

C^* -algebras: structure and classification

David Kerr

The theory of C^* -algebras traces its origins back to the development of quantum mechanics and it has evolved into a large and highly active field of mathematics. Much of the progress over the last couple of decades has been driven by an ambitious program of classification launched by George A. Elliott in the 1980s, and just recently this project has succeeded in achieving one of its central goals in an unexpectedly dramatic fashion. This Snapshot aims to recount some of the fundamental ideas at play.

1 Structure and classification in mathematics

In the 17th and 18th centuries, philosophers like René Descartes (1596–1650) and Immanuel Kant (1724–1804) signaled a profound transformation in the way that mathematics was being understood and created, away from the idea of uncovering a universe of static truths and towards a more dynamical conception based on the notions of operation, function, and relation, and driven by an ethic of radical and prodigious invention. One need only think of the emergence of group theory, which expresses the shift in emphasis from object to relation in its starkest form and consequently has come to underpin much of modern mathematics and theoretical physics, in addition to having become an important subject in its own right. Even geometry could now be described as the study of invariance under groups of transformations like rotations and reflections. In

another direction, one can also think of the way in which analysis was electrified by the discovery by Georg Cantor (1845–1918) of a hierarchy of infinities and the construction of startling examples such as that by Karl Weierstrass (1815–1897) of a continuous but nowhere differentiable function.

These developments were intimately tied up with the crisis in logic and set theory at the end of the 19th century that ultimately paved the way for the practice of mathematics as we know it today, where invention is firmly grounded in strict and clearly formulated standards of rigour and logical precision. Together with inspiration from advances in physics such as quantum mechanics and general relativity, these new standards set the stage in the 20th century for the creation of whole new theories built upon the foundation of a few simple axioms. The axiomatic approach to defining mathematical objects led naturally to questions of *structure* and *classification* (that is, how much variety is there among the objects which model a given set of axioms, and what kind of structural differences express this variety?). These two terms taken together as a pair represent a distinctly post-19th century paradigm that continues to guide much of the mathematical activity in our own century.

Two remarkable achievements which epitomize this new paradigm are the classification of finite simple groups^[1] and Perelman’s proof of Thurston’s geometrization conjecture in topology. Another concerns the Elliott classification program in the theory of C^* -algebras, which is less well known but has been just as dramatic. In fact, a recent flourish of progress has now brought the central chapter in this classification program to a spectacular climax, capping off several decades of vigorous activity by a large international group of researchers. This is a story with many surprises and twists, and over the course of the following pages I will aim to recount some of the main ideas that have propelled it.

2 What is a C^* -algebra?

A common way of constructing a mathematical theory is to build things up from simple components. The axiomatic viewpoint reverses this picture, so that we start from a mathematical structure, whether defined in an explicit way or via some abstract properties, and try to gain some better understanding of it by breaking it down into simpler and more manageable components. To see how this applies to the world of C^* -algebras, let us first try to get a rough sense of what C^* -algebras are and where they come from, and then we will delve into some concrete examples which will motivate and illustrate many of the concepts and techniques at play in their structure and classification.

[1] For an introduction to the classification of finite simple groups see Snapshot 5/2016 *Symmetry and characters of finite groups* by Eugenio Giannelli and Jay Taylor.

C^* -algebras are historically linked to the development of quantum mechanics through the groundbreaking work of John von Neumann (1903–1957) in the late 1920s, and since the 1960s they have served as a natural mathematical framework for quantum field theory. Over the last several decades the internal study of C^* -algebras has evolved into a large and intensely active subject in its own right, with fertile connections to many other areas of mathematics, including topology, number theory, dynamical systems, graph theory, differential geometry, and random matrix theory.

As the name suggests, C^* -algebras are first of all algebras, which means that we can perform operations like addition, multiplication, and scaling by a constant, but their power derives from the fact that they have extra “analytic” structure which interacts with the algebraic operations, especially the multiplication, in a rather rigid way that leads to a rich structure theory. In practice most examples of C^* -algebras are constructed from combinatorial or topological data, or from dynamical systems like the irrational rotations that we will encounter below.

The elements of a C^* -algebra can be concretely viewed as continuous linear transformations of a particular type of vector space \mathcal{H} over the complex numbers, called a “Hilbert space”, which comes equipped with the geometric notions of angle and distance. These transformations are called *bounded operators*, and a C^* -algebra A is a collection of such operators on a fixed Hilbert space \mathcal{H} such that

1. A is an algebra in the sense that we can add, scale, and multiply (here meaning compose) operators inside of A ,
2. there is an operation $*$ on elements of A which is its own inverse (such a function is called an *involution*), and,
3. we can take limits inside of A in a very strong sense which is much like uniform convergence for a sequence of functions on \mathbb{R} , in which the limit function must be close to its approximants at every point.

The idea of being close everywhere in some parameter space is a fundamental property of C^* -algebras and usually means that some challengingly tight maneuvering has to be done when making approximations, as we will attempt to illustrate below. The “ C ” in “ C^* -algebra” refers to the fact that the algebra is *closed* under taking limits as in (iii), which means that the limit operators must also be elements of the C^* -algebra.


One can think of C^* -algebras as far-reaching generalizations of the complex numbers. Readers who have studied some complex analysis will be familiar with the principle that many of the pathologies and mysteries in the study of functions of a real variable vanish when passing to functions of a complex variable. This has to do with the special mixture of geometry and algebra in the complex numbers. For example, rotation of the complex plane by an angle θ can be implemented in an algebraic way via multiplication by $e^{i\theta} = \cos(\theta) + i \sin(\theta)$,

while the “size” or “length” of a complex number z , viewed as a vector in the plane, can be computed algebraically as the positive real number $\sqrt{\bar{z}z}$, where $\bar{z} = a - ib$ is the complex conjugate of the complex number $z = a + ib$. In fact, the involution in (ii) can be thought of as a generalization of the operation of taking the complex conjugate. It is precisely this combination of geometry, algebra, and the concept of positivity that is built into the basic fabric of a C^* -algebra, and indeed C^* -algebras also admit an abstract definition that is very reminiscent of the complex numbers (see Chapter 1 of [3] or Section 2.1 of [11]) but flexible enough to allow for a vast array of infinite-dimensional phenomena.

Finite-dimensional C^* -algebras are nothing but algebras of matrices (more precisely, they are direct sums of the form $M_{n_1} \oplus \cdots \oplus M_{n_k}$ where M_n denotes the algebra of all $n \times n$ matrices), and are thus completely understood. Such matrix algebras turn out to be the key to unlocking the structure of large classes of naturally arising infinite-dimensional C^* -algebras via a process of approximation. We will now proceed to describe how this approximation works, beginning with permutations as a simple but suggestive model.

3 Permutations as models for approximation

Consider the finite set $X = \{1, \dots, d\}$ where d is a large positive integer and let $T : X \rightarrow X$ be the map which cyclically permutes the elements in X (which we will refer to and imagine as *points*), so that $T(k) = k + 1$ for $k = 1, \dots, d - 1$ and $T(d) = 1$ (such a T is referred to as a *transformation* of X):



Suppose now that we are given another positive integer n , possibly much smaller than d , and let us suppose for simplicity that n divides d , that is, $d = kn$ for some integer k . Let $T' : X \rightarrow X$ be the map which for each $j = 1, \dots, k$ cyclically permutes the subset $\{jn + 1, jn + 2, \dots, (j + 1)n\}$, as illustrated below in the case $d = 20$ and $n = 4$:



Then T' is a union of k cyclic permutations of sets of size n , and it agrees with T everywhere except at the points $n, 2n, \dots, kn$, of which there are k . Thus T' disagrees with T on a subset of X of proportional size $k/d = 1/n$, which will be very small if n is very large.

Now if n does not necessarily happen to divide d , then one can still try to define a map T' as before by taking cyclic permutations of consecutive segments of length n , but there will be some remaining segment at the end containing at most $n - 1$ points, which we can then permute amongst themselves in order to round out the definition of our updated T' , as illustrated below when $d = 18$ and $n = 5$:



Then the number of points at which T' disagrees with T will be no greater than the (possibly nonintegral) quantity $1/n + (n - 1)/d$, the first term coming from the permutations of the segments of length n and the second from the remainder. We therefore see that if n is very large and d is very, very large (that is, large compared with n) then T' and T are approximately equal in the sense that they agree at all but a proportionally small number of points. This kind of play at two or more different scales is a common theme in C^* -algebra theory and in many other areas of analysis.

The point is that we can now view T approximately via T' as consisting of a single permutation of a relatively small set with size n not depending on d (for a given tolerance in the approximation) but repeated with multiplicity k . This suggests that we might be able to keep n and the tolerance in the approximation fixed but replace $\{1, \dots, d\}$ with an *infinite* set. This idea lies at the root of matrix approximation in C^* -algebra theory, as we will see, but before jumping to that let us illustrate with a basic but very important example of a transformation of an infinite set.

Let \mathbb{T} be the unit circle in \mathbb{R}^2 and θ an irrational number between 0 and 1, and let T be the transformation of \mathbb{T} which rotates by the angle $2\pi\theta$ radians. Given

an $\varepsilon > 0$ and a natural number n , one can show that there is a finite disjoint union of arcs, which we'll call C , such that the images $TC, T^2C, \dots, T^n C$ of C under n iterations of T (that is, the images of C under rotation by the angles $2\pi\theta, 4\pi\theta, \dots, n\pi\theta$) are pairwise disjoint and the lengths of the arcs making up the complement of these images sum to less than ε . This is a bit tricky to do by hand (the reader might wish to try, especially in the case that n is large compared to $1/\theta$), but it is a straightforward consequence of a general result known as the *Rokhlin lemma*. We thereby obtain much the same kind of approximation as before with prescribed period n . Note however that this time we cannot exactly cycle back to the beginning, since the images of C under rotation by $2\pi\theta$ can never coincide with C itself in view of the irrationality of θ , although $T^n C$ will mostly overlap with C .

The rotation T generates a C^* -algebra called an *irrational rotation algebra*, which is constructed by taking the algebra $C(\mathbb{T})$ of continuous complex-valued functions on \mathbb{T} and enlarging it by the addition of an abstract element u which implements the rotation via conjugation, so that ufu^{-1} for a function $f \in C(\mathbb{T})$ is equal to the composition of f with the rotation of \mathbb{T} by $2\pi\theta$. Some technical work is required to turn this into a fully-fledged C^* -algebra, but the construction, called the “crossed product”, is a natural one that also applies to actions of general groups. Irrational rotation algebras are ubiquitous prototypes in the field of noncommutative geometry [2] and were also an early impetus for the development of C^* -algebra classification theory, as we will discuss below.

4 Approximation in matrix algebras

Let us return now to the permutation T of $\{1, \dots, d\}$ and its approximant T' as in the first paragraph of the previous section. We can translate this kind of approximation into C^* -algebraic language by expressing the maps T and T' as what are called permutation matrices. We replace our set $\{1, \dots, d\}$ with the d -dimensional complex vector space \mathbb{C}^d , with an element $j \in \{1, \dots, d\}$ corresponding to the standard basis vector $e_j = (0, \dots, 0, 1, 0, \dots, 0)$ where the 1 occurs in the j th coordinate. For every positive integer m we define R_m to be the $m \times m$ permutation matrix

$$\begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}.$$

Note that the action of this matrix is to send the vector (x_1, \dots, x_n) to $(x_n, x_1, \dots, x_{n-1})$. Then T corresponds to the matrix R_d , which we'll write as

A_T , while T' corresponds to the block diagonal permutation matrix

$$A_{T'} = \begin{bmatrix} R_n & 0 & \cdots & 0 \\ 0 & R_n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & R_n \end{bmatrix}.$$

Even though we've passed to matrices, we can still accommodate our original set of points in this picture by identifying the integer j with the matrix whose entries are all 0 except at the j th position down the diagonal, where the entry is 1. In this picture the diagonal matrices can be viewed as functions on X , so that we are effectively replacing points with functions. One can now check that if D is a diagonal matrix then $A_T D A_T^{-1}$ is again a diagonal matrix whose entries are the same as those of D except that they have been permuted according to T . This is similarly the case for $A_{T'}$ and T' . We can thereby incorporate *all* information about our space X and the transformations T and T' into a common framework, namely the algebra of $d \times d$ matrices, which we denote as before by M_d . One can verify that the subalgebra of M_d generated by A_T and the diagonal matrices is all of M_d , while the subalgebra \mathcal{A} of M_d generated by $A_{T'}$ and the diagonal matrices is much smaller, consisting of all matrices which decompose into diagonal blocks of size $n \times n$ (assuming for simplicity that n divides d). This subalgebra \mathcal{A} contains the diagonal matrices, and it also approximately contains the matrix A_T coming from the original cyclic permutation T , since it contains $A_{T'}$, which is approximately equal to A_T .

If one only needs \mathcal{A} to approximately contain a certain selection A_1, \dots, A_k of diagonal matrices in addition to A_T , then one can typically replace \mathcal{A} with a smaller subalgebra given by what is called an “embedding”, where a certain number of copies of each of them are distributed as blocks down the diagonal in M_d (in order to express this map in such a clean form, it is necessary to do some coordinate shuffling in \mathbb{C}^d , which amounts to conjugating all of the matrices in M_d by a fixed permutation matrix). What is new here is the phenomenon of multiplicity, whose simplest illustration occurs in the case $k = 1$, $n = 2$, and $d = 4$, where this embedding is given by

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \mapsto \quad \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} = \begin{bmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & c & d \end{bmatrix},$$

and if one shuffles coordinates then one can also express this by


$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} aI & bI \\ cI & dI \end{bmatrix} = \begin{bmatrix} a & 0 & b & 0 \\ 0 & a & 0 & b \\ c & 0 & d & 0 \\ 0 & c & 0 & d \end{bmatrix}$$

where I is the 2×2 identity matrix $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. One can start to compose these kinds of embeddings in an iterative manner, increasing the matrix sizes at every stage, and even continue the process indefinitely in order to obtain an infinite-dimensional “limit” object known as an *approximately finite-dimensional C*-algebra*, or *AF-algebra* for short. In the 1970s a whole calculus was developed to classify the different algebras that can arise in this manner, first by Bratteli in terms of combinatorial data attached to the successive embeddings and then by Elliott in the language of groups (K -theory). Elliott’s classification of AF-algebras set the stage for the program of classifying much broader classes of C*-algebras defined in a more abstract way by a less rigid matrix approximation property. This project came quite early on to include the irrational rotation algebras, which can be shown not to be AF-algebras. In order to describe and appreciate how the classification program has ultimately panned out, we will need to pursue our discussion of matrices a little further.


5 C*-algebra approximation and Berg’s technique

Although our discussion above veered into AF-algebras, it turns out that the kind of approximation that we have been discussing, beginning with permutations, is in fact the wrong one for C*-algebras. Here we can appreciate the essential dividing line between C*-algebras and their cousins the von Neumann algebras, which are built around approximations which allow us to completely ignore small parts of the space or of the matrix algebra. The more stringent demands of C*-approximation force us to leverage additional symmetry within the algebra itself and not merely in the dynamics around which the algebra may be built. Changing a permutation at a single point turns out to be far too drastic an operation in the C*-algebra world.


A more subtle trick that works very well to give these “extra good” matrix approximations for the irrational rotation algebras is Berg’s technique, which we will try to give a flavour of using a few pictures. Think back again to our cyclic permutation T of the set $\{1, \dots, d\}$, which we now illustrate in two rows:



If we wish to chop up this into two smaller permutations it suffices to swap the images of a pair of points as follows:



As mentioned, this is already too drastic for the purposes of C^* -approximation, so we do it more gradually by introducing extra “ghost” points that allow for a gentle interchange of the two streams:



It is precisely within the matrix algebra M_d that we can make these points

concrete by using 2×2 rotation matrices of the form

$$e^{-i\theta} \begin{bmatrix} \cos \theta & i \sin \theta \\ i \sin \theta & \cos \theta \end{bmatrix}$$

with θ varying in discrete but small increments in the range from 0 to $\pi/2$, which begins with the identity matrix $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and ends up at the coordinate flip $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Use of the complex number i here ensures that these matrices will preserve lengths when acting on vectors in \mathbb{C}^2 . The pairs of vertically aligned points in the middle portion of the diagram, viewed as standard basis vectors in \mathbb{C}^d on which our matrices act, each get “twisted” inside the amplified space \mathbb{C}^d , with the amount of twisting gradually increasing as we move along the streams (one should view the red part of the diagram three-dimensionally as a 180-degree twist in a double helix, so that the two lines do not actually cross but rather rotate around each other). Each intermediate twisted “ghost” pair is a nontrivial linear combination of the standard basis vectors corresponding to the original pair. The two streams then interchange at the end of the twisting, producing a pair of cyclic permutations.

Now if two angles θ_1 and θ_2 are close together then the matrices

$$e^{-i\theta_1} \begin{bmatrix} \cos \theta_1 & i \sin \theta_1 \\ i \sin \theta_1 & \cos \theta_1 \end{bmatrix} \quad \text{and} \quad e^{-i\theta_2} \begin{bmatrix} \cos \theta_2 & i \sin \theta_2 \\ i \sin \theta_2 & \cos \theta_2 \end{bmatrix}$$

will be close together entry by entry, without exception, which is consistent with the requirements of C^* -algebra approximation. The above procedure of lining up the two streams and slowly twisting thus softens the brutal chopping operation we began with to something much more C^* -friendly.

While Berg’s technique and related methods have been successfully employed to unravel the structure of many C^* -algebras arising from transformations, like the irrational rotation algebras (see [4] and Chapter VI of [3]) and the crossed products of minimal transformations of the Cantor set [13], their use becomes severely limited if the algebra is built up not from a single transformation but from many transformations, because in this case there are too many “directions” to be able to pair off points and twist them as before.

One alternative approach that has turned out to be surprisingly effective was initiated by Winter in the late 1990s and turns on the concept of *dimensionality*, not in the sense that some C^* -algebras are finite-dimensional and others infinite-dimensional as vector spaces or algebras, but rather in a more inherently topological sense that is in line with the notion that an interval is one-dimensional and a sphere is two-dimensional. Some overlapping is still permitted (and usually even required) as in Berg’s technique, but one no longer needs to resolve these overlaps by any kind of twisting, and in addition the matrix models are now allowed to be distorted. There are some variations in the technical

implementation of this idea, but the most important for our present discussion is due to Winter and Zacharias, who referred to the optimal bound in the amount of “overlap” as the *nuclear dimension* of the C^* -algebra [19].

Nuclear dimension is a highly flexible tool that applies in a tractable way to large classes of examples. The drawback however is that it does not give us as direct an access to the structure of the C^* -algebra due to the possible distortion in and interference between the matrix models that appear in its definition. The point now is to rely instead on a general classification theorem to do all of the heavy lifting in producing a complete structural picture of the C^* -algebra based on genuine matricial approximation, with nuclear dimension being precisely the device that will clear the path leading us there.

6 Nuclear dimension

As we discussed in Section 4, given an even positive integer d the embedding $A \mapsto \text{diag}(A, \dots, A)$ from M_2 into M_d can alternatively be described, by shuffling coordinates, as

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} aI & bI \\ cI & dI \end{bmatrix}$$

where I is the $(d/2) \times (d/2)$ identity matrix. Now if we replace I by an arbitrary $(d/2) \times (d/2)$ matrix H then

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} aH & bH \\ cH & dH \end{bmatrix}$$

still defines an embedding $\varphi : M_2 \rightarrow M_d$, but it will no longer be multiplicative in general, that is, we might have $\varphi(AB) \neq \varphi(A)\varphi(B)$ for some A and B , which means that the image might no longer be a matrix algebra. However, a vestige of multiplicativity remains: one can check that $\varphi(A)\varphi(B) = 0$ whenever $AB = 0$, in which case we call φ an *order-zero* map. Winter had the insight that for a great many purposes it enough to perform approximations using finite collections of such order-zero maps, with the images of these maps being allowed to overlap or “interfere” with each other in any way as long as there is a uniform bound on the number of maps which are used at every scale of approximation. This leads to the notion of *nuclear dimension*. For the record, we give the precise definition below. More details and an explanation of all of the terms can be found in [19]. What is important to note is that the definition is *local* in nature, which means that one only needs to check the condition on finitely many elements in the algebra at a time.

Definition 6.1. *The nuclear dimension of a C^* -algebra A is the least integer $d \geq 0$ such that for every finite set $\Omega \subseteq A$ and $\varepsilon > 0$ there are finite-dimensional C^* -algebras B_0, \dots, B_d and completely positive linear maps*

$$A \xrightarrow{\varphi} B_0 \oplus \dots \oplus B_d \xrightarrow{\psi} A$$

such that φ is contractive, $\psi|_{B_i}$ is contractive and order zero for each $i = 0, \dots, d$, and

$$\|\psi \circ \varphi(a) - a\| < \varepsilon$$

for every $a \in \Omega$. If no such d exists then we define the nuclear dimension to be ∞ .


The definition of nuclear dimension turns out to be compatible enough with dynamical phenomena that one can compute the value, or at least derive useful estimates, for many crossed products (see [15, 9]). The irrational rotation algebras can be shown for example to have nuclear dimension exactly one. A breakthrough by Szabó in [15] showed that the crossed products of a large class of \mathbb{Z}^n -actions have finite nuclear dimension, taking us well outside the domain where Berg's technique is applicable. Szabó's results have since been extended in [16].

The point now is that the periodic approximations that we needed before to construct our matrix algebras can still be used to construct the matrix approximations in the definition of nuclear dimension (note that the algebras B_0, \dots, B_d , being finite-dimensional, are all direct sums of matrix algebras) but they no longer need to cycle back on themselves, even approximately. However, these segments must overlap with each other to some degree in order to be able to produce the matrix approximations B_0, \dots, B_d , and the extent of the overlapping is what controls the parameter d . These are wonderfully supple conditions compared to the usual rigid demands of C^* -algebra theory, and since they don't involve in any essential way the kind of past-and-future directionality that we relied on to carry out Berg's technique, they end up applying to actions of a wide variety of groups. For more details see Section 8 of [19] or Section 8 of [9].

Dynamics aside, nuclear dimension is also closely related to the ordinary notion of dimension in topology. Consider for example the unit interval $[0, 1]$, which we can cover using two collections of disks with diameter as small as we wish such that the disks in each subcollection are pairwise disjoint, as illustrated below by using different colours (empty and hatched) to indicate the subcollection to which each disk belongs:



The d -cube $[0, 1]^d$ can be similarly covered by balls but now grouped into $d + 1$ subcollections each of which is disjoint, as illustrated below in the case $d = 2$:



The (*covering*) *dimension* is then defined to be one less than the smallest number of subcollections each of which is disjoint. We thus allow balls to overlap, just not within each subcollection. A subcollection being disjoint is the analogue of a map being order-zero, and in fact this correspondence goes beyond mere analogy as one can use this picture of covering dimension to show that the C^* -algebra of continuous functions on $[0, 1]^d$ has nuclear dimension d . This requires the notion of partition of unity, which allows one to build functions over individual balls in a coherent fashion.

7 Structure and classification of C^* -algebras

We will now finally formulate the classification theorem. By an *invariant* we mean a collection of data which is associated to a mathematical object (here, a C^* -algebra) in some structural way. A prototypical example is the dimension

of a vector space. One can picture an invariant as a kind of skeleton which provides us enough information to be able to completely flesh out the original object.

To say that a class of C^* -algebras is *classified* by an invariant means that two C^* -algebras in the class are equivalent (technically, “isomorphic”) if and only if they have the same invariant. Of course, for such a classification to have any utility the invariant should be much simpler than the C^* -algebra itself. Here we are working with the *Elliott invariant*, about which we will not say anything in detail except that it consists of ordered K -theory paired with traces. Let us now state the theorem and then follow up with some explanation. The theorem is a culmination of several decades of effort by many people and was clinched in recent work of Gong–Lin–Niu [8], Elliott–Gong–Lin–Niu [5], and Tikuisis–White–Winter [17]. Although this Snapshot has been implicitly concerned with C^* -algebras that are *tracial*, which is an incompressibility property satisfied by irrational rotation algebras and many other crossed products and is the case targeted by the papers just cited, the statement below also incorporates an earlier classification of Kirchberg and Phillips from the 1990s in the “compressible” purely infinite case.

Theorem 7.1. *The class of infinite-dimensional simple separable unital C^* -algebras satisfying the UCT and having finite nuclear dimension is classified by the Elliott invariant.*

The conditions of simplicity, separability, and unitalness are all ambient assumptions. The first means that the C^* -algebra is indecomposable (it is a kind of “atom” in the theory), the second that the C^* -algebra is sufficiently set-theoretically simple, and the third that the C^* -algebra behaves like a compact topological space, so that the kinds of issues typically associated with unboundedness do not arise. The UCT (universal coefficient theorem) is a technical assumption that allows one to leverage a version of K -theory for pairs of C^* -algebras and might very well hold automatically, and at least is known to hold in all examples of interest. Thus the only truly operative hypothesis is finite nuclear dimension, and it is a necessary hypothesis since it can fail to hold in many cases, in particular for some crossed products of transformations of infinite-dimensional spaces [18, 6] (these crossed products are nevertheless *nuclear* in the sense that matrix approximations as in the definition of nuclear dimension always exist, just not with a uniform bound on the d ; crossed products of actions of free groups, on the other hand, frequently fail even to be nuclear).

The conciseness of the theorem statement obscures the fact that there is a whole cast of characters standing behind it, including the concepts of \mathcal{L} -stability, strict comparison, decomposition rank, and tracial rank, all of which play a significant technical role in the proof. That nuclear dimension would in fact play the starring role was not even clear until shortly before the result was established.

I might also mention that it is a prominent and still not fully resolved conjecture of Toms and Winter that for simple separable infinite-dimensional nuclear C^* -algebras the properties of finite nuclear dimension, \mathcal{L} -stability, and strict comparison are all equivalent.

It is worth noting that the contributions of Gong–Lin–Niu and Elliott–Gong–Lin–Niu on the one hand and of Tikuisis–White–Winter on the other are of a completely different nature. The latter was a singular and surprising advance that showed that a large class of C^* -algebras, much larger than previously known, satisfy an external finite-dimensional approximation property called quasidiagonality, which can be extremely difficult to verify in individual examples. The former on the other hand was an extraordinarily labour-intensive achievement (totalling some 300 pages and building on thousands of pages of earlier work) that established the classification on the basis of hypotheses that could be pared down to finite nuclear dimension once one knew the quasidiagonality result. It is absolutely remarkable that these two pieces of the puzzle, disparate in technical content but precisely complementary in their outcome, came together at about just the same moment in 2015. Another crucial earlier contribution that should be mentioned here is a result of Sato, White, and Winter [14] that clarified the relation between quasidiagonality and nuclear dimension.

Returning now to the statement of the classification theorem itself, the question that immediately arises is: how do we compute the Elliott invariant? It can be done in some special cases (see for example the beautiful Giordano–Putnam–Skau theory of crossed products of minimal transformations of the Cantor set [7]) but in general is a difficult business. One line of attack in the case of group C^* -algebras and general types of crossed products is mapped out by the Baum–Connes conjecture, which has spawned an extensive theory with fascinating connections to topology and coarse geometry [12, 9].

Notwithstanding the problem of actually computing the invariant, the theorem still provides us an extraordinary dividend in structural understanding. Just as the spaces \mathbb{R}^n for $n = 1, 2, \dots$ are models for all of the possible finite dimensions among real vector spaces, it is known that, in the tracial case, the Elliott invariant is already exhausted among C^* -algebras which are limits of nested sequences of algebras of matrices whose entries are not necessarily complex numbers (like in the case of AF-algebras) but, more generally, continuous functions on some space like a torus or a sphere. One concludes that all of the tracial C^* -algebras appearing in the theorem can be expressed in this rigid asymptotic matricial form, all as a consequence of the relatively soft and amenable property of finite nuclear dimension. Such a stark structural picture is simply impossible to obtain in a bare-hands way for a great many examples, including crossed products of various kinds of group actions. This is a supreme

example of how the modern paradigms of structure and classification are, in the end, two sides of the same coin.

What next? The world of C^* -algebras extends to horizons far beyond finite nuclear dimension, and this wider C^* -universe remains largely uncharted. For a glimpse of what lurks at the outer boundaries of the subject the reader may wish to take a look at the striking recent work of Breuillard, Kalantar, Kennedy, and Ozawa on the simplicity of group C^* -algebras [10, 1].

Image credits

All figures created by the author.

References

- [1] E. Breuillard, M. Kalantar, M. Kennedy, and N. Ozawa, *C^* -simplicity and the unique trace property for discrete groups*, Publications mathématiques de l’IHÉS **126** (2017), 35–71.
- [2] A. Connes, *Noncommutative Geometry*, Academic Press, Inc., 1994.
- [3] K. R. Davidson, *C^* -Algebras by Example*, Fields Institute Monographs, vol. 6, American Mathematical Society, 1996.
- [4] G. A. Elliott and D. E. Evans, *The structure of the irrational rotation C^* -algebra*, Annals of Mathematics **2** (1993), 477–501.
- [5] G. Elliott, G. Gong, H. Lin, and Z. Niu, *On the classification of simple amenable C^* -algebras with finite decomposition rank, ii.*, Arxiv preprint arXiv:1507.03437.
- [6] J. Giol and D. Kerr, *Subshifts and perforation*, Journal für die reine und angewandte Mathematik **639** (2010), 107–119.
- [7] T. Giordano, I. F. Putnam, and C. F. Skau, *Topological orbit equivalence and C^* -crossed products*, Journal für die reine und angewandte Mathematik **469** (1995), 51–111.
- [8] G. Gong, H. Lin, and Z. Niu, *Classification of finite simple amenable \mathcal{L} -stable C^* -algebras*, Arxiv preprint arXiv:1501.00135.
- [9] E. Guentner, R. Willett, and G. Yu, *Dynamic asymptotic dimension: relation to dynamics, topology, coarse geometry, and C^* -algebras*, Mathematische Annalen **367** (2017), 785–829.

- [10] M. Kalantar and M. Kennedy, *Boundaries of reduced C^* -algebras of discrete groups*, Journal für die reine und angewandte Mathematik **727** (2017), 247–267.
- [11] G. J. Murphy, *C^* -Algebras and Operator Theory*, Academic Press, Inc., 1990.
- [12] P. W. Nowak and G. Yu, *Large Scale Geometry*, EMS Textbooks in Mathematics, European Mathematical Society, 2012.
- [13] I. F. Putnam, *On the topological stable rank of certain transformation group C^* -algebras*, Ergodic Theory and Dynamical Systems **10** (1990), 197–207.
- [14] Y. Sato, S. White, and W. Winter, *Nuclear dimension and \mathcal{L} -stability*, Inventiones mathematicae **202** (2015), 893–921.
- [15] G. Szabó, *The Rokhlin dimension of topological \mathbb{Z}^m -actions*, Proceedings of the London Mathematical Society **110** (2015), 673–694.
- [16] G. Szabó, J. Zacharias, and J. Wu, *Rokhlin dimension for actions of residually finite groups*, Ergodic Theory and Dynamical Systems **39** (2017), 2248–2304.
- [17] A. Tikuisis, S. White, and W. Winter, *Quasidiagonality of nuclear C^* -algebras*, Annals of Mathematics **185** (2017), 229–284.
- [18] A. S. Toms, *On the classification problem for nuclear C^* -algebras*, Annals of Mathematics **167** (2008), 1029–1044.
- [19] W. Winter and J. Zacharias, *The nuclear dimension of C^* -algebras*, Advances in Mathematics **224** (2010), 461–498.

David Kerr is a professor of mathematics
at the University of Münster.

License
Creative Commons BY-SA 4.0

Mathematical subjects
Analysis

DOI
10.14760/SNAP-2021-002-EN

Connections to other fields
Physics

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

ISSN 2626-1995

Junior Editor
Sara Munday
junior-editors@mfo.de

Senior Editor
Sophia Jahns
senior-editor@mfo.de

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9–11
77709 Oberwolfach
Germany

Director
Gerhard Huisken



Mathematisches
Forschungsinstitut
Oberwolfach



IMAGINARY
open mathematics