COMPANION NOTES FOR THE WINTER SCHOOL COURSE "THE RADIUS OF COMPARISON OF C*-ALGEBRA CROSSED PRODUCTS BY MINIMAL HOMEOMORPHISMS"

N. CHRISTOPHER PHILLIPS

This set of notes is a companion to my Winter School Course "The Radius of Comparison of C*-Algebra Crossed Products by Minimal Homeomorphisms", 14–15 December 2020, at Seoul National University. It contains some basic definitions and theorems, most of which will be presented in the lectures, but a few of which are background which will be assumed. Even for many of the the ones which will be presented in the lectures, more detail and background is given here. On the other hand, the main results the course leads up to are not stated here.

I suggest reading through the relevant sections of these notes before the corresponding lectures, and then referring back to them as needed and rereading them after the corresponding lectures.

These notes are still rough, and are incomplete: they only go roughly through the first two of the three lectures. They have also not been carefully proofread.

1. Banach algebras and C*-algebras

Begin by recalling the definition of a Banach algebra.

Definition 1.1. A *Banach algebra* is a Banach space A over $\mathbb{C}$ which is also an algebra over $\mathbb{C}$, such that $\|ab\| \leq \|a\|\|b\|$ for all $a, b \in A$. If A is a unital algebra, we also require $\|1\| = 1$.

Definition 1.2. A *Banach *-algebra* is a Banach algebra A over $\mathbb{C}$ equipped with an operation $a \mapsto a^*$, the *adjoint*, satisfying:

(1) $(\lambda a + \mu b)^* = \overline{\lambda} a^* + \overline{\mu} b^*$ for all $a, b \in A$ and all $\lambda, \mu \in \mathbb{C}$.
(2) $(ab)^* = b^* a^*$ for all $a, b \in A$.
(3) $a^{**} = a$ for all $a \in A$.
(4) $\|a^*\| = \|a\|$ for all $a \in A$.

If A is unital then it follows that $1^* = 1$.

Definition 1.3. A *C*-algebra* is a Banach *-algebra A such that $\|a^* a\| = \|a\|^2$ for all $a \in A$.

It is automatic that $\|a^* a\| \leq \|a\|^2$, so we are really only ruling out $\|a^* a\| < \|a\|^2$. We (informally) call this condition the "C* condition". Despite initial appearances, it is *very* restrictive and *very* powerful.

Example 1.4. The basic examples are:

Date: 13 December 2020.
Partially supported by the Simons Foundation Collaboration Grant for Mathematicians #587103.

(1) $M_n = M_n(\mathbb{C})$, with the norm gotten by identifying it with operators on $l^2(\{1, 2, \ldots, n\})$, and adjoint given by conjugate transpose: if $a = (a_{j,k})_{j,k=1}^n$ then $(a^*)_{j,k} = \overline{a_{k,j}}$.
(2) $C(X)$, the algebra of all continuous complex valued functions on a compact Hausdorff space X, with $\|f\| = \sup_{x \in X} |f(x)|$ and $f^*(x) = \overline{f(x)}$.
(3) The algebra $L(H)$ of all bounded operators on a Hilbert space H, with the operator norm and with a^* determined by $\langle a^*\xi, \eta\rangle = \langle \xi, a\eta\rangle$ for all $\xi, \eta \in H$. (Part (1) is a special case of this.)
(4) Any subalgebra $A \subset L(H)$ which is closed in the norm topology and self-adjoint (closed under the adjoint, that is, $a \in A$ implies $a^* \in A$).

Theorem 1.5. Every C*-algebra is isomorphic to a selfadoint subalgebra of $L(H)$ for some Hilbert space H, as in Example 1.4(4).

An isomorphism is an algebra isomorphism which preserves the norm and the adjoint. For C*-algebras, the requirement that the norm be preserved turns out to be redundant. More generally, we define a homomorphism $\varphi\colon A \to B$ of C*-algebras to be an algebra homomorphism which also preserves the adjoint. It is a striking early theorem in the subject that, without adding any continuity assumptions, homomorphisms of C*-algebras are automatically contractive ($\|\varphi(a)\| \leq \|a\|$ for all a) and have closed range. Moreover, injective homomorphisms are isometric.

Theorem 1.6. Every commutative unital C*-algebra is isomorphic to $C(X)$ for some compact Hausdorff space X, as in Example 1.4(2). Moreover, for compact Hausdorff spaces X and Y, there is a bijection from continuous maps $h\colon Y \to X$ to unital homomorphisms $\varphi\colon C(Y) \to C(X)$, given by assigning to h the homomorphism $\varphi(f) = f \circ h$ for $f \in C(Y)$.

Given a commutative unital C*-algebra A, the corresponding space is gotten as the maximal ideal space of A.

Remark 1.7. In Theorem 1.6, one actually has a contravariant category equivalence from compact Hausdorff spaces and continuous maps to commutative unital C*-algebras and unital homomorphisms. This gives license to thinking of not necessarily commutative C*-algebras as "noncommutative topological spaces", an idea which has motivated a vast amount of interesting work on C*-algebras but which doesn't apply to general Banach algebras.

Example 1.8. Some basic nonexamples are:

(1) The algebra of upper triangular $n \times n$ matrices over $\mathbb{C}$. (This algebra is not selfadjoint.)
(2) The disk algebra, the set of all continuous functions on the closed unit disk in $\mathbb{C}$ whose restriction to the interior is holomorphic. (There is an adjoint, namely $f^*(z) = \overline{f(\overline{z})}$, but it behaves very badly. See the example after Theorem 2.2.)
(3) A group algebra $L^1(G)$ for a locally compact group G (if G is discrete, $l^1(G)$), with convolution, L^1 norm, and adjoint $a^*(g) = \Delta(g)^{-1}\overline{a(g^{-1})}$. This algebra is fairly close to being a C*-algebra, but the norm must be "corrected" and $L^1(G)$ must then be completed in the new norm. The result is the group C*-algebra $C^*(G)$.

(4) M_n with the norm gotten by identifying it with operators on $l^p(\{1, 2, \ldots, n\})$ with fixed $p \in [1, \infty) \setminus \{2\}$. The adjoint is still continuous, but the norm is wrong.

To illustrate Example 1.8(3), take $G = \mathbb{Z}$. The algebra $l^1(\mathbb{Z})$ can be identified with the Banach algebra of functions on the circle S^1 with absolutely convergent Fourier series, and $C^*(\mathbb{Z})$ can be identified with the Banach algebra $C(S^1)$ of all continuous functions on the circle.

2. Spectrum, functional calculus, and order

The material here is not entirely in the order in which it must be proved in the development of the subject.

We recall the definition of the spectrum of an element of a Banach algebra.

Definition 2.1. Let A be a unital Banach algebra, and let $a \in A$. Then the *spectrum* $\mathrm{sp}(a)$ is defined by

$$\mathrm{sp}(a) = \big\{\lambda \in \mathbb{C} \colon \lambda - a \text{ is not invertible}\big\}.$$

Here $\lambda - a$ is a (standard) abbreviation for $\lambda \cdot 1 - a$. The set $\mathrm{sp}(a)$ is a nonempty compact subset of $\mathbb{C}$.

If $A = M_n$, then $\mathrm{sp}(a)$ is the set of eigenvalues of a. Even in $L(H)$ for a Hilbert space H, things can be much more complicated. The unilateral shift is injective, even isometric, but not invertible because it is not surjective. Multiplication by the function $f(\lambda) = \lambda$ on $L^2([0, 1])$ (using Lebesgue measure) has no eigenvalues, but every $\lambda \in [0, 1]$ is an "approximate eigenvalue" and is in the spectrum.

Theorem 2.2. Let A be a unital C*-algebra, and let $a \in A$ be selfadjoint. Then $\mathrm{sp}(a) \subset \mathbb{R}$.

This theorem is false in Banach *-algebras not satisfying the C* condition. In the disk algebra, Example 1.8(2), the function $f(z) = z$ satisfies $f^* = f$, but $\mathrm{sp}(f)$ is the closed unit disk.

Theorem 2.3. Let A be a unital C*-algebra, and let $a \in A$. Then the following are equivalent:

(1) a is selfadjoint and $\mathrm{sp}(a) \subset [0, \infty)$.
(2) There is $b \in A$ such that $a = b^*b$.
(3) There is a selfadjoint element $b \in A$ such that $a = b^2$.
(4) There is a selfadjoint element $b \in A$ such that $a = b^2$ and $\mathrm{sp}(b) \subset [0, \infty)$.
(5) For every representation $\pi \colon A \to L(H)$ of A on a Hilbert space H, for all $\xi \in L(H)$ one has $\langle \pi(a)\xi, \xi \rangle \geq 0$.
(6) There is an injective representation $\pi \colon A \to L(H)$ of A on a Hilbert space H such that for all $\xi \in L(H)$ one has $\langle \pi(a)\xi, \xi \rangle \geq 0$.

Representations are of course required to preserve the adjoint. The condition $\langle \pi(a)\xi, \xi \rangle \geq 0$ for all $\xi \in H$ is, of course, the usual definition of positivity of an operator on a Hilbert space.

The element b in Theorem 2.3(4) is unique, and is called $a^{1/2}$.

Definition 2.4. Let A be a unital C*-algebra, and let $a \in A$. We say that a is *positive* if a is selfadjoint and satisfies the conditions in Theorem 2.3. We write $a \geq 0$.

If $b, c \in A$ are selfadjoint, we further write $b \leq c$ to mean $c - b \geq 0$.

When one represents A faithfully on a Hilbert space, Definition 2.4 corresponds to the usual definitions of positivity and order for operators on a Hilbert space.

A function $f \in C(X)$ is selfadjoint if and only if $f(x) \in \mathbb{R}$ for all $x \in X$, and positive if and only if $f(x) \geq 0$ for all $x \in X$. So real valued functions f and g satisfy $f \leq g$ if and only if $f(x) \leq g(x)$ for all $x \in X$. In general, though, the order is not as well behaved as the example of $C(X)$ suggests. It is not true that $0 \leq a \leq b$ implies $a^2 \leq b^2$; there are counterexamples in M_2. There is no such thing as $a \wedge b$ or $a \vee b$ unless a and b commute. It is true that if $x \leq y$ then $a^*xa \leq a^*ya$. (Proof: if $y - x = t^*t$ then $a^*ya - a^*xa = (ta)^*(ta)$.) But it is not true that if $a \geq 0$ and $0 \leq x \leq y$ then $x^{1/2}ax^{1/2} \leq y^{1/2}ay^{1/2}$. Nor is anything similar true.

Let A be a unital C*-algebra, and let $a \in A$ be normal, that is, a commutes with a^*. Then the norm closed subalgebra B generated by a, a^*, and 1 is a commutative unital C*-algebra, and therefore isomorphic to $C(X)$ for some compact Hausdorff space X. It turns out that one can identify X with $\mathrm{sp}(a)$, giving an isomorphism $C(\mathrm{sp}(a)) \to B \subset A$, which is written $f \mapsto f(a)$. This map agrees with holomorphic functional calculus when f is the restriction to $\mathrm{sp}(a)$ of a holomorphic function on a neighborhood of $\mathrm{sp}(a)$. In particular, in holomorphic functional calculus for a normal element a of a C*-algebra, only the values of f on $\mathrm{sp}(a)$ matter. Formally, we have the following theorem.

Theorem 2.5. Let A be a unital C*-algebra, and let $a \in A$ be normal. Then there exists a unique unital homomorphism $\varphi\colon C(\mathrm{sp}(a)) \to A$ such that, if $h \in C(\mathrm{sp}(a))$ is the function $h(\lambda) = \lambda$ for all $\lambda \in \mathrm{sp}(a)$, then $\varphi(h) = a$.

If $\mathrm{sp}(a) \subset [0, \infty)$, $\alpha > 0$, and $f(\lambda) = \lambda^\alpha$, then $f(a)$ is written a^α, and analogously for other functions written with similar notation. The element $a^{1/2}$ from this notation agrees with the notation $a^{1/2}$ after Theorem 2.3.

3. Crossed products

For simplicity, we restrict almost entirely to discrete groups, although crossed products are defined for actions of locally compact groups. We also usually assume the group is amenable. This class of of groups contains finite groups and $\mathbb{Z}$ In fact, all abelian groups, and even all solvable groups, are amenable. This class does not contain the free group F_2 on two generators or $\mathrm{SL}_2(\mathbb{Z})$.

Definition 3.1. Let G be a group, taken with the discrete topology, and let A be a C*-algebra. An *action of G on A* is a group homomorphism $\alpha\colon G \to \mathrm{Aut}(A)$, usually written $g \mapsto \alpha_g$.

If G is not discrete, one adds the continuity condition: for every $a \in A$, the function $g \mapsto \alpha_g(a)$, from G to A, is norm continuous.

Definition 3.2. Let G be a group, taken with the discrete topology, let X be a compact Hausdorff space, and let $(g, x) \mapsto gx$ be a continuous action of G on X. The induced action of G on $C(X)$, say α, is given by $\alpha_g(f)(x) = f(g^{-1}x)$ for $g \in G$, $f \in C(X)$, and $x \in X$.

When A is unital and G is discrete, the crossed product $C^*(G, A, \alpha)$ is the universal unital C*-algebra generated by a copy of A and a unitary representation $g \mapsto u_g$ of G which implements the action, in the sense that $u_g a u_g^* = \alpha_g(a)$ for all $g \in G$ and $a \in A$.

If, in addition, G is amenable, it can be obtained as follows. Choose any unital faithful representation $\pi_0 \colon A \to L(H_0)$ on a Hilbert space H_0. (It is a theorem that the choice of π_0 doesn't matter.) Set $H = l^2(G, H_0)$, the Hilbert space of all functions $\xi \colon G \to H_0$ such that $\sum_{g \in G} \|\xi(g)\|^2 < \infty$. For $g \in G$, define a unitary $u_g \in L(H)$ by

$$(u_g \xi)(h) = \xi(g^{-1}h).$$

for $h \in G$. Define a representation π of A on H as follows. For $a \in A$ and $g \in G$, set

$$\big(\pi(a)\xi\big)(h) = \pi_0\big(\alpha_{h^{-1}}(a)\big)\big(\xi(h)\big).$$

Then take $C^*(G, A, \alpha)$ to be the C*-subalgebra of $L(H)$ generated by

$$\big\{u_g \colon g \in G\big\} \cup \big\{\pi(a) \colon a \in A\big\}.$$

(If G is not amenable, this algebra is the reduced crossed product $C^*_{\mathrm{r}}(G, A, \alpha)$.)

When α is understood, it can be omitted from the notation, and we write $C^*(G, A)$. If $A = C(X)$, and the action arises as in Definition 3.2, we sometimes write $C^*(G, X)$. For the action of $\mathbb{Z}$ generated by $\alpha \in \mathrm{Aut}(A)$, we write $C^*(\mathbb{Z}, A, \alpha)$, and for the action of $\mathbb{Z}$ generated by a homeomorphism $h \colon X \to X$, we write $C^*(\mathbb{Z}, X, h)$.

Example 3.3. If G acts trivially on $\mathbb{C}$, the crossed product is the group C*-algebra $C^*(G)$. If G is abelian, it is isomorphic to the algebra of continuous functions on the Pontryagin dual $\widehat{G}$ of G.

Example 3.4. Let G be finite, and let G act on $C(G)$ via the translation action of G on G. Call this action α. Set $n = \mathrm{card}(G)$. Then the crossed product $C^*(G,\, C(G),\, \alpha)$ is isomorphic to M_n.

We prove this by exhibiting a family of matrix units in $C^*(G,\, C(G),\, \alpha)$ which spans $C^*(G,\, C(G),\, \alpha)$ and which is indexed by G.

Let $\delta_g \in C(G)$ be the function $\chi_{\{g\}}$. Then $\alpha_g(\delta_h) = \delta_{gh}$ for $g, h \in G$. Also, $\mathrm{span}\big(\{\delta_g \colon g \in G\}\big) = C(G)$. For $g, h \in G$, the element

$$v_{g,h} = \delta_g u_{gh^{-1}}$$

is in $C^*(G,\, C(G),\, \alpha)$. Moreover, for $g_1, h_1, g_2, h_2 \in G$, we have

$$\begin{aligned} v_{g_1,h_1} v_{g_2,h_2} &= \delta_{g_1} u_{g_1 h_1^{-1}} \delta_{g_2} u_{g_2 h_2^{-1}} \\ &= \delta_{g_1} \alpha_{g_1 h_1^{-1}}(\delta_{g_2}) u_{g_1 h_1^{-1}} u_{g_2 h_2^{-1}} = \delta_{g_1} \delta_{g_1 h_1^{-1} g_2} u_{g_1 h_1^{-1} g_2 h_2^{-1}}. \end{aligned}$$

Thus, if $g_2 \neq h_1$, the answer is zero, while if $g_2 = h_1$, the answer is v_{g_1,h_2}. Similarly, $v_{g,h}^* = v_{h,g}$. That is, the elements $v_{g,h}$ satisfy the relations for a system of matrix units indexed by G. It is easily checked that $\mathrm{span}\big(\{v_{g,h} \colon g, h \in G\}\big)$ is the algebra generated by $\{u_g \colon g \in G\}$ and $C(G)$. Since this algebra is finite dimensional, it is complete, and this algebra is all of $C^*(G,\, C(G),\, \alpha)$.

We describe imporant special classes of actions of groups on spaces.

Definition 3.5. Let G be a group, let X be a set, and let $(g, x) \mapsto gx \colon G \times X \to X$ be an action of G on X. The action is *free* if whenever $g \in G$ and $x \in X$ satisfy $gx = x$, then $g = 1$.

Equivalently, every nontrivial group element moves every point.

Definition 3.6. Let G be a group, let X be a set, and let $(g,x) \mapsto gx \colon G \times X \to X$ be an action of G on X. For $x \in X$, the *orbit* (or *G-orbit*) of x is the set $Gx = \{gx \colon g \in G\}$.

Definition 3.7. Let G be a group, let X be a topological space, and let $(g,x) \mapsto gx \colon G \times X \to X$ be a continuous action of G on X. The action is *minimal* if the only closed G-invariant subsets of X are $\varnothing$ and X. If $G = \mathbb{Z}$ and the action is generated by a homeomorphism h, we say that h is a *minimal homeomorphism.*

It is equivalent to say that the only open G-invariant subsets of X are $\varnothing$ and X. It is also equivalent to say that every orbit in X is dense.

Theorem 3.8. Let X be a compact Hausdorff space, let G be a discrete amenable group, and let $(g,x) \mapsto gx \colon G \times X \to X$ be a continuous action of G on X which is free and minimal. Then $C^*(G,X)$ is simple.

Freeness is slight overkill; something called "essential freeness" is good enough. Amenability can be dropped if one uses the reduced crossed product. The proof requires work.

If the action is not minimal, it is easy to see that simplicity fails. Let $U \subset X$ be a nontrivial open G-invariant subset. Then $C_0(U)$ is a nontrivial G-invariant ideal in $C(X)$, and one can show that this implies that $C^*(G,U)$ is a nontrivial ideal in $C^*(G,X)$.

Crossed products by minimal homeomorphisms sometimes contain interesting information about the homeomorphisms they come from.

Theorem 3.9 (Giordano-Putnam-Skau)**.** Let X be the Cantor set, and let $h_1, h_2 \colon X \to X$ be minimal homeomorphisms. Then $C^*(\mathbb{Z}, X, h_1) \cong C^*(\mathbb{Z}, X, h_2)$ if and only if h_1 and h_2 are what is called "strongly orbit equivalent".

We don't define strong orbit equivalence here. It is a stronger condition than orbit equivalence, which is the condition that there exist a homeomorphism $g \colon X \to X$ which sends orbits of h_1 to orbits of h_2, that is,

$$g\big(\big\{h_1^n(x) \colon n \in \mathbb{Z}\big\}\big) = \big\{h_2^n(x) \colon n \in \mathbb{Z}\big\}$$

for all $x \in X$.

However, sometimes crossed products preserve very little information. For example, let $n_1, n_2 \in \mathbb{Z}_{>0}$, and let

$$h_1 \colon S^{2n_1+1} \to S^{2n_1+1} \qquad \text{and} \qquad h_2 \colon S^{2n_2+1} \to S^{2n_2+1}$$

be minimal homeomorphisms. Then $C^*(\mathbb{Z},\, S^{2n_1+1},\, h_1) \cong C^*(\mathbb{Z},\, S^{2n_2+1},\, h_2)$ if and only if h_1 and h_2 have isomorphic simplexes of invariant Borel probability measures. In particular, this can happen even if $n_1 \neq n_2$.

4. Dimension and mean dimension

We will use covering dimension $\dim(X)$ of a compact metric space X. It is defined for a general topological space, although only useful under additional conditions. If M is a compact manifold, $\dim(M)$ is the usual dimension of M. If X is a finite CW complex, $\dim(X)$ is the usual dimension of the highest dimensional cell of X. In particular, if X is the disjoint union of finitely many compact manifolds, then $\dim(X)$ is the usual dimension of the manifold in the disjoint union with the largest dimension; the rest are ignored.

For compact metric spaces, covering dimension agrees with the large and small inductive dimensions, as well as several other versions of dimension, but even in this context bad behavior is possible. For example, we always have $\dim(X \times Y) \leq \dim(X) + \dim(Y)$, but $\dim(X \times Y) < \dim(X) + \dim(Y)$ is possible, even when $X = Y$.

Notation 4.1. Let X be a compact Hausdorff space. We write $\mathrm{Cov}(X)$ for the set of all finite open covers of X.

Definition 4.2. Let X be a compact Hausdorff space, and let $\mathcal{U}$ be a finite open cover of X. The *order* $\mathrm{ord}(\mathcal{U})$ of $\mathcal{U}$ is the least number $n \in \mathbb{Z}_{>0}$ such that the intersection of any $n+2$ distinct elements of $\mathcal{U}$ is empty.

That is, $\mathrm{ord}(\mathcal{U})$ is the largest $n \in \mathbb{Z}_{>0}$ such that there are $n+1$ distinct sets in $\mathcal{U}$ whose intersection is not empty. An alternative formulation is

$$\mathrm{ord}(\mathcal{U}) = -1 + \sup_{x \in X} \sum_{U \in \mathcal{U}} \chi_U(x).$$

The normalization is chosen so that if $\mathcal{U}$ is cover of X by disjoint open sets, and $X \neq \varnothing$, then $\mathrm{ord}(\mathcal{U}) = 0$: the intersection of any two distinct sets in $\mathcal{U}$ is empty, but the sets themselves need not be empty.

Definition 4.3. Let X be a compact Hausdorff space, and let $\mathcal{U}$ and $\mathcal{V}$ be finite open covers of X. Then $\mathcal{V}$ *refines* $\mathcal{U}$, written $\mathcal{V} \prec \mathcal{U}$, if for every $V \in \mathcal{V}$ there is $U \in \mathcal{U}$ such that $V \subset U$.

That is, every set in $\mathcal{V}$ is contained in some set in $\mathcal{U}$.

Definition 4.4. Let X be a compact Hausdorff space, and let $\mathcal{U}$ be a finite open cover of X. We define the *dimension* $\mathcal{D}(\mathcal{U})$ of $\mathcal{U}$ by

$$\mathcal{D}(\mathcal{U}) = \inf\big(\big\{\, \mathrm{ord}(\mathcal{V}) \colon \mathcal{V} \in \mathrm{Cov}(X) \text{ and } \mathcal{V} \prec \mathcal{U} \big\}\big).$$

That is, $\mathcal{D}(\mathcal{U})$ is the least possible order of a finite open cover which refines $\mathcal{U}$.

Definition 4.5. Let X be a nonempty compact Hausdorff space. The *covering dimension* $\dim(X)$ is

$$\dim(X) = \sup\big(\big\{\mathcal{D}(\mathcal{U}) \colon \mathcal{U} \in \mathrm{Cov}(X)\big\}\big).$$

By convention, $\dim(\varnothing) = -1$.

That is, $\dim(X)$ is the supremum of $\mathcal{D}(\mathcal{U})$ over all finite open covers $\mathcal{U}$ of X.

Now let $h \colon X \to X$ be a homeomorphism. We care mainly about minimal homeomorphisms, but minimality is not needed in the definitions. The notion of mean dimension generalizes to actions of arbitrary countable amenable groups, but we don't give the definitions here.

Definition 4.6. Let X be a compact Hausdorff space, and let $\mathcal{U}$ and $\mathcal{V}$ be two finite open covers of X. Then the *join* $\mathcal{U} \vee \mathcal{V}$ of $\mathcal{U}$ and $\mathcal{V}$ is

$$\mathcal{U} \vee \mathcal{V} = \big\{U \cap V \colon U \in \mathcal{U} \text{ and } V \in \mathcal{V}\big\}.$$

Definition 4.7. Let X be a compact Hausdorff space, let $\mathcal{U}$ be a finite open cover of X, and let $h \colon X \to X$ be a homeomorphism. We define

$$h(\mathcal{U}) = \big\{h(U) \colon U \in \mathcal{U}\big\}.$$

Definition 4.8. Let X be a compact metric space and let $h\colon X \to X$ be a homeomorphism. Then the *mean dimension* of h is

$$\operatorname{mdim}(h) = \sup_{\mathcal{U} \in \operatorname{Cov}(X)} \lim_{n \to \infty} \frac{\mathcal{D}\big(\mathcal{U} \vee h^{-1}(\mathcal{U}) \vee \cdots \vee h^{-n+1}(\mathcal{U})\big)}{n}.$$

There are n covers in the join in the definition. We don't discuss why the limit exists.

The mean dimension attempts to measure the following. Assume X is infinite dimensional. (Otherwise, one immediately sees that $\operatorname{mdim}(h) = 0$.) Start with some finite open cover of X. It "sees only finitely many of the dimensions of X". For example, if X is a shift space $K^{\mathbb{Z}}$ (Definition 4.9 below), then any open set is nontrivial in only finitely many coordinates and so any finite open cover "sees" only finitely many coordinates. Now join this cover with its images under powers of the homeomorphism. The result still "sees only finitely many of the dimensions of X", but it may see more of them. Mean dimension tells you "how many more of the dimensions of X" you see with every additional power of h. We give a somewhat less vague description after Theorem 4.10.

More formally, the definition is designed to make Theorem 4.10 below true.

Definition 4.9. Let K be a set. The *shift* $h_K\colon K^{\mathbb{Z}} \to K^{\mathbb{Z}}$ is the bijection given by $h_K(x)_k = x_{k+1}$ for $x = (x_k)_{k \in \mathbb{Z}} \in K^{\mathbb{Z}}$ and $k \in \mathbb{Z}$.

Theorem 4.10. Let $d \in \mathbb{Z}_{>0}$, let K be a finite simplicial complex with $\dim(K) = d$ (such as $K = [0,1]^d$), and let h_K be as in Definition 4.9. Then $\operatorname{mdim}(h_K) = d$.

To understand the result heuristically, for $n \in \mathbb{Z}$ let $p_n\colon K^{\mathbb{Z}} \to K$ be the projection on the n-th coordinate. In Definition 4.8 consider a finite open cover $\mathcal{U}_0$ of K, and form the finite open cover of $K^{\mathbb{Z}}$ given by

$$\mathcal{U} = \big\{p_0^{-1}(U)\colon U \in \mathcal{U}_0\big\}.$$

Then the cover $\mathcal{U} \vee h_K^{-1}(\mathcal{U}) \vee \cdots \vee h_K^{-n+1}(\mathcal{U})$ sees only n of the coordinates in $K^{\mathbb{Z}}$, so that

$$\mathcal{D}\big(\mathcal{U} \vee h^{-1}(\mathcal{U}) \vee \cdots \vee h^{-n+1}(\mathcal{U})\big) \leq \dim(K^n) \leq n \dim(K).$$

The proof of $\operatorname{mdim}(h_K) \leq d$ in Theorem 4.10 requires only one modification of this idea, namely that the original cover $\mathcal{U}$ must be allowed to depend on an arbitrary finite number of coordinates rather than just one.

One does not expect $\operatorname{mdim}(h_K) = \dim(K)$ for general compact metric spaces K, because of the possibility of having $\dim(K^n) < n \dim(K)$.

Shifts are certainly not minimal. For example, every constant sequence is a fixed point, so is an invariant closed subset.

Theorem 4.11 (Dou)**.** Let G be a countable infinite amenable group and let P be a polyhedron. Suppose $0 \leq \beta < \dim(P)$. Then there is a closed shift invariant subset $X \subset P^G$ such that the restriction of the shift to X is minimal and has mean dimension equal to β.

In particular, G can be $\mathbb{Z}$.

The subshifts in this construction have additional useful properties.

It is not known whether there can be a minimal subshift with mean dimension equal to $\dim(P)$.

5. Projections, vector bundles, and comparison

Definition 5.1. Let A be a C*-algebra. An element $p \in A$ is a *projection* if $p^2 = p$ and $p^* = p$.

Thus, a projection is a selfadjoint idempotent. In $L(H)$, the projections are exactly the orthogonal projections onto closed subspaces. One can show that $p \leq q$ if and only if $pq = p$, and it follows that $p \leq q$ if and only if the range of p is contained in the range of q.

In $C(X)$, the projections are exactly the characteristic functions of compact open subsets of X. Moreover, $p \leq q$ if and only if the compact open subset corresponding to p is a subset of the one corresponding to q.

Definition 5.2. Let A be a C*-algebra, and let $p, q \in A$ be projections. We say that p and q are *Murray-von Neumann equivalent,* written $p \sim_{\mathrm{MvN}} q$, if there is $s \in A$ such that $s^*s = p$ and $ss^* = q$.

Murray-von Neumann equivalence is an equivalence relation. (Other relations, namely homotopy and unitary equivalence, are also used, but we won't need them here.)

In $L(H)$, the element s is an isometric isomorphism from the range of p to the range of q. Thus, $p \sim_{\mathrm{MvN}} q$ if and only if the ranges of p and q have the same dimension as Hilbert spaces, that is, they have orthonormal bases of the same cardinality.

In $C(X)$, $p \sim_{\mathrm{MvN}} q$ if and only if $p = q$.

In a general C*-algebra A, think of $p \sim_{\mathrm{MvN}} q$ as meaning that the ranges of p and q are "isomorphic via an isomorphism in A".

Definition 5.3. Let A be a C*-algebra, and let $p, q \in A$ be projections. We say that p *is Murray-von Neumann subequivalent to* q, written $p \precsim_{\mathrm{MvN}} q$, if there is a projection $e \in A$ such that $p \sim_{\mathrm{MvN}} e$ and $e \leq q$.

One can check that this is the same as saying that there is $s \in A$ such that $s^*s = p$ and $ss^* \leq q$. In $L(H)$, this happens if and only if the range of p is isomorphic to a subspace of the range of q.

The direct sum of subspaces corresponds to a "direct sum" of projections, but one may need to pass to a larger algebra. If $p \in M_m(A)$, the $m \times m$ matrices over A, and $q \in M_n(A)$, then we take

$$p \oplus q = \mathrm{diag}(p, q) = \begin{pmatrix} p & 0 \\ 0 & q \end{pmatrix} \in M_{m+n}(A).$$

To make things fit together, we introduce the following notation.

Notation 5.4. Let A be a C*-algebra. We let $M_\infty(A)$ denote the *-algebra of all infinite matrices over A such that all but finitely many entries are zero. We regard $M_n(A)$ as embedded in $M_\infty(A)$ via

$$a \mapsto \begin{pmatrix} a & 0 \\ 0 & 0 \end{pmatrix},$$

that is, as the upper left corner.

Then $M_\infty(A) = \bigcup_{n=1}^{\infty} M_n(A)$. Warning: $M_\infty(A)$ is not a C*-algebra, because it is not complete. (The completion is $K \otimes A$.)

Murray-von Neumann equivalence and Murray-von Neumann subequivalence still make sense in $M_\infty(A)$, and we define $V(A)$ to be the set of Murray-von Neumann subequivalence classes of projections in $M_\infty(A)$.

The direct sum, as above, turns out to be well defined on $V(A)$.

We now turn to vector bundles. Their importance for us is that, on the one hand, a vector bundle over X give rise to a projection in $M_\infty(C(X))$ (see Theorem 5.15 below), and, on the other hand, vector bundles arise naturally in several ways and have a well developed theory.

A (complex) vector bundle E over a topological space X consists of an assignment of a finite dimensional complex vector space E_x to each point $x \in X$, in such a way that the vector spaces E_x "fit together reasonably", and moreover in such a way that $E = \coprod_{x \in X} E_x$ is a reasonable topological space and is "locally trivial". The vector space E_x is called the *fiber* of E over the point x. We list some examples before giving the formal definition.

Example 5.5. A vector bundle over a one point space is just a finite dimensional complex vector space.

Example 5.6. If X is any topological space, then $E = \mathbb{C}^n \times X$, with the product topology and $E_x = \mathbb{C}^n \times \{x\}$ for $x \in X$, is a vector bundle.

A vector bundle isomorphic to one of these is called *trivial*.

The next example is for those who have seen differentiable manifolds. If the manifold is real (the usual situation), one gets a real vector bundle rather than a complex vector bundle. A complex manifold gives a complex vector bundle.

Example 5.7. Let M be a smooth manifold. Then the tangent bundle TM, with fiber over $x \in M$ taken to be $T_x M$, the tangent space of M at x, is a vector bundle over M.

Example 5.8. The open Möbius strip can be regarded as a real vector bundle over its "center line", which is a copy of S^1. It locally looks like the product of an open subset of S^1 with, say, $(-1, 1)$. Topologically, $(-1, 1)$ is the same as $\mathbb{R}$, so think of the open Möbius strip as locally looking like the product of an open subset of S^1 with $\mathbb{R}$. In this way, one gets fibers which are one dimensional real vector spaces.

The Möbius strip bundle is locally trivial but is not trivial.

Example 5.9. Let $\mathbb{C}P^n$ be the n dimensional complex projective space, thought of as the set of complex lines in $\mathbb{C}^{n+1}$. It is the quotient space $(\mathbb{C}^{n+1} \setminus \{0\})/(\mathbb{C} \setminus \{0\})$ coming from the action of the multiplicative group $\mathbb{C} \setminus \{0\}$ on $\mathbb{C}^{n+1} \setminus \{0\}$ by scalar multiplication. Equivalently, it is the quotient space of the action of S^1 on the unit sphere in $\mathbb{C}^{n+1}$ by scalar multiplication. There is a "tautological bundle" whose fiber over the point $L \in \mathbb{C}P^n$ (thought of as the set of complex lines in $\mathbb{C}^{n+1}$) is exactly the complex vector space L. Formally, it is

$$\big\{(\xi, L) \in \mathbb{C}^{n+1} \times \mathbb{C}P^n \colon \xi \in L\big\},$$

with the relative product topology.

This particular vector bundle plays a key role here.

Example 5.10. Let $k, n \in \mathbb{Z}_{\geq 0}$ with $k \leq n$. The (complex) Grassmannian $\mathrm{Gr}_k(\mathbb{C}^n)$ is the set of all k dimensional subspaces of $\mathbb{C}^n$, with a suitable topology (gotten from a realization as a quotient of the open subset of $(\mathbb{C}^n)^k$ consisting of linearly

independent k-tuples). There is a complex vector bundle E over $\mathrm{Gr}_k(\mathbb{C}^n)$, whose fiber over V, regarded as a point in $\mathrm{Gr}_k(\mathbb{C}^n)$, is V, regarded as a k dimensional complex vector space.

Now we give the definition. Part (2) asserts, in view of Example 5.6, that E is locally (on the base) isomorphic to a trivial vector bundle, that is, locally trivial. Requiring local triviality eliminates the need to require continuity of addition, scalar multiplication, etc., as well as the need to require that q be topologically a quotient map.

Definition 5.11. Let X be a topological space. A *vector bundle* over X is a topological space E equipped with a continuous surjective map $q\colon E \to X$ and with a finite dimensional complex vector space structure on $E_x = q^{-1}(\{x\})$ for every $x \in X$, with the following local triviality property. For every $x \in X$ there exist an open set $U \subset X$ containing x, an integer $n \in \mathbb{Z}_{\geq 0}$, and a homeomorphism $b\colon \mathbb{C}^n \times U \to q^{-1}(U)$ such that:

(1) $b(\mathbb{C}^n \times \{y\}) = E_y$ for all $y \in U$.
(2) $b|_{\mathbb{C}^n \times \{y\}}\colon \mathbb{C}^n \times \{y\} \to E_y$ is an isomorphism of complex vector spaces for all $y \in U$.

We describe several operations on vector bundles.

Definition 5.12. Let X be a topological space, and let E_1 and E_2 be vector bundles over X, with projection maps $q_1\colon E_1 \to X$ and $q_2\colon E_2 \to X$. The *direct sum* $E_1 \oplus E_2$ is the vector bundle

$$E_1 \oplus E_2 = \big\{(\xi_1, \xi_2) \in E_1 \times E_2\colon q_1(\xi_1) = q_2(\xi_2)\big\},$$

with the restricted product topology. The fiber $(E_1 \oplus E_2)_x = (E_1)_x \oplus (E_2)_x$ is given the direct sum vector space structure.

We are making the obvious identification $(E_1)_x \times (E_2)_x = (E_1)_x \oplus (E_2)_x$, using vector space structure on the right. One must, of course, prove that $E_1 \oplus E_2$ is locally trivial, but this is easy.

Definition 5.13. Let X be a topological space, and let E and F be vector bundles over X. A *homomorphism* from E to F is a continuous map $a\colon E \to F$ such that $a(E_x) \subset F_x$ for all $x \in X$, and such that $a|_{E_x}\colon E_x \to F_x$ is linear for all $x \in X$. The map a is an *isomorphism* if, in addition, it is bijective.

One can show that if a is bijective, then a^{-1} is continuous.

Definition 5.14. Let X be a compact Hausdorff space, and let E be a vector bundle over X. Then a subset $F \subset E$ is a *subbundle* of E if F is closed and is a vector bundle with the obvious restriction of the vector bundle structure on E.

The following result is known as Swan's Theorem. We will not use the parts about finitely generated projective modules, except in the description of the idea of the proof.

Theorem 5.15. For compact metric spaces X, there are natural bijections between the following sets:

(1) $V(C(X))$, the set of Murray-von Neumann equivalence classes of projections in $M_\infty(C(X))$.
(2) The set of isomorphism classes of vector bundles over X.

(3) The set of isomorphism classes of finitely generated projective modules over $C(X)$.

If $p, q \in M_\infty(A)$ are projections, E and F are corresponding vector bundles, and M and N are corresponding finitely generated projective modules, then the vector bundle corresponding to $p \oplus q$ can be taken to be $E \oplus F$, and the finitely generated projective module corresponding to $p \oplus q$ can be taken to be $M \oplus N$.

Moreover, if $p, q \in M_\infty(A)$ are projections, E and F are corresponding vector bundles, and M and N are corresponding finitely generated projective modules, then the following are equivalent:

(4) $p \precsim_{\mathrm{MvN}} q$.
(5) E is isomorphic to a subbundle of F.
(6) M is isomorphic to a submodule of N.

The proof is not hard, but is a bit long. If E is a vector bundle with structure map $q \colon E \to X$, choose as corresponding finitely generated projective module the set $\Gamma(E)$ of continuous sections of E, that is, continuous functions $s \colon X \to E$ such that $s(x) \in E_x$ for all $x \in X$. If $p \in M_n(A)$ is a projection, regard p as a (right) A-module map $A^n \to A^n$ via the usual formula for applying a matrix to a vector, and choose as corresponding finitely generated projective right module the range of p.

When is one vector bundle isomorphic to a subbundle of another? To keep the discussion simple, suppose that the fibers of E all have the same dimension, and the fibers of F all have the same dimension. Refer to these dimensions as the ranks $\operatorname{rank}(E)$ and $\operatorname{rank}(F)$. If X is a point, then E and F are just vector spaces of dimensions $\operatorname{rank}(E)$ and $\operatorname{rank}(F)$, and E is isomorphic to a subbundle of F if and only if $\operatorname{rank}(E) \leq \operatorname{rank}(F)$. The existence of nontrivial vector bundles shows that this is too good to be true in general. One might hope that it would suffice to require $\operatorname{rank}(E) < \operatorname{rank}(F)$, but this is also not good enough. The gap might need to be quite large; how large depends on the space X. The best possible theorem of its kind is as follows.

Theorem 5.16. Let X be a compact metric space. Let E and E be vector bundles over X. Suppose that for every $x \in X$, we have

$$\dim(F_x) \geq \dim(E_x) + \frac{1}{2}\big(\dim(X) - 1\big).$$

Then E is isomorphic to a subbundle of F.

The factor $\frac{1}{2}$ appears because we are using complex vector bundles. The factor $\frac{1}{2}$ in the generalized Phillips-Toms conjecture has exactly the same origin.

Theorem 5.16 can be used to get a lower bound on $\dim(X)$. It doesn't determine $\dim(X)$, even approximately. For example, if X is contractible (such as the Hilbert cube $[0,1]^{\mathbb{Z}}$), then every vector bundle over X is trivial. Thus, for vector bundles over the Hilbert cube, $\operatorname{rank}(E) \leq \operatorname{rank}(F)$ implies that E is isomorphic to a subbundle of F, even though the Hilbert cube is infinite dimensional. To extract better information, one must look at subspaces, and this requires that one use C*-algebra elements more general than projections.

Department of Mathematics, University of Oregon, Eugene OR 97403-1222, USA.