

By Proposition 4.4 we therefore now have a lift $g_1 : A_1 \to B_2$. By repeating the same argument with the labels associated to A and B interchanged (and fixing some indices, and using Φ^{-1}), we also get an order and scale preserving map $\psi_1 : B_2 \to A_3$ with a lift $f_1 : B_2 \to A_3$. The final problem is that the triangle

$$\begin{tikzcd} A_1 \arrow[r, "\alpha_{1*}"] \arrow[dr, "g_1"'] & A_2 \arrow[r, "\alpha_{2*}"] & A_3 \\ & B_2 \arrow[ur, "f_1"'] & \end{tikzcd} \tag{4.1}$$

need not actually commute, and so our lift was slightly premature. In order to rectify this observe that by construction the maps ϕ_1 and ψ_1 satisfy the equations $\Phi \circ \hat{\alpha}_{1*} = \hat{\beta}_{2*} \circ \phi_1$ and $\Phi^{-1} \circ \hat{\beta}_{2*} = \hat{\alpha}_{3*} \circ \psi_1$, and therefore we have

$$\hat{\alpha}_{3*} \circ (\psi_1 \circ \phi_1) = \Phi^{-1} \circ \hat{\beta}_{2*} \circ \phi_1 = \Phi^{-1} \circ \Phi \circ \hat{\alpha}_{1*} = \hat{\alpha}_{1*}.$$

This equality factors through a finite-dimensional algebra A_k, in that there is some $k \in \mathbf{N}$ for which we have

$$\alpha_{3\to k*} \circ (\psi_1 \circ \phi_1) = \alpha_{1\to k*}$$

with $\alpha_{j\to j'*}$ the morphism which is just the composition $\alpha_{j'-1*} \circ \alpha_{j'-2*} \circ \cdots \circ \alpha_{j*}$. We can now "fix" the morphism ψ_1 by post-composing it with $\alpha_{3\to k*}$, so without loss of generality we can assume that $k = 3$ and the desired commutativity relation $\psi_1 \circ \phi_1 = \alpha_{1\to 3*}$ is satisfied.

Now considering the lifts f_1 and g_1, we still do not have commutativity of the triangle (4.1), but we do have $(f_1 \circ g_1)_* = \alpha_{1\to 3*}$, which is almost what we desire. The proof is completed by Lemma 4.5, which provides that the composite $f_1 \circ g_1$ must differ from $\alpha_{1\to 3}$ by post-composition with a map which is conjugation by a unitary. Redefining f_1 to include this post-composition, the triangle (4.1) commutes, and by induction we obtain an entire infinite family to which Lemma 4.8 applies.

Now suppose that Φ is not scale preserving. Then it is easily seen that $\Xi(A \otimes \mathscr{K}(H))$ is all of $K_0^+(A \otimes \mathscr{K}(H))$ (and similarly for B), and thus by Proposition 3.9 we have an isomorphism $\hat{\Phi} : K_0(A \otimes \mathscr{K}(H)) \to K_0(B \otimes \mathscr{K}(H))$ which is order and (automatically) scale preserving. The previous part of the theorem now applies to $\hat{\Phi}$ (we must use continuity to insert tensor products with $\mathscr{K}(H)$ everywhere in the above diagrams), and this completes the proof. $\square$

There is a generalisation of this theory to nonunital C^*-algebras, and those scaled ordered groups which arise as the extended K_0 groups of AF-algebras have also been completely classified. Much work has been done attempting to extend the classifying power of this construction to more general settings than AF-algebras, but this has happened in general with little success (for example, a counterexample to direct extension of this classification to all of the C^*-algebras is provided by $C_0(\mathbf{R}^2)$, given our discussion of its K-theory in the previous section). Comprehensive accounts of these attempts can be found in [15, 2].

5 Higher K-groups

In this section we rapidly introduce the K_1 group and summarise its properties, before establishing the (expected, given the topological case) isomorphism between K_1 and $K_0\Sigma$ and defining the higher K-groups. We conclude by showing how this map is a special case of the very important *index map*, and we also easily obtain split exactness of the K_n functors as a consequence of our work.

5.1 The functor K_1

In comparison to the K_0 functor, the functor K_1 is much easier to define—no Grothendieck completion is required, and there are no potential problems due to the absence of a unit:

Definition 5.1. Let A be a C^*-algebra. As in the case of the family of algebras $\mathbf{M}_n(A)$, we can take a direct limit to from the C^*-algebra

$$\mathrm{GL}_\infty^+(A) = \varinjlim_n \mathrm{GL}_n^+(A)$$

where we include $T \in \mathrm{GL}_n^+(A)$ into $\mathrm{GL}_{n+1}^+(A)$ by $T \mapsto \begin{bmatrix} T & 0 \\ 0 & 1 \end{bmatrix}$. Then the group $K_1(A)$ is defined by

$$K_1(A) = \mathrm{GL}_\infty^+(A) / \mathrm{GL}_\infty^+(A)_0,$$

for $\mathrm{GL}_\infty^+(A)_0$ the path component of $\mathrm{GL}_\infty^+(A)$ containing the identity. (In particular, $\mathrm{GL}_\infty^+(A)$ is a topological group under the induced topology.) Elements of $K_1(A)$ are then represented by particular invertible elements of $\mathbf{M}_n(A^+)$ (for some n) which can be viewed as elements of $\mathrm{GL}_\infty^+(A)$ by direct sum with the "infinite identity matrix". Of great utility is the fact that under this quotient we can assume that the representatives are themselves unitary.

It is an elementary fact about the paths in GL_n^+ that the operation

$$[T] + [S] \mapsto [TS]$$

is a commutative group multiplication (see [19, 2] for the details).

In order to get somewhat of a handle on this definition, we will first calculate $K_1(\mathbf{C})$. Indeed, let $U \in \mathrm{GL}_n^+(A)$ be unitary. Then in fact $U = V + I_n$ for $V \in \mathbf{M}_n(A)$ unitary and I_n the $n \times n$ identity matrix in $\mathbf{M}_n(A^+)$ defined above. We can then define a continuous map $f_t : [0,1] \times \mathbf{M}_n(A) \to \mathbf{M}_n(A)$ by $f_t = e^{t \ln V}$ by employing the Borel functional calculus (which is actually available for any von Neumann algebra other than $\mathbf{C}$ as well). Thus we have a path connecting V and thus U to the respective identity matrices of their groups. This shows that $\mathrm{GL}_\infty^+(A)$ consists of only a single path component, and hence $K_1(\mathbf{C}) \cong 0$. This is a prime example of large machinery in analysis (in this case the functional calculus) solving problems in the more algebraic setting of K-theory.

In fact, $K_1(A)$ can be defined using the kernel of a map from $K_1(A^+)$ in the same way that $K_0(A)$ was defined in terms of $GV(A^+) \cong K_0(A^+)$. However, this turns out to always give the same definition because $K_1(A) \cong K_1(A^+)$ always due to the (forthcoming) properties of K_1 given below combined with the fact that $K_1(\mathbf{C}) \cong 0$.

Furthermore, despite the materially different definition of K_1 in comparison to K_0, the four fundamental properties of Subsection 3.3 which hold for K_0 also hold for K_1 by modifying the matrix computations in the proofs; thus, we conclude this subsection by only quickly listing these properties (without proof) below. As a consequence of continuity in particular, we obtain $K_1(A) \cong 0$ whenever A is almost-finite—this follows because the argument above for $K_1(\mathbf{C})$ extends to matrix algebras $K_1(\mathbf{M}_n(\mathbf{C}))$, and therefore any almost-finite dimensional C^*-algebra as well.

Proposition 5.2 (Half-exactness)**.** *Let*

$$0 \longrightarrow A \xrightarrow{\;f\;} B \xrightarrow{\;g\;} C \longrightarrow 0$$

be an exact sequence of C^*-algebras. Then the sequence

$$K_1(A) \xrightarrow{\;f_*\;} K_1(B) \xrightarrow{\;g_*\;} K_1(C)$$

is also exact.

Proposition 5.3 (Continuity)**.** *The functor K_1 is continuous, in that it preserves direct limits of C^*-algebras.*

Proposition 5.4 (Stability)**.** *If A is a C^*-algebra, then $K_0(A \otimes \mathcal{K}(H)) \cong K_0(A)$, so that K_1 is invariant under stabilisation (tensoring with $\mathcal{K}(H)$).*

Proposition 5.5 (Homotopy invariance)**.** *Homotopic maps $f, g : A \to B$ induce the identical maps $f_*, g_* : K_1(A) \to K_1(B)$.*

5.2 The suspension map

In complete analogy with the topological case, we desire (and can find) a *suspension map* (actually, an isomorphism) $\theta_A : K_1(A) \to K_0(\Sigma A)$ for each C^*-algebra A. The corresponding isomorphism the other way (i.e. $K_1\Sigma \to K_0$) is called the *Bott map*, and will be defined in the next section.

Perhaps because the suspension map (along with the more general *index map* to be defined below) is not induced by the action of a functor, the construction of each θ_A is technically very challenging and requires a great deal of unappealing matrix calculations. Indeed, in [19] Wegge-Olsen advises in such cases to "...fasten seat belts and in particular to keep an eye on the *sizes* of the rather overwhelming number of matrices involved." Here we avoid this prospect and instead prefer to simply give the definitions and general idea of the construction.

Theorem 5.6. *There is a natural transformation $\theta : K_1 \to K_0\Sigma$.*

Proof sketch. Fix a C^*-algebra A. There are five major steps:

1. Given a representative $T \in \mathrm{GL}_n^+(A)$ of an element of $K_1(A)$, we must construct an element of $K_0(\Sigma A)$—at the very least, we need a loop of projections on (a matrix algebra of) A. The idea is to set
$$S = \begin{bmatrix} T & 0 \\ 0 & T^{-1} \end{bmatrix} \in \mathrm{GL}_{2n}^+(A)$$
and note that because multiplication by elementary matrices can be performed via continuous homotopy, there is a path $R_t : [0,1] \times \mathrm{GL}_{2n}^+(A) \to \mathrm{GL}_{2n}^+(A)$ from S to I_{2n}. In order to make this into a loop of projections, we conjugate the projection defined by
$$\hat{I}_n = \begin{bmatrix} I_n & 0 \\ 0 & 0 \end{bmatrix} \in \mathrm{GL}_{2n}^+(A).$$
That is, we define $\hat{R}_t : [0,1] \times \mathrm{GL}_{2n}^+(A) \to \mathrm{GL}_{2n}^+(A)$ by $\hat{R}_t = R_t \hat{I}_n R_t^{-1}$. Then
$$\hat{R}_0 = S\hat{I}_n S^{-1} = \begin{bmatrix} TT^{-1} & 0 \\ 0 & 0 \end{bmatrix} = \hat{I}_n,$$
and $\hat{R}_1 = I_{2n}\hat{I}_n I_{2n} = \hat{I}_n$, so we have a loop. Moreover, we have
$$\rho_A(\hat{R}_t) = \rho_A(R_t)\rho_A(\hat{I}_n)\rho_A(R_t^{-1}) = \rho_A(\hat{I}_n) = \hat{I}_n$$
for all $t \in [0,1]$, and so $\hat{R}_t$ has the claimed codomain.

Thus $\hat{R}_t$ is in $GV(A^+)$, but the problem is that it need not *a priori* be in $K_0(A)$. Fortunately, by construction the difference $\hat{R}_t - \hat{I}_n$ has entries in A, and so we can safely define $\theta_A : K_1(A) \to K_0(\Sigma A)$ by

$$\theta_A([u]) = [\hat{R}_t] - [\hat{I}_n].$$

As would appear likely after making so many arbitrary choices in the definition of a function, checking well-definedness (and then injectivity and surjectivity!) is a great deal of work. This requires a tremendous amount of matrix computation; we do not do this here, instead only listing our objectives if we were to proceed.

2. We are required to show that this definition of θ_A is independent of the homotopy class of u as chosen above, and it will be very difficult to show that θ_A is a morphism of groups unless we show also that $\theta_A([u])$ is independent of the homotopy we took.

3. To show injectivity we must explicitly construct a path between representatives u and v of $[u]$ and $[v]$ respectively, which have the same image under θ_A.

4. The proof of surjectivity is a difficult matrix calculation relying on the fact that every element of $K_0(A)$ is represented by a difference $[p] - [I_n]$ as in Proposition 3.5.

5. Finally, showing naturality is not too difficult—we just expand the definitions.

See [2, 19] for an elaboration of these final steps (note however that the "index map" defined below is a generalisation of this construction). □

Given the content of Theorem 5.6, we are immediately motivated to make the following (partially alternate/re-)definition.

Definition 5.7. For each $n \in \mathbf{N}$, we define the functor K_n by (Σ^n is just n-fold application of the suspension functor)

$$K_n = K_0 \Sigma^n.$$

5.3 The index map

On the path to establishing the six-term short exact sequence, we devote the remainder of this section to at least defining the connecting map $\partial : K_1(A/B) \to K_0(B)$ permitting us to obtain the exact sequence

$$\cdots \xrightarrow{\partial} K_1(B) \to K_1(A) \to K_1(A/B) \xrightarrow{\partial} K_0(B) \to K_0(A) \to K_0(A/B) \tag{5.1}$$

for each C^*-algebra A and ideal B of A. In particular, each "run" of three groups in the sequence is exact by half-exactness of the functor K_j (for $f \in \{0,1\}$), and the role of ∂ is to connect these maps together (via suspension, we will be able to obtain all of the required connecting maps given just $\partial : K_1(B/A) \to K_0(B)$).

We will at least define ∂. Indeed, it is not too difficult to see how one could arrive at the following definition by generalising their successful attempt at defining θ above.

Definition 5.8. Let A be a C^*-algebra and let B be an ideal of A. For each $[U] \in K_1(A/B)$ (i.e. with $U \in \mathrm{GL}_n^+(A/B)$ for some $n \in \mathbf{N}$), find a $V \in \mathrm{GL}_{2n}^+(A)$ which projects to the matrix

$$\begin{bmatrix} U & 0 \\ 0 & U^{-1} \end{bmatrix} \in \mathrm{GL}_{2n}^+(A/B).$$

Then the *index map* $\partial : K_1(A/B) \to K_0(A)$ is defined to act on $[U]$ by

$$\partial([U]) = [V\hat{I}_n V^{-1}] - [\hat{I}_n]$$

for $\hat{I}_n \in GL_{2n}^+(A)$ defined in blocks by

$$\hat{I}_n = \begin{bmatrix} I_n & 0 \\ 0 & 0 \end{bmatrix}.$$

In the special case that $B = \mathscr{B}(H)$ and $A = \mathscr{K}(H)$, we get an index map $\partial : K_0(\mathscr{B}(H)/\mathscr{K}(H)) \to K_0(\mathscr{K}(H))$. Composing with the canonical isomorphism $K_0(\mathscr{K}(H)) \to \mathbf{Z}$ of Subsection 3.2 gives a way of computing an integer associated to element of the Calkin algebra of H. In fact, this corruption of ∂ computes the famous *Fredholm index* ([3] provides a solid introduction), and thus confers the index map its name.

Theorem 5.9 (Index map existence)**.** *The map $\partial : K_1(B/A) \to K_0(B)$ of Definition* 5.8 *is well-defined, and the resulting ∂ makes* (5.1) *into a long exact sequence.*

Moreover, using Lemma 3.6 *to obtain a map a $\partial : K_1(C) \to K_0(B)$ for general short exact sequences $0 \to A \to B \to C \to 0$, the family of maps called ∂ define a natural transformation from K_1 to K_0 in the sense that every commutative diagram*

$$\begin{tikzcd} 0 \arrow[r] & A \arrow[r] \arrow[d] & B \arrow[r] \arrow[d] & C \arrow[r] \arrow[d] & 0 \\ 0 \arrow[r] & A' \arrow[r] & B' \arrow[r] & C' \arrow[r] & 0 \end{tikzcd}$$

of exact sequences gives rise to a commutative and exact diagram

$$\begin{tikzcd} K_1(A) \arrow[r] \arrow[d] & K_1(B) \arrow[r] \arrow[d] & K_1(C) \arrow[r, "\partial"] \arrow[d] & K_0(A) \arrow[r] \arrow[d] & K_0(B) \arrow[r] \arrow[d] & K_0(C) \arrow[d] \\ K_1(A') \arrow[r] & K_1(B') \arrow[r] & K_1(C') \arrow[r, "\partial"] & K_0(A') \arrow[r] & K_0(B') \arrow[r] & K_0(C') \end{tikzcd} . \tag{5.2}$$

The proof of the well-definedness of ∂ is very difficult, and is rivalled in length only by the proof of the Bott Periodicity theorem. One must show

- that ∂ does not depend on any of the data (lifts and representatives) used to define it, and then
- one must then actually check that it makes the sequence (5.1) exact (neither part of which is easy).

One then inspects the definitions of the maps in (5.2) in order to directly check exactness and commutativity, and for naturality only the square involving ∂ need be checked (see [2], or for a more comprehensive account [19]).

As consolation, we will at least show (as promised) how Theorem 5.9 gives an easy proof of Theorem 5.6, and therefore that the latter is a special case of the former.

Corollary 5.9.1 (Reproof of Theorem 5.6)**.** *There is a natural isomorphism $\theta : K_1 \to K_0\Sigma$.*

Proof. Let A be a C^*-algebras. By Proposition 3.14 there is a short exact sequence

$$0 \longrightarrow \Sigma A \longrightarrow CA \longrightarrow A \longrightarrow 0 .$$

Therefore Theorem 5.9 provides[7] that the sequence

$$K_1(\Sigma A) \longrightarrow K_1(CA) \longrightarrow K_1(A) \xrightarrow{\partial} K_0(\Sigma A) \longrightarrow K_0(CA) \longrightarrow K_0(A)$$

is exact. But cones are contractible by Lemma 3.13, and therefore $\partial : K_1(A) \to K_0(\Sigma A)$ is actually an isomorphism. The fact that $\theta = \partial$ is a natural transformation then completes the proof. □

Finally, we will use the additional generality of the index map ∂ (over the suspension map) in order to obtain a useful result (in its own right), on the split exactness of the K-group functors. It will also prove useful in slightly simplifying the task of proving the Bott Periodicity theorem.

[7] Once again, we use the general result of Lemma 3.6 in order to generalise an exact sequence obtained from an actual honest quotient of C^*-algebras.

Proposition 5.10 (Split exactness of K_n)**.** *The functor K_0 is split exact. Therefore, the functor K_n is split exact for every $n \in \mathbf{N}$.*

Proof. A split exact sequence

$$0 \longrightarrow A \longrightarrow B \longleftrightarrow C \longrightarrow 0$$

by Theorem 5.9 combined with Lemma 3.6 gives rise to a split exact sequence

$$K_1(A) \longrightarrow K_1(B) \longleftrightarrow K_1(C) \xrightarrow{\partial} K_0(A) \longrightarrow K_0(B) \longleftrightarrow K_0(C) \ .$$

This forces the connecting map ∂ to be zero, and the rightmost map must be surjective because it has a section, splitting it. Therefore the rightmost three terms break off into a split exact sequence

$$0 \longrightarrow K_0(A) \longrightarrow K_0(B) \longleftrightarrow K_0(C) \longrightarrow 0 \ ,$$

as desired. The fact that there is a natural isomorphism $K_n \cong K_{n-1}\Sigma$ (or, depending on one's definitions, an equality) then completes the proof. □

6 Bott Periodicity, and a "Thom isomorphism"

Given the connection between topological and C^*-algebraic K-theory (which is made particularly concrete in Section 7), it should be expected that we have a Bott Periodicity theorem for C^*-algebras—and especially in at least the commutative case! As it turns out, we do, and we also have a "Thom isomorphism" which is actually much more closely related to Bott Periodicity (it is a high-powered generalisation) than the Thom isomorphism in the topological case.

6.1 Traditional Bott Periodicity and the 6-term exact sequence

Indeed, the obvious analog of the celebrated Bott Periodicity holds in the case of general C^*-algebras. The proof of the Bott Periodicity theorem is notoriously intricate, even if possible to formulate in quite elementary terms, and therefore cannot possibly be included here. Instead, we overview the map which induces the Bott Periodicity and introduce related theorems known as the P-V sequence and the Connes "Thom" isomorphism, these latter two results being specific to C^*-algebraic K-theory. We conclude the section by noting how both Bott Periodicity and the P-V sequence follow from the Connes "Thom" isomorphism (which, despite the name, is a vast generalisation of Bott Periodicity and not the Thom isomorphism!).

Theorem 6.1 (Bott Periodicity)**.** *There is a natural isomorphism $\beta : K_0 \to K_1\Sigma$. Therefore, for every C^*-algebra A we have $K_n(A) \cong K_n(\Sigma^2 A)$ for every $n \in \mathbf{N}$ (by Theorem 5.6).*

Proof sketch. Let A be a C^*-algebra. The crux of the proof is define the *Bott map* $\beta_A : K_0(A) \to K_1(\Sigma A)$, for which it can then be shown quite easily that β_A defines a component of a natural transformation. One then performs the quite arduous task of verifying that each map β_A is actually an isomorphism. We primarily follow the exposition of [2], and also [19], in which the additional details may be found.

We will at least give the definition of the Bott map. Fix a projection $p \in \mathbf{M}_\infty(A^+)$ with some dimension n. Then we have a loop in $\mathrm{GL}_n(A^+)$ defined by

$$s_p(z) = (z-1)p + I_n = (z-1)p + \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & 1 \end{bmatrix}.$$

The map s_p is invertibly-valued (and thus we have a companion map s_p^{-1} valued in the inverse of s_p), and has $s_p(1)$ the identity. It therefore represents a class in $K_1(\Sigma A)$. The Bott map is then defined simply by

$$\beta_A([p]-[q]) = [s_p s_q^{-1}].$$

This is well defined as a function because (by using compactness of S^1 to verify continuity) homotopic projections p and p' give rise to homotopic maps s_p and $s_{p'}$, and in particular this shows additivity. The map β_A also obviously preserves the identity.

We have actually defined the components of a natural transformation because for any map $f : A \to B$ we have that

$$(f_* \circ \beta_A)([p]-[q]) = [f(s_p s_q^*)] = [f(s_p) f(s_q)^*] = [s_{f(p)} s_{f(q)}^*] = (\beta_B \circ f_*)([p]-[q]).$$

We reduce the proof to the case of unital A by using a variant of the short five lemma applied to the diagram (with the C^*-algebra A arbitrary)

$$\begin{tikzcd}
0 \arrow[r] & K_0(A) \arrow[r] \arrow[d, "\beta_A"] & K_0(A^+) \arrow[r] \arrow[d, "\beta_{A^+}"] & K_0(\mathbf{C}) \arrow[r] \arrow[d, "\beta_{\mathbf{C}}"] & 0 \\
0 \arrow[r] & K_1(\Sigma A) \arrow[r] & K_1(\Sigma A^+) \arrow[r] & K_1(\Sigma \mathbf{C}) \arrow[r] & 0
\end{tikzcd}$$

i.e. we need only show that the two rightmost vertical arrows are isomorphisms, given that the rows are exact by split exactness of the functors K_0 and K_1 (this is Proposition 5.10).

The actual body of the proof can then proceed in many ways. Of note is the fact that Atiyah's famous method for the case of topological K-theory actually generalises to the "noncommutative case"; we can show both injectivity and surjectivity by approximating arbitrary loops as Laurent polynomials and then simply polynomials... see [18] or [2] for the familiar details.

This method is truly "K-theoretic", in the sense that it downplays the C^*-structure. Proofs using honestly-noncommutative machinery include Cuntz's proof via the unilateral shift operator and Toeplitz algebras (see e.g. [11]), later proofs by Atiyah [1] and Phillips [12] also using Toeplitz operators, and Nest, Natsume, and Elliott's proof in [6] via the study of particular so-called *crossed products* of C^*-algebras. □

It follows immediately that the infinite exact sequence (5.1) wraps around on itself, giving rise to the famous six-term exact sequence of the K-theory of C^*-algebras.

Theorem 6.2. *Let $0 \to A \to B \to C \to 0$ be an exact sequence of C^*-algebras. Then the following six term sequence is exact.*

$$\begin{tikzcd}
K_0(A) \arrow[r] & K_0(B) \arrow[r] & K_0(C) \arrow[d, "\partial"] \\
K_1(C) \arrow[u, "\partial"] & K_1(B) \arrow[l] & K_1(A) \arrow[l]
\end{tikzcd}$$

In particular, the map $\partial : K_0(C) \to K_1(A)$ is the dotted arrow in the diagram

$$\begin{tikzcd}
K_1(\Sigma C) \arrow[r, "\partial"] & K_0(\Sigma A) \\
K_0(C) \arrow[u, "\beta_C"] \arrow[r, dotted] & K_1(A) \arrow[u, "\theta_B"']
\end{tikzcd}$$

with the top row coming from the long exact sequence associated to the short exact sequence defined in Lemma 3.15.

The connecting map $\partial : K_0(C) \to K_1(A)$ is called the *exponential map*, because of the special case where A is an ideal of B and $C = B/A$. There, given a difference $[p]-[I_n]$ in $K_0(B/A)$ the map ∂ acts by finding $q \in \mathbf{M}_\infty(B^+)$ such that $\rho_B(q) = p$, and then uses the the functional calculus to emit $\partial([p]-[I_n]) = [e^{i2\pi q}]$.

6.2 Surjectivity of K_0 and K_1

The establishment of the six-term exact sequence above (finally!) provides the opportunity for a first neat application which combines all of the machinery we have developed thus far; we will show that the functors K_0 and K_1 are surjective on the abelian groups. The problem essentially reduces to some elementary group theory, after we establish the following definitions and pair of easy lemmas.

Definition 6.3. The *mapping cylinder* for a morphism $f : A \to B$ of C^*-algebras is

$$Y_f = \{x \oplus s : A \oplus C([0,1] \to B) : s(1) = f(x)\}.$$

The *mapping cone* is similarly defined by

$$C_f = \{x \oplus s \in A \oplus CB : s(1) = f(x)\},$$

and is obviously a subset of the corresponding mapping cylinder.

Lemma 6.4. *Given a morphism $f : A \to B$ of C^*-algebras, we have that A is a deformation retract of the mapping cylinder Y_f. Therefore $K_n(A) = K_n(Y_f)$ for all $n \in \mathbf{Z}$.*

Proof. The idea is just to "crush" the function part $C([0,1] \to B)$ of Y_f. Indeed, the map $g_t : [0,1] \times Y_f \to Y_f$ defined by

$$g_t(x \oplus s) = x \oplus ((1-t)s + tf(x))$$

suffices. □

Lemma 6.5. *Given a morphism $f : A \to B$ of C^*-algebras, the sequence*

$$0 \longrightarrow C_f \xrightarrow{\subset} Y_f \xrightarrow{\mathrm{ev}_0} B \longrightarrow 0$$

with $\mathrm{ev}_0 : Y_f \to B$ *defined by* $x \oplus s \mapsto s(0)$ *is exact.*

Proof. If $x \oplus s \in C_f$ then in particular $s(0) = 0$ by hypothesis, and therefore $\mathrm{ev}_0(x \oplus s) = 0$. Conversely, if $\mathrm{ev}_0(x \oplus s) = s(0) = 0$, then given that $x \oplus s \in Y_f$, this is precisely the condition that $x \oplus s \in C_f$. □

Proposition 6.6. *Let G be an abelian group. Then there exists a C^*-algebra G^* such that $K_0(G^*) \cong 0$ and $K_1(G^*) \cong G$.*

Proof. We will essentially transfer the canonical free resolution/presentation of G to the K-theory world via the six-term exact sequence (the main idea is credited to Higson and Brown, and can be found in [19]). First, consider the short exact sequence

$$0 \longrightarrow \ker\pi \hookrightarrow \mathbf{Z}\{G\} \xrightarrow{\pi} G \longrightarrow 0$$

where $\pi : \mathbf{Z}\{G\} \to G$ is the canonical map from the free abelian group on the elements of G to G, which is completely defined by specifying for each $n \in \mathbf{Z}$ and $g \in G$ that $\pi(ng)$ is the n-fold sum of g with itself. Now, $\ker\pi$ is a subgroup of a free abelian group, and therefore is itself free. It will be problematic for us that the inclusion $\ker\pi \hookrightarrow \mathbf{Z}\{G\}$ can give rise to linear combinations involving negative coefficients, and so we must actually refine this short exact sequence to a general one

$$0 \longrightarrow \mathbf{Z}\{I\} \xrightarrow{\phi} \mathbf{Z}\{J\} \xrightarrow{\psi} G \longrightarrow 0$$

for some indexing sets I and J, so that each $i \in I$ is sent via ϕ to a positive linear combination in $\mathbf{Z}\{J\}$. This is always possible by "hiding the negatives in the maps ψ"; we form an exact sequence $0 \to \ker\pi \oplus \mathbf{Z}\{G\} \xrightarrow{\phi} \mathbf{Z}\{G\} \oplus \mathbf{Z}\{G\} \xrightarrow{\psi} G \to 0$ by defining the ψ of $x \oplus y$ to take the difference $\pi(x) - \pi(y)$, and defining[8] ϕ to ensure that the sequence remains exact. For each $i \in I$ there are thus nonnegative integers $k_{i,j}$ such that $\phi(i) = \sum_{j \in J} k_{i,j} j$.

[8] This can be achieved for example by defining ϕ on $x \oplus y$ by sending the "positive part" of the inclusion of $x \in \ker\pi$ into $\mathbf{Z}\{G\}$ to the first copy of $\mathbf{Z}\{G\}$ and the "negative part" to the second copy, and sending the y part unchanged to both components (when x and y are both nonzero we sum).

The idea is to replicate the free groups $\mathbf{Z}\{I\}$ and $\mathbf{Z}\{J\}$ with copies of the compact operators; we define $A = \bigoplus_{i\in I}\mathscr{K}(H)$ and $B = \bigoplus_{j\in J}\mathscr{K}(H)$, so that $K_0(A) \cong \bigoplus_{i\in I}\mathbf{Z}$ and $K_0(B) \cong \bigoplus_{j\in J}\mathbf{Z}$, while $K_1(A) \cong K_1(B) \cong 0$ (by continuity). We then have a map $f : A \to B \otimes \mathscr{K}(H)$ which sends x in the copy of $\mathscr{K}(H)$ with index $i \in I$ to $x \otimes I_{k_{i,j}}$ in each copy of $\mathscr{K}(H) \otimes \mathscr{K}(H)$ in B respectively indexed by $j \in J$ (only finitely many of these $k_{i,j}$'s are nonzero at a time). Note that here we use I_n to denote a rank n projection in $\mathscr{K}(H)$. Now because $[x \otimes I_{k_{i,j}}]$ is sent to $k_{i,j}[x]$ under the isomorphism $\chi : K_0(B \otimes \mathscr{K}(H)) \cong K_0(B)$, the map f is such that $\chi \circ K_0(f) = \phi$.

By Lemma 6.5 and then Theorem 6.2 we thus have a six-term exact sequence

$$\begin{tikzcd}
K_0(C_f) \arrow[r] & K_0(A) \arrow[r] & K_0(B) \arrow[d, "\partial"] \\
0 \arrow[u, "\partial"] & 0 \arrow[l] & K_1(C_f) \arrow[l]
\end{tikzcd},$$

where we have used Lemma 6.4 to replace $K_j(Y_f)$ with $K_j(A)$ for $j \in \{0,1\}$. But the map $K_0(A) \to K_0(B)$ is essentially just ϕ, which is injective. We therefore have an exact sequence

$$0 \longrightarrow K_0(A) \cong \mathbf{Z}\{I\} \xrightarrow{\ \phi\ } \mathbf{Z}\{J\} \cong K_0(B) \xrightarrow{\ \partial\ } K_1(C_f) \longrightarrow 0$$

which proves that $K_1(C_f) \cong G$! Letting $G^* = C_f$ then completes the proof. □

The general statement of the surjectivity of K_0 and K_1 is now an easy consequence of Bott Periodicity.

Theorem 6.7. *Let G and G' be abelian groups. Then there exists a C^*-algebra A such that $K_0(A) \cong G$ and $K_1(A) \cong G'$.*

Proof. By Proposition 6.6 there exist C^*-algebras B and B' so that $K_1(B) \cong G$, $K_1(B') \cong G'$, and $K_0(B) \cong K_0(B') \cong 0$. Therefore, setting $A = \Sigma B \oplus B'$ we have

$$K_0(A) \cong K_1(B) \oplus K_0(B') \cong G \quad \text{and} \quad K_1(A) \cong K_0(B) \oplus K_1(B') \cong G',$$

as desired. □

6.3 Generalisations: Connes' "Thom" Isomorphism and the P-V sequence

That our remarks concluding the "proof" of Bott periodicity referenced crossed products of C^*-algebras is quite topical, for Nest, Natsume, and Elliott's method can be generalised to give a proof of the so called Connes "Thom" Isomorphism. The purpose of this following development is to sketch the statement of this result, and show how the other famous Pimsner-Voiculescu (P-V) sequence follows as a corollary. These results stand on their own in sharp contrast to the results of topological K-theory, where they have no directly analogous statement. In order to introduce them, we must first establish the following firmly C^*-algebraic construction;

Definition 6.8. Let A be a C^*-algebra, and let $\mathrm{Aut}(A)$ be the C^*-automorphism group of A equipped with the topology induced by pointwise convergence. Given a continuous group homomorphism $T : G \to \mathrm{Aut}(A)$ from any locally compact group G (we will only actually care about the cases $G = \mathbf{Z}$ and $G = \mathbf{R}$), we can then form the *crossed product* $A \ltimes_T G$ through a slightly elaborate construction (for which he have neither the time nor necessity to fully explain):

- The idea is to consider pairs of representations of A and G on the same Hilbert space H. Pairs $(\rho : A \to \mathscr{B}(H), \sigma : G \to \mathscr{B}(H))$ which are compatible in the sense that

$$\sigma(g)\rho(x)\sigma(g)^* = \rho(\phi_g(x))$$

  for every $g \in G$ and $x \in A$ are called *covariant representations of ϕ*.

- Because G is a locally compact group, to each covariant representation is associated a certain normed associative algebra with involution called the *convolution algebra* $C_{(\rho,\sigma)}$ for the automorphism ϕ. This is the set $C_c(G \to A)$ of compactly supported continuous functions $G \to A$ with a with a norm defined by a convolution integral (which we will not describe here). This norm is submultiplicative, and so the completion of each convolution algebra is a C^*-algebra.
- We can define another norm on $C_c(G \to A)$ by taking the supremum over all of the norms on the convolution algebras associated to each covariant representation, and again under the completion we obtain a C^*-algebra.
- The resulting object actually satisfies a universal property, and is called $A \ltimes_\phi G$, *the crossed product of A with G by ϕ.*

Of note is the fact that the crossed product of an abelian C^*-algebra with an abelian group need not be abelian, which is a significant obstacle to the transfer of the results of this subsection to topological K-theory. On the other hand, when $A = C_0(X)$ (for X a compact space) actions $T : G \to \mathrm{Aut}(A)$ define parameterised families of continuously-varying homeomorphisms of X (by Gelfand duality), and so the analogy is far from completely broken.

The primary result on the calculation of the K-theories of crossed products is the following celebrated theorem.

Theorem 6.9 (Connes "Thom" Isomorphism)**.** *Let A be a C^*-algebra and let $T : \mathbf{R} \to \mathrm{Aut}(A)$ be a continuous group homomorphism. Then*

$$K_n(A \ltimes_T \mathbf{R}) \cong K_{n+1}(A)$$

for every $n \in \mathbf{N}$.

This is a "Thom isomorphism" only in the sense that a version of the topological Thom isomorphism gives for a k-dimensional vector bundle $E \to X$ that $K^n(E) \cong K^{n+k}(X)$ for every $n \in \mathbf{N}$ (in $K^n(E)$ we consider E a locally compact base space in its own right). One method of proof of Theorem 6.9 proceeds (quite intuitively) by carefully deforming an arbitrary $\mathbf{R}$ action into the trivial one, but this is not quite as easy as it sounds—see 10.9 of [2] for one possible elaboration of this idea.

As a trivial example of the immense power of this result, consider the trivial action $T : \mathbf{R} \to \mathrm{Aut}(A)$ on a C^*-algebra A defined by $T(r) = \mathrm{id}_A$ for all $r \in \mathbf{R}$. Then (by the definition of the convolution product)

$$A \ltimes_T \mathbf{R} \cong A \otimes C_0(\mathbf{R}) \cong \Sigma A,$$

in which case Theorem 6.9 is exactly the statement of the Bott Periodicity theorem!

Now let $\phi \in \mathrm{Aut}(A)$ for A a unital C^*-algebra. Then observe that we automatically have an action of $\mathbf{Z}$ on A in the sense that the map $T : \mathbf{Z} \to \mathrm{Aut}(A)$ defined by $T(j) = \phi^j$ is automatically continuous (ϕ^j is just j-fold composition, and makes sense for negative j because ϕ is invertible). Thus, we are free to write $A \ltimes_\phi \mathbf{Z}$ without any ambiguity. Crossed products with $\mathbf{Z}$ are thus at least as easy to obtain as C^*-automorphisms of A, and hence technical tools which deal with them are of great utility. The Pimsner-Voiculescu (P-V) exact sequence is a cornerstone in this regard, and was established prior to Connes' result during the study of the K-theory of so-called irrational rotation algebras by the aforenamed. As another example of the power of the Connes isomorphism, and to highlight the P-V sequence, we conclude the section by showing how the latter follows from the former.

Theorem 6.10 (Pimsner-Voiculescu exact sequence)**.** *Let A be a unital C^*-algebra, and let $\phi \in \mathrm{Aut}(A)$. Then the following six-term sequence is exact everywhere.*

$$\begin{tikzcd}
K_1(A) \arrow[r, "\iota_*"] & K_1(A \ltimes_\phi \mathbf{Z}) \arrow[r] & K_0(A) \arrow[d, "\mathrm{id}-\phi_*"] \\
K_1(A) \arrow[u, "\mathrm{id}-\phi_*"] & K_0(A \ltimes_\phi \mathbf{Z}) \arrow[l] & K_0(A) \arrow[l, "\iota_*"]
\end{tikzcd}$$

The map $\iota : A \to A \rtimes_\phi \mathbf{Z}$ is canonically provided be the definition of the crossed product. The unlabelled connecting maps were originally constructed from an extension

$$0 \to A \otimes \mathscr{K}(H) \to \mathscr{T}(A,\phi) \to A \rtimes_\phi \mathbf{Z} \to 0$$

with $\mathscr{T}(A,\phi)$ called a "Toeplitz extension" (see [13]).

Proof sketch. We only explain how to obtain the six-term exact sequence. In fact, this result follows almost immediately from Theorem 6.9 by applying Green's result [8] on Morita equivalence. We will instead elect to navigate via relative elementary means as in [2], through the *mapping torus of* ϕ defined by

$$M_\phi = \{f : [0,1] \to A : f(1) = \phi(f(0))\}.$$

This is a C^*-algebra.

Now, there is an exact sequence

$$0 \longrightarrow \Sigma A \xrightarrow{\ \subset\ } M_\phi \xrightarrow{\ \mathrm{ev}_0\ } A \longrightarrow 0$$

with the first nonzero map the inclusion (if $f(0) = f(1) = 0$, then $f(1) = 0 = \phi(0) = \phi(f(0))$), and the second is the evaluation-at-0 map $f \mapsto f(0)$. Indeed, it is trivial that $\mathrm{ev}_0(\Sigma A) = \{0\}$, and conversely if $\mathrm{ev}_0(f) = f(0) = 0$ subject to the condition that $f(1) = \phi(f(0)) = \phi(0) = 0$, then $f \in \Sigma A$. This then gives rise to a six-term exact sequence (by Theorem 6.2)

$$\begin{tikzcd} K_1(A) \arrow[r] & K_0(M_\phi) \arrow[r] & K_0(A) \arrow[d, "\delta"] \\ K_1(A) \arrow[u, "\delta"] & K_1(M_\phi) \arrow[l] & K_0(A) \arrow[l] \end{tikzcd} . \tag{6.1}$$

The trick is to use M_ϕ to bootstrap the induced $\mathbf{Z}$-action into an induced $\mathbf{R}$-action, for which Theorem 6.9 will then apply. Indeed, define $\hat{\phi} : \mathbf{R} \to \mathrm{Aut}(M_\phi)$ by (we consider the zero-fold composition of a functor with itself to be the identity, and for invertible maps define negative integral exponents in the obvious way)

$$\hat{\phi}_t(f : [0,1] \to A) = \Big(s \mapsto (\phi^{\lfloor t \rfloor} \circ f)(s + t - \lfloor t \rfloor)\Big).$$

That this map is continuous is true by the definition of the mapping torus M_ϕ; indeed, we have $\phi^n(f(1)) = \phi^{n-1}(f(0))$ for every $f \in M_\phi$ and $n \in \mathbf{N}$.

We can now state and apply[9] (a special case of) a famous theorem—the Takai duality—which transports a similar result of Takesaki for von Neumann algebras to the case of C^*-algebras.

Theorem 6.11 (Takai duality for mapping cones)**.** *Let $\phi \in \mathrm{Aut}(A)$, and $\hat{\phi} : \mathbf{R} \to M_\phi$ the induced action as described above. Then*

$$M_\phi \rtimes_{\hat{\phi}} \mathbf{R} \cong (A \rtimes_\phi \mathbf{Z}) \otimes \mathscr{K}(H).$$

Applying the K_j functor to both sides of this isomorphism and then using stability, we have

$$K_j(M_\phi \rtimes_{\hat{\phi}} \mathbf{R}) \cong K_j((A \rtimes_\phi \mathbf{Z}) \otimes \mathscr{K}(H)) \cong K_j(A \rtimes_\phi \mathbf{Z}).$$

But then by Theorem 6.9 we obtain

$$K_j(M_\phi \rtimes_{\hat{\phi}} \mathbf{R}) \cong K_{j+1}(M_\phi).$$

By Bott periodicity the six term exact sequence of (6.1) becomes

$$\begin{tikzcd} K_1(A) \arrow[r] & K_1(A \rtimes_\phi \mathbf{Z}) \arrow[r] & K_0(A) \arrow[d, "\delta"] \\ K_1(A) \arrow[u, "\delta"] & K_0(A \rtimes_\phi \mathbf{Z}) \arrow[l] & K_0(A) \arrow[l] \end{tikzcd} ,$$

[9] Because we have left the definition of a crossed product somewhat imprecise, we must take the conclusion of Theorem 6.11 at face value—despite the fact that is a significant specialisation of the general result, and (relatedly) is easier to prove.

as desired. One finally sees that $\partial = \mathrm{id} - \phi_*$ by a direct calculation, which essentially exploits the fact that $e^{a+b} = e^a e^b$ in the explicit (special case) formula for the exponential map ([2] has the details). □

By generalising the proof of exactness of the P-V sequence, we can also obtain analogous exact sequences for other crossed products (e.g. with $\mathbf{Z}_n$), but the study of such crossed products is a rich subject in its own right and we would not be able to give it the respect it deserves in the short time we have to spend. Indeed, much is yet to be known; see [2] for a survey of the state-of-the-art.

We end this section by noting the existence of many additional results in C^*-algebraic K-theory which we have not yet covered and *do* have direct analogs in the topological case; for example, we have a Künneth theorem for an extremely large class of C^*-algebras, but once again we must end the discussion here.

7 The connection to topological K-theory

In rounding out our exposition, we conclude by making rigorous the link between topological K-theory and the C^*-algebraic one. Historically, the following results were developed essentially completely in reverse order, and inspired the entire generalisation of topological K-theory to the C^*-algebra case in the first place!

7.1 Phase one: Reinterpreting topological K-theory

If there is to be any connection between topological K-theory and the K-theory of C^*-algebras heretofore developed, we must find a way to relate the two theories' principal objects of study. That is, we must find a way to relate topological objects associated to a base space X (in this case, complex vector bundles) to some kind of algebraic objects associated to a C^*-algebra. Furthermore, there is only one sensible C^*-algebra to choose at that; $C_0(X)$.

Let $\mathrm{Vect}(X)$ denote the category of finite-dimensional complex vector bundles over the base space X (not isomorphism classes). We first introduce the following algebraic definitions.

Definition 7.1. Let M be a module over a ring R.

- If there exists another R-module N such that $M \oplus N$ is free, then M is said to be *projective*.
- If there is a finite subset of M such that the smallest[10] R-submodule of M containing this subset is all of M, then M is said to be *finitely generated*. Equivalently, M is finitely generated if for some $n \in \mathbf{N}$ there is a surjective module map $R^n \to M$.

Let R-Mod denote the category of modules over a ring R, and let $\mathrm{ProjFin}(R)$ denote the full subcategory of R-Mod given by those R-modules which are finitely generated and projective. We are now ready to establish the first pieces of the equivalence;

Proposition 7.2. *There is a functor* $\Gamma : \mathrm{Vect}(X) \to C_0(X)\text{-Mod}$. *If in addition X is compact, then in fact Γ maps into* $\mathrm{ProjFin}(C_0(X))$.

Proof. The point is to observe that the sections of a vector bundle are a module over $C_0(X)$. The ring $C_0(X)$ acts just by pointwise multiplication, $\Gamma(X)$ is an abelian group in its own right, and these structures are compatible. Maps of $E \to F$ of vector bundles send sections of E to sections of F (and this happens functorially, because function composition is associative). This makes Γ into at least a functor from $\mathrm{Vect}(X)$ to $C_0(X)$-modules.

Now fix a vector bundle $p : E \to X$ for X compact. It is an elementary theorem of the theory of vector bundles[11] (for instance, Proposition 1.4 of [9]) that for X compact, there exists another vector bundle $p' : F \to X$ such that $E \oplus F$ is isomorphic to a trivial bundle $\mathbf{C}^n \times X$ (for some $n \in \mathbf{N}$). Because the section functor Γ clearly respects direct sums, it follows that

$$\Gamma(E) \oplus \Gamma(F) \cong \Gamma(E \oplus F) \cong \Gamma(\mathbf{C}^n \times X) \cong \Gamma(\mathbf{C} \times X)^n \cong C_0(X)^n.$$

[10] Concretely, given a set of R-modules each containing a common set S, we can take their intersection.

[11] It is a curious fact that this result is essentially trivial in the noncommutative case (it is effectively the content of Proposition 3.5).

This shows that both $\Gamma(E)$ is a projective module (it is a direct summand of the free module $C_0(X)^n$), and is finitely generated (the n generators of $C_0(X)^n$ project down to $\Gamma(E)$, and must still generate). $\square$

The greater miracle is that there is a "partner" functor the other way!

Proposition 7.3. *For X compact, there is a functor $\Delta : \mathrm{ProjFin}(C_0(X)) \to \mathrm{Vect}(X)$.*

Proof. Let $M \in \mathrm{ProjFin}(C_0(X))$ be arbitrary. By assumption, there exists an $n \in \mathbf{N}$ and a $C_0(X)$-module N such that there is an isomorphism $\phi : C_0(X)^n \to M \oplus N$. We have the projection $\pi : M \oplus N \to M$ and injection $\iota : M \to M \oplus \mathbf{0} \to M \oplus N$ module maps, and thus can form a diagram

$$\begin{tikzcd} C_0(X)^n \arrow[r, "\phi"] \arrow[rrrr, bend right, dotted, "P"'] & M \oplus N \arrow[r, "\pi"] & M \arrow[r, "\iota"] & M \oplus N \arrow[r, "\phi^{-1}"] & C_0(X)^n \end{tikzcd} .$$

The composite $P = \phi^{-1} \circ \iota \circ \pi \circ \phi$ is in particular an idempotent module map $C_0(X)^n \to C_0(X)^n$. For each $x \in X$, we thus have a (linear) projection $P_x : \mathbf{C}^n \to \mathbf{C}^n$ induced by P. In particular, we have a continuous function $P : X \to \mathbf{M}_n(\mathbf{C})$.

The idea is to define

$$\Delta(M) = \{(x, v) \in X \times \mathbf{C}^n : v \in \operatorname{im} P_x\},$$

equipping $\Delta(M)$ with the subspace topology. We claim that the projection $p : \Delta(M) \to X$ defined by $p(x, v) = x$ gives a vector bundle. Now, each fibre $p^{-1}(x)$ is certainly a vector space of dimension $\operatorname{rank} P_x$, and so to show $\Delta(M)$ is a vector bundle it remains to show that we have local trivialisations.

Fix $x \in X$. Then the fibre $p^{-1}(x)$ has some basis $\{v_1, \dots, v_k\}$ for some $k \le n$, and by the definition of $\Delta(M)$ we have that $P_x v_j = v_j$ for every $1 \le j \le k$. By extending this set to a basis of $\mathbf{R}^n$, we can then form the "characteristic matrix"

$$\chi_y = \left[P_y v_1, \dots, P_y v_k, v_{k+1}, \dots v_n \right]$$

for each $y \in X$. Because P is continuous, we have a continuous map $\chi : X \to \mathbf{M}_n(X)$. The key is that by constructing χ in this way we have ensured that χ_x is of full rank. The determinant function is continuous, and therefore there exists a neighbourhood U of x such that χ has full rank on all of U. But this in particular implies that the maps $\psi_j(y) = (y, P_y v_j) : X \to \mathbf{C}^n$ give a linearly independent set $\{\psi_1(y), \dots, \psi_k(y)\}$ for all $y \in U$ (these are just the first k columns of χ_y). Therefore we have constructed k linearly-independent local sections at $x \in X$, from which a local trivialisation is readily obtained.

Hence $\Delta(M)$ is a vector bundle over X for every $M \in \mathrm{ProjFin}(C_0(X))$. Let $f : N \to M$ be a map of finitely generated projective $C_0(X)$-modules, and let $P : X \to \mathbf{M}_n(\mathbf{C})$ and $Q : X \to \mathbf{M}_m(\mathbf{C})$ be the continuous maps associated to N and M respectively in the above construction. We have a commutative diagram (with isomorphisms $\phi : C_0(X)^n \to N \oplus N'$ and $\psi : C_0(X)^m \to M \oplus M'$)

$$\begin{tikzcd} C_0(X)^n \arrow[r, "\pi\circ\phi"] \arrow[rr, bend left, "P"] \arrow[drr, bend right, dotted, "\hat{f}"'] & N \arrow[r, "\phi^{-1}\circ\iota"] \arrow[d, "f"] & C_0(X)^n & & \\ & M \arrow[r, "\psi^{-1}\circ\iota"'] \arrow[rr, bend left, "\mathrm{id}"] & C_0(X)^m \arrow[r, "\pi\circ\psi"'] \arrow[rr, bend right, "Q"'] & M \arrow[r, "\psi^{-1}\circ\iota"'] & C_0(X)^m \end{tikzcd} ,$$

whence we obtain a module map $\hat{f} : C_0(X)^n \to C_0(X)^m$ and hence continuous map $\hat{f} : X \to \mathbf{M}_{n\times m}$. This defines the data of a map $f_* : \Delta(N) \to \Delta(M)$ by

$$f_*(x, v) = (x, \hat{f}(x) v),$$

which is clearly continuous and fibrewise linear, and in particular $\hat{f}(x)v \in \operatorname{im} Q$ by the commutativity of the diagram ($\pi \circ \psi$ is a left inverse of $\psi^{-1} \circ \iota$). Thus we define $\Delta(f) = f_*$, and this is obviously[12] functorial as is seen by adding another "rung" to the diagram. $\square$

These two pieces assemble to give the following theorem of Swan [17].

Theorem 7.4 (Serre–Swan, analytic version)**.** *The functors* Γ *and* Δ *together determine an equivalence of categories* $\operatorname{ProjFin}(C_0(X)) \cong \operatorname{Vect}(X)$ *for any compact space* X.

Proof. First let $M \in \operatorname{ProjFin}(C_0(X))$. Then by definition, we have (using the notation of the previous propositions)

$$\begin{aligned}
\Gamma(\Delta(M)) &= \left\{s : X \to \Delta(M) : p \circ s = \mathrm{id}\right\} \\
&= \{s : X \to X \times \mathbf{C}^n : s(x) \in \{x\} \times \operatorname{im} P_x \quad \forall x \in X\} \\
&\cong \{s : X \to \mathbf{C}^n : s(x) \in \operatorname{im} P_x \quad \forall x \in X\} \\
&= \{s : X \to \mathbf{C}^n : s \in \operatorname{im} P\} \\
&= \operatorname{im} P.
\end{aligned}$$

In fact, $\operatorname{im} P$ is canonically isomorphic to M, due to the fact that the following diagram commutes (without the dotted arrow)

$$\begin{tikzcd}
M \arrow[r, "\phi^{-1}\circ\iota"] \arrow[rr, bend right, "\mathrm{id}"'] & C_0(X)^n \arrow[r, "\pi\circ\phi"] \arrow[rr, bend left, "P"] & M \arrow[r, "\phi^{-1}\circ\iota"] & C_0(X)^n \arrow[l, dotted, bend left, "\pi\circ\phi"]
\end{tikzcd} ,$$

combined with the fact that the dotted arrow is a left inverse of $\psi^{-1} \circ \iota$.

Now let $p : E \to X$ be a vector bundle, and let $f : E \oplus F \to X \times \mathbf{C}^n$ be an isomorphism of vector bundles. Then again from the definitions we have

$$\Delta(\Gamma(E)) = \{(x, v) \in X \times \mathbf{C}^n : v \in \operatorname{im} P_x\}.$$

When computing Δ of $\Gamma(E)$, the map $P : \Gamma(X \times \mathbf{C}^n) \to \Gamma(X \times \mathbf{C}^n)$ applies the isomorphism f, everywhere zeroes the F component of the direct sum, and then applies the isomorphism f^{-1}. Therefore, for fixed $x \in X$, the set $\operatorname{im} P_x$ consists of exactly the pairs $(x, v) \in X \times \mathbf{C}^n$ which come from evaluating the image of the composite

$$\Gamma(E) \longrightarrow \Gamma(E) \oplus \Gamma(F) \longrightarrow \Gamma(X \times \mathbf{C}^n)$$

at x. These are exactly the vectors in the E summand of $E \oplus F \to X \times \mathbf{C}^n$, and a canonical isomorphism realising this is immediately obtained from the isomorphism f.

That we have natural isomorphisms $\Delta\Gamma \to \mathbf{1}$ and $\Gamma\Delta \to \mathbf{1}$ follows immediately from the fact that we found isomorphisms in each case by only making canonical choices (subject to the caveat of the previous footnote), and hence this completes the proof. $\square$

7.2 Phase two: Reinterpreting C^*-algebraic K-theory

We now turn to building the second half of the bridge, from the side of the C^*-algebra theory.

Definition 7.5. Let A be a unital C^*-algebra. Then the monoid $\hat{V}(A)$ is defined to be the set of isomorphism classes of finitely generated projective modules over A (addition is just the direct sum of A-modules).

[12] Actually, there is a minor subtlety here; we need to ensure that for each finitely generated projective module M over $C_0(X)$ we choose a *specific*, and *fixed*, isomorphism $\phi : C_0(X)^n \to M \oplus M'$, before defining Δ on the morphisms using all of the same choices. If one is being principled from the category theoretic perspective, one should really define projectivity of a module not as a condition, but as the extra data of such an isomorphism.

Proposition 7.6. *If A is any unital C^*-algebra, then the monoids $V(A)$ and $\hat{V}(A)$ are isomorphic.*

Proof. We will explicitly write down an isomorphism of monoids. Fix $[p] \in V(A)$ for p a projection of dimension n. Then $C_0(X)^n$ splits as a direct sum

$$C_0(X)^n \cong pC_0(X)^n \oplus (1-p)C_0(X)^n.$$

As above, this proves that $pC_0(X)^n$ is a finitely generated projective $C_0(X)$-module. Performing this construction with equivalent projections obviously gives rise to isomorphic modules (by Lemma 3.1), and therefore we can safely set

$$\Phi([p]) = [pC_0(X)^n].$$

By the definition of the sum of classes of projections this map is also additive, and because it also preserves identities it is therefore is a monoid homomorphism. It is also obvious that Φ has trivial kernel, and is surjective because every finitely generated projective $C_0(X)$-module M gives rise to an isomorphism $\phi : C_0(X)^n \to M \oplus N$ (for some $n \in \mathbf{N}$ and other $C_0(X)$-module N). Indeed, composing with the canonical map $M \oplus N \to M$ and then inclusion $M \to N \oplus M \to C_0(X)^n$ (as we often did above) then gives a projection p for which $pC_0(X)^n \cong M$. □

We have actually just proved that the "C^*-algebraic K_0-group" of every C^*-algebra is the same as the algebraic K_0-group of its underlying unital ring. Curiously, the same is not true for the first K-group as well (see II.6.13 of [10])—the latter is the quotient of $\mathrm{GL}_\infty^+(A)$ by its commutator subgroup! The equivalence also disperses some hope for a multiplicative structure on the C^*-algebraic K_0 group, as there is no generally sensible tensor product for modules over noncommutative rings (discrepancies such as this one explain why topological K-theory continues to be studied in its own right).

Now let $\widetilde{\mathrm{Vect}}(X)$ denote the quotient of $\mathrm{Vect}(X)$ by taking isomorphism classes, so that (recalling the definition of topological K-theory) $K^0(X) = G\widetilde{\mathrm{Vect}}(X)$ for G the Grothendieck group completion functor.

Theorem 7.7. *The functors $K_0 C_0$ and K^0 are isomorphic on the compact spaces. Therefore, by the additional fact that $K_1 \cong K_0\Sigma$, the topological and C^*-algebraic K-theories of compact spaces are identical.*

Proof. We have done the hard work already.

Let X be a compact space, and so by Gelfand duality $C_0(X)$ is unital. By Theorem 7.4, we have an isomorphism $\widetilde{\mathrm{Vect}}(X) \cong \hat{V}(C_0(X))$. By Proposition 7.6 we also have $V(C_0(X)) \cong \hat{V}(C_0(X))$. Both of these isomorphisms pass to isomorphisms in the Grothendieck group completion, and therefore by Proposition 3.4 we have

$$K_0(C_0(X)) \cong GV(C_0(X)) \cong G\hat{V}(C_0(X)) \cong G\widetilde{\mathrm{Vect}}(X) = K^0(X),$$

as desired. Because of the equivalences involved, one easily checks that we actually have an isomorphism of functors. □

Thus, we get isomorphic K-groups! The contravariant-ness of the Gelfand duality functor "exactly cancels" the contravariant-ness of topological K-theory, and everything works out. Incidentally, Theorem 7.7 provides a proof of Bott periodicity in the topological case as well (given the proof in the more general C^*-algebraic setting, or perhaps a proof of the Connes "Thom" isomorphism!).

8 Greatest Hits™: Volume 2

The constructions and theorems of C^*-algebraic K-theory which we have detailed here are only just the beginning. In our brief survey we have stopped just short of the general theory of extensions of C^*-algebras, for which there is a rich array of results in the literature. Indeed, our K-theory functors K_0 and K_1 appear as a special case in Kasparov's KK-theory, along with the corresponding functors in K-homology which we have not even attempted to mention (and there is something of a pairing between these!). In the KK-theoretic setting, Cuntz has for example proved that our K_0 and K_1 functors are uniquely determined (on a suitable but large class of C^*-algebras, up to some "boundary conditions") by some of the properties which we have established; namely of being half exact, homotopy invariant, and stable.

There is also a general Universal Coefficient Theorem and Künneth Formula in KK-theory, which recover all of the corresponding K-theoretic results upon specialisation. The setting of KK-theory is a natural place for generalisations of index theorems to be proved; much work has been done generalising the Atiyah–Singer index theorem, for example. Because KK-theory theory can fail to have a generalised six-term sequence in all circumstances, the special case of E-theory has been pioneered by Connes and Higson. The E-theory can be thought of as a version of KK-theory for which one always has six-term exact sequences; as a consequence, more algebraic and category-theoretic methods typically apply. The additional structure present in E-theory also gives one reason to say that they are not just doing noncommutative topology, but actually *noncommutative stable homotopy theory*!

This concludes our account of some first steps into the theory of noncommutative vector bundles; the K-theory of C^*-algebras. The reader is wished the very best for their continuation of this journey, and it is hoped that the road ahead shall not prove too treacherous to navigate!—the paths carved by [19], [16], and (if one is particularly confident) [2], are all excellent ones to venture down next.

References

[1] M F Atiyah. "Algebraic topology and operators in Hilbert space". *Lectures in Modern Analysis and Applications I*. Springer, 1969, pp. 101–121.

[2] Bruce Blackadar. *K-theory for operator algebras*. Vol. 5. Cambridge University Press, 1998.

[3] John B Conway. *A course in functional analysis*. Vol. 96. Springer Science & Business Media, 2013.

[4] Joachim Cuntz. "K-theory for certain C*-algebras". *Annals of Mathematics* (1981), pp. 181–197.

[5] George A Elliott. "On the classification of inductive limits of sequences of semisimple finite-dimensional algebras". *Journal of Algebra* 38.1 (1976), pp. 29–44.

[6] George A Elliott, Toshikazu Natsume, and Ryszard Nest. "The Heisenberg group and K-theory". *K-theory* 7.5 (1993), pp. 409–428.

[7] Peter A Fillmore. *A user's guide to operator algebras*. Vol. 14. Wiley-Interscience, 1996.

[8] Philip Green. "C-algebras of transformation groups with smooth orbit space". *Pacific Journal of Mathematics* 72.1 (1977), pp. 71–97.

[9] Allen Hatcher. "Vector Bundles and K-Theory" (2018). Online notes.

[10] M Karoubi. "K-theory, an introduction Spriger-Verlag". *Berlin, Heidelberg, New York* (1978).

[11] Gerald J Murphy. *C*-algebras and operator theory*. Academic press, 2014.

[12] N Christopher Phillips. "The Toeplitz operator proof of noncommutative Bott periodicity". *Journal of the Australian Mathematical Society* 53.2 (1992), pp. 229–251.

[13] Mihai Pimsner and Dan Voiculescu. "Exact sequences for K-groups and Ext-groups of certain cross-product C*-algebras". *Journal of Operator Theory* (1980), pp. 93–118.

[14] Ian Putnam. "Lecture Notes on C-algebras" (2016).

[15] Mikael Rørdam. *Structure and classification of C*-algebras*. Citeseer, 2006.

[16] Mikael Rørdam, Flemming Larsen, Flemming Larsen, and N Laustsen. *An Introduction to K-theory for C*-algebras*. Vol. 49. Cambridge University Press, 2000.

[17] Richard G Swan. "Vector bundles and projective modules". *Transactions of the American Mathematical Society* 105.2 (1962), pp. 264–277.

[18] Joseph L Taylor. *Banach algebras and topology*. University of Utah, 1973.

[19] N. E. Wegge-Olsen. "K-theory and C-star-algebras: A friendly approach" (1993).