Friendly introduction to the concepts of noncommutative geometry

Notes pour le séminaire "Philosophie et Physique", Laboratoire SPHERE, Paris 7

Fabien Besnard, EPF

May 24, 2013

1 Introduction

Noncommutative geometry is such an impressive field that giving an introductory talk to it is quite daunting. Many threads lead to it and it has ramifications in many branches of mathematics, most of them I have meagre knowledge about. So instead of rushing through so many subject, counting on speed to conceal my incompetence, I will follow only one thread, the one that starts with quantum physics, where it all began, and comes back to it. Moreover, I will keep as much a low-tech profile as possible, so that everything should be easy to follow with an undergrad level in math. In fact, if you have ever diagonalized a matrix, you should be able to keep up. (I'm overselling a little bit, since I will need bounded operators, but if you don't know what a bounded operator is, you can imagine it to be an infinite matrix and it should be OK. You will have to know a bit of topology too, but very basic.)

Of course there are many things that I will leave out. I will feel no remorse for sweeping technicalities under the rug, since you can learn about them in the texts given in reference. On the other hand it is certainly a pity that I will not say a word about beautiful and important subjects such as K-theory, cyclic co/homology, and the deep noncommutative generalization of the Atiyah-Singer index theorem. To give an idea of the crime, let me say that these issues cover at least half of the [Red Book], and are certainly indispensable to understand the big picture. If you want to learn about this you are urged to consult the nice introductory texts [Kh 04].

A word of caution before you proceed : even though our trip will start with quantum mechanics and end with the standard model of particle physics, noncommutative geometry is not quantum in the same way quantum mechanics is. The reason is simple : noncommutative geometry is not about mechanics, it's about space. So its "quantumness" applies inside the configuration space, not between configuration and impulsion variables like in quantum mechanics. Take a look at figures 4 and 5. They illustrate the difference between commutative and noncommutative geometry. The same figures could be used to illustrate the difference between classical and quantum mechanics, but only if we understand that the space in this case is the *phase space*. Of course we may hope that one day the whole of physics will be reduced to some kind of noncommutative geometry, so that the distinction will disappear. But the for the moment it is still a dream...

2 c-numbers and q-numbers

We know that for Dirac, going from classical to quantum mechanics basically amounted to replace c-numbers by q-numbers, that is commuting numbers with non-commuting numbers. (The c can mean commuting, classical or complex !)

But what is a q-number ? This is essentially a matrix, either of finite or infinite dimension. To give a rigorous meaning to the latter case, we must interpret matrices as representing operators on Hilbert spaces, of either finite or infinite dimension.

So what is the analogy between a complex number and an (bounded) operator ? Complex numbers have three operations defined on them : addition, multiplication, and complex conjugation (which we denote by $*$). The following properties hold for these operations :

1. $(\mathbb{C}, +)$ is an additive group,
2. $1.a = a.1 = a$,
3. $(ab)c = a(bc)$,
4. $a(b+c) = ab + ac$, $(a+b)c = ab + bc$,
5. $(a^*)^* = a$ (complex conjugation is involutive),
6. $(a+b)^* = a^* + b^*$,
7. $(ab)^* = b^* a^*$.

Of course they satisfy another property, which is commutativity : $ab = ba$. Now we see at once that operators do satisfy all the above properties if we interpret complex conjugation as adjunction. They also satisfy another property, which we did not see in the case of $\mathbb{C}$ because it was trivial : they form a linear space over $\mathbb{C}$, and the multiplication with a complex number satisfies the following rules :

- $\lambda.(ab) = (\lambda.a)b = a(\lambda.b)$,
- $(\lambda.a)^* = \bar{\lambda} a^*$.

We sum up all these properties by saying that both $\mathbb{C}$ and $B(H)$ are examples of unital $*$-algebras over $\mathbb{C}$. For the moment, this all seem to be very formal, but it starts to be interesting when we single out several kinds of c-numbers which are of particular interest :

- a is a real number iff $a^* = a$,
- a is nonnegative iff a the square of a real number. Alternatively, $a = b^*b$ for some $b \in \mathbb{C}$.
- a is unitary iff $aa^* = a^*a = 1$.

Now what is interesting is that these notions continue to make sense in $B(H)$, and that, moreover, the spectrum of real operator is real, the spectrum of nonnegative operator is nonnegative, and the spectrum of a unitary operator is a subset of the unit circle. (We think that seeing things this way is enlightening even for a beginner in linear algebra, think for a example of the parallel that can be drawn between the polar decomposition of a complex number and a matrix).

3 c-functions, q-functions, and C^*-algebras

However, the analogy so far is a bit misleading. The reason is that c-numbers have another crucial property which operators lack completely : they are all invertible (except 0). This is why matrices are not really analogs of complex numbers, they are analogs of complex *functions*.

Indeed, the set $\mathcal{F}(X, \mathbb{C})$ of functions from an arbitrary space X to $\mathbb{C}$ comes equipped with a natural structure of commutative unital $*$-algebra, with pointwise laws and unit $\mathbf{1}$, the constant function with value 1. Now the most important fact is that there is an algebraic way to express that some complex number λ is a value taken by the function a : it is so iff $a - \lambda.\mathbf{1}$ is not invertible. This is so important that we put it in a box :

$$\forall a \in \mathcal{F}(X, \mathbb{C}),\ \sigma(a) = a(X)$$

The interpretation of the eigenvalues of an operator as the values taken by a function is particularly obvious if we take $X = \{1, \dots, n\}$ a finite set, in which case $\mathcal{F}(X, \mathbb{C})$ can be identified with the diagonal matrices in $\mathrm{M}_n(\mathbb{C})$.

There is one last crucial ingredient which is still missing on our algebras of c/q-numbers : some notion of topology (on the algebra as well as on the space itself, both being related). The technically simpler setting is that of continuous functions on compact Hausdorff spaces, so from now on we consider a compact space X, and we want to "quantize" the algebra $\mathcal{C}(X) = \mathcal{C}(X, \mathbb{C})$ of continuous complex functions on X, that is to say we want to do the following :

- Identify which extra structure the algebra $\mathcal{C}(X)$ gains thanks to the the topology on X.

- Express algebraically the fact that $\mathcal{C}(X)$ is an algebra of continous functions on a compact hausdorff space.
- Suppress the commutativity and define a structure of "continuous q-functions".

The extra-structure is easy to identify : it is the uniform norm $\|a\| = \sup_X |a(x)|$. This norm posseses certain obvious properties with respect to the product and involution :

1. $\|ab\| \leq \|a\|\|b\|$,
2. $\|a^*\| = \|a\|$,
3. $\|aa^*\| = \|a\|\|a^*\|$.

Moreover $\mathcal{C}(X)$ is complete with respect to this norm. Now the uniform norm of a has an obvious algebraic meaning in the case of $\mathcal{C}(X)$: it is the largest eigenvalue of a in modulus, i.e. $\|a\| = \rho(a)$, the spectral radius. Clearly the spectral radius continues to make sense for instance in $B(H)$, but it is not a norm anymore. However a little yoga can cure this inconvenience. We can also write $\|a\| = \sqrt{\rho(|a|^2)} = \sqrt{\rho(a^*a)}$. This time the mapping $a \mapsto \sqrt{\rho(a^*a)}$ is a norm on $B(H)$, which coincides with the operator norm. Moreover it also satisfies the three properties above and $B(H)$ is complete with respect to this norm.

All this is fine but how can we be sure we haven't forgotten an important property ? This can be done by proving a reconstruction result. In order to state it, let us give a definition.

Definition 1 *A (unital) C^*-algebra A is :*

1. *a (unital) algebra,*
2. *with a map $* : A \to A$ such that*
   (a) *it is anti-linear,*
   (b) *it is an involution,*
   (c) $(ab)^* = b^*a^*$.
3. *with a norm such that*
   (a) $\|ab\| \leq \|a\|\|b\|$,
   (b) $\|aa^*\| = \|a\|\|a^*\|$,
   (c) *A is complete with respect to $\| \, \|$.*

Note that we have removed the axiom that $\|a\| = \|a^*\|$. The reason is that it can be derived from the other ones.

Now we have already 3 examples of C^*-algebras :

- $\mathcal{C}(X)$, which is commutative,

- $M_n(\mathbb{C})$,
- more generally $B(H)$, the bounded operators on a Hilbert space.

With a little thought we see that a norm closed $*$-subalgebra of $B(H)$ will also be automatically a C^*-algebra. The famous theorems of Gelfand and Naimark precisely show that :

- There is no other commutative C^*-algebras besides the $\mathcal{C}(X)$ (Commutative GN theorem).
- There is no other C^*-algebras besides the C^*-subalgebras of $B(H)$ (Noncommutative GN theorem).

The two theorems of Gelfand and Naimark give a strong rationale for the motto "operators are like noncommutative functions", and in fact it is even a little more precise : bounded operators are like noncommutative continuous functions. Moreover we can see a deep parallel with quantum mechanics, which completely justifies Dirac's point of view (with the restriction that we deal with functions, not just numbers). This is not surprising then that C^*-algebras have often been proposed as the correct setting for quantum mechanics and even quantum field theory (even though the simplest examples of quantum systems already involve unbounded operators, there are ways to deal with this problem). (By the way the C of C^*-algebra comes from "continuous".)

We will have to say a little more later about the two theorems of Gelfand and Naimark and the way in which they are proven. But for the moment let us continue to explore the analogy between operators and functions.

When $a : X \longrightarrow \mathbb{C}$ is a continuous function, and $f : U \longrightarrow \mathbb{C}$ is a continuous function defined on a subset U of $\mathbb{C}$ which contains $a(X) = \sigma(a)$, then $f \circ a$ is defined and belongs to $\mathcal{F}(X, \mathbb{C})$. Now when a belongs to an arbitrary unital C^*-algebra, and provided $aa^* = a^*a$ (a is *normal*), one can also "compose " a with any continuous f defined on $\sigma(a)$ (in fact this is an application of the commutative GN theorem), that is to say one can uniquely define $f(a)$, and one has $\sigma(f(a)) = f(\sigma(a))$. This is called *continuous functional calculus*. When we restrict to the case of a hermitian operator in $B(H)$, one can even define $f(a)$ for an arbitrary Borel function on $\sigma(a)$.

A philosophical digression: The commutative Gelfand-Naimark theorem has the precise technical meaning of a dual equivalence of categories. It means that all that can be said using the language of compact Hausdorff spaces and continuous functions can be exactly translated into the algebraic language of commutative C^*-algebras and $*$-morphisms. There are other such duality theorems in mathematics, bearing the names of Stone, Pontryagin, Tannaka-Krein, etc. Now suppose for the sake of the argument that someone come up with a "theory of everything" in which the fundamental entities are compact spaces. We could immediately translate it into a different, but strictly equivalent theory using commutative C^*-algebras instead. Wouldn't this pose a serious problem for the issue of ontology ? There would be no unique answer to the famous question "what is the world made of ?" We might draw the conclusion that either this question will never receive a unique answer, or that a theory of everything should admit no duality. Or maybe that an answer can be given in terms of equivalence classes of categories, whatever this means...

4 Spaces associated with a C^*-algebra

According to what we said in section 3 it is possible to recover the space X given the commutative C^*-algebra $\mathcal{C}(X)$. How can we do ? The trick is as simple as it is deep : it is to write

$$f(x) = x(f)$$

It means that we should not see points as things functions act upon, but things which act on functions. When we see in this way a point x as acting on functions, it is customary to write it $\hat{x}$.

The very definitions of the involution and of the laws of addition and multiplication now read :

$$\hat{x}(\alpha f + \beta g) = \alpha\hat{x}(f) + \beta\hat{x}(g); \qquad \hat{x}(fg) = \hat{x}(f)\hat{x}(g); \qquad \hat{x}(f^*) = \hat{x}(f)^*$$

It means that $\hat{x}$ is a morphism of $*$-algebra from $\mathcal{C}(X)$ to $\mathbb{C}$. Since $\mathcal{C}(X)$ contains the constant function $\mathbf{1}$, we see that it is also unital ($\hat{x}(\mathbf{1}) = 1$), and in particular it is non-zero (this is equivalent thanks to the identity $\hat{x}(\mathbf{1}^2) = \hat{x}(\mathbf{1})^2$).

Non-zero morphisms of $*$-algebra with values in $\mathbb{C}$ are called *characters*. If we write $\mathcal{X}(\mathcal{C}(X))$ the set of characters of $\mathcal{C}(X)$, we have just defined a map $x \mapsto \hat{x}$ from X to $\mathcal{X}(\mathcal{C}(X))$. It turns out that this map is bijective. It is injective because continuous functions separate the points of X. Let us quickly explain why it is surjective. Take a character χ on $\mathcal{C}(X)$. Then its kernel is an ideal of $\mathcal{C}(X)$, and since the image of χ is one-dimensional, it must be a maximal ideal. Now it is a well-known (and interesting !) exercice in analysis that every maximal ideal of $\mathcal{C}(X)$ is of the form $I_x = \{f | f(x) = 0\}$ for some $x \in X$. This x is easily seen to be the antecedent of χ.

Thus we can identify X with the character space of $\mathcal{C}(X)$ as sets, but in fact we can also do it as topological spaces if we equip $\mathcal{X}(\mathcal{C}(X))$ with the topology of pointwise convergence.

Now that we have characterized algebraically the notion of points, we can apply it to noncommutative C^*-algebras. For instance take $A = M_n(\mathbb{C})$. Then if χ were a character on A, its kernel would be a two-sided ideal of A. But matrix algebras are simple, hence the ideal would be trivial, which is impossible for dimension reasons as soon as $n \geq 2$. Thus, matrix algebras have no "points" ! This is a quite general situation for noncommutative algebras : they tend to have very few characters, if at all. We will see below that there are other notions of points, which agree in the commutative case, and which generalize better. But we must not be naive : no such notion is going to give us an unambiguous generalization (first of all because there are several of them !).

We must admit that "noncommutative space" is most of the time a *metaphoric expression*, and when we speak of noncommutative functions, we must not think that these functions are always defined on some kind of space.

To find another notion of space, we have to do a little yoga.

A representation of a C^*-algebra is a $*$-morphism $\pi : A \to B(H)$, where H is a Hilbert space. It is called *irreducible* if the only closed invariant subspaces of H under the action of A are trivial. For the sake of simplicity consider a finite dimensional representation. That is, each element of A is represented by a matrix $\pi(a) \in M_n(\mathbb{C})$. The representation will be reducible if there is a block matrix decomposition such that each $\pi(a)$ is block diagonal. "Schur's lemma" tells us that a representation is irreducible iff the only operators which commute with $\pi(A)$ are scalar multiples of Id_H. (This is quite obvious in the finite dimensional case.)

Two representation $\pi_1 : A \to B(H_1)$ and $\pi_2 : A \to B(H_2)$ are said to be *unitary equivalent* iff there exists a unitary transformation $U : H_1 \to H_2$ (i.e. a linear isomorphism preserving scalar products) such that

$$\pi_1(a) = U^{-1}\pi_2(a)U$$

for any $a \in A$.

The set of unitary equivalence classes of non-zero irreducible representations of A is called the *spectrum* (or sometimes the structure space) of A, and is denoted by $\hat{A}$.

The reason why I bother you with these definitions, is that a 1-dimensional representation (which is necessarily irreducible and alone in its equivalence class) is the same thing as a character. Moreover, when A is commutative, $\pi(A)$ commutes with itself, hence, by Schur's lemma, irreducible representations of commutative algebras are necessarily 1-dimensional[1]. Thus, $\mathcal{X}(\mathcal{C}(X)) = \widehat{\mathcal{C}(X)} = X$. Now the last equality is an identification of topological spaces, and it means that we must look for a way to put a topology on $\hat{A}$ such that $\hat{\mathcal{C}}(X)$ is canonically homeomorphic to X.

We will go back to this subject in a minute, for we need first to introduce another kind of space that we can attach to a C^*-algebra : its *primitive spectrum*. It is the set of ideals of A which are the kernels of irreducible representations. It is quite natural to consider this space, which in the commutative case is just the space of maximal ideals (it's a simple exercise to prove that).

Now there is an obvious surjective map $p : \hat{A} \to \mathrm{Prim}(A)$ which is $\pi \mapsto \ker \pi$. The topology we put on $\hat{A}$ is the pullback of the Jacobson topology by this surjection (for a friendly introduction to the Jacobson topology, see [Landi]). In general (for noncommutative algebras), both spaces are somewhat pathological. They are compact (for unital algebras), but generally not Hausdorff. The primitive spectrum is always T_0, but the spectrum is not (that is, two points may have exactly the same neighbourood system), unless when p is a bijection (in which case both spaces are homeomorphic).

[1] From then on I will always assume that representations are non-zero.

Finally there is a third notion of space.

For any unital C^*-algebra A, let a *state* be a linear functional $\phi : A \to \mathbb{C}$ such that :

1. ϕ is positive, that is $\phi(a^*a) \geq 0$ for any $a \in A$,
2. $\phi(1) = 1$.

This notion evidently comes from quantum mechanics. The following example will make the connection transparent. Let A be a C^*-subalgebra of $B(H)$ for some Hilbert space A (remember that it is not really a restriction). Let ψ be a unit vector in H, and let $\phi_\psi : A \to \mathbb{C}$ be $\phi_\psi(a) = \langle \psi, a\psi \rangle$. This is obviously a state on A (called a vector state), which in quantum mechanics gives the expectation value of a when a is hermitian and the system is in the state (in the quantum mechanical sense) ψ.

Clearly, any convex combination of states is again a state. Thus, the *state space* $\mathcal{S}(A)$ is a convex set (a technical note : the set of states is not closed in infinite dimension. The state space is then defined as the closure of the set of states). Given the $*$-weak topology, it can be shown to be compact (by the Banach-Alaoglu theorem).

Extreme points of $\mathcal{S}(A)$, that is, states which are not convex combinations of two other states are called *pure states*. The *pure state space* will be denoted $\mathcal{PS}(A)$.

Let us look at the simplest example imaginable : $A = \mathbb{C}^n$ is the set of complex functions on the discrete space $X = \{1; \ldots; n\}$. It is easily seen that a state on a A is of the form $f \mapsto \phi(f)$ where

$$\phi(f) = p_1 f(1) + \ldots + p_n f(n)$$

with $p_i \geq 0$ and $\sum p_i = 1$. That is, ϕ is the expectation of f for the probability measure $\mu = \sum p_i \delta_i$. Pure states are obviously the Dirac masses δ_i.

This example can be generalized to any compact space X : states are expectations under (Borel) probability measures, thus they can be identified with probability measures on X, and pure states correspond to Dirac masses.

Since Dirac masses are multiplicative : $\delta_x(fg) = f(x)g(x) = \delta_x(f)\delta_x(g)$, we see that pure states of commutative C^*-algebras are the same thing as characters, and thus can be identified with the points of the underlying space.

Now there is a very important tool in the theory of C^*-algebra, which is called the GNS construction, and that we won't describe here. It will suffice to say that it gives a surjective map from states to representations, and that irreducible representations are exactly the ones which are constructed out of pure states. Thus we have a diagram :

$$P(A) \longrightarrow \hat{A} \longrightarrow \mathrm{Prim}(A)$$

where each arrow is surjective. Moreover, the topologies are such that the maps are continuous and open. These are sufficient conditions to ensure that $\hat{A}$ is a topological quotient of $P(A)$, and $\mathrm{Prim}(A)$ of $\hat{A}$.

Here $P(A)$ is the set of pure states. Thus its closure is $\mathcal{PS}(A)$, which is compact and Hausdorff. Each space has its own kind of "pathology": $P(A)$ is generally not closed, and tends to be too big, $\hat{A}$ and $\mathrm{Prim}(A)$ are generally not Hausdorff, but are compact. They are sometimes homeomorphic, but sometimes $\mathrm{Prim}(A)$ is very small while $\hat{A}$ is quiet complicated (see the examples below).

While the general attitude is to give $\hat{A}$ the privilige of representing "the" space associated to A, let me stress that it is the pure state space which will enable us to introduce metric structures.

For each $[\pi] \in \hat{A}$, its pre-image under GNS is isomorphic to set of vector states associated to the representation $\pi : A \to H$, hence it is the projective space $P(H)$.

Note that the character space $\mathcal{X}(A)$ can be seen as included in each one of this space (in the middle one as the set of 1-dimensional irreps, and in the last as a subspace of the space of maximal ideals). And of course $P(A) = \hat{A} = \mathrm{Prim}(A) = \mathcal{X}(A)$ when A is commutative.

Some examples :

- For $A = M_n(\mathbb{C})$, $P(A) = \mathbb{P}(\mathbb{C}^n)$ and $\hat{A} = \{[\mathrm{Id}]\}$, $\mathrm{Prim}(A) = \{0\}$.
- Any finite dimensional C^*-algebra A_f is a direct sum of matrix algebra. Since direct sums of the algebras correspond to disjoint unions of the different kinds of spaces associated to it, we see that $P(A_f)$ is a disjoint union of projective space, each component of which shrinks to a single point in $\hat{A}_f \simeq \mathrm{Prim}(A_f)$.
- For $A = \mathcal{C}(M, M_n(\mathbb{C}))$, with M compact, one has $P(A) \simeq M \times \mathbb{P}(\mathbb{C}^n)$, and $\hat{A} \simeq \mathrm{Prim}(A) \simeq M$. (Note : this is a consequence of the fact that A and $\mathcal{C}(M)$ are strongly Morita equivalent).
- For $A = B(H)$, with H infinite dimensional and separable, there are two kinds of pure states : vector states and singular states, the latter being null on the set $K(H)$ of compact operators (see the section below for the definition). Irreducible representations also falls into two groups : the first consists of the single (up to unitary equivalence) representation $\mathrm{Id} : B(H) \to B(H)$, the second, which corresponds to singular states via the GNS construction, is formed by thoses irreps which vanish on K. It can be shown that each of these representations is non-separable, and that there are 2^c equivalence classes of them, where c is the power of the continuum ([K-R], p 750-757). Thus $\hat{A}$ has cardinal 2^c. However, $\mathrm{Prim}(B(H))$ has only two elements : $\{0\}$ and $K(H)$. Indeed, $K(X)$ is the only maximal norm closed two-sided ideal of $B(H)$, and since the kernel of a representation is a norm closed two-sided ideal, if $\pi(K) = \{0\}$ then $\ker \pi$ must be equal to $K(H)$.

We see that the space of pure states is a "blown up" version of the spectrum. This gives us the insight that what was just a point in the commutative case can have a structure in noncommutative geometry. We will continue to explore this idea below with noncommutative quotients and metrics.

5 Infinitesimals and quantized calculus

5.1 Infinitesimals as compact operators

One of the striking features of the analogy between functions and operators is that there is a natural notion of infinitesimal. Connes' claim that this notion of infinitesimal is superior to the one of nonstandard analysis has been the starting point of an ongoing dispute...

Let us recall that an operator $a \in B(H)$ is said to be *compact* if it is the norm limit of a sequence of finite rank operators (operators with a finite dimensional image). Alternatively, an operator is compact iff, given any $\epsilon > 0$, there always exists a finite dimensional $V \subset H$ (of sufficiently high dimension), such that $\|a_{|V^\perp}\| < \epsilon$. If we accept the idea that finite dimensional subspaces count for nothing, then we understand that a compact operator is negligible in some way.

We will write $K(H)$ for the set of compact operators. If H is finite dimensional we obviously have $K(H) = B(H) \simeq M_n(\mathbb{C})$.

In the rest of the section we suppose that H is infinite dimensional. In that case, a compact operator a cannot be inversible, hence 0 always belongs to its spectrum. The other elements of the spectrum are eigenvalues of finite multiplicity, and the spectrum of a is either finite (for finite rank operators) or countable with 0 as its only accumulation point. The spectrum of $|a| = \sqrt{a^*a}$ can then be arranged as a non-increasing sequence $\mu_n(a)$ which converges to 0 as $n \to \infty$ (it will be stationnary for finite rank operators).

The fastest $\mu_n(a)$ converges to 0, the smallest $|a|$ will be on the orthogonal of the finite-dimensional eigenspace W_n spanned by $(\mu_0(a), \ldots, \mu_n(a))$. So it is reasonable to say that a is an infinitesimal of order $\alpha > 0$ if $\mu_n(a) = O(n^{-\alpha})$ (one can adopt the convention that any bounded operator is of order 0). That this definition is sensible is confirmed by the following facts :

- if a_i is an infinitesimal of order α_i, for $i = 1, 2$, then $a_1 a_2$ is an infinitesimal of order $\alpha_1 + \alpha_2$,
- if a is an infinitesimal of order α, and b is any bounder operator, then ab and ba are infinitesimals of order α. (One should think of $f(x)dx$ where f is a "normal " function and dx an order 1 infinitesimal.)
- One can define a noncommutative integral $⨍$ such that order 1 infinitesimal can be integrated, and higher order infinitesimal have a null integral.

We won't describe in detail this noncommutative integral. Let us just say that the usual trace of an order one infinitesimal is (at most) logarithmically divergent, and that there is a way (the Dixmier trace) to extract the coefficient of the logarithmic divergence, even though the partial sums $\frac{1}{\ln N}\sum_{n=0}^{n-1}\mu_n(a)$ are not necessarily convergent. The existence of the Dixmier trace depends on the axiom of choice, and thus cannot be explicitly constructed. Moreover its value would depend on some linear form ω on the space of bounded sequences, satisfying some criteria. However, in concrete situtations, the value of the Dixmier trace $\mathrm{Tr}_\omega(a)$ does not depend on ω (it can be explicitly computed). When it is the case we say that a is *measurable* (there is an explicit criterion for a being measurable, see [Red Book]). When a is measurable, we write $⨍ a$ for its Dixmier trace relative to any ω, and we call it the noncommutative integral of a.

The noncommutative integral is positive, linear, neglects infinitesimal of order > 1 and is a trace : $⨍ ab = ⨍ ba$ for any $b \in B(H)$.

Some arguments of Connes in favor of the superiority of this notion of infinitesimal over the one from nonstandard analysis :

- a compact operator can be exhibited while a nonstandard infinitesimal cannot,
- random variables with discrete or continuous ranges can coexist in $B(H)$.

Let us give an example (from [Co 95]). Let Ω be a disk and ask the question : what is the probability of a dart hitting a given point ? In standard probability theory, it is of course 0, a result that some people find unsatisfactory. In nonstandard analysis, I presume that the answer would be a non zero infinitesimal. In Connes' formalism also, the result is such. Let $H = L^2(\Omega, d\mu)$ where $d\mu$ is the Lebesgue measure on Ω. Then the answer of Connes is Δ^{-1}, where Δ is the Dirichlet Laplacian (Laplacian accepting only functions satisfying the Dirichlet boundary condition $u = 0$ on $\partial\Omega$) in Ω. Any continuous function f on Ω can be seen as an operator in H acting by multiplication and one has the remarkable result

$$⨍ f\Delta^{-1} = \int_\Omega f d\mu$$

Since the Laplacian is the square of the Dirac operator (when there is no curvature), the formula above gives some weight to the insight that "$ds = D^{-1}$" in NCG. (Here $d\mu = dxdy \leftrightarrow D^{-2}$. In general $|D|^{-p}$ will correspond to the volume form, where p is the dimension.)

5.2 Noncommutative differentiation or "quantized calculus"

We will now quickly explain the noncommutative notion of differentiation. We take an algebra $\mathcal{A}$ of would-be noncommutative differentiable functions. A priori we cannot expect every continuous functions to be differentiable so that we only require that $\mathcal{A}$ is a $*$-algebra, not a C^*-algebra. We will also need a representation π of $\mathcal{A}$ on some Hilbert space H.

In a way which is quite similar to the replacement of Poisson brackets with commutators in quantum mechanics, one defines the differential df of a "q-function" f by a commutator :

$$df = [F, f]$$

Note that this automatically satisfies the Leibniz rule. We also require that F is a self-adjoint densely defined operator such that $F^2 = 1$. The requirement that $F^2 = 1$ ensures that $[F, [F, f]]_+ = 0$, which corresponds to $d(df) = 0$ (d acts on df by anticommutator). We could justify the requirement that $F^* = F$ by asking that the differential of a real f is still real, but for this it would be a good thing to put an i in front of F in the definition. Some authors do, but we will stick with the conventions of [Red Book] in order to ease the comparison, and we will live with the fact that if f is real, then df is imaginary.

Naturally, we also require that $[F, f] \in K(H)$ for every $f \in A$.

Now you certainly want some examples.

Example 1 (a discrete space) : Let us consider $X = \{1; 2\}$ and $\mathcal{A} = \mathcal{C}(X)$. As we already saw, $\mathcal{A}$ can be represented on $H = \mathbb{C}^2$ by diagonal matrices. This representation being faithful (injective) we won't bother with the π's. So an element of $\mathcal{A}$ is of the form $a = \begin{pmatrix} a(1) & 0 \\ 0 & a(2) \end{pmatrix}$. Since the diagonal elements of F will drop out of commutators with diagonal matrices, they serve no purpose and we can suppose them to be 0. You can check that the general solution to $F = F^*$, $F^2 = 1$ is then $F = \begin{pmatrix} 0 & e^{i\alpha} \\ e^{-i\alpha} & 0 \end{pmatrix}$. The choice of an α can be understood as the choice of a differential structure on X. We choose $\alpha = 0$ and we calculate $[F, a]$:

$$da = [F, a] = (a(2) - a(1)) \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

Thus the noncommutative differential here boils down to a finite difference operator times a constant matrix (which depends on the choice of α). Note for later purpose that the spectrum of this matrix is $\{-i; +i\}$. (Remark that $\sigma(da) = i\{a(1) - a(2); a(2) - a(1)\}$. Up to the infamous factor of i which is missing in the definition of F, it is a most natural result !).

Example 2 (a fractal space) : we consider the middle-third Cantor set K. Remember the construction of K : one starts with $[0; 1]$ and removes the open middle third $]1/3; 2/3[$, then one does the same to each of the two remaining closed intervals, etc. The compact set K is what is left in the end. It can be described as the set of points of $[0; 1]$ which can be written entirely with 0's and 2's in base 3. The end points of removed intervals are among those, since they have infinite ternary representations composed of 2's and 0's, like $1/3 = 0.02222\ldots$ (in base 3). Let us call E the set of such end points. These points can be arranged in pairs $(e_k^-; e_k^+)_{k\in\mathbb{N}}$ where $]e_k^-; e_k^+[$ is a removed interval (with the exception of 0 and 1, but we make a pair of these two).

We let $H = \ell^2(E)$ and $A = \mathcal{C}(K)$. The elements of A act on H by multiplication :

$$\forall f \in A, \forall \xi \in H,\ f.\xi(e) = f(e)\xi(e), \forall e \in E$$

Consider the sequences δ_k^+ and δ_k^- which take the value 1 at e_k^+ (resp. e_k^-) and 0 elsewhere. They span the 2-dimensional space V_k. We define the operator

F_k on V_k by the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ as in the previous example. Finally we set $F = \bigoplus_{k\in\mathbb{N}} F_k$. This operator exchanges the values a sequence in H takes on $e_k^{\pm}$.

Using on first example on each V_k we see that

$$[F, f] = \bigoplus_{k\in\mathbb{N}} (f(e_k^+ - f(e_k^-)) \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{1}$$

Hence $[F, f]$ contains the information of the discrete jumps of f on each gap of K.

Now let us use (1) on the particularly simple function $x : K \to \mathbb{C}$ defined by $t \mapsto t$, that is, the inclusion function of K into $\mathbb{C}$. We get

$$dx := [F, x] = \bigoplus_k (e_k^+ - e_k^-) \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} := \bigoplus_k \ell_k \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{2}$$

Since $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ has eigenvalues $\pm i$, one has

$$|dx| = \bigoplus_k \ell_k I_2 \tag{3}$$

Hence the spectrum of $|dx|$ is 1 (coming from the exceptional couple $(0;1)$), $1/3$, $1/9$, ..., and the multiplicity of $1/3^k$ is $2 \times 2^{k-1}$, the first factor coming from the matrix I_2, and the factor 2^{k-1} coming from the fact that there are 2^{k-1} couple of endpoints separated by an interval of length $1/3^k$. Incidentally, one sees that $|dx|$ is a compact operator. Now if we decreasingly order the eigenvalues of $|dx|$, then $\mu_n(|dx|) = 3^{-k}$ where k is the largest integer such that $1 + 1 + 2 + \ldots + 2^{k-1} = 2^k \leq n$ (the first 1 comes from the multiplicity of 2 of the eigenvalue 1, which is exceptional, and the formula works for $n \geq 1$, which is enough for our purpose).

Hence $k = E(\frac{\log n}{\log 2})$. Thus $\mu_n(|dx|) \sim (\frac{1}{n})^{\log 3/\log 2}$, showing that $|dx|$ is an infinitesimal of order > 1. To have an infinitesimal of order 1 we must take $|dx|^p$, with $p = \frac{\log 2}{\log 3}$. Now in this case it is easy to compute the Dixmier trace. Since $(1/3^j)^p = 1/2^j$ and the multiplicity of $(1/3)^j$ is precisely 2^j, one has $\sum_{j=0}^{n-1} \mu_j(|dx|^p) = 1 + 1 + 2^{-k}(n - 2^k)$, by simple counting. Hence we get

$$\frac{1}{\log n}(\frac{\log n}{\log 2} + 1) \leq \frac{1}{\log n} \sum_{j=0}^{n-1} \mu_j(|dx|^p) < \frac{1}{\log n}(\frac{\log n}{\log 2} + 3)$$

Thus $\frac{1}{\log n} \sum_{j=0}^{n-1} \mu_j(|dx|^p)$ converges. This is the simplest instance in which an operator is measurable (i.e. its Dixmier trace does not depend on ω), and one has $\mathrm{Tr}_\omega(|dx|^p) = \frac{1}{\log 2}$. One can show more generally that for any $f \in \mathcal{C}(K)$, one has

$$\mathrm{Tr}_\omega(f|dx|^p) = \frac{1}{\log 2} \int f d\Lambda$$

where Λ is the Hausdorff measure on K. Thus one recovers the Hausdorff dimension and the Hausdorff measure of K (the latter up to a constant).

Of course this is a particularly simple case, but it can be vastly generalized. Let J be the Julia set of the sequence of iterations of the complex polynomial z^2+c. (Recall that J is the boundary of the set of initial values for which the sequence is bounded. This is typically a fractal). For c small enough, J, which is of Hausdorff dimension $p \in]0;1[$, is the image of a continuous map $Z : S^1 \to \mathbb{C}$.

It is proven in [Red Book] that if we set $H = L^2(S^1)$, $A = \mathcal{C}(S^1)$, and F is the Hilbert transform (that one can define by $Fe_n = sgn(n)e_n$ on $e_n(\theta) = e^{in\theta}$), then

$$f \mapsto ⨍ f(Z)|[F,Z]|^p \tag{4}$$

is a state on the algebra $\mathcal{C}(J)$ which is just the integration with respect to the Hausdorff measure.

Remark : one can check that on $\mathbb{R}$: $[F,f]\psi(t) = \int k(s,t)\psi(s)ds$ where $k(s,t) = \frac{1}{i\pi}\frac{f(s)-f(t)}{s-t}$. On S^1 : use the Cayley transform.

Thanks to these examples we see that the so-called quantized calculus has interesting features, even in the commutative case :

1. It applies as well to discrete, fractal, and smooth spaces,
2. df can make sense even if f is not differentiable,
3. $|df|^p$ makes obvious sense by continuous functional calculus, even for non-integer p.

6 Noncommutative geometry

6.1 The Dirac operator

With $\hbar = c = 1$, the relativistic dispersion formula $E^2 - p^2 = m^2$ for a particle of mass m can be quantized into the Klein-Gordon equation

$$(-\frac{\partial^2}{\partial t^2} + \triangle)\Psi = m^2\Psi$$

This equation for was obtained by Schrödinger (even before his eponymous equation), but it poses several problems, one of which being of the second order. Dirac realized that in order to get a relativistic wave equation of the first order, one had to somehow take the square root of the operator $-\frac{\partial^2}{\partial t^2} + \triangle$. One then looks for an equation of the form

$$(i\not{\partial} - m)\Psi = (i\gamma^\mu\partial_\mu - m)\Psi = 0$$

where γ^μ are formal symbols for the moment. Acting with the conjugate $-i\not{\partial} - m$ we immediately see that in order to recover the Klein-Gordon equation one must have

$$\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu,\nu}$$

where $\{.;.\}$ is the anti-commutator and $g^{\mu,\nu}$ is the Minkowski metric. Thus the γ^μ cannot be usual numbers and the smallest algebra containing such objects turns out to be $M_4(\mathbb{C})$. One popular matrix representation is

$$\gamma^0 = \begin{pmatrix} \sigma_0 & 0 \\ 0 & -\sigma_0 \end{pmatrix}, \gamma^i = \begin{pmatrix} 0 & \sigma_i \\ -\sigma_i & 0 \end{pmatrix}, \ i = 1, 2, 3$$

where $\sigma_0 = I_2$, $\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $\sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, $\sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ are the Pauli spin matrices. The elements of the four dimensional space on which the γ matrices act are called *Dirac spinors*. (Remark : one has to put a $-i$ in front of $\gamma^{1,2,3}$ in euclidean signature.)

The Dirac operators which appear in noncommutative geometry are abstract versions of the Dirac operators which had been used in *riemannian* (as opposed to Lorentzian) geometry by Atiyah and Singer. These can be defined on spin manifolds M, that is, oriented riemannian manifolds on which one can construct a spinor bundle $S \to M$. (For those not familiar with vector bundles, it is a smooth and "locally trivial" way of gluing copies S_x of the space of spinors at each point $x \in M$.)

Using a local orthonormal frame (e_i) one will be able to write like this :

$$D\Psi = \sum_{i=1}^{n} e_i \cdot \nabla_{e_i} \Psi$$

where $\cdot$ is the Clifford multiplication between vectors and spinors (this is a way to see e_i as an operator on spinors, thus the e_i in this formula replace the gamma matrices), and ∇ is the canonical lift of the Levi-Civita connection on the spinor bundle. There is a canonical hermitian scalar product on spinors, thus one can define the Hilbert space $L^2(S)$, and D will be a densily defined symmetric operator on this Hilbert space. All of this still works when one replaces ∇ with ∇^A with A a connection one-form. With this generalized Dirac operator, one can prove the Lichnerowicz formula $D^2 = \triangle + R/4$ where R is the scalar curvature.

If M is compact, the spectrum of $|D|$ is discrete, and if we order the eigenvalues as an increasing sequence λ_n, then $\lambda_n \sim n^{-1/p}$, where p =dim(M) by Weyl's spectral theorem.

One of the very striking aspects of Connes NCG is that the abstract Dirac operator will allow one to define at the same time the metric and the differential structure in the noncommutative setting (the two aspects can be partially decoupled by writing the polar decomposition $D = |D|F$. The phase F of the Dirac will give the differential structure, while in some cases the metric only depends on $|D|$, though not always, see below). That the Dirac operator is interwoven with the metric is apparent already in the commutation relations followed by the gamma matrices (which generalize in the Clifford algebra). There is also a very heuristic argument that I would like to be able to make more precise, but unfortunately I cannot : the euclidean line element ds satisfied $ds^2 = dx^2 + dy^2 + \ldots$. Now the Laplacian is $\frac{\partial^2}{\partial x^2} + \ldots$, so somehow (this is the very slippery part), $ds^2 \leftrightarrow 1/\triangle$. Now since $DD^* = D^2 = \triangle$, one has $|D| \leftrightarrow 1/ds$. Indeed we will see that the line element ds corresponds to $|D|^{-1}$ in noncommutative geometry.

6.2 Spectral triples

Definition 2 *(Connes, 1994) A spectral triple is a triple $(A, \mathcal{H}, D)$ where*

1. *A is a $*$-subalgebra of $\mathcal{B}(\mathcal{H})$,*
2. *D a densely defined self-adjoint operator on $\mathcal{H}$,*
3. *D has compact resolvent,*
4. *$\forall \Psi \in Dom(D)$, $\forall a \in A$, $a\Psi \in Dom(D)$ and $[D, a] \in B(H)$.*

How do we understand these axioms ? The norm closure $\bar{A}$ will be a C^*-algebra which is to be thought of the algebra of continuous functions on some virtual noncommutative space. The algebra A itself will represent elements which are more regular than being just continuous (they are at least Lipschitz).

Let us explain condition 3. Technically it says that $(D - \lambda)^{-1}$ is compact for some λ not in the spectrum. It will be perhaps better understood as a condition on the spectrum (which is equivalent) : D has a real discrete spectrum made of eigenvalues with finite multiplicities such that $|\lambda_n| \to \infty$ (in infinite dimension). Compare with Weyl's spectral theorem. Note that a pair (H, D) such as in the definition is completely characterized (up to unitary equivalence) by the list of eigenvalues of D with multiplicities.

Remark : this is really the definition of an "odd" spectral triple. We will not enter into the discussion of even spectral triples.

To understand these axioms we need to check them on some examples, and this we will do below. But first let us see how we can define some notion of metric. This is given by the following formula, for any two states ϕ, ψ on $\bar{A}$:

$$d_C(\phi, \psi) = \sup_{a \in A}\{|\phi(a) - \psi(a)|, \ \|[D, a]\| \leq 1\}$$

This gives a (not always bounded) distance on $\mathcal{S}(A)$, which we call the Connes distance. We will rather consider its restriction on the state of pure states, which corresponds to the points of the manifold in the commutative case.

Example 1 : the canonical triple over the circle. We take $H = L^2(S^1)$, $A = \mathcal{C}^\infty(S^1)$, $D = \frac{1}{i}\frac{d}{d\theta}$. A function $f \in A$ acts H by multiplication. For any $f \in A$, and for any $\Psi \in H$ such that Ψ is derivable,

$$[D, f]\Psi = \frac{1}{i}(\frac{d}{d\theta}(f\Psi) - f\frac{d\Psi}{d\theta}) = \frac{1}{i}\frac{df}{d\theta}\Psi$$

Thus $[D, f]$ acts as the multiplication by the bounded functions $-if'(\theta)$ on a dense subset of H, and can thus be extended to an element of $B(H)$.

Finally, if $\mathcal{F}$ is the Fourier transform $\mathcal{F} : H \to \ell^2(\mathbb{Z})$, then $\mathcal{F}^{-1}D\mathcal{F}$ is just the diagonal multiplication by $(n)_{n\in\mathbb{Z}}$ which has all the required spectral properties for proving that D has compact resolvent.

Let us now look at Connes distance. Using Gelfand transform, we identify two pure states on $\mathcal{C}(S^1)$ with two points p, q on the circle and we want to compute their distances using Connes formula. First of all, we can suppose with no loss of generality that p is the origin of arguments and that the principal argument of q is $\theta \in [0; \pi]$. We identify A with periodic $\mathcal{C}^\infty$ functions on $[0; 2\pi]$. We have $\|[D, f]\| = \|\frac{df}{d\theta}\|$, so the condition $\|[D, f]\| \leq 1$ together with the mean value theorem means that $|f(\theta) - f(0)| \leq \theta$. Now take the function whose graph is a tent with $f(0) = 0 = f(2\pi)$ and $f(\pi) = \pi$. It is not $\mathcal{C}^\infty$, but we can approximate it by a $\mathcal{C}^\infty$ function as well as we wish. With this f the supremum is reached (or rather approximated as close as we wish) and we see that $d(p, q) = \theta$, which is precisely the geodesic distance on the circle.

In fact, this example can be vastly generalized.

Example 2 : the canonical triple over a manifold. Let M be a spin manifold, $H = L^2(S)$ be the Hilbert space of L^2 sections of the spinor bundle over M, and D is the Dirac operator associated to the Levi-Civita connection.

In some coordinates one finds that $[D, f]\Psi = (\gamma^\mu \partial_\mu f)\Psi$. So once again $[D, .]$ is a sort of differentiation.

When restricted to pure states, Connes distance gives back the geodesic distance on M. This is a relatively simple though very important result :

Theorem 1 *Let (A, H, D) be the canonical triple over a manifold. Then*

$$d(p, q) = \sup_{a\in A}\{|a(p) - a(q)|, \ \|[D, a]\| \leq 1\}$$

is the geodesic distance between p and q.

Remark : if $n = 1$ the spinor bundle has dimension $2^{[n/2]} = 1$, so that example 1 is really a particular case of example 2.

Example 3 : finite dimensional commutative triple. Now we look at what is probably the simplest possible example. We take $H = \mathbb{C}^2$, $A = \mathbb{C} \oplus \mathbb{C} = Diag(M_2(\mathbb{C}))$. Looking at our axioms we see that there is really no restriction on D. However, if we separate D into an off-diagonal and a diagonal part : $D = \Delta_1 + \Delta_2$, we see that the diagonal part disappears in the commutators. Hence we suppose that $D = \begin{pmatrix} 0 & m \\ \bar{m} & 0 \end{pmatrix} = |m|F$, with $F = \begin{pmatrix} 0 & e^{i\alpha} \\ e^{-i\alpha} & 0 \end{pmatrix}$. The commutator $[D, a]$, with $a = diag(a_1, a_2)$ is $(a_2 - a_1)|m| \begin{pmatrix} 0 & e^{i\alpha} \\ -e^{-i\alpha} & 0 \end{pmatrix}$.

Now let us compute Connes distance. There are two pure states, which we call p_1 and p_2, and which are none other than the two points of our commutative space : $p_i(a) = a(i) = a_i$. The norm of the matrix $|m| \begin{pmatrix} 0 & e^{i\alpha} \\ -e^{-i\alpha} & 0 \end{pmatrix}$ is $|m|$, hence the condition on a boils down to $|a_1 - a_2| \leq |m|^{-1}$. We thus have to take the supremum of $|a_1 - a_2|$ subject to the condition $|a_1 - a_2| \leq |m|^{-1}$. The result is obviously $|m|^{-1}$.

We see now that the commutator $[D, a]$ is really a discrete derivative :

$(a_1 - a_2)|m| = \frac{a_1 - a_2}{d(p_1, p_2)}$, times a constant matrix which depends on "the differential structure" F.

Note that if we do not remove the diagonal part of D, the distance will still only depend on the modulus of Δ_1, whereas the modulus of $|D|$ will also depend on Δ_2. We see that we cannot always clearly read a decoupling of the metric and differential structure from the polar decomposition of D.

Remark : It has recently been shown [Ch-Iv 07] that the distance of any compact metric space can be recovered thanks to an ad hoc generalization of this 2-dimensional spectral triple. We see that even though the spectral distance is intended to be a noncommutative generalization of the geodesic distance on manifolds, it naturally applies, even in the commutative case, to much more general kinds of spaces.

Example 4 : finite dimensional noncommutative triple $A = M_2(\mathbb{C})$ acting on $H = \mathbb{C}^2$. There is still no restriction on D which can be any hermitian matrix. Up to some unitary transformation we can thus suppose that $D = diag(d_1, d_2)$. We can also remove $\frac{1}{2}\mathrm{Tr}(D)I_2$ without changing the commutators, thus we will suppose that $D = d \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ with $d > 0$.

Then if $M = (m_{ij})$ one finds $[D, M] = 2d \begin{pmatrix} 0 & m_{12} \\ -m_{21} & 0 \end{pmatrix}$.

Recall that a pure state on $M_2(\mathbb{C})$ is always a vector state of the form $\omega_\xi(a) = \langle \xi | a(\xi) \rangle$ with $\|\xi\| = 1$. Since ω_ξ really depends only on the direction $[\xi]$, the state space of A is here $\mathbb{P}^1(\mathbb{C})$. It turns out that we can visualize elements of $\mathbb{P}^1(\mathbb{C})$ thanks to the Hopf fibration : to $\xi = \begin{pmatrix} \xi_1 \\ \xi_2 \end{pmatrix}$ we associate $h(\xi) = (2\Re(\xi_1\bar{\xi}_2), 2\Im(\xi_1\bar{\xi}_2), |\xi_1^2| - |\xi_2|^2)$ which turns out to be an homeomorphism between $\mathbb{P}^1(\mathbb{C})$ and S^2 (when S^2 is viewed this way it is sometimes called the Bloch sphere in physics). Note that orthogonal vectors in $\mathbb{C}^2$ are mapped to antipodal points in S^2. Now one can compute [I-K-M 01] that :

- If $h(\alpha)$ and $h(\beta)$ lie at the same altitude in S^2, then $d(\omega_\alpha, \omega_\beta) = \frac{\sqrt{1 - |\langle \alpha | \beta \rangle|^2}}{d} = \frac{1}{d}\delta(h(\alpha), h(\beta))$ where δ is the euclidean distance in $\mathbb{R}^3$,
- else $d(\omega_\alpha, \omega_\beta) = \infty$.

In case you wonder why the north pole/south pole axis is singled out by this distance, it is because we have supposed at the start that D was diagonal. In the general case, the eigendirections of D would give two antipodal points in S^2 which would then play the role of the poles.

We can heuristically describe the situation as such : if you look with "commutative eyes" you see only one point, because the structure space of $M_2(\mathbb{C})$ has only one point. But the spectral triple endows this point with an internal

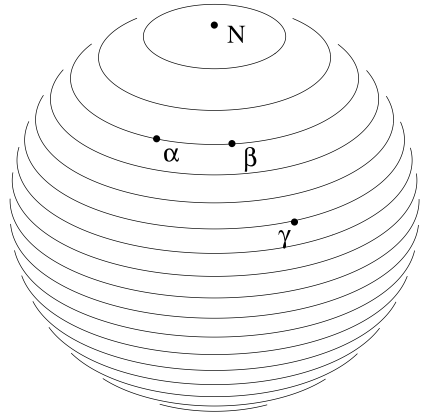

Figure 1: Distances between pure states in example 4.

structure. In this case you can see the structure if you blow-up the point, passing from the structure space to the pure state space. The structure consist of a length scale given by d and a differential structure given by the foliation of S^2 by parallel circles.

Remark : We can even interpret the result on $[D, M]$. At first sight it doesn't make sense to derive a function defined at a single point. But we can see this point as a noncommutative quotient of a two-points sets (see the appendix), and the matrix M giving the values of the function on the arrows of the graph of the equivalence relation. This is very heuristic, and I'm afraid I can't be much more specific than that.

6.3 Noncommutative manifolds

An important issue is to characterize the spectral triples which are the canonical triples over manifolds. Commutativity of the algebra is an obvious requirement, but we already see thanks to example 3 that it does not suffice. Thanks to Weyl's spectral theorem we can express the dimension of the manifold in a spectral way by supposing that the characteristic values of the resolvent of D are of the order of $n^{-1/p}$ with p some natural integer.

But there are other assumptions which are quite technical, and which ensure that we will recover a manifold M such that the algebra A we started from is exactly the algebra of smooth functions, that D will be an order one differentiable operator, and that M will be orientable (see [Co 08] for details). In order to recover also the spin structure, and be sure that the Dirac operator corresponds to some metric on the manifold, one also supposes the existence of a real structure J.

Let us collectivelly call these assumptions the *axioms of commutative spin* *geometry*. Under these assumptions, it is possible to prove a reconstruction theorem. It is now natural to call spectral triples satisfying all these axioms except commutativity *noncommutative spin geometries*.

Remark : if you want to learn more about that issue our advice is that learn first about the axioms of spin geometries in section 10.5 of [GB-V-F]. In this book (chapter 11), a reconstruction theorem is stated and proven under the additional hypothesis that we know from the beginning that $A = \mathcal{C}^\infty(M)$. At the end of this chapter, the authors allude to the fact that this quite strong hypothesis can in fact be removed and replaced by "A is commutative". However, the paper they refer to for a proof has flaws. A correct proof has now been given in [Co 08], to which one should refer instead.

6.4 The standard model

The symmetry group of the Einstein-Hilbert Lagragian is Diff(X) where X is the spacetime manifold. On the other hand, the one of the Standard Model Lagrangian is $G_{SM}(X) = \mathcal{C}^\infty(X, U(1) \times SU(2) \times SU(3))$ (at least when the bundle is trivial, e.g. if X is contractible). Thus the full symmetry group of gravity minimally coupled to the Standard Model is the semi-direct product $G(X) = G_{SM}(X) \rtimes \text{Diff}(X)$. A geometrization of the SM, in the spirit of Kaluza and Klein, would be to find a higher dimensional space Y such that Diff$(Y) = G(X)$. However, it is impossible, since the diffeo group of an ordinary manifold cannot be a semi-direct product. However it is possible to find a spectral triple with that property. Moreover, and this is remarkable, this spectral triple satisfies the axioms of noncommutative spin geometries.

Let us be more specific. The diffeo group of X is canonically isomorphic to the automorphism group of $\mathcal{C}^\infty(X)$ (by the map $\phi \mapsto (f \mapsto f \circ \phi^{-1})$. Now there exists a noncommutative algebra $\mathcal{A}$, which is of the form $\mathcal{C}^\infty(X) \otimes \mathcal{A}_F$, with $\mathcal{A}_F$ a finite dimensional C^*-algebra which we will describe below, such that Aut$(\mathcal{A})$ is exactly the desired symmetry group. It is worth noting that the inner automorphisms of $\mathcal{A}$ correspond to what are called internal symmetries in physics, whereas the outer automorphisms, that is to say the quotient Aut$\mathcal{A}/Inn(\mathcal{A})$ correspond to diffeomorphisms.

In a nutshell, the usual gauge theory, which postulates that Y has the structure of a bundle over X and reduces the symmetries to the ones which preserve this structure, is replaced by a theory which is closer in spirit to General Relativity, allowing all symmetries. Attaching a noncommutative discrete space to each point of X, instead of a commutative smooth one, permits to interpret internal symmetries as "noncommutative diffeomorphisms in the discrete direction", if one wants.

Let us now briefly discuss this spectral triple. For simplicity's sake, we describe only the first model, which first appeared in the begging of the nineties in the works of Alain Connes, and is described in full details in [Ch-Co 96]. A better model, which takes neutrino mixing into account appeared in [Co 06] and independently in [Ba 07].

The finite dimensional algebra $\mathcal{A}_F$ is $\mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$. (Remark : in this talk we only considered complex C^*-algebras for simplicity, but we see that we really need real algebras such as $\mathbb{H}$.) This algebra acts on a finite dimensional space H_F which has the set of elementary fermions as basis (=the list of quarks and leptons of all generations, taking into account chirality and the antiparticles.)

Specifically, $A = \mathcal{C}^\infty(M) \otimes A_F$, $H = L^2(M,S) \otimes H_F$, $D = D_M \otimes 1 + \gamma_5 \otimes D_F$.

The Dirac operator D_F for the finite part encodes the Yukawa couplings, that is the masses of the fermions and the CKM mixing matrix.

The gauge bosons appear thanks to "inner fluctuations of the metric". It means that we replace D by a fluctuated Dirac operator $D_A = D + A + \epsilon' JAJ^*$, where A is a general gauge potential of the form $A = \sum_i a_i[D, b_i]$, $a_i, b_i \in A$, and ϵ' is a sign which is given by the real K-theoretic dimension of the algebra (here it is equal to 1). If one thinks of D as a fixed origin, it is quite normal that the degrees of freedom of the Dirac operator, which is a dynamical variable of the theory, are expressed in the one-form A.

Note that the action of the inner automorphism $\alpha_u : a \mapsto uau^*$ on D_A is to replace A by $A^u = uAu^* + u[D, u^*]$.

Let us now focus on the bosonic part of the action for this model. It comes from the *spectral action principle*, put forward by Connes and Chamseddine. This principle stipulates that the action defining the theory should be invariant by all automorphisms of the algebra A (in fact this is not exactly true, there is a little subtlely : these automorphisms have to be implementable by unitary operators in H. For the case of the canonical triple over a manifold, it is always the case.) This principle generalizes the principle of general covariance (or if one prefers, diff-invariance) on which GR is based, to which it reduces in the case of the canonical triple. Since the only piece of data which is necessarily invariant in this way are the eigenvalues of the Dirac operator, we guess that the action will involve only these eigenvalues. Indeed, it is given by

$$S(D_A) = \mathrm{Tr}(f(D_A^2/\Lambda^2)) \tag{5}$$

Here Λ is an energy scale which plays the role of a cut-off parameter, and f is a smooth approximation of the characteristic function of $[0;1]$. Thus, in essence, S just counts the eigenvalues of D_A^2 which are below the cut-off.

Note the "tour-de-force" : it is only an effective theory, with a cut-off parameter. However, the truncation is made in a way which preserves the symmetries, which is quite remarkable.

Just for the record, let us mention that the fermionic part of the action is the maybe more familiar $S_F(\Psi, A) = \langle \Psi, D_A \Psi \rangle$.

The punchline is the following : using this action on this particular spectral triple one recovers the usual Lagragian of the standard model minimally coupled to (euclidean) gravity. There are some bonuses which are outlined below :

- The standard model spectral triple satisfies the axioms of a noncommutative spin geometry, which is not a trivial requirement. For instance, models of GUT do not. (Recently Connes and Chamseddine have shown that if one drops the first-order condition, one is naturally led to the Pati-Salam model.) Also, at least 2 generations of particles are required (sadly this prediction comes 60 years too late !).

- The Higgs field is not an input : it appears naturally as the component of the connection 1-form in the discrete direction
- The elementary forces, including *classical euclidean* gravity are unified.
- The bosonic part of the action is pure gravity.
- The old model which I have described above had a problem of fermion doubling. It turned out that the new model in which neutrino masses appear also cures this problem. The cure comes from giving the finite algebra a real K-theoretic dimension of 6 modulo 8 instead of 0, which changes some signs.

On the side of open problems :

- For the time being we can deal only with riemannian (as opposed to lorentzian) manifolds. (A first step in the Lorentzian direction has already been taken in [Ba 07], allowing the understanding of the correct change of signs in the spinorial chessboard. But full-fledged Lorentzian noncommutative geometry is still a work in progress : see the introduction of [Be 09] for a very short review, or [Fr 11] for a thorough one.)
- There is no quantum gravity effect.
- In fact there is not even quantum mechanics : it gives a classical field theory which must be quantized afterwards.
- It is an effective theory, with a cut-off parameter : by no way a fundamental theory.
- The finite algebra is an input. Recently Connes has put forward some assumptions under which this algebra turns out to be an almost unique and thus natural choice (see [Ch-Co 08]). Of course one can always discuss the naturalness of the assumptions, but it is still an important conceptual progress.
- With this model one can compute the Higgs mass under the big desert hypothesis. Since the predicted mass is already ruled out experimentally, we come to the conclusion that either the model or the BDH is wrong. On the other hand one gets a retrodiction of the top quark mass which agrees with experiment to 10 % : this can maybe be due to the fermionic action being simpler and thus less sensitive to new physics than the bosonic action.

7 Conclusion

Noncommutative geometry is characterized by a remarkable interplay between mathematics and physics. This interplay is not at all easy to analyze: I think that there is some work here for the philosopher. Maybe we can say, as a first intuitive approach, that the emergence of noncommutative geometry can be seen as the impact on mathematics of the slow assimilation of our discovery of the quantum world. All the concepts that we have long abstracted from the classical world (points, lines, geometry, topology, ...) get replaced by subtler and deeper ones, which turn out also to be more effective. Surely we could not stay forever pondering the commutative shadow of a bigger world. A step had to be made into the noncommutative realm, and it first happened in physics. On the other hand, our brains which grew in the commutative shadow are always in need of geometric pictures. The effectiveness of the geometric intuition outside its scope of application is for me a source of amazement.

Of course we are very far from a complete theory. For the purely mathematical part, the theory is still not entirely formalized, and some important conjectures are still open (like the Baum-Connes conjecture). As for the application to physics, it should be apparent from what we have seen above that some radically new ideas are needed in order to obtain a theory which include quantum gravity and make correct predictions. This promises even more far-reaching consequences on mathematics and on our fundamental concepts about the physical world.

A Building noncommutative spaces

A.1 Noncommutative quotients

This section comes mostly from [Kh 04], to which we refer for a more careful treatment of the convergence problems which we here sweep under the rug.

There are several ways of defining a groupoid. The most visual one is to say that it is a directed graph with the following properties :

1. Every time there is an arrow f from a to b and an arrow g from b to c, there is an arrow from a to c which is an arrow $g \circ f$ from a to c (composition of arrows).

2. The composition is associative.

3. For every vertex a there is a unique arrow from a to itself which is called Id_a.

4. Every arrow has an inverse.

In other words a groupoid is a (small) category in which every arrow is an isomorphism.

Let us explain in which way groupoids are generalized groups. For a group G we can draw the following graph : there is only one vertex, and each arrow represent an element of the group, the composition being just the group multiplication. Thus a groupoid is a sort of group with many local units instead of a global one, and a partially defined multiplication. (Remark: we can dispense with the vertices and use the local units as a substitute, since there is a bijection between local units and vertices. In this way we work only with a bunch of arrows.)

There are two fundamental examples of groupoids :

- An equivalence relation, when seen as a graph, is obviously a groupoid.
- Let us consider a group G acting on a set X. We take X as the set of vertices, an we put an arrow labeled by $g \in G$ between x and y iff $y = g.x$. This is called the *transformation groupoid* of the group action. Note that the loops at a vertex v correspond to the stabilizer (or isotropy subgroup) of v.

In fact the two examples are related : given a group action we have the equivalence relation whose classes are G-orbits. Conversely, at least at the level of discrete spaces, one can always consider an equivalence relation to be given by a product group $G_1 \times G_2 \times \ldots$ where G_i is a group acting transitively on the i-th class (for instance the group of permutation of this class).

Just as groups can be used to define group algebras, groupoids give rise to groupoid algebras thanks to the convolution product. Let $\mathcal{G}$ be a groupoid and $\mathbb{C}\mathcal{G}$ be the set of functions from $\mathcal{G}$ to $\mathbb{C}$ with finite support (we consider here only discrete groupoids for the sake of simplicity). Define $f * g$ by

$$f * g(\gamma) = \sum_{\gamma_1 \circ \gamma_2 = \gamma} f(\gamma_1) g(\gamma_2) \tag{6}$$

If we identify $\gamma \in \mathcal{G}$ with $\mathbf{1}_\gamma \in \mathbb{C}\mathcal{G}$, we see that the convolution product is just an extension to linear combinations of the composition product of the groupoid (with the convention that $\gamma_1 * \gamma_2 = 0$ if the arrows are not composable). In particular the groupoid algebra is unital iff the groupoid has a finite number of vertices and $1 = \sum_{v \text{ vertex}} \mathrm{Id}_v$.

It is time now for some concrete examples. Let $V = \{1; 2; 3\}$. On V we consider :

- The equivalence relation $1 \sim 2$. Call the corresponding groupoid $\mathcal{G}_1$.
- The action of the group $G = \{e; \tau\}$ of permutation of V which leave 3 fixed. Call the corresponding groupoid $\mathcal{G}_2$.

We thus obtain two groupoids which are almost the same : the only difference is that the second one has two loops on 3 instead of one. Of course $\sim$ is the equivalence relation whose classes are the orbit under the action of G.

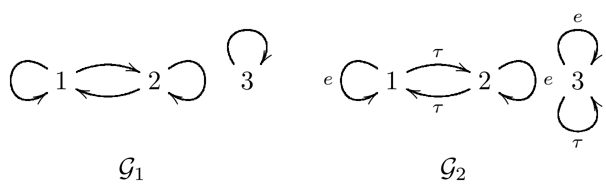

Figure 2: Two groupoids.

Let us see what the two groupoid algebras look like. Obviously, product of arrows from the two different components of the graphs is zero. So we have an algebra which is a direct sum of the form $A \oplus B$. Let us concentrate on the component of the two first vertices.

The equivalence relation $\sim$ is then the set of all pairs (i, j) with $i, j = 1, 2$. Let us associate with (i, j) the elementary matrix $E_{ij} \in M_2(\mathbb{C})$. We immediately see that the product $(i, j)(k, l) = \delta_{jk}(i, l)$ in $\mathbb{C}\mathcal{G}_{1,2}$ translates into the matrix product. Using linearity we see that $A \simeq M_2(\mathbb{C})$.

Now we turn our attention to the second component. For $i = 1$, it is obvious that $B = \mathbb{C}$. For $i = 2$, using obvious notations, let us associate $e_3 \mapsto \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $\tau_3 \mapsto \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. We can verify that it extends to an isomorphism of algebra $B \simeq \mathrm{Diag}(M_2(\mathbb{C})) \simeq \mathbb{C} \oplus \mathbb{C}$. Thus we get

$$\mathbb{C}\mathcal{G}_1 \simeq M_2(\mathbb{C}) \oplus \mathbb{C}; \qquad \mathbb{C}\mathcal{G}_2 \simeq M_2(\mathbb{C}) \oplus \mathbb{C} \oplus \mathbb{C}$$

The idea of noncommutative quotient is that the groupoid C^*-algebra is a replacement for the algebra $\mathcal{C}(V/G)$. Here we see that $\mathbb{C}\mathcal{G}$ contains more information than $\mathcal{C}(V/G)$ which is $\mathbb{C}^2$ in both cases : the groupoid algebra "remembers" the number of ways in which two points were equivalent. We see also that the structure space of $\mathbb{C}\mathcal{G}_1$ is homeomorphic to $V/\sim$ (a two-points discrete space), while $\widehat{\mathbb{C}\mathcal{G}_2}$ has an additional point : this is because the action of G on V is not free.

Relation between classical and noncommutative quotient in the good cases :

Theorem 2 *(Rieffel) If G acts freely and properly on a compact Haussdorf space X, then $\mathcal{C}(X/G)$ and $\mathcal{C}(X) \rtimes_r G$ are strongly Morita equivalent.*

Let us look at an example showing the interest of the construction. Let $X = [0; 1] \times \{a; b\}$ be the disjoint union of two unit segments. We define the equivalence relation by $(x, a) \sim (x, b)$ for $x \in]0; 1[$. The quotient space is the set $\{[x] := [(x, a)] = [(x, b)] | x \in]0; 1[\} \cup \{0_a; 0_b; 1_a; 1_b\}$ where $0_a = [(0; a)]$ etc. A continuous function of $X/\sim$ is of the form $[x] \mapsto f(x)$ where f is a continous function on X such that $f(x, a) = f(x, b)$ for $x \in]0; 1[$. Taking the limit $x \to 0/1$ we see that $f(0, a) = f(0, b)$, thus $f(0_a) = f(0_b)$, the same for $1_a, 1_b$. Hence

$\mathcal{C}(X/\sim) = \mathcal{C}([0;1])$. The extra end points are not seen by the continuous functions on the quotient : they just see a segment with two endpoints.

Figure 3: A bad quotient.

On the contrary, it is easy to see that the groupoid algebra is

$$\{f \in \mathcal{C}([0;1], M_2(\mathbb{C})) | f(0) \text{ and } f(1) \text{ are diagonal}\}$$

Indeed we can see it at the algebraic level if we replace $[0;1]$ by a finite discrete space $V = \{1; \ldots; n\}$. As the example 1 above, the groupoid algebra will be $\mathbb{C} \oplus \mathbb{C} \oplus \bigoplus_{i=2}^{n-1} M_2(\mathbb{C}) \oplus \mathbb{C} \oplus \mathbb{C}$ which is isomorphic to $\{f : V \to M_2(\mathbb{C}) | f(0), f(1)$ diagonal$\}$. Careful considerations of topology permit to extend this to the continous case.

Remark : The noncommutative quotients builded above would really merit to be called "quantum quotient". Indeed, the way points are identified classically is via an equivalence relation, but here they are made equivalent by *allowing every complex linear combinations* of them. This is just the superposition principle applied to points of a space !

A.2 The noncommutative torus

Let $\mathbb{T}$ be $\mathbb{R}^2/\mathbb{Z}^2$. Let $\theta \in \mathbb{R} \setminus \mathbb{Q}$. The *Kronecker foliation* is the partition of $\mathbb{T}$ into (the image in the quotient of) straight lines of slope θ (such a line is called an irrational helix of the torus). Let $\sim$ be the equivalence relation on $\mathbb{T}$ defined by $x \simeq y$ iff x and y lie on the same leaf of the Kronecker foliation ($\Leftrightarrow$ the slope of the line from x to y is θ). Let us call $X = \mathbb{T}/\simeq$.

We can see at once that the quotient topology on X is the trivial topology. An open set U of X has an inverse image by p which is a saturated open set under $\simeq$. However such a saturated open set is either $\emptyset$ or $\mathbb{T}$. Thus the only continuous functions on X are the constant ones. This topology, or if one wants, the algebra of continuous functions, does not distinguishes X from a one-point space. This is true even of measurable functions. The reason is the *ergodic* property of the foliation. It is easy to see it by restricting the foliation to a transverse S^1, for instance the set $(\alpha; 0) \in (\mathbb{R}/\mathbb{Z}) \times \{0\}$. We have $(\alpha; 0) \simeq (\alpha'; 0)$ iff $\alpha' - \alpha \in \theta\mathbb{Z}$, thus we see that $X = \mathbb{T}/\simeq \approx S^1/\sim$ where $\sim$ is the equivalence relation defined by the action of the rotation of angle θ. But one can prove the ergodic property :

$$\frac{1}{2\pi} \int_0^{2\pi} f(e^{i\phi}) d\phi = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} f(e^{i\alpha + k\theta})$$

for any integrable function f on S^1, for almost every α. Now if f is defined on X, it means that it is constant on each orbit, hence the RHS is equal to $f(e^{i\alpha})$. Since the LHS is a constant, we see that an integrable function is constant on every orbit iff it is constant almost everywhere on the circle. In fact this is true also for functions which are just measurable.

Since the usual continuous functions do not give any useful information about X, the trick will be to use the philosophy of noncommutative quotients sketched above. Hence, we will replace the useless algebra $\mathcal{C}(X) = \mathbb{C}$ by the groupoid algebra A_θ of the equivalence relation $\simeq$ on $\mathbb{T}^2$, or equivalently, of $\sim$ on S^1. The elements of A_θ are functions on the graph of $\sim$, thus are of the form $(\alpha, \alpha') \mapsto a(\alpha, \alpha')$, with $\alpha \sim \alpha' \Leftrightarrow \exists n \in \mathbb{Z}, \alpha' = \alpha + n\theta$. Hence they can also be viewed as functions $(\alpha, n) \mapsto a(\alpha, n)$ from $S^1 \times \mathbb{Z}$ to $\mathbb{C}$, which are continuous in the first variable. These functions are multiplied with the convolution product (6) which here takes the form :

$$a \star b(\alpha, n) = \sum_{k\in\mathbb{Z}} a(\alpha, k) b(\alpha + k\theta, n - k) \tag{7}$$

Note that the $*$ operation which is normally $a^*(\alpha, \alpha') = \bar{a}(\alpha', \alpha)$ (matrix adjoint) becomes in the representation we are using :

$$a^*(\alpha, n) = \bar{a}(\alpha + n\theta, -n) \tag{8}$$

Let us define the following elements of A_θ :

$$U(\alpha, n) = 1 \text{ if } n = 1, U(\alpha, n) = 0 \text{ otherwise} \tag{9}$$

$$V(\alpha, n) = e^{2i\pi\alpha} \text{ if } n = 0, V(\alpha, n) = 0 \text{ otherwise} \tag{10}$$

It is then easy to check using (7) and (8) that (dropping the $\star$ to simplify the notations) :

$$\begin{array}{rcl} UU^* = U^*U = 1 & ; & VV^* = V^*V = 1 \\ UV & = & e^{2i\pi\theta} VU \end{array} \tag{11}$$

Where 1 means the unit of A_θ which is the function $(\alpha, n) \mapsto \delta_{n,0}$.

One also checks that $U^n = \delta_{n,0}$ for every $n \in \mathbb{Z}$, so that each element of A_θ can be written

$$f = \sum_{n\in\mathbb{Z}} f_n U^n$$

where $f_n \in \mathcal{C}(S^1)$. Since $V^k(\alpha, n) = e^{2ik\pi\alpha}\delta_{n,0}$ for all $n \in \mathbb{Z}$, one has by Fourier $f_n = \sum_{k\in\mathbb{Z}} f_{n,k} V^k$. Thus

$$f = \sum_{n,k\in\mathbb{Z}} f_{n,k} V^k U^n \tag{12}$$

We have been kind of sloppy about convergence up to now. Let us just mention that if we take $(f_{n,k})$ to be a sequence in the Schwartz space of rapid decay, we get a $*$-subalgebra $\mathcal{A}_\theta$ which is just $\mathcal{C}^\infty(\mathbb{T})$ in the commutative case. Moreover, $\|f\| := \sup|f_{n,k}|$ is a norm on $\mathcal{A}_\theta$ and A_θ is its completion with respect to this norm. Thus A_θ consists of guys like (12) with $|f_{n,k}|$ bounded. On a purely algebraic level, A_θ is the universal C^*-algebra generated by elements satisfying (11).

It should be clear why this beast is called "a noncommutative torus" (one should add "algebra"). In the case $\theta = 0$ one has the universal C^*-algebra generated by two commuting unitaries U and V. From this algebraic setting we will recover the usual torus, as an exercise. We know from the GN theorem that $A_0 = \mathcal{C}(K)$ where K is the character space of A_0, so all we have to do is prove that $K = \mathbb{T}^2$. However, thanks to the universal property of A_0, in order to define a character χ, it suffices to define it on U and V. The only restriction on $\chi(U)$ comes from the relation $UU^* = 1$ which forces $\chi(U)$ to satisfy $|\chi(U)|^2 = 1$, and similarly for $\chi(V)$.

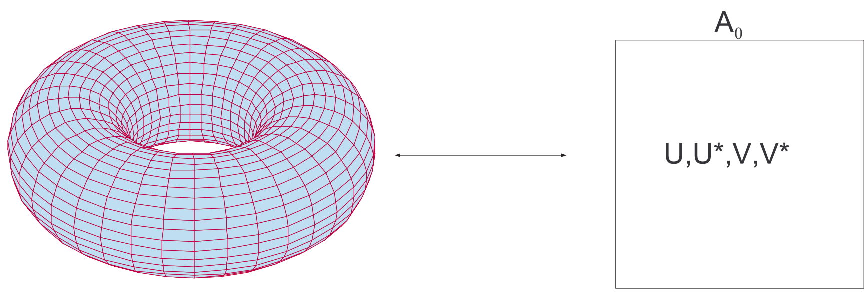

Figure 4: The commutative torus, which canonically corresponds to its algebra of continuous functions.

Thus $\Psi : K \to \mathbb{T}^2$, $\Psi(\chi) = (\chi(U), \chi(V))$ is a bijection. There remains to prove that it is also an homeomorphism. Since we are dealing with compact metrizable spaces, it will suffice to prove that Ψ respects convergence of sequences (hence is continous, hence is an homeomorphism). Now $\chi_n \to \chi$ means $\chi_n(a) \to \chi(a)$ for every $a \in A_0$, in particular $\chi_n(U) \to \chi(U)$ and $\chi_n(V) \to \chi(V)$. Thus $\chi_n \to \chi \Rightarrow \Psi(\chi_n) \to \Psi(\chi)$.

Final remarks :

- For θ rational we obtain an algebra which is strongly Morita equivalent to the algebra of a commutative torus.
- We have restricted our attention from the equivalence relation given by the Kronecker foliation on $\mathbb{T}^2$ to the one given by the iteration of an irrational rotation on S^1, but what if we carry on doing the noncommutative quotient with the torus instead ? Well, in this case we end up with an algebra B_θ which is $B_\theta = A_\theta \otimes K(H)$.

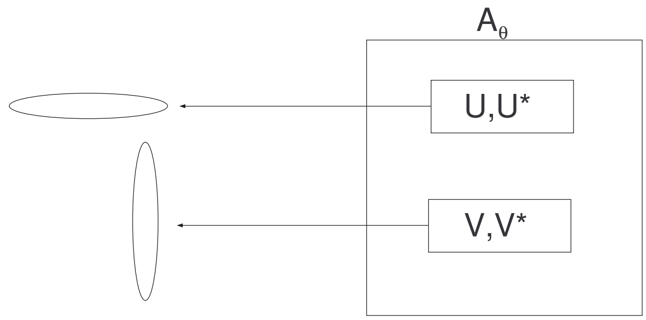

Figure 5: The noncommutative torus algebra : it contains two commutative sub-algebras which have a circle as spectrum, but the two circles cannot be integrated into a torus because the two-subalgebras do not commute.

- The noncommutative torus (and higher dimensional NC tori), have attracted much attention since its introduction, and many things are known about it. However, to my knowledge, its state space has not yet been determined.

References

[Red Book] A. Connes, *Noncommutative Geometry*, Academic Press (1994)

[Landi] G. Landi, *An Introduction to Noncommutative Spaces and Their Geometries*, Springer LNP (1997) hep-th/9701078

[K-R] R. Kadison, J. Ringrose, *Fundamentals of the Theory of Operator Algebras*, AMS (1997)

[Co 95] A. Connes, *Noncommutative Geometry and Reality*, Journal of Mathematical Physics (1995)

[Ch-Iv 07] E. Christensen, C. Ivan, *Sums of Two Dimensional Spectral Triples*, Math. Scand. **100** (2007), 35-60

[I-K-M 01] B. Iochum, T. Krajewski, P. Martinetti, *Distances in Finite Spaces from Noncommutative Geometry*, Journ. of Geom. Phys. **37** (2001), 100-125, hep-th/9912217

[Co 08] A. Connes, *On the Spectral Characterization of Manifolds*, abs/0810.2088

[GB-V-F] J. Gracia-Bondía, J. Varilly, H. Figueroa, *Elements of Noncommutative Geometry*, Birkhäuser (2000)

[Ch-Co 96] A. Chamseddine, A. Connes, *The spectral action principle*, Commun. Math. Phys. **186** (1996) 737-750 hep-th/9606001

[Co 06] A. Connes, *Noncommutative geometry and the standard model with neutrino mixing*, JHEP 0611:081 (2006) hep-th/0608226

[Ba 07] J. Barrett, *Lorentzian version of the noncommutative geometry of the standard model of particle physics*, Journal of mathematical physics (2007) hep-th/0608221

[Fr 11] N. Franco, *Lorentzian approach to noncommutative geometry*, PhD Thesis abs/1108.0592 (2011)

[Be 09] F. Besnard, *A noncommutative view on topology and order*, J. Geom. Phys. **59**, pp. 861–875 (2009) abs/0804.3551

[Ch-Co 08] A. Chamseddine, A. Connes, *Why the standard model ?*, J. Geom. Phys. **58**, pp. 38–47 (2008), abs/0706.3688

[Kh 04] M. Khalkhali, *Very basic noncommutative geometry* math/0408416 (Extended version : *Basic noncommutative geometry*, EMS (2009)