

Clifford Algebras and Bott Periodicity

Dexter Chua

1	Clifford Algebras and Representations	1
2	Fredholm Operators and K-theory	7
3	Proving Bott Periodicity I	13
4	Proving Bott Periodicity II	19
5	Coda: The ABS Orientation	23

In [2], Atiyah, Bott and Shapiro calculated certain groups A_k associated to real Clifford algebra representations, and observed that they were isomorphic to $KO^{-k}(\ast)$. The same can be said for complex K -theory, but a sequence of groups alternating between \mathbb{Z} and 0 is less impressive. In their paper, they constructed a map $A_k \rightarrow KO^{-k}(\ast)$, and *using their knowledge of $KO^{-k}(\ast)$* , they showed this is in fact an isomorphism. It wasn't a particularly exciting proof, since both sides have a ring structure that the map respects, and so they only had to check the map does the right thing to the handful of generators.

Later, in [3], Atiyah and Singer found a *good* reason why they had to be isomorphic. Roughly, the idea is that KO^{-k} is represented by some group of Clifford algebra representation homomorphisms, and it is not too difficult to show that π_0 of this space is isomorphic to A_k . The goal of these notes is to work through these results and conclude the Bott periodicity theorem.

If the reader finds the details lacking, they may refer to [3] (which does not provide a very detailed treatment either, but the intersection of details missed out should be minimal).

In these notes, “homotopy equivalence” will mean “weak homotopy equivalence”. However, it follows from results of Milnor [7] that the spaces involved are CW complexes, so it doesn't actually matter.

1 CLIFFORD ALGEBRAS AND REPRESENTATIONS

We begin by defining Clifford algebras.

Definition. Let V be a vector space over \mathbb{R} and $Q : V \rightarrow \mathbb{R}$ a quadratic form. The *Clifford algebra* $C(V; Q)$ is defined by the quotient

$$C(V; Q) = \frac{\bigoplus_{k=0}^{\infty} V^{\otimes k}}{(x \otimes x - Q(x))}.$$

Example. Let Q_k be the standard negative-definite quadratic form on \mathbb{R}^k , so that

$$Q(x_1, \dots, x_k) = -\sum x_i^2.$$

We define $C_k = C(\mathbb{R}^k; Q_k)$ and $C'_k = C(\mathbb{R}^k; -Q_k)$. These are the Clifford algebras we are interested in.

Observe that our Clifford algebras are non-commutative algebras. For many purposes, it is convenient to treat them as \mathbb{Z}_2 -graded algebras*, with V being in degree 1. In particular, it gives the correct tensor product.

Recall that if A, B are \mathbb{Z}_2 -graded algebras, then the *graded tensor product* $A \hat{\otimes} B$ is defined by the algebra with underlying vector space $A \otimes B$ and multiplication

$$(x \otimes y) \cdot (v \otimes w) = (-1)^{|y||v|} xv \otimes yw.$$

Lemma. If $V = V_1 \oplus V_2$ and $Q(x + y) = Q_1(x) + Q_2(y)$ for all $x \in V_1$ and $y \in V_2$, then

$$C(V; Q) = C(V_1; Q_1) \hat{\otimes} C(V_2; Q_2). \quad \square$$

Here the graded product is needed because we want

$$Q(a + b) = (a \otimes 1 + 1 \otimes b)^2 = a^2 \otimes 1 + 1 \otimes b^2 = Q_1(a) + Q_2(b).$$

From this lemma, we immediately deduce that

Corollary. We have

$$C_k = \underbrace{C_1 \hat{\otimes} \cdots \hat{\otimes} C_1}_{k \text{ times}}.$$

Explicitly, C_k is generated by e_1, \dots, e_k subject to the relations

$$e_i^2 = -1, \quad e_i e_j + e_j e_i = 0.$$

In particular, $\dim C_k = 2^k$, and the dimension in each degree is 2^{k-1} . Similar statements can be made for C'_k .

Ultimately, we wish to understand the representations of C_k . To do so, it is convenient to explicitly identify the Clifford algebras C_k . While graded tensor products are the “morally correct” things to consider, they are not computationally helpful. Instead, we want to rephrase everything in terms of ordinary tensor products. We have the following result:

* Here \mathbb{Z}_2 always means the integers mod 2; we will make no use of the 2-adics in these notes.

Proposition. We have

$$\begin{aligned} C_{k+2} &\cong C'_k \otimes_{\mathbb{R}} C_2 \\ C'_{k+2} &\cong C_k \otimes_{\mathbb{R}} C'_2 \end{aligned}$$

Proof. Let e'_1, \dots, e'_k be the generators of C'_k . Consider the map of algebras $C'_{k+2} \rightarrow C_k \otimes_{\mathbb{R}} C'_2$ defined by

$$\begin{aligned} e'_1 &\mapsto 1 \otimes e'_1 \\ e'_2 &\mapsto 1 \otimes e'_2 \\ e'_i &\mapsto e_{i-2} \otimes e'_1 e'_2 \text{ for } i \geq 3. \end{aligned}$$

One checks that this is well-defined, and surjects onto the generators of $C_k \otimes_{\mathbb{R}} C'_2$. So this map is surjective, and since both sides have the same dimension, this is an isomorphism. The other isomorphism is obtained by swapping e'_i with e_i . \square

Using this, we are now in a position to explicitly describe the Clifford algebras inductively. We first identify the Clifford algebras in low dimensions explicitly.

Notation. For any algebra R , we write $R(n)$ for the algebra of $n \times n$ matrices with entries in R .

Proposition.

$$\begin{aligned} C_1 &\cong \mathbb{C} & C'_1 &\cong \mathbb{R} \oplus \mathbb{R} \\ C_2 &\cong \mathbb{H} & C'_2 &\cong \mathbb{R}(2). \end{aligned}$$

Proof. The descriptions of C_1 and C_2 are essentially by definition. The identification of C'_1 is obtained by picking basis vectors $\frac{1 \pm e'_1}{2}$. The identification of C'_2 is obtained by setting

$$e'_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad e'_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix}. \quad \square$$

We will also need the following observations:

Lemma. We have

$$\mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} \cong \mathbb{C} \oplus \mathbb{C}$$

and

$$\begin{aligned} \mathbb{H} \otimes_{\mathbb{R}} \mathbb{C} &\cong \text{End}_{\mathbb{C}}(\mathbb{H}) \cong \mathbb{C}(2) & \mathbb{H} \otimes_{\mathbb{R}} \mathbb{H} &\cong \text{End}_{\mathbb{R}}(\mathbb{H}) \cong \mathbb{R}(4) \\ p \otimes z &\mapsto (q \mapsto zq\bar{p}) & p \otimes r &\mapsto (q \mapsto pq\bar{r}) \end{aligned} \quad \square$$

We can then write down the table of Clifford algebras

k	C_k	C'_k
1	\mathbb{C}	$\mathbb{R} \oplus \mathbb{R}$
2	\mathbb{H}	$\mathbb{R}(2)$
3	$\mathbb{H} \oplus \mathbb{H}$	$\mathbb{C}(2)$
4	$\mathbb{H}(2)$	$\mathbb{H}(2)$
5	$\mathbb{C}(4)$	$\mathbb{H}(2) \oplus \mathbb{H}(2)$
6	$\mathbb{R}(8)$	$\mathbb{H}(4)$
7	$\mathbb{R}(8) \oplus \mathbb{R}(8)$	$\mathbb{C}(8)$
8	$\mathbb{R}(16)$	$\mathbb{R}(16)$

Moreover, $C_{k+8} = C_k \otimes_{\mathbb{R}} \mathbb{R}(16)$.

Note that all our Clifford algebras are semi-simple, and most of them are in fact simple. One might think the representations are therefore quite uninteresting, but that would be wrong.

Definition. Let A be an \mathbb{Z}_2 -graded algebra. We define $M(A)$ to be the Grothendieck group of graded representations of A , and $N(A)$ be the Grothendieck group of ungraded representations of A . We also write A^0 for the degree zero part of A .

Lemma. We have an isomorphism of groups

$$\begin{aligned}
M(C_k) &\cong N(C_k^0) \\
M^0 \oplus M^1 &\mapsto M^0 \\
C_k \otimes_{C_k^0} M &\leftarrow M.
\end{aligned}$$

□

Moreover, we can also easily identify C_k^0 :

Lemma. We have an isomorphism of algebras

$$\begin{aligned}
\phi : C_{k-1} &\cong C_k^0 \\
e_i &\mapsto e_i e_k.
\end{aligned}$$

This immediately allows us to write down the Grothendieck group of representations of C_k , as well as the dimensions of the representations.

k	C_k	$M(C_k)$	$\dim_{\mathbb{R}}(M^0)$
1	\mathbb{C}	\mathbb{Z}	1
2	\mathbb{H}	\mathbb{Z}	2
3	$\mathbb{H} \oplus \mathbb{H}$	\mathbb{Z}	4
4	$\mathbb{H}(2)$	$\mathbb{Z} \oplus \mathbb{Z}$	4
5	$\mathbb{C}(4)$	\mathbb{Z}	8
6	$\mathbb{R}(8)$	\mathbb{Z}	8
7	$\mathbb{R}(8) \oplus \mathbb{R}(8)$	\mathbb{Z}	8
8	$\mathbb{R}(16)$	$\mathbb{Z} \oplus \mathbb{Z}$	8

Now there is always an inclusion $i : C_k \rightarrow C_{k+1}$, which induces a map $i^* : M(C_{k+1}) \rightarrow M(C_k)$.

Definition. We define $A_k = M(C_k)/i^*M(C_{k+1})$.

When $k \neq 4n$, computing A_k is trivial, since there is only one irreducible representation. Hence we know what the image of i^* is simply by looking at the dimension of the pulled back representation. We claim that $A_{4n} = \mathbb{Z}$ for all n .

One can prove this using some more direct arguments, but it is convenient for us to prove this by understanding the module structure better. Given a graded C_k module $M = M^0 \oplus M^1$, we can obtain a new C_k module $M^* = M^1 \oplus M^0$. If $k \neq 4n$, then for purely dimensional reasons, we must have $M \cong M^*$.

Proposition. Let $k = 4n$, and let M and N be the irreducible C_k -modules. Then $M = N^*$ and $N = M^*$.

Proof. We wish to reduce this problem to a problem of ungraded modules.

Note that in general, for an algebra A and an automorphism α , if M is an A -module, then we can construct another A -module M^α whose underlying set is the same, but the action is given by

$$x \cdot m = \alpha(x)m.$$

With this in mind, consider the automorphisms

$$\begin{aligned} \alpha : C_k &\mapsto C_k & \beta : C_{k-1} &\mapsto C_{k-1} \\ x &\mapsto e_k x e_k^{-1} & x &\mapsto (-1)^{|x|} x. \end{aligned}$$

Check that these are related by the commutative diagram

$$\begin{array}{ccc} C_{k-1} & \xrightarrow{\phi} & C_k \\ \downarrow \beta & & \downarrow \alpha \\ C_{k-1} & \xrightarrow{\phi} & C_k. \end{array}$$

Since multiplication by e_k gives an isomorphism of vector spaces $M^0 \cong M^1$, we know that $M^* \cong M^\alpha$. Hence under the correspondence $M(C_k) \cong N(C_{k-1})$, the operation $M \mapsto M^*$ corresponds to $M^0 \mapsto (M^0)^\beta$.

Thus, it suffices to show that β swaps the two irreducible ungraded representations of C_{k-1} . Since $C_{k-1} \cong C_{k-2} \oplus C_{k-2}$, we want to find the central idempotents that project to the two ideals, and show that β swaps the two. We can just write this down.

Set

$$w = e_1 e_2 \cdots e_{4n-1}.$$

Then 1 and w are in the center, and $w^2 = 1$. So the desired idempotents are $\frac{1 \pm w}{2}$, and $\beta(w) = -w$. \square

Corollary. $A_{4n} \cong \mathbb{Z}$.

Proof. Let x, y be the two irreducible modules in $M(C_{4n})$, and z the irreducible module in $M(C_{4n+1})$. Then $z^* = z$, so $(i^*z)^* = i^*z^* = i^*z$. So we must have $i(z) = x + y$. \square

This w will be useful later on. Observe that, when restricted to degree 0 components, w acts as 1 on one of the irreducible representations and -1 on the other. So if we are given some representation of C_{4n} , if we want to know how many of each irrep we have got, we simply have to take the degree 0 component and count the eigenvectors of $\phi(w)$.

We can now tabulate

k	C_k	$M(C_k)$	$\dim_{\mathbb{R}}(M^0)$	A_k
1	\mathbb{C}	\mathbb{Z}	1	\mathbb{Z}_2
2	\mathbb{H}	\mathbb{Z}	2	\mathbb{Z}_2
3	$\mathbb{H} \oplus \mathbb{H}$	\mathbb{Z}	4	0
4	$\mathbb{H}(2)$	$\mathbb{Z} \oplus \mathbb{Z}$	4	\mathbb{Z}
5	$\mathbb{C}(4)$	\mathbb{Z}	8	0
6	$\mathbb{R}(8)$	\mathbb{Z}	8	0
7	$\mathbb{R}(8) \oplus \mathbb{R}(8)$	\mathbb{Z}	8	0
8	$\mathbb{R}(16)$	$\mathbb{Z} \oplus \mathbb{Z}$	8	\mathbb{Z}

Note that we have $C_{k+8} = C_k \otimes_{\mathbb{R}} \mathbb{R}(16)$, so we see that A_k is 8-periodic. In fact, more is true.

Observe that if $M \in M(C_k)$ and $N \in M(C_m)$, then since $C_{k+m} \cong C_k \hat{\otimes} C_m$, we can view $M \hat{\otimes} N$ as an C_{k+m} module. Then this turns $M_* = M(C_*)$ into a ring, and since the image of $i_* : M_* \rightarrow M_*$ is an ideal ($M \cdot i^*N = i^*(M \cdot N)$), this implies A_* is a ring as well.

Theorem (Bott Periodicity). Let λ be the irreducible module of C_8 . Then multiplication by λ induces an isomorphism $M(C_k) \rightarrow M(C_{k+8})$, hence an isomorphism $A_k \cong A_{k+8}$.

Proof. This is trivial except in the case $k = 4n$, by dimension counting. If $M(C_{4n})$ is generated by the irreps x, y , then $\lambda \cdot x$ is one of the irreps in $M(C_{4n+8})$, and then

$$\lambda \cdot y = \lambda \cdot x^* = (\lambda \cdot x)^*$$

is the other irrep. \square

This is interesting, since we observe that we have $A_k \cong KO^{-k}(*)$, and after a bit of checking, this is in fact an isomorphism of rings! Of course, it would be terrible to just observe that we know what both sides are, and that they are the same. We should prove that they are the same without knowledge of $KO^{-k}(*)$, and this would give us both Bott periodicity, as well as explicit computations of $KO^{-k}(*)$.

2 FREDHOLM OPERATORS AND K -THEORY

It is difficult to directly connect Clifford algebras to K -theory. Instead, we have to find an alternative model for K -theory. One such model is given by the space of Fredholm operators.

2.1 BASIC PROPERTIES OF FREDHOLM OPERATORS

Definition. Let H be a Hilbert space. A bounded linear operator $T : H \rightarrow H$ is *Fredholm* if $\ker T$ and $\operatorname{coker} T$ are both finite-dimensional.

Exercise. T induces an isomorphism between $\ker T^\perp$ and $\operatorname{im} T$. In particular, $\operatorname{im} T$ is closed and (if $\ker T \neq 0$) 0 is an isolated point in the spectrum of T .

Definition. The *index* of a Fredholm operator $T : H \rightarrow H'$ is

$$\operatorname{idx} T = \dim \ker T - \dim \operatorname{coker} T.$$

Example. If $H = H'$ is finite-dimensional, then all operators are Fredholm with index 0 , by the rank-nullity theorem.

Example. Let $\operatorname{shift} : \ell^2 \rightarrow \ell^2$ be the map

$$\operatorname{shift}(x_1, x_2, x_3, \dots) = (x_2, x_3, x_4, \dots).$$

Then this has index 1 . The adjoint of shift , written $\operatorname{shift}^{-1}$, is given by

$$\operatorname{shift}^{-1}(x_1, x_2, x_3, \dots) = (0, x_1, x_2, \dots).$$

which has index -1 .

One can similarly see that $\operatorname{idx}(\operatorname{shift}^k) = k$ for all $k \in \mathbb{Z}$. So in particular, every integer is the index of some operator.

The fact that the index of shift^k is the negative of the index of its adjoint is not a coincidence.

Exercise. Let $T : H \rightarrow H$ be Fredholm. Then so is T^* , and $\operatorname{idx} T^* = -\operatorname{idx} T$.

We shall prove some basic properties of Fredholm operators. They all follow from the following result:

Lemma. Consider the commutative diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & H & \longrightarrow & H' & \longrightarrow & H'' & \longrightarrow & 0 \\ & & \downarrow T & & \downarrow S & & \downarrow R & & \\ 0 & \longrightarrow & H & \longrightarrow & H' & \longrightarrow & H'' & \longrightarrow & 0 \end{array}$$

If the rows are short exact, and any two of T , S and R are Fredholm, then so is the third, and

$$\operatorname{idx} S = \operatorname{idx} T + \operatorname{idx} R.$$

Proof. Apply the snake lemma to obtain a long exact sequence

$$0 \longrightarrow \ker T \longrightarrow \ker S \longrightarrow \ker R \longrightarrow \operatorname{coker} T \longrightarrow \operatorname{coker} S \longrightarrow \operatorname{coker} R \longrightarrow 0. \quad \square$$

Corollary. If $T : H \rightarrow H$ and $S : H' \rightarrow H'$ are Fredholm, then so is $T \oplus S : H \oplus H' \rightarrow H \oplus H'$, and we have

$$\operatorname{idx}(T \oplus S) = \operatorname{idx} T + \operatorname{idx} S. \quad \square$$

Of course, we could have proved this directly.

Corollary. Let $T : H \rightarrow H'$ and $S : H' \rightarrow H''$ be Fredholm. Then ST is Fredholm, and

$$\operatorname{idx}(ST) = \operatorname{idx} T + \operatorname{idx} S.$$

Proof.

$$\begin{array}{ccccccccc} 0 & \longrightarrow & H & \xrightarrow{(I,T)} & H \oplus H' & \xrightarrow{T-I} & H' & \longrightarrow & 0 \\ & & \downarrow T & & \downarrow ST \oplus I & & \downarrow S & & \\ 0 & \longrightarrow & H' & \xrightarrow{(S,I)} & H'' \oplus H' & \xrightarrow{I-S} & H'' & \longrightarrow & 0 \end{array} \quad \square$$

Corollary. Let $I : H \rightarrow H$ be the identity and $K : H \rightarrow H$ be a compact operator. Then $I + K$ is Fredholm and has index 0.

Proof. First consider the case where K has finite-dimensional image. Then consider the short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \operatorname{im} K & \longrightarrow & H & \longrightarrow & \operatorname{coker} K & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & \operatorname{im} K & \longrightarrow & H & \longrightarrow & \operatorname{coker} K & \longrightarrow & 0 \end{array}$$

where the vertical maps are all restrictions of $I + K$. One sees that the maps are well-defined.

Now note that the left-hand map is one between finite-dimensional vector spaces, so is Fredholm with index 0. On the other hand, the right-hand map is in fact the identity, since we quotiented out by the image of K , and is in particular Fredholm with index 0. So we are done.

For general compact K , we can approximate K arbitrarily well by finite-rank operators[†]. So pick K' such that K' has finite rank and $\|K - K'\| < 1$. Thus $I + (K - K')$ is invertible, and we can write

$$I + K = (I + (K - K'))(I + (I + K - K')^{-1}K').$$

[†] The image of the unit ball is totally bounded, so there is a finite-dimensional subspace that approximates it well. Then compose K with the projection to this subspace. Note that this is false for general Banach spaces [5].

The first term is invertible, hence Fredholm with index 0. The second is I plus something of finite rank, so is Fredholm with index 0. So the general result follows. \square

Theorem. Let T be a bounded linear operator. Then T is Fredholm iff there exists $S : H \rightarrow H$ such that $TS - I$ and $ST - I$ are compact.

Proof.

(\Rightarrow) Note that T gives an isomorphism between $(\ker T)^\perp$ and $\operatorname{im} T$. So we can make sense of $(T|_{\operatorname{im} T})^{-1} : \operatorname{im} T \rightarrow (\ker T)^\perp$. Then we can take S to be

$$H \cong \operatorname{im} T \oplus (\operatorname{im} T)^\perp \xrightarrow{(T|_{\operatorname{im} T})^{-1} \oplus 0} (\ker T)^\perp \oplus \ker T \cong H .$$

Then $ST - I$ and $TS - I$ have finite rank, as they vanish on $(\ker T)^\perp$ and $\operatorname{im} T$ respectively.

(\Leftarrow) Using the previous corollary, we know ST and TS are Fredholm. Since ST is Fredholm, $\ker T$ is finite-dimensional. Since TS is Fredholm, $\operatorname{coker} T$ is finite-dimensional. So we are done. \square

Corollary. If T is Fredholm and K is compact, then $T + K$ is also Fredholm.

Proof. We can use the same S , since the algebra of compact operators is a two-sided ideal in $\operatorname{End}(H)$. \square

2.2 KUIPER'S THEOREM

We shall take a short break from Fredholm operators, and prove the following extremely important theorem:

Theorem (Kuiper, 1964, [6]). Let H be a separable infinite-dimensional (real or complex) Hilbert space. Then for any compact space X , we have $[X, \operatorname{GL}(H)] = 0$. In particular, $\operatorname{GL}(H)$ is weakly contractible.

To prove this, we first establish a useful “move”.

Lemma. Let X be a fixed space, and $S, T : X \rightarrow \operatorname{GL}(H)$ continuous functions. Then the maps

$$\begin{pmatrix} ST & 0 \\ 0 & I \end{pmatrix}, \begin{pmatrix} T & 0 \\ 0 & S \end{pmatrix} : X \rightarrow \operatorname{GL}(H \oplus H)$$

are homotopic as maps $X \rightarrow \operatorname{GL}(H \oplus H)$.

The mental picture we should have in mind is that we are allowed to “rotate” S from the top-left corner to the bottom-right corner. The proof, unsurprisingly, is a literal rotation.

Lemma. Suppose $H = H_0 \oplus H_1$, $f(x) : H \rightarrow H$, and $g(x) : H_0 \rightarrow H$ are such that $f(x)(H_0) \perp g(x)(H_0)$ for all x and $g(x)$ is an isomorphism onto its image. Then f is homotopic to a map \tilde{f} such that $\tilde{f}(x)|_{H_0} = g(x)$.

Proof. For each x , decompose $H = f(x)H_0 \oplus g(x)H_0 \oplus H_x$ for some H_x . Then the map

$$\begin{aligned} \varphi_x : f(x)H_0 \oplus g(x)H_0 \oplus H_x &\rightarrow f(x)H_0 \oplus g(x)H_0 \oplus H_x \\ f(x)a \oplus g(x)b \oplus c &\mapsto f(x)b \oplus (-g(x)a) \oplus c. \end{aligned}$$

is homotopic to the identity. Indeed, we can achieve this if $f(x)$ and $g(x)$ weren't there by a rotation, and we just have to conjugate that homotopy by $f(x) \oplus g(x) \oplus I_{H_x}$. Then take $\tilde{f}(x) = \varphi_x^{-1}f(x)$. \square

Thus, if we can decompose $H = H_0 \oplus H_1 \oplus H_2$ such that $f(x)(H_0) \perp H_2$, then the theorem follows by performing two swaps — first homotope f so that $f(x)(H_0) = H_2$, and then swap the image back to the identity on H_0 .

Proposition. We may assume the image of f is contained in a finite-dimensional subspace of $\text{End}(H)$.

Proof. This is a special case of the following more general fact: Let V be a normed vector space and $U \subseteq V$ an open set. Suppose x_1, \dots, x_n are points in U and $\varepsilon_i > 0$ are such that $B(x_i, 3\varepsilon_i) \subseteq U$. Then there is a deformation retract of $U_* = \bigcup_{i=1}^n B(x_i, \varepsilon_i)$ into the simplicial complex spanned by x_1, \dots, x_n . Indeed, pick a partition of unity

$$\begin{aligned} \psi_j(x) &= \max(\varepsilon_j - \|x - x_j\|, 0), \\ \phi_j(z) &= \frac{\psi_j(z)}{\sum_{k=1}^n \psi_k(z)}. \end{aligned}$$

Then we can use the deformation

$$g_t(x) = (1-t)x + t \sum_{j=1}^n \phi_j(x)x_j,$$

and the hypothesis ensures this remains in $\text{GL}(H)$. \square

Proposition. There exists an orthogonal decomposition $H = H_1 \oplus H_2 \oplus H_3$ such that $f(x)(H_1) \perp H_3$ for all x , and H_1 and H_3 are both infinite-dimensional.

Proof. Pick a vector a_1 , put it in H_1 . The span of all $f(x)(a_1)$ is finite-dimensional, so we can put finitely many vectors in H_2 so that $f(x)(a_1) \in H_1 \oplus H_2$ for all x . Pick any vector orthogonal to the vectors we have chosen and put it in H_3 .

Next, pick a_2 that is orthogonal to everything chosen so far, and also so that $f(x)(a_2)$ is orthogonal to everything chosen so far for all x . This is possible since these conditions only exclude a finite-dimensional subspace. Keep going on countably many times, and afterwards if there is anything left, put it in H_2 . \square

2.3 THE ATIYAH–JÄNICH THEOREM

Equipped with Kuiper’s theorem, we can finally prove what we wanted to.

Theorem (Atiyah, Jänich). Let H be an infinite-dimensional real Hilbert space, and \mathcal{F} the space of all Fredholm operators on H . Then for every compact space X , there is a natural isomorphism

$$\text{idx} : [X, \mathcal{F}] \rightarrow KO(X).$$

The same holds in the complex case with KO replaced by KU .

For $f : X \rightarrow \mathcal{F}$, we would like to define $\text{idx}(f)$ by setting the fiber at each $x \in X$ to be the formal difference between $\ker f(x)$ and $\text{coker } f(x)$. In general, this does not give a vector bundle, but it is a general fact (which we shall not prove) that we can homotope f so that it does, and the difference $[\ker f] - [\text{coker } f]$ is independent of the choice of homotopy. In particular, it depends only on the homotopy class of f .

Proof (cf. [1]). First observe the following trivial lemma:

Lemma. Let $f : M \rightarrow G$ be a surjective homomorphism from a monoid to an abelian group with trivial kernel, i.e. $f^{-1}(\{0\}) = \{0\}$. Then f is in fact injective, hence an isomorphism. \square

So it suffices to show our map is surjective and has trivial kernel.

- To show the kernel is trivial, by Kuiper’s theorem, it suffices to show that every map $T : X \rightarrow \mathcal{F}$ with index 0 is homotopic to a map with image in $\text{GL}(H)$.

If $\text{idx } T = 0$, we can homotope T so that $[\ker T] = [\text{coker } T]$. Thus, by definition of the K -groups, we can find some large m such that

$$\ker T \oplus (X \times \mathbb{R}^m) \cong \text{coker } T \oplus (X \times \mathbb{R}^m).$$

Thus, by homotoping T to vanish further on a trivial subbundle isomorphic to $X \times \mathbb{R}^m$, we may assume that there is an isomorphism $\phi : \ker T \rightarrow \text{coker } T \cong (\text{im } T)^\perp$. Then we have a homotopy

$$(t\phi, T|_{(\ker P_n T)^\perp}) : \ker P_n T \oplus (\ker P_n T)^\perp \rightarrow H$$

going from T to an isomorphism. Then we are done.

- To show surjectivity, since every element in $K(X)$ is of the form $[E] - [X \times \mathbb{R}^k]$, it suffices to show that we can produce maps with index $[E]$ and $-[X \times \mathbb{R}^k]$.

The case of $-[X \times \mathbb{R}^k]$ is easy — take the map T to constantly be shift^{-k} .

To construct the one with $[E]$, first pick some bundle F such that $E \oplus F \cong X \times \mathbb{R}^N$ for some N .

Now consider the bundle $(E \oplus F) \oplus (E \oplus F) \oplus (E \oplus F) \oplus \dots$, which is isomorphic to $X \times H$. Then take the Fredholm operator to be

$$(e_1, f_1, e_2, f_2, e_3, f_3, \dots) \mapsto (e_2, f_1, e_3, f_2, e_4, f_3, \dots).$$

3 PROVING BOTT PERIODICITY I

3.1 THE LAST PART OF THE PROOF

We can now return to Clifford algebras and Bott periodicity. The idea of Atiyah and Singer [3] was to extend the Atiyah–Jänich theorem to provide explicit descriptions of the loop spaces $\Omega^k \mathcal{F}(H)$. It turns out we can interpret this space (up to homotopy equivalence, of course) as Fredholm homomorphisms of Clifford algebra representations.

Note that C_k naturally comes with an involution defined by $e_i^* = -e_i$, and we shall require our representations to respect this. Let $H = H^0 \oplus H^1$ be a \mathbb{Z}_2 -graded Hilbert space with a simultaneous graded action of all C_k , i.e. there are bounded linear operators J_1, J_2, \dots of degree 1 such that

$$J_i J_j + J_j J_i = 0, \quad J_i^2 = -1, \quad J_i^* = -J_i,$$

for all $i \neq j$. We will focus on skew-Hermitian operators. Since we have a graded Hilbert space, we require our skew-Hermitian operators to have degree 1. So that B takes the form

$$\begin{pmatrix} 0 & -T^* \\ T & 0 \end{pmatrix}.$$

Definition. We let

$$\begin{aligned} \hat{\mathcal{F}}(H) &= \{B \in \mathcal{F}(H) : B \text{ is skew-Hermitian}\} \\ \mathcal{F}^k(H) &= \{B \in \hat{\mathcal{F}} : BJ_i + J_i B = 0 \text{ for } i = 1, \dots, k\} \end{aligned}$$

After picking an arbitrary isomorphism $H^0 \cong H^1$ of vector spaces, sending $B \in \hat{\mathcal{F}}(H) = \mathcal{F}^0(H)$ to $T \in \mathcal{F}(H^0)$ as above gives a homeomorphism $\mathcal{F}(H^0) \rightarrow \mathcal{F}^0(H)$. Hence \mathcal{F}^0 represents $KO = KO^0$.

It turns out these $\mathcal{F}^k(H)$ are not what we want. To understand this, suppose $B \in \mathcal{F}^k(H)$ is unitary (which we can achieve by restricting to $\ker B^\perp$ and then rescaling), so that $B^2 = -B^*B = -1$. Then B “acts as” J_{k+1} on H , and this gives H a new C_{k+1} action. If $k \not\equiv -1 \pmod{4}$, then there is a unique irreducible C_{k+1} module, hence the structure of H as a C_{k+1} -module is completely determined. However, if $k \equiv -1 \pmod{4}$, then we want to ensure H has infinitely many copies of each irreducible module, for it to be well-behaved.

Recall that to count how many copies of each each irrep we’ve got, we are supposed to count the eigenvalues of $w = e_1 e_2 \cdots e_k$ after turning H^0 to a C_k -module. The process of turning H^0 into a C_k -module involves the inclusion $C_k \hookrightarrow C_{k+1}$, which sends e_i to $e_i e_{k+1}$. In our case, $e_{k+1} = B$, and $e_i = J_i$ for $i = 1, \dots, k$. Thus, we are counting the eigenvalues of

$$w = (J_1 B)(J_2 B)(J_3 B) \cdots (J_k B).$$

Since $k \equiv -1 \pmod{4}$ and $B^2 = -1$, this is equal to

$$w = J_1 J_2 \cdots J_k B.$$

We then want to require that the eigenvalues ± 1 to both have infinite multiplicity. In general, for an arbitrary B , it is not unitary, or even injective. The desired “niceness” property is now that when we restrict B to the $\ker B^\perp$, and then rescale B so that it is unitary, the multiplicity of ± 1 in w are both infinite. Equivalently, without doing the restriction and rescaling business, we want $J_1 J_2 \cdots J_k B$ to have infinitely many eigenvalues of each sign, counted with multiplicity.

Definition. We say a self-adjoint operator is *essentially definite* if all but finitely many of its eigenvalues are of the same sign.

Definition. If $k \not\equiv -1 \pmod{4}$, we define $\mathcal{F}_*^k(H) = \mathcal{F}^k(H)$. Otherwise, we define

$$\mathcal{F}_*^k(H) = \{B \in \mathcal{F}^k(H) : J_1 J_2 \cdots J_k B|_{H^0} \text{ is not essentially definite}\}.$$

Observe that these form a component of $\mathcal{F}^k(H)$.

The main theorem of Atiyah and Singer is the following:

Theorem A. For $k \geq 1$, there is a homotopy equivalence

$$\mathcal{F}_*^k(H) \rightarrow \Omega(\mathcal{F}_*^{k-1}(H)).$$

Thus, combined with the Atiyah–Jänich theorem, we know that $\mathcal{F}_*^k(\mathcal{H})$ represents KO^{-k} .

We first see how this proves that $KO^{-k}(\ast) = A_k$, and in particular implies Bott periodicity.

Theorem (Bott periodicity). There is a natural[‡] homotopy equivalence $\mathcal{F}_*^{k+8} \simeq \mathcal{F}_*^k$.

Proof. Let M be an irreducible C_8 -module. Then there is an isomorphism

$$\begin{aligned} \mathcal{F}_*^k(H) &\cong \mathcal{F}_*^{k+8}(H \hat{\otimes} M) \\ A &\mapsto A \otimes I, \end{aligned}$$

and there is a contractible space of C_{k+8} -module isomorphisms $H \cong H \hat{\otimes} M$. \square

Theorem (Computation of $KO^{-k}(\ast)$). The map $\text{id}_X : \mathcal{F}_*^k(H) \rightarrow A_k$ defined by $A \mapsto [\ker A]$ is continuous, and induces a bijection $\pi_0(\mathcal{F}_*^k) \rightarrow A_k$.

This in fact gives us an isomorphism of *rings*.

[‡] More precisely, there is a natural contractible space of such homotopy equivalences.

Proof. For convenience, we drop the (H) in $F_*^k(H)$ when it is clear.

We leave the $k = 0$ case for the reader. Note that $C_0 = \mathbb{R}$ is concentrated in degree 0, and it has two irreducible graded representations given by $0 \oplus \mathbb{R}$ and $\mathbb{R} \oplus 0$.

For $k > 0$, we first show that the kernel map is continuous, so that it factors through $\pi_0(\mathcal{F}_*^k)$. Let $B \in \mathcal{F}_*^k$. Note that B^2 is self-adjoint and negative, since

$$\langle B^2 x, x \rangle = -\langle Bx, Bx \rangle \leq 0.$$

Hence we know that $\sigma(B^2) \subseteq (-\infty, 0]$. In fact, since B is Fredholm, we know 0 is an isolated point in the spectrum, and by scaling B , we may assume

$$\sigma(B^2) \subseteq (-\infty, -1] \cup \{0\}.$$

Since the spectrum depends continuously on B^2 , we can pick a small neighbourhood of B in $\hat{\mathcal{F}}_*^k$ such that whenever C is in the neighbourhood,

$$\sigma(C^2) \subseteq (-\infty, -1 + \varepsilon] \cup [-\varepsilon, 0],$$

and further that $\|B^2 - C^2\| < \varepsilon < \frac{1}{2}$.

Let E be the spectral projection to $[-\varepsilon, 0]$. We claim that E is isomorphic to $\ker B$. Indeed, consider the orthogonal projection P from E to $\ker B$.

- P is injective. If not, suppose $x \in E \cap (\ker B)^\perp$, and wlog assume $\|x\| = 1$. Since $\sigma(B^2|_{\ker B^\perp}) \subseteq (-\infty, -1]$, we have

$$|\langle (B^2 - C^2)x, x \rangle| = |\langle B^2 x, x \rangle - \langle C^2 x, x \rangle| \geq 1 - \varepsilon,$$

a contradiction.

- P is surjective. If not, suppose $x \in \ker B \cap E^\perp$, and again assume $\|x\| = 1$. Then

$$\langle (B^2 - C^2)x, x \rangle = -\langle C^2 x, x \rangle \geq 1 - \varepsilon,$$

a contradiction.

Now B^2 and C^2 commute with C_k , so the orthogonal projection is in fact a C_k -module isomorphism. We write $E = \ker C \oplus \ker C^\perp$. Then

$$\text{idx } B - \text{idx } C = [\ker C^\perp].$$

But the restriction of C to $\ker C^\perp$ is non-singular and skew-Hermitian, so we can use the action of C to turn $\ker C^\perp$ into a C_{k+1} module. Morally, we should be able to just “scale” C , and the correct thing to write down is

$$J_{k+1} = C(-C^2)^{-1/2}.$$

Surjectivity is trivial (for $[M] \in A_k$, pick an isomorphism $H \cong H \oplus M$ and use $J_{k+1} \oplus 0 \in \mathcal{F}_*^k(H \oplus M)$).

Injectivity follows from the following sequence of observations:

Claim. Any $B \in \mathcal{F}_*^k$ can be deformed so that it is unitary on the complement of $\ker B$.

By restricting to the complement of $\ker B$, we may assume $\ker B$ is trivial. Note that the operator $B(-B^2)^{-1/2}$ is unitary. So we can use the path $B(t - (1-t)B^2)^{-1/2}$.

Claim. Let $B \in \mathcal{F}_*^k$, and let V be a C_{k+1} module. Then we can deform B to B' so that $\ker B' = \ker B \oplus V$ as a C_k -module.

We may assume B is unitary on the complement of $\ker B$. Since $B \in \mathcal{F}_*^k$, we know each C_{k+1} module appears as a direct summand of $\ker B^\perp$ when B acts as J_{k+1} . So we can decompose

$$H = \ker B \oplus V \oplus \text{remaining},$$

and then linearly scale B to vanish on V .

Thus, if $B, C \in \mathcal{F}_*^k$ are such that $[\ker B] = [\ker C]$, then we can deform B and C so that their kernels are in fact isomorphic.

Claim. If $B, C \in \mathcal{F}_*^k$ are such that $\ker B \cong \ker C$ as C_k -modules, then there is a C_k -module isomorphism $T : H \rightarrow H$ such that $TBT^{-1} = C$.

Again wlog assume B and C are unitary on the complement, so that they act as J_{k+1} . Then the complements are isomorphic as C_{k+1} modules, since both have infinitely many copies of each irrep. Hence there must be a C_k -module isomorphism that sends B to C .

We can then conclude the theorem if we can find a path from T to the identity.

Claim. The group of C_k -module isomorphisms $H \rightarrow H$ is connected.

By Schur's lemma, this group is isomorphic to $\text{GL}(H)$, hence is in fact contractible. \square

3.2 THE FIRST PART OF THE PROOF

We can now begin to prove Theorem A. In the remainder of this section, we shall reduce the problem to something easier via a chain of rather straightforward homotopy equivalences, and then in the next section, we shall do some work to finish it off.

The representation theory of Clifford algebras was easier to study with graded modules, but analysis is easier on ungraded things. For $k \geq 1$, we let

$$\tilde{\mathcal{F}}^k(H^0) = \{B \in \mathcal{F}(H^0) : B^* = -B, BJ_i + J_i B = 0 \text{ for } i = 1, \dots, k-1\},$$

and whenever $k \equiv -1 \pmod{4}$, we set

$$\tilde{\mathcal{F}}_*^k(H^0) = \{B \in \tilde{\mathcal{F}}_*^k(H^0) : J_1 J_2 \cdots J_k B \text{ is not essentially definite}\}.$$

Otherwise, we set $\tilde{\mathcal{F}}_*^k(H^0) = \tilde{\mathcal{F}}^k(H^0)$. Then $B \mapsto BJ_k$ gives a homeomorphism $\tilde{\mathcal{F}}_*^k(H^0) \cong \mathcal{F}_*^k(H)$. Similarly, we set $\tilde{\mathcal{F}}_*^0(H^0) = \mathcal{F}(H^0)$. Thus, from now on, we will write \mathcal{F}_*^k to mean $\tilde{\mathcal{F}}_*^k(H^0)$.

In this notation, we want to prove

Theorem A'. For $k \geq 1$, there is a homotopy equivalence

$$\mathcal{F}_*^k \rightarrow \Omega(\mathcal{F}_*^{k-1}).$$

Note that our loops start from J_{k-1} to $-J_{k-1}$, but that doesn't matter.

To prove the theorem, it is convenient to consider the subspace of unitary elements in our operator algebras. We make the following huge list of definitions. Let $\mathcal{A} = \text{End}(H)$, and U denote the unitary elements in a Banach algebra. Then define

$$\begin{aligned} \mathcal{A}^k &= \{T \in \mathcal{A} : T^* = -T, TJ_i + J_i T = 0 \text{ for } i = 1, \dots, k-1\} \\ \mathcal{K} &= \text{compact operators in } \mathcal{A} & \mathcal{K}^k &= \mathcal{A} \cap \mathcal{A}^k \\ \mathcal{B} &= \mathcal{A}/\mathcal{K} & \mathcal{B}^k &= \mathcal{A}^k/\mathcal{K}^k \\ \mathcal{G} &= \mathcal{B}^\times & \mathcal{G}^k &= \mathcal{G} \cap \mathcal{B}^k & \mathcal{G}_*^k &= \mathcal{G}^k \cap p(\mathcal{F}_*^k) \\ G_*^k &= \mathcal{G}_*^k \cap U & F_*^k &= p^{-1}G_*^k \cap U \\ \mathcal{L} &= \mathcal{A}^\times = \text{GL}(H) & L &= \mathcal{L} \cap U & L^k &= \mathcal{L} \cap \mathcal{A}^k \\ \Omega_k &= \{T \in \mathcal{A}^k : T^*T = 1, T \equiv J_{k-1} \pmod{\mathcal{K}}\} \end{aligned}$$

where $p : \mathcal{A} \rightarrow \mathcal{A}/\mathcal{K}$ is the natural projection (recall that $\mathcal{F} = p^{-1}(\mathcal{G})$).

We need a theorem of Bartle and Graves.

Theorem (Bartle–Graves, [4]). Let U, B be Banach spaces and $\pi : U \rightarrow B$ be a bounded linear surjection. Then π has a (not necessarily linear) section.

4 PROVING BOTT PERIODICITY II

It remains to show that the green and red arrows are homotopy equivalences. Note that the loop space involves only the identity component. Let L_{\blacksquare}^{k-1} and G_{\blacksquare}^{k-1} denote the components of L^{k-1} and G^{k-1} respectively that contain J_{k-1} . Then the fact that the green arrows are homotopy equivalences follow from the following two lemmas:

Lemma. We have a fiber sequence

$$\Omega_{k-1} \rightarrow L_{\blacksquare}^{k-1} \rightarrow G_{\blacksquare}^{k-1},$$

Lemma. L_{\blacksquare}^{k-1} is contractible.

Proof of first lemma. Let L_k be the subgroup of L consisting of matrices that commute with J_1, \dots, J_k .

First observe that $L_{\blacksquare}^{k-1} \rightarrow G_{\blacksquare}^{k-1}$ has a local section by restricting the local section of $\mathcal{A} \rightarrow \mathcal{A}/\mathcal{K}$, averaging, and then applying unitary retract. So $L_{\blacksquare}^{k-1} \rightarrow G_{\blacksquare}^{k-1}$ is a fiber bundle. So it suffices to show that the fiber through J_{k-1} is Ω_{k-1} . It is clear that it is given by $\Omega_{k-1} \cap L_{\blacksquare}^{k-1}$. So we want to show that $\Omega_{k-1} \subseteq L_{\blacksquare}^{k-1}$.

Since L_{k-1} is connected (and in fact contractible, by Schur's lemma and Kuiper's theorem), it suffices to show that if $J \in \Omega_{k-1}$, then J is conjugate to J_{k-1} by an element in L_{k-1} . Equivalently, we want to show that when J acts as J_{k-1} , there are infinitely many copies of each irrep.

We only have to consider the case when C_{k-1} is not simple. In this case, we know the projections onto the subspaces spanned by each irrep are given by $\frac{1 \pm w}{2}$. Thus, the image is finite dimensional iff $\frac{1 \pm w}{2}$ is compact. But the $w(J) \equiv w(J_{k-1})$ modulo compact operators. Since $w(J_{k-1})$ is not compact, the same is true for $w(J)$. \square

Proof of second lemma. We claim that there is a fiber sequence

$$L_k \rightarrow L_{k-1} \rightarrow L_{\blacksquare}^{k-1}.$$

This would imply the result, since both L_k and L_{k-1} are contractible.

The second map is given the action $A \mapsto AJ_{k-1}A^{-1}$, and since L_{k-1} is connected, the image lies in L_{\blacksquare}^{k-1} . It follows from [4, Theorem 6] that $L_{k-1} \rightarrow L_{\blacksquare}^{k-1}$ has a local section, which implies that the orbit of J_{k-1} is open. Since the same would be true for the orbits of all other elements in L^{k-1} under the conjugation action, the orbit of J_{k-1} is closed, hence a component, i.e. it is L_{\blacksquare}^{k-1} . Since the stabilizer of J_{k-1} is L_k , we would be done. \square

It remains to show that the red arrow is a homotopy equivalence. To do so, we first translate the spaces involved by J_{k-1} , so that the map between them is simply given by exponentiation. Define

$$\tilde{F}_*^k = F_*^k J_{k-1}, \quad \tilde{\Omega}_{k-1} = J_{k-1} \Omega_{k-1}.$$

Then we can characterize \tilde{F}_*^k as the operators B such that

1. B is Fredholm and skew-Hermitian;
2. B commutes with J_1, \dots, J_{k-2} and anti-commutes with J_{k-1} ;
3. $\|B\| = 1$ and $BB^* = 1$ modulo compact operators.
4. If $k \equiv -1 \pmod{4}$, then $J_1 \cdots J_{k-2}B$ is not essentially definite.

Similarly, $\tilde{\Omega}_{k-1}$ consists of operators S such that


1. S is orthogonal;
2. S commutes with J_1, \dots, J_{k-2} and $(J_{k-1}S)^2 = -1$
3. $S \equiv -1$ modulo compact operators.

We now want to show that $\exp \pi : \tilde{F}_*^k \rightarrow \tilde{\Omega}_{k-1}$ is a homotopy equivalence. Dropping the tildes, we consider filtrations on F_*^k and Ω_{k-1} as follows: We define $\Omega_{k-1}(n)$ to be those operators T such that $I + T$ has rank $\leq n$, and $F_*^k(n) = (\exp \pi)^{-1}(\Omega_{k-1}(n))$.

Lemma. For any m and $a \in \Omega_{k-1}(m), b \in F_*^k(m)$, the natural inclusion induces bijections

$$\begin{aligned} \lim_{n \rightarrow \infty} \pi_k(\Omega_{k-1}(n), a) &\rightarrow \pi_k(\Omega_{k-1}, a) \\ \lim_{n \rightarrow \infty} \pi_k(F_*^k(n), b) &\rightarrow \pi_k(F_*^k, b). \end{aligned}$$

Proof. We only do the case of Ω_{k-1} . The other is similar. Given an element $T \in \Omega_{k-1}$, its spectrum looks roughly like this:



Here orthogonality forces the eigenvalues to lie on the unit circle, and the fact that $I + T$ is compact means -1 is the only point in the essential spectrum.

Thus, given any angle θ , we can apply a spectral deformation to shrink all the points that are at most θ away from -1 to -1 , and what is left is something in $\Omega_{k-1}(n)$ for some n , since there are only finitely many eigenvalues at an angle $> \theta$ away from -1 . By picking a spectral deformation that maps the unit circle to the unit circle and invariant under complex conjugation, this ensures the operator stays orthogonal and real. If we further choose it so that it is odd, then this preserves the set of operators that commute and anti-commute with it. So we stay inside Ω_{k-1} .

In fact, given any compact subset of Ω_{k-1} , applying such a spectral deformation will send it to something lying in $\Omega_{k-1}(n)$ for some large n . Thus, given any $a \in \Omega_{k-1}(m)$, we can pick θ so that all the eigenvalues of T are either -1 or $> \theta$ away from -1 . Then the associated spectral deformation fixes a , but retracts any compact subset (and in particular, any image of maps from spheres) into one of the $\Omega_{k-1}(n)$.

In the case of $F_k^*(m)$, the spectrum instead looks like



□

Now it suffices to show that $\exp \pi : F_*^k(m) \rightarrow \Omega_{k-1}(m)$ is a homotopy equivalence for all m . For convenience, set $H^k(m) = F_*^k(m) \setminus F_*^k(m-1)$ and $D^k(m) = \Omega_{k-1}(m) \setminus \Omega_{k-1}(m-1)$. We shall show that $\exp \pi : H^k(m) \rightarrow D^k(m)$ is a fiber bundle with contractible fiber for all m , and then use the following standard result in homotopy theory:

Lemma. Let (Y, Y_1, Y_2) and (Z, Z_1, Z_2) be excisive triads, and $f : Y \rightarrow Z$ a map of triads. If

$$f|_{Y_1} : Y_1 \rightarrow Z_1, \quad f|_{Y_2} : Y_2 \rightarrow Z_2, \quad f|_{Y_1 \cap Y_2} \rightarrow Z_1 \cap Z_2$$

are all weak homotopy equivalences, then so is f .

□

We shall take


$$\begin{aligned} Y &= F_*^k(m) & Y_1 &= F_*^k(m-1) \\ Z &= \Omega_{k-1}(m) & Z_1 &= \Omega_{k-1}(m-1). \end{aligned}$$

By a simple spectral deformation argument, we can show that $F_*^k(n-1)$ is the deformation retract of some open neighbourhood of $F_*^k(n)$ (add to $F_*^k(n-1)$ those operators whose eigenvalue closest to -1 has small real part, then retract those to -1), and similarly for $\Omega_{k-1}(n-1)$ in a compatible way. We can then Y_2 and Z_2 to be these neighbourhoods.

By induction, we can conclude that $f|_{Y_1}$ is a homotopy equivalence. Then the fact that $\exp \pi$ restricts to a fiber bundle with contractible fiber shows that the other two maps are homotopy equivalences.

Theorem. $\exp \pi : H^k(m) \rightarrow D^k(m)$ is a fiber bundle with contractible fiber.

Proof. Visually, the exponential map does something like this:



Essentially, the “information lost” by passing from $T \in H^k(m)$ to $\exp \pi T \in D^k(m)$ is how the spectrum is distributed between $\pm i$.

To make this concrete, observe that since $A \in D(n)$ is such that $I + A$ has constant rank n , it follows that $\ker(I + A)$ is a Hilbert space subbundle \mathcal{H} of $D(n) \times H$ with orthogonal complement \mathcal{H}^\perp an n -dimensional vector bundle.

Given $T \in H^k(m)$, it acts on \mathcal{H} and \mathcal{H}^\perp disjointly. Its action on \mathcal{H}^\perp is completely determined by $\exp \pi T$ (by linear algebra), and the action of T on \mathcal{H} is orthogonal, skew-Hermitian, etc. Thus, local triviality of \mathcal{H} and \mathcal{H}^\perp implies $\exp \pi : H^k(m) \rightarrow D^k(m)$ is a fiber bundle whose fiber consists of operators $T : H' \rightarrow H'$ (where H' is any fiber of \mathcal{H}) that are

1. Orthogonal and square -1 ;
2. Commuting with J_1, \dots, J_{k-2} and anti-commuting with J_{k-1} ;
3. If $k \equiv -1 \pmod{4}$, then $J_1 \cdots J_{k-2} T$ is not essentially definite.

Equivalently, if we replace B by $B J_{k-1}$ (thus reversing the previous shift), this is the number of choices of J_k such that each irrep appears with infinite multiplicity. Thus, this is given by the quotient of the space of unitary C_{k-1} -automorphisms of H' by the centralizer of J_k , i.e. by the space of C_k -automorphisms of H' . Both of these are contractible. So we are done. \square

5 CODA: THE ABS ORIENTATION

Suppose $\pi : V \rightarrow X$ is a rank k vector bundle with a Euclidean metric. Then taking the Clifford algebra with respect to (the negation of) the metric fiberwise gives us a C_k -bundle $C(V)$ over X . Let $M(V)$ be the Grothendieck group of $C(V)$ -modules. If $E = E^0 \oplus E^1$ is a graded $C(V)$ -module, then the pullbacks π^*E_1 and π^*E_0 are isomorphic over $V^\#$, the complement of the zero section, where the isomorphism at $v \in V^\#$ is given by multiplication by v (with inverse a multiple of v , since $v^2 = -\|v\|$). This thus gives us an element of $\widetilde{KO}(V, V^\#) \cong \widetilde{KO}(X^V)$, where X^V is the Thom space. In other words, we have a canonical map

$$M(V) \rightarrow \widetilde{KO}(X^V).$$

Suppose E came from a $C(V \oplus 1)$ -module, then by a linear deformation, we can use a vector in the trivial factor as the isomorphism between π^*E_1 and π^*E_2 instead, and this extends to an isomorphism defined on all of V . So the resulting class is trivial in $\widetilde{KO}(X^V)$. Moreover, this map is natural with respect to products, namely if $V \rightarrow X$ and $W \rightarrow Y$ are two bundles, then the following diagram commutes:

$$\begin{array}{ccc} A(V) \otimes A(W) & \longrightarrow & A(V \oplus W) \\ \downarrow & & \downarrow \\ \widetilde{KO}(X^V) \otimes \widetilde{KO}(Y^W) & \longrightarrow & \widetilde{KO}(X \times Y^{V \oplus W}). \end{array}$$

In particular, taking $X = *$, we obtain *ring* homomorphisms

$$A_* \rightarrow \sum_{k \geq 0} KO^{-k}(*).$$

In fact, this is the same map as the one we constructed above, which is just a matter of tracing through all the constructions we have done.

Let $\text{GSpin}(k)$ be the even invertible elements of C_k that preserve $\mathbb{R}^k \subseteq C_k$ under the conjugation action, and let $\text{Spin}(k)$ consist of those elements of “spinor norm” 1. The natural action on \mathbb{R}^k begets a map $\text{Spin}(k) \rightarrow \text{SO}(k)$ that is a non-trivial double cover. Crucially, $\text{Spin}(k)$ acts as automorphisms of the whole of C_k .

Thus, V has the structure of a $\text{Spin}(k)$ -bundle, so that there is a principal $\text{Spin}(k)$ -bundle $P \rightarrow X$ with $V = P \times_{\text{Spin}(k)} \mathbb{R}^k$, then $C(V) = P \times_{\text{Spin}(k)} C_k$ and there is a natural map $A_k \rightarrow A(V)$ that sends an C_k module M to $P \times_{\text{Spin}(k)} M$. This gives a long composite

$$A_k \rightarrow \widetilde{KO}(X^V).$$

In particular, if $k = 8m$, and λ is the generator of A_{8m} , then by naturality, the image of λ^m restricts to a generator of $\widetilde{KO}(S^{8m})$. So this is a Thom class of the bundle V . Stabilizing, we get

Theorem. Every spin bundle is naturally KO -oriented.

REFERENCES

- [1] M. F. Atiyah and D. W. Anderson. *K-theory*. Advanced Books Classics Series. Avalon Publishing, 1994.
- [2] M. F. Atiyah, R. Bott, and A. Shapiro. Clifford modules. *Topology*, 3(Supplement 1):3 – 38, 1964.
- [3] M. F. Atiyah and I. M. Singer. Index theory for skew-adjoint fredholm operators. *Publications Mathématiques de l’IHÉS*, 37:5–26, 1969.
- [4] R. G. Bartle and L. M. Graves. Mappings between function spaces. *Transactions of the American Mathematical Society*, 72(3):400–413, 1952.
- [5] P. Enflo. A counterexample to the approximation problem in banach spaces. *Acta Math.*, 130:309–317, 1973.
- [6] N. H. Kuiper. The homotopy type of the unitary group of hilbert space. *Topology*, 3(1):19 – 30, 1965.
- [7] J. W. Milnor. On spaces having the homotopy type of a cw-complex. *Transactions of the American Mathematical Society*, 90(2):272–280, 1959.