

A Presentation of Certain New Trends in Noncommutative Geometry

Réamonn Ó Buachalla

Thesis submitted for MSc by Research
National University of Ireland, Cork
2006

Department of Mathematics
Faculty of Science

Head of Department: Prof. J. Berndt
Supervisors: Prof. G. J. Murphy, Dr S. Wills

Contents

1	Motivation	7
1.1	The Gelfand–Naimark Theorem	7
1.1.1	C^* -algebras	7
1.1.2	The Gelfand Transform	11
1.1.3	Gelfand–Naimark Theorem	13
1.1.4	The Algebra–Space Correspondence	14
1.2	Noncommutative Topology	15
1.2.1	Some Noncommutative Generalisations	16
1.3	Vector Bundles and Projective Modules	23
1.3.1	Vector Bundles	23
1.3.2	Standard results	27
1.3.3	Finite Projective Modules	28
1.3.4	Serre–Swan Theorem	29
1.4	Von Neumann Algebras	33
1.4.1	Noncommutative Measure Theory	34
1.5	Summary	37
2	Differential Calculi	38
2.1	The de Rham Calculus	39
2.1.1	The Exterior Derivative	40
2.2	Differential calculi	41
2.2.1	The Universal Calculus	43
2.3	Derivations-Based Differential Calculi	45
3	Cyclic Cohomology and Quantum Groups	49
3.1	Cyclic Cohomology	49

3.1.1	The Chern Character	49
3.1.2	Traces and Cycles	52
3.1.3	(Co)Chain Complex (Co)Homology	55
3.1.4	Hochschild (Co)Homology	58
3.1.5	Cyclic (Co)Homology	62
3.1.6	The Chern–Connes Character Maps	65
3.2	Compact Quantum Groups	66
3.2.1	Compact Quantum Groups	66
3.2.2	Differential Calculi over Quantum Groups	74
3.2.3	Twisted Cyclic Cohomology	77
3.2.4	Twisted Hochschild Homology and Dimension Drop	79
4	Dirac Operators	81
4.1	Euclidean and Geometric Dirac Operators	83
4.1.1	Clifford Algebras	83
4.1.2	Euclidean Dirac Operators	85
4.1.3	Geometric Dirac Operators	86
4.1.4	Spin Manifolds and Dirac Operators	91
4.1.5	Properties of the Atiyah–Singer–Dirac Operator	95
4.2	Spectral Triples	97
4.2.1	The Noncommutative Riemannian Integral	99
4.2.2	Connes’ State Space Metric	100
4.3	Dirac Operators and Quantum Groups	102
4.3.1	The Dirac and Laplace Operators	102
4.3.2	The Hodge Decomposition	104
4.3.3	The Hodge Operator	105
4.3.4	A Dirac Operator on Woronowicz’s Calculus	106
5	Fuzzy Physics	107
5.1	Quantum Field Theory	107
5.2	Noncommutative Regularization	108
5.2.1	The Fuzzy Sphere	109
5.2.2	Fuzzy Coadjoint Orbits	112
5.2.3	Fuzzy Physics	114

6	Compact Quantum Metric Spaces	116
6.1	Compact Quantum Metric Spaces	117
6.1.1	Noncommutative Metrics and the State Space	117
6.1.2	Lipschitz Seminorms	118
6.1.3	Order-unit Spaces	119
6.1.4	Compact Quantum Metric Spaces	121
6.1.5	Examples	122
6.1.6	Spectral Triples	124
6.2	Quantum Gromov–Hausdorff distance	125
6.2.1	Gromov–Hausdorff Distance	125
6.2.2	Quantum Gromov–Hausdorff	126
6.2.3	Examples	129
6.3	Matrix Algebras Converging to the Sphere	129
6.3.1	The Berezin Covariant Transform	131
6.3.2	The Berezin Contravariant Transform	133
6.3.3	Estimating the QGH Distance	139
6.3.4	Matrix Algebras Converging to the Sphere	140
6.4	Matricial Gromov–Hausdorff Distance	140

Introduction

Einstein was always rather hostile to quantum mechanics. How can one understand this? I think it is very easy to understand, because Einstein had proceeded on different lines, lines of pure geometry. He had been developing geometrical theories and had achieved enormous success. It is only natural that he should think that further problems of physics should be solved by further development of geometrical ideas. How to have $a \times b$ not equal to $b \times a$ is something that does not fit in very well with geometric ideas; hence his hostility to it. P. A. M. Dirac

The development of quantum mechanics in the first half of the twentieth century completely revolutionized classical physics. In retrospect, its effect on many areas of mathematics has been no less profound. The mathematical formalism of quantum mechanics was constructed immediately after its birth by John von Neumann. This formulation in turn gave birth to the theory of operator algebras. In the following years the work of mathematicians such as Israil Gelfand and Mark Naimark showed that by concentrating on the function algebra of a space, rather than on the space itself, many familiar mathematical objects in topology and measure theory could be understood as ‘commutative versions’ of operator algebra-structures. The study of the noncommutative versions of these structures then became loosely known as *noncommutative topology* and *noncommutative measure theory* respectively. In the first chapter of this thesis we shall present the work of von Neumann, Gelfand, Naimark and others in this area.

From the middle of the twentieth century on, geometers like Grothendieck, Atiyah, Hirzebruch and Bott began to have great success using algebraic formulations of geometric concepts that were expressed in terms of the function algebras of spaces. *K*-theory, which will be discussed in Chapter 3, is one important example. Later, it was realised that these algebraic formulations are well defined for any algebra, and that the structures involved had important roles to play in the operator theory. We cite algebraic *K*-theory, and Brown, Fillmore, and Douglas’ *K*-homology as primary examples. However, the first person to make a serious attempt to interpret all of this work as a noncommutative version of differential geometry was a French mathematician named Alain Connes. Connes first came to prominence

in the 1970s and he was awarded the Fields medal in 1982 for his work on von Neumann algebras. Connes advanced the mathematics that existed on the borderlines between differential geometry and operator algebras further than anyone had previously imagined possible. In 1994 he published a book called *Noncommutative Geometry* [12] that gave an expository account of his work up to that time. It was an expanded translation of a book he had written in French some years earlier called *Géométrie Noncommutative*. Noncommutative Geometry was hailed as a ‘milestone for mathematics’ by Connes’ fellow Fields medalist Vaughan Jones.

Connes developed noncommutative geometry to deal with certain spaces called singular spaces. These arise naturally in many problems in classical mathematics (usually as quotient spaces) but they are badly behaved from the point of view of the classical tools of mathematics, such as measure theory, topology and differential geometry. For example, the space, as a topological space, may not be Hausdorff, or its natural topology may even be the coarse one; consequently, the tools of topology are effectively useless for the study of the problem at hand. The idea of noncommutative geometry is to replace such a space by a canonically corresponding noncommutative C^* -algebra and to tackle the problem by means of the formidable tools available in noncommutative geometry. This approach has had enormous success in the last two decades.

We give the following simple example: Let X be a compact Hausdorff space and let G be a discrete group acting upon it. If the action is sufficiently complicated then the quotient topology on X/G can fail to separate orbits, and in extreme cases X/G can even have the indiscrete topology. Thus, the traditional tools of mathematics will be of little use in its examination. However, the noncommutative algebra $C(X) \rtimes G$, the crossed product of G with $C(X)$, is a powerful tool for the study of X/G .

In the third and fourth chapters of this thesis we shall present the fundamentals of Connes’ work.

In Connes’ wake, noncommutative geometry has become an extremely active area of mathematics. Applications have been found in fields as diverse as particle physics and the study of the Riemann hypothesis. At the start of the decade a group of European universities and research institutes across seven countries formed an alliance for the purposes of co-operation and collaboration in research in the theory of operator algebras and noncommutative geometry; the alliance was called the *European Union Operator Algebras Network*. The primary aim of this thesis is to examine some of the areas currently being explored by the Irish based members of the network. The Irish institutions in question are the National University of Ireland, Cork, and the Dublin Institute for Advanced Studies (DIAS). In Chapters 3 and 4 we shall present some of the work being pursued in Cork.

Specifically, we shall present the work being done on the interaction between Connes' noncommutative geometry and compact quantum groups. Compact quantum groups are a noncommutative generalisation of compact groups formulated by the Polish mathematician and physicist Stanislaw Lech Woronowicz. The relationship between these two theories is troublesome and ill understood. However, as a research area it is showing great promise and is at present the subject of very active investigation.

In the fifth chapter we shall present an overview of some of the work being done in Dublin. We shall focus on the efforts being made there to use noncommutative methods in the renormalisation of quantum field theory. This area is also a very active area of research; it is known as *fuzzy physics*. We shall not, however, present a detailed account of the DIAS work. Its highly physical nature is not well suited to a pure mathematical treatment. Instead, an exposition of Marc Rieffel's compact quantum metric space theory will be given in the sixth chapter. Compact quantum metric spaces were formulated by Rieffel with the specific intention of providing a sound mathematical framework in which to discuss fuzzy physics. Indeed, the papers of DIAS members are often cited by Rieffel when he discusses compact quantum metric spaces.

Chapter 1

Motivation

As explained in the introduction, noncommutative geometry is based upon the fact that there exist a number of correspondences between basic mathematical structures and the commutative versions of certain operator algebraic structures. In this chapter we shall present the three prototypical examples of these correspondences: the Gelfand–Naimark Theorem, the Serre–Swan Theorem, and the characterisation of commutative von Neumann Algebras. We shall also show how these results naturally lead us to define ‘noncommutative versions’ of the mathematical structures in question.

1.1 The Gelfand–Naimark Theorem

The Gelfand–Naimark Theorem is often regarded as the founding theorem of noncommutative geometry. Of all the results that we shall present in this chapter it was the first to be formulated, and, from our point of view, it is the most important. It is the theorem that motivated mathematicians to consider the idea of a ‘noncommutative generalisation’ of locally compact Hausdorff spaces.

The theorem, and indeed most of noncommutative geometry, is expressed in the language of C^* -algebras. Thus, we shall begin with an exposition of the basics of these algebras, and then progress to a proof of the theorem itself.

1.1.1 C^* -algebras

The motivating example of a C^* -algebra is $B(H)$ the normed algebra of bounded linear operators on a Hilbert space H . The algebra operations of $B(H)$ are defined pointwise and its norm, which is called the *operator norm*, is defined by $\|A\| = \sup\{\|A(x)\| : \|x\| \leq 1\}$, $A \in B(H)$. With respect to these definitions,

$B(H)$ is a *Banach algebra*; that is, an algebra with a complete submultiplicative norm. Completeness of the norm is a standard result in Hilbert space theory, and submultiplicativity follows from the fact that

$$\|AB\| = \sup_{\|x\| \leq 1} \|AB(x)\| \leq \sup_{\|x\| \leq 1} (\|A\| \|B(x)\|) = \|A\| \|B\|,$$

for all $A, B \in B(H)$.

A central feature of Hilbert space operator theory is the fact that for every $A \in B(H)$, there exists an operator $A^* \in B(H)$, called the *adjoint* of A , such that

$$\langle Ax, y \rangle = \langle x, A^*y \rangle, \quad \text{for all } x, y \in H.$$

As is well known, and easily verified, $(A + \lambda B)^* = A^* + \bar{\lambda}B^*$, $A^{**} = A$, and $(AB)^* = B^*A^*$, for all $A, B \in B(H)$, $\lambda \in \mathbf{C}$. Another well known, and very important, equation that involves an operator and its adjoint is

$$\|A^*A\| = \|A\|^2, \quad \text{for all } A \in B(H).$$

Now that we have reviewed the relevant features of Hilbert space operator theory, we are ready to begin generalising. A **-algebra* is an algebra A , together with a mapping

$$* : A \rightarrow A, \quad a \mapsto a^*,$$

called an *algebra involution*, such that for all $a, b \in A$, $\lambda \in \mathbf{C}$;

1. $(a + \lambda b)^* = a^* + \bar{\lambda}b^*$ (*conjugate-linearity*),
2. $a^{**} = a$ (*involutivity*),
3. $(ab)^* = b^*a^*$ (*anti-multiplicativity*).

If A and B are two *-algebras and φ is an algebra homomorphism from A to B , then φ is called a **-algebra homomorphism* if $\varphi(a^*) = \varphi(a)^*$, for all $a \in A$. If the homomorphism is bijective, then it is called a **-isomorphism*. If $a \in A$ such that $a^* = a$, then a is called *self-adjoint*; we denote the subset of self-adjoint elements of A by A_{sa} . If X is a subset of A such that $X^* = \{x^* : x \in X\} = X$, then X is called a *self-adjoint subset* of A .

Definition 1.1.1. A *C^* -algebra* \mathcal{A} is a *-algebra that is also a Banach algebra, such that

$$\|a^*a\| = \|a\|^2, \quad \text{for all } a \in \mathcal{A}.$$

From our comments above we see that $B(H)$ is a C^* -algebra for every Hilbert space H . The simplest example of a C^* -algebra is the complex numbers \mathbf{C} . Any closed self-adjoint subalgebra of a C^* -algebra is clearly also a C^* -algebra.

One of the most important examples is $C_0(X)$, the algebra of continuous functions that vanish at infinity on a locally compact Hausdorff space X . (We recall that for a function $f \in C_0(X)$, to *vanish at infinity* means that for each $\varepsilon > 0$ there exists a compact subset $K \subseteq X$ such that $|f(x)| < \varepsilon$, for all $x \notin K$.) If we endow $C_0(X)$ with the standard supremum norm and define

$$f^*(x) = \overline{f(x)}, \quad \text{for all } x \in X,$$

then the conditions for a C^* -algebra are clearly fulfilled.

This algebra is one of the standard examples of a C^* -algebra that is not necessarily unital. Notice that if X is non-compact, then $1 \notin C_0(X)$. On the other hand, if X is compact, then $C_0(X) = C(X)$, the algebra of continuous functions on X , and so it is unital.

We should also notice that $C_0(X)$ is a commutative C^* -algebra. In fact, as we shall see later, all commutative C^* -algebras are of this form.

The Universal Representation

As we have seen, C^* -algebras generalise the algebras of bounded linear operators on Hilbert spaces. A *representation* of a C^* -algebra \mathcal{A} is a pair (U, H) where H is a Hilbert space and $U : \mathcal{A} \rightarrow B(H)$ is a $*$ -algebra homomorphism. If U is injective, then the representation is called *faithful*. What is interesting to note is that for any C^* -algebra \mathcal{A} there exists a distinguished faithful representation of \mathcal{A} called its *universal representation*. Thus, every C^* -algebra is isometrically $*$ -isomorphic to a closed $*$ -subalgebra of $B(H)$, for some Hilbert space H . This result is due to Gelfand and Naimark, for further details see [80].

Unitisation

If \mathcal{A} is a non-unital C^* -algebra, then it often proves useful to embed it into a unital C^* -algebra using a process known as *unitisation*: One starts with the linear space $\mathcal{A} \oplus \mathbf{C}$ and defines a multiplication on it by setting

$$(a, \lambda).(b, \mu) = (ab + \lambda b + \mu a, \lambda\mu), \tag{1.1}$$

and an involution by setting

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

One then defines a norm on it by setting

$$\|(a, \lambda)\| = \sup\{\|ab + \lambda b\| : b \in \mathcal{A}, \|b\| \leq 1\}.$$

It is not too hard to show that this norm makes it a C^* -algebra. (As we shall see later, this norm is necessarily unique.) The element $(0, 1)$ clearly acts as a unit. We denote this new unital C^* -algebra by $\tilde{\mathcal{A}}$, and embed \mathcal{A} into it in the canonical manner. It is easily seen to be an isometric embedding.

Spectrum of an Element

Let \mathcal{A} be a unital C^* -algebra and $a \in \mathcal{A}$. We define the *spectrum* of a to be the set

$$\sigma(a) = \{\lambda \in \mathbf{C} : (\lambda 1 - a) \notin \text{Inv}(\mathcal{A})\},$$

where $\text{Inv}(\mathcal{A})$ is the set of invertible elements of \mathcal{A} . If a is an element of a non-unital algebra then we define its *spectrum* to be the set

$$\sigma(a) = \{\lambda \in \mathbf{C} : (\lambda 1 - a) \notin \text{Inv}(\tilde{\mathcal{A}})\}.$$

Three basic facts about the spectrum of an element are that: $\sigma(a) \neq \emptyset$, for any $a \in \mathcal{A}$; if $a = a^*$, then $\sigma(a) \subseteq \mathbf{R}$; and if $\lambda \in \sigma(a)$, then $|\lambda| \leq \|a\|$. (These results hold for the unital or non-unital definitions of the spectrum.)

We define the *spectral radius* of an element a to be

$$r(a) = \sup\{|\lambda| : \lambda \in \sigma(a)\};$$

note that this is a purely algebraic definition and takes no account of the algebra's norm. It is a well known result of Beurling that, for any C^* -algebra \mathcal{A} ,

$$r(a) = \lim_{n \rightarrow \infty} \|a^n\|^{\frac{1}{n}}, \quad \text{for all } a \in \mathcal{A}.$$

Now, if $a \in \mathcal{A}$ is self-adjoint, then $\|a^2\| = \|a^*a\| = \|a\|^2$. Thus,

$$r(a) = \lim_{n \rightarrow \infty} \|a^n\|^{\frac{1}{n}} = \lim_{n \rightarrow \infty} \|a^{2^n}\|^{\frac{1}{2^n}} = \lim_{n \rightarrow \infty} \|a\| = \|a\|. \quad (1.2)$$

This result will be of use to us in our proof of the Gelfand–Naimark Theorem.

As a more immediate application, we can use it to show that there is at most one norm on $*$ -algebra making it a C^* -algebra. Suppose that $\|\cdot\|_1$ and $\|\cdot\|_2$ are two norms on a $*$ -algebra \mathcal{A} making a C^* -algebra. Then they must be equal since

$$\|a\|_1^2 = \|aa^*\|_1 = r(aa^*) = \|aa^*\|_2 = \|a\|_2^2,$$

for all $a \in \mathcal{A}$. A useful consequence of this fact is that any $*$ -isomorphism is an isometric mapping.

1.1.2 The Gelfand Transform

A *character* φ on an algebra A is a non-zero algebra homomorphism from \mathcal{A} to \mathbf{C} ; that is, a non-zero linear functional φ such that

$$\varphi(ab) = \varphi(a)\varphi(b), \quad \text{for all } a, b \in A.$$

We denote the set of all characters on A by $\Omega(A)$.

Let \mathcal{A} be a unital C^* -algebra and $\varphi \in \Omega(\mathcal{A})$, then $\varphi(a) = \varphi(a1) = \varphi(a)\varphi(1)$, for all $a \in \mathcal{A}$, and so $\varphi(1) = 1$. If $a \in \text{Inv}(\mathcal{A})$, then $1 = \varphi(aa^{-1}) = \varphi(a)\varphi(a^{-1})$. Thus, if $\varphi(a) = 0$, a cannot be invertible. The fact that $\varphi(\varphi(a) - a) = 0$ then implies that $(\varphi(a) - a) \notin \text{Inv}(\mathcal{A})$. Thus, $\varphi(a) \in \sigma(a)$, for all $a \in \mathcal{A}$, and for all $\varphi \in \Omega(\mathcal{A})$. If we now recall that $|\lambda| \leq \|a\|$, for all $\lambda \in \sigma(a)$, then we see that

$$\|\varphi\| = \sup\{|\varphi(a)| : \|a\| \leq 1, a \in \mathcal{A}\} \leq 1. \quad (1.3)$$

(In fact, in the unital case $\|\varphi\| = 1$, since $\varphi(1) = 1$.) It follows that each $\varphi \in \Omega(\mathcal{A})$ is norm continuous, and that $\Omega(\mathcal{A})$ is contained in \mathcal{A}_1^* ; where \mathcal{A}^* is the space of bounded linear functionals on \mathcal{A} , and \mathcal{A}_1^* is the closed unit ball of \mathcal{A}^* .

Another fact about characters is that they are *Hermitian*, that is, $\varphi(a^*) = \overline{\varphi(a)}$, for all $\varphi \in \Omega(\mathcal{A})$. This is shown as follows: every element a of \mathcal{A} can be written uniquely in the form $a = a_1 + ia_2$, where a_1 and a_2 are the two self-adjoint elements

$$a_1 = \frac{1}{2}(a + a^*), \quad a_2 = \frac{1}{2i}(a - a^*).$$

Clearly,

$$\varphi(a^*) = \varphi((a_1 + ia_2)^*) = \varphi(a_1 - ia_2) = \varphi(a_1) - i\varphi(a_2).$$

Then, since $\varphi(a_i) \in \sigma(a_i)$, and $\sigma(a_i) \subset \mathbf{R}$, for $i = 1, 2$, we have that

$$\varphi(a_1) - i\varphi(a_2) = \overline{\varphi(a_1) + i\varphi(a_2)} = \overline{\varphi(a)}.$$

Gelfand Topology

We now find it convenient to endow $\Omega(\mathcal{A})$ with a topology different from the norm topology. We do this by putting the weak* topology on \mathcal{A}^* , and restricting it to $\Omega(\mathcal{A})$. Recall that the *weak* topology* on \mathcal{A}^* is the weakest topology with respect to which all maps of the form

$$\widehat{a} : \mathcal{A}^* \rightarrow \mathbf{C}, \quad \varphi \mapsto \varphi(a), \quad a \in \mathcal{A}, \quad (1.4)$$

are continuous. Of course, we may alternatively describe it as the weakest topology with respect to which a net $\{\varphi_\lambda\}_\lambda$ in \mathcal{A}^* converges to $\varphi \in \mathcal{A}^*$ if, and only if, $\widehat{a}(\varphi_\lambda) \rightarrow \widehat{a}(\varphi)$; or equivalently, if, and only if, $\varphi_\lambda(a) \rightarrow \varphi(a)$, for all $a \in \mathcal{A}$.

When $\Omega(\mathcal{A})$ is endowed with this topology we call it the *spectrum* of \mathcal{A} . We call the topology itself the *Gelfand topology*. We should note that the spectrum is a Hausdorff space since the weak* topology is a Hausdorff topology.

With respect to the weak* topology $\Omega(\mathcal{A}) \cup \{0\}$ is closed in \mathcal{A}^* . To see this take a net $\{\varphi_\lambda\}_\lambda$ in $\Omega(\mathcal{A}) \cup \{0\}$ that converges to $\varphi \in \mathcal{A}^*$. Now, φ_λ is bounded for each λ , and so

$$\varphi(ab) = \lim_\lambda \varphi_\lambda(ab) = \lim_\lambda \varphi_\lambda(a)\varphi_\lambda(b) = \varphi(a)\varphi(b),$$

for all $a, b \in \mathcal{A}$. Thus, if φ is non-zero, then it is a character. Either way however, φ is contained in $\Omega(\mathcal{A}) \cup \{0\}$, and so $\Omega(\mathcal{A}) \cup \{0\}$ is closed.

If \mathcal{A} is unital, then taking $\{\varphi_\lambda\}_\lambda$ and φ as above, we have that

$$\varphi(1) = \lim_\lambda \varphi_\lambda(1) = 1 \neq 0.$$

Hence the zero functional lies outside the closure of $\Omega(\mathcal{A})$, and so $\Omega(\mathcal{A})$ is closed. We should now note that equation (1.3) implies that $\Omega(\mathcal{A}) \cup \{0\}$ is contained in \mathcal{A}_1^* , which is weak* compact by the Banach–Alaoglu Theorem. Hence, $\Omega(\mathcal{A})$ is compact in the unital case, and locally compact in the non-unital case.

Gelfand Transform

Define a mapping Γ , called the *Gelfand transform*, by setting

$$\Gamma : \mathcal{A} \rightarrow C_0(\Omega(\mathcal{A})), \quad a \mapsto \widehat{a};$$

where by \widehat{a} we now mean the mapping defined in (1.4) with a domain restricted to $\Omega(\mathcal{A})$. The image of \mathcal{A} under the Gelfand transform is contained in $C_0(\Omega(\mathcal{A}))$. To see this, choose an arbitrary $\varepsilon > 0$ and consider the set

$$\Omega_{\widehat{a}, \varepsilon} = \{\varphi \in \Omega(\mathcal{A}) : |\widehat{a}(\varphi)| \geq \varepsilon\}.$$

Using an argument similar to the one above, we can show that this set is weak* closed in \mathcal{A}_1^* . Thus, by the Banach–Alaoglu Theorem, it is compact. Now, for $\varphi \notin \Omega_{\widehat{a}, \varepsilon}$ we have $|\widehat{a}(\varphi)| < \varepsilon$, and so $\widehat{a} \in C_0(\Omega(\mathcal{A}))$. However, if \mathcal{A} is unital, then this tells us nothing new since $\Omega(\mathcal{A})$ will be compact and $C_0(\Omega(\mathcal{A}))$ will be equal to $C(\Omega(\mathcal{A}))$.

Gelfand Transform and the Spectral Radius

We saw earlier that if $\varphi \in \Omega(\mathcal{A})$, then $\varphi(a) \in \sigma(a)$, for all $a \in \mathcal{A}$. If \mathcal{A} is assumed to be commutative and unital, then it can be shown that all elements of $\sigma(a)$ are of this form; that is,

$$\sigma(a) = \{\varphi(a) \mid \varphi \in \Omega(\mathcal{A})\}.$$

In the non-unital case we almost have the same result: if \mathcal{A} is a non-unital commutative C^* -algebra, then

$$\sigma(a) = \{\varphi(a) \mid \varphi \in \Omega(\mathcal{A})\} \cup \{0\}.$$

These two results are important because they give us the equation

$$r(a) = \max\{|\lambda| : \lambda \in \sigma(a)\} = \max\{|\lambda| : \lambda \in \widehat{a}(\Omega(\mathcal{A}))\} = \|\widehat{a}\|.$$

Hence, for any commutative C^* -algebra \mathcal{A} , we have that

$$r(a) = \|\widehat{a}\|, \quad \text{for all } a \in \mathcal{A}. \quad (1.5)$$

It is instructive to note that this is the point at which the requirement of commutativity (which is needed for the Gelfand–Naimark theorem to hold) enters our discussion of C^* -algebras.

1.1.3 Gelfand–Naimark Theorem

The following result is of great importance. It allows us to completely characterise commutative C^* -algebras. It first appeared in a paper [37] of Gelfand and Naimark in 1943, and it has since become the principle theorem motivating noncommutative geometry.

Theorem 1.1.2 *Let \mathcal{A} be an abelian C^* -algebra. The Gelfand transform*

$$\Gamma : \mathcal{A} \rightarrow C_0(\Omega(\mathcal{A})), \quad a \mapsto \widehat{a},$$

is an isometric $$ -isomorphism.*

Proof. That Γ is a homomorphism is clear from

$$\Gamma(a + \lambda b)(\varphi) = \widehat{a + \lambda b}(\varphi) = \varphi(a) + \lambda\varphi(b) = \Gamma(a)(\varphi) + \lambda\Gamma(b)(\varphi),$$

and

$$\Gamma(ab)(\varphi) = \widehat{ab}(\varphi) = \varphi(ab) = \varphi(a)\varphi(b) = (\Gamma(a)\Gamma(b))(\varphi).$$

Hence, $\Gamma(\mathcal{A})$ forms a subalgebra of $C_0(\Omega(\mathcal{A}))$.

If $\varphi, \psi \in \Omega(\mathcal{A})$ and $\varphi \neq \psi$, then there exists an $a \in \mathcal{A}$ such that $\varphi(a) \neq \psi(a)$. Hence, $\widehat{a}(\varphi) \neq \widehat{a}(\psi)$, and so $\Gamma(\mathcal{A})$ separates points of $\Omega(\mathcal{A})$.

If $\varphi \in \Omega(\mathcal{A})$, then by definition φ is non-zero. Therefore, there exists an $a \in \mathcal{A}$ such that $\varphi(a) \neq 0$, and so $\widehat{a}(\varphi) \neq 0$. Hence, for each point of the spectrum there exists a function in $\Gamma(\mathcal{A})$ that does not vanish there.

It is seen that $\Gamma(\mathcal{A})$ is closed under conjugation from

$$\widehat{a}^*(\varphi) = \overline{\widehat{a}(\varphi)} = \overline{\varphi(a)} = \varphi(a^*) = \widehat{a^*}(\varphi).$$

We shall now show that Γ is an isometry. Firstly, we note that since aa^* is self-adjoint

$$\|aa^*\| = r(aa^*) = \|\widehat{aa^*}\|.$$

Then since $\Gamma(\mathcal{A})$ is a subset of the C^* -algebra $C_0(\Omega(\mathcal{A}))$, we have that $\|\widehat{a}\|^2 = \|\widehat{a^*a}\|$. It follows that

$$\|\widehat{a}\|^2 = \|\widehat{a^*a}\| = \|\widehat{aa^*}\| = \|aa^*\| = \|a\|^2.$$

Therefore, $\|a\| = \|\widehat{a}\|$ and Γ is an isometry. This implies that $\Gamma(\mathcal{A})$ is complete, and therefore closed in $C_0(\Omega(\mathcal{A}))$. Since $\Omega(\mathcal{A})$ is locally compact, we can invoke the Stone–Weierstrass Theorem and conclude that $\Gamma(\mathcal{A}) = C_0(\Omega(\mathcal{A}))$. \square

Corollary 1.1.3 *For \mathcal{A} a unital abelian C^* -algebra, the Gelfand transform*

$$\mathcal{A} \rightarrow C(\Omega(\mathcal{A})), \quad a \mapsto \widehat{a},$$

is an isometric $$ -isomorphism*

1.1.4 The Algebra-Space Correspondence

Following the definition of a C^* -algebra, we saw that if X is a locally compact Hausdorff space, then $C_0(X)$ is a commutative C^* -algebra. Therefore, every locally compact Hausdorff space is naturally associated with an abelian C^* -algebra. Conversely, the Gelfand–Naimark Theorem showed us that every abelian C^* -algebra \mathcal{A} is naturally associated with a locally compact Hausdorff space, namely $\Omega(\mathcal{A})$. If we could show that these associations were inverse to each other, then we would have established a one-to-one correspondence between spaces and algebras. In fact, this is a direct consequence of the following theorem.

Theorem 1.1.4 *If X is a locally compact Hausdorff space, then X is homeomorphic to $\Omega(C_0(X))$.*

Proof. We shall prove this result in the compact case only and we refer the interested reader to the first chapter [114].

Let us define the mapping

$$\mathcal{F} : X \rightarrow \Omega(C(X)), \quad x \mapsto \mathcal{F}_x,$$

by setting $\mathcal{F}_x(f) = f(x)$, for all $f \in C(X)$. It is clear that \mathcal{F}_x is a character, for all $x \in X$.

If $x_1 \neq x_2$ in X , then it follows from Urysohn's Lemma that there exists an f in $C(X)$ such that $f(x_1) = 0$ and $f(x_2) = 1$. This shows that \mathcal{F} is injective.

To show that \mathcal{F} is surjective, take any $\psi \in \Omega(C(X))$ and consider the ideal

$$I = \ker(\psi) = \{f \in C(X) : \psi(f) = 0\}.$$

We shall show that there exists an $x_0 \in X$ such that $f(x_0) = 0$, for all $f \in I$. If this is not the case, then for each $x \in X$, there is an $f_x \in I$ such that $f_x(x) \neq 0$. The continuity of f implies that each x has an open neighborhood U_x on which f_x is non-vanishing. By the compactness of X , there exist x_1, \dots, x_n in X such that $X = \bigcup_{k=1}^n U_{x_k}$. Let

$$f(x) = \sum_{k=1}^n |f_{x_k}(x)|^2, \quad x \in X.$$

Clearly f is non-vanishing on X , and so it is invertible in $C(X)$. This implies that $\psi(f) \neq 0$. On the other hand, since ψ is multiplicative,

$$\psi(f) = \sum_{k=1}^n \psi(f_{x_k})\psi(\overline{f_{x_k}}) = 0$$

This contradiction shows that there must exist some $x_0 \in X$ such that $f(x_0) = 0$, for all $f \in I$.

Now, if f is an arbitrary element in $C(X)$, then $f - \psi(f)$ is in I . Thus,

$$(f - \psi(f))(x_0) = f(x_0) - \psi(f) = 0,$$

or equivalently $\psi(f) = f(x_0)$. Therefore, $\psi = \mathcal{F}(x_0)$, and so \mathcal{F} is surjective.

If $(x_\lambda)_\lambda$ is a net in X that converges to x , then $f(x_\lambda) \rightarrow f(x)$, for every $f \in C(X)$. This is equivalent to saying that, $\psi_{x_\lambda}(f) \rightarrow \psi_x(f)$, for all $f \in C(X)$; or that $\widehat{f}(\psi_{x_\lambda}) \rightarrow \widehat{f}(\psi_x)$, for all $f \in C(X)$. Hence, $\psi_{x_\lambda} \rightarrow \psi_x$ with respect to the weak* topology, and so \mathcal{F} is a continuous mapping.

We have now shown that \mathcal{F} is bijective continuous function from the compact Hausdorff space X to the compact Hausdorff space $\Omega(C(X))$. Therefore, we can conclude that it is a homeomorphism. \square

1.2 Noncommutative Topology

We shall now stop and reflect on what we have established: we have shown that there is a one-to-one correspondence between commutative C^* -algebras and locally compact Hausdorff spaces. Thus, the function algebra of a space contains all

the information about that space. This means that nothing would be lost if we were to study the algebra alone and ‘forget’ about the space. This approach is common in other areas of mathematics, most notably algebraic geometry. The idea behind noncommutative topology is to take it one step further: loosely speaking, noncommutative topology views compact Hausdorff spaces as special commutative examples of general C^* -algebras, and studies them in this context. The subject can be described as the investigation of those C^* -algebraic structures that correspond to topological structures when the algebra is commutative.

As often happens when one works in greater generality, results that were previously complex or technical become quite straightforward when this approach is used. It has allowed formerly ‘unsolvable’ problems in topology to be solved. Also, the application of our topological intuition to noncommutative C^* -algebras has helped in the discovery of new algebraic results. Quite often, the algebraic structures that generalise topological structures are much richer and have features with no classical counterparts.

However, in all of this one does notice a lack of duality. While algebraic properties of commutative C^* -algebras correspond to geometric properties, no such correspondence exists for noncommutative algebras. This prompts us to consider the possibility that, in a very loose sense, every noncommutative C^* -algebra could be viewed as the *function algebra of some type of ‘noncommutative’ or ‘quantum space’*. This is the basic heuristic principle upon which most of the vocabulary of noncommutative topology is based.

Quantum spaces are imagined to be a type of generalised set, and those spaces that have a classical point set representation are considered to be special case. Noncommutative topology is then thought of as the investigation of these quantum spaces through their function algebras.

It must be stressed, however, that quantum spaces only exist as an intuitive tool based upon an analogy. They do not have any kind of proper definition. It is only when one ventures into the literature of physics that the concept gains any concrete form.

1.2.1 Some Noncommutative Generalisations

Now that we have presented the general philosophy behind noncommutative topology we can move on and explore some more concrete aspects of the subject. In this section we shall present some simple examples of generalisations of topological properties and structures to the noncommutative setting.

1. As we saw above, points in a space X are in one-to-one correspondence with

characters on the C^* -algebra $C_0(X)$. Therefore, we view characters as the appropriate generalisation of points to the noncommutative case:

$$\text{points} \longrightarrow \text{characters.}$$

However, it must be noted that the non-emptiness of the spectrum of a noncommutative C^* -algebra is not guaranteed. The easiest example of such a C^* -algebra is $M_n(\mathbf{C})$, for $n \geq 2$. It is well known $M_n(\mathbf{C})$ has no proper ideals for $n \geq 2$. However, if φ is a character on any C^* -algebra \mathcal{A} , then its kernel is clearly a proper two-sided proper ideal of \mathcal{A} (in fact, it is a maximal ideal). Therefore, the spectrum of $M_n(\mathbf{C})$ must be empty.

For this reason the notion of a point will be of little use in the noncommutative world. In fact, it is often written that ‘quantum spaces are pointless spaces’.

2. Let \mathcal{F} be a homeomorphism from a locally compact Hausdorff space X to itself. Consider the map

$$\Phi : \mathcal{A} \rightarrow \mathcal{A}, \quad f \mapsto f \circ \mathcal{F},$$

where $\mathcal{A} = C(X)$. A little thought will verify that it is a $*$ -isomorphism. Thus, we can associate a $*$ -isomorphism to each homeomorphism.

On the other hand, let Φ be an algebra $*$ -isomorphism from \mathcal{A} to itself, and consider the mapping

$$\mathcal{F} : \Omega(\mathcal{A}) \rightarrow \Omega(\mathcal{A}), \quad \varphi \mapsto \varphi \circ \Phi.$$

With the aim of establishing the continuity of \mathcal{F} , we shall consider a net $(\varphi_\lambda)_\lambda$ in $\Omega(\mathcal{A})$ that converges to φ . Since $\Omega(\mathcal{A})$ is endowed with the weak* topology, $\varphi_\lambda(a) \rightarrow \varphi(a)$, for all $a \in \mathcal{A}$. Now, $\mathcal{F}(\varphi_\lambda)(a) = \varphi_\lambda(\Phi(a))$, and since $\varphi_\lambda(\Phi(a)) \rightarrow \varphi(\Phi(a))$ and $\varphi(\Phi(a)) = \mathcal{F}(\varphi)(a)$, \mathcal{F} must be continuous.

Again, it is straightforward to show that \mathcal{F} is bijective. Using an argument similar to that above, we can also show that \mathcal{F}^{-1} is continuous, and so \mathcal{F} is a homeomorphism.

Because of the Gelfand–Naimark Theorem \mathcal{F} can also be considered as a homeomorphism from X to X . Thus, to each $*$ -isomorphism we can associate a homeomorphism.

A little extra work will verify that these two associations are inverse to each other. Hence, we have established a one-to-one correspondence between homeomorphisms and $*$ -isomorphisms. This motivates our next generalisation:

homeomorphisms \longrightarrow *-isomorphisms.

3. If $f \in C_0(X)$, for some compact Hausdorff space X , then clearly $\lambda \in \sigma(f)$ if, and only if, $f(x) = \lambda$, for some $x \in X$. (In the locally compact case $\sigma(f) = \text{im}(f) \cup \{0\}$.) Thus, the spectrum of an element of a C^* -algebra is a generalisation of the image of a function:

image of a function \longrightarrow spectrum of a element.

This motivates us to define a *positive element* of a C^* -algebra \mathcal{A} to be a self-adjoint element with positive spectrum; we denote that $a \in \mathcal{A}$ is positive by writing $a \geq 0$, and we denote the set of positive elements of \mathcal{A} by \mathcal{A}_+ . The self-adjointness requirement is necessary because there may exist non-self-adjoint elements of \mathcal{A} with positive spectrum. One merely needs to look in $M_2(\mathbf{C})$ for an example.

It is pleasing to note that, just as in the classical case, every positive element is of the form aa^* , for some $a \in \mathcal{A}$. Note that this implies that when $\mathcal{A} = B(H)$, for some Hilbert space H , then the positive elements of \mathcal{A} will coincide with the positive operators on H .

We now record this generalisation:

positive function \longrightarrow positive element.

4. Let X be a compact Hausdorff space, and let μ be a regular complex Borel measure on X . The scalar-valued function λ on $C_0(X)$, defined by $\lambda(f) = \int_X f d\mu$, is clearly bounded and linear by the properties of the integral. Therefore, it is an element of $C_0(X)^*$. The following well known theorem shows us that every bounded linear functional arises in this way.

Theorem 1.2.1 (Riesz Representation Theorem) *Let X be a locally compact Hausdorff space. For all $\lambda \in C_0(X)^*$ there exists a unique regular complex Borel measure μ on X such that*

$$\lambda(f) = \int_X f d\mu, \quad f \in C(X).$$

Thus, there exists a one-to-one correspondence between the elements of the dual space of $C_0(X)$ and the regular complex Borel measures on X . This motivates our next generalisation:

regular complex Borel measures \longrightarrow bounded linear functionals.

5. As we noted earlier, if X is a compact Hausdorff space, then $C_0(X) = C(X)$ is a unital algebra; and if X is a non-compact Hausdorff space, then $C_0(X)$ is a non-unital algebra.

As is well known, we may compactify a non-compact space by adding to it a point at infinity. We denote this new space by X_∞ . The algebra $C_0(X_\infty) = C(X_\infty)$ is then a unital algebra. If we unitise $C_0(X)$, then we also get a unital algebra $\widetilde{C_0(X)}$. What is interesting about this is that $C(X_\infty)$ is isometrically $*$ -isomorphic to $\widetilde{C_0(X)}$. An obvious isomorphism is

$$C(X_\infty) \rightarrow \widetilde{C_0(X)}, \quad f \mapsto ((f - f(\infty)1)|_X, f(\infty)).$$

These observations motivate the following generalisation:

$$\begin{array}{ll} \text{compact spaces} & \longrightarrow \text{unital } C^*\text{-algebras,} \\ \text{one point compactification} & \longrightarrow \text{unitisation.} \end{array}$$

6. Each closed subset K of a compact Hausdorff space X is a compact Hausdorff space. Through our algebra-space correspondence, X and K are associated to the C^* -algebra of continuous functions defined upon them. With the aim of generalising closed sets to the algebra setting, we shall examine the relationship between these two C^* -algebras.

Let K be a closed subset of X and write

$$I = \{f \in C(X) : f|_K = 0\}. \tag{1.6}$$

Clearly I is a closed $*$ -ideal of $C(X)$; that is, a closed self-adjoint ideal. This implies that $C(X)/I$ is well defined as a $*$ -algebra. We can define a norm on it by

$$\|f + I\| = \inf_{h \in I} \|f + h\|_\infty.$$

It is a standard result that when $C(X)/I$ is endowed with this norm it is a Banach algebra. In fact, as can be routinely verified, it is a C^* -algebra.

Let us now consider the mapping

$$\pi : C(X) \rightarrow C(K), \quad f \mapsto f|_K.$$

For any $g \in C(K)$, the Tietze extension theorem implies that there exists an $f \in C(X)$, such that f extends g . Thus, π is surjective. Since π is clearly a $*$ -algebra homomorphism with kernel equal to I , it induces a $*$ -isomorphism from $C(X)/I$ to $C(K)$.

If we could show that all the closed $*$ -ideals of $C(X)$ were of the same form as (1.6), then we would have a one-to-one correspondence between quotient algebras and closed subsets. The following lemma shows exactly this (for a proof see the first chapter of [114]).

Lemma 1.2.2 *Let X be a compact Hausdorff space and let I be a closed $*$ -ideal of $C(X)$. Then there exists a closed subset K of X such that*

$$I = \{f \in C(X) : f|_K = 0\}.$$

Using a similar line of argument one can also show that there is a bijective correspondence between the open sets of X and the ideals in $C(X)$.

This gives us our next generalisations:

$$\begin{array}{ll} \text{closed sets of a compact space} & \longrightarrow \text{quotients of unital } C^*\text{-algebras,} \\ \text{open sets of a compact space} & \longrightarrow \text{ideals of unital } C^*\text{-algebras.} \end{array}$$

7. With the aid of a simple definition and a standard result we can generalise connectedness.

A *projection* in a C^* -algebra \mathcal{A} is an element $a \in \mathcal{A}$ such that

$$a^2 = a^* = a.$$

An $*$ -algebra A is called *projectionless* if the only projections it contains are 0, and 1 if A is unital.

Let X be a connected space. If $p \in C(X)$ is a projection, then for all $x \in X$, $(p(x))^2 = p(x)$, so $p(x)$ is equal to 0 or 1. Since X is connected, $X = p^{-1}\{0\} \cup p^{-1}\{1\}$ cannot be a disconnection of X , therefore p is equal to 0 or 1.

Conversely, if $C(X)$ is projectionless, then X must be connected because a non-trivial projection can easily be defined on an unconnected space. This motivates the following generalisation:

$$\text{connected compact space} \longrightarrow \text{projectionless unital } C^*\text{-algebra.}$$

8. **Theorem 1.2.3** *If a compact Hausdorff space X is metrisable, then $C(X)$ is separable.*

Proof. Let X be metrisable with a metric d , and denote

$$B_r(x) = \{y : d(x, y) < r, y \in X\}, \quad \text{for } r > 0.$$

This means that $\mathcal{C}_r = \{B_r(x) : x \in X\}$ is an open cover of X . Since X is compact, \mathcal{C}_r has a finite subcover \mathcal{C}'_r , for all $r \geq 0$. The family $\mathcal{C} = \bigcup_{n=1}^{\infty} \mathcal{C}'_{\frac{1}{n}}$ will then form a countable base for the topology of X .

Let x be an element of $B_1 \in \mathcal{C}$. Since X is a compact Hausdorff space, it is easy to see that there exists a $B_2 \in \mathcal{C}$, such that

$$x \in B_2 \subset \overline{B_2} \subset B_1. \quad (1.7)$$

Thus, the countable family

$$\mathcal{D} = \{(B_1, B_2) \in \mathcal{C} \times \mathcal{C} : \overline{B_2} \subset B_1\},$$

is non-empty. For each $(B_1, B_2) \in \mathcal{D}$, Urysohn's Lemma guarantees the existence of a function $f_{B_1, B_2} \in C(X)$ satisfying

$$f_{B_1, B_2}(X \setminus B_1) = \{0\}, \quad \text{and} \quad f_{B_1, B_2}(\overline{B_2}) = \{1\}.$$

We write

$$\mathcal{F} = \{f_{B_1, B_2} : (B_1, B_2) \in \mathcal{D}\}.$$

If $x \neq y$ in X , then there clearly exists $(B_1, B_2) \in \mathcal{D}$ such that $x \in B_2$, and $y \in X \setminus B_1$. Since

$$f_{B_1, B_2}(x) = 1 \neq 0 = f_{B_1, B_2}(y),$$

\mathcal{F} must separate the points of X . Also, it is *non-vanishing*; that is, for any $x \in X$, there must exist an $f \in \mathcal{F}$ such that $f(x) \neq 0$.

Let \mathcal{A} be the smallest algebra that contains $\mathcal{F} \cup \mathcal{F}^*$. It is clearly a non-vanishing self-adjoint algebra that separates points. Hence, by the Stone-Weierstrass Theorem, it is dense in $C(X)$. Clearly, there exists a countable subset of \mathcal{A} that is dense in \mathcal{A} . Take, for example, the smallest algebra over \mathbf{Q} that contains $\mathcal{F} \cup \mathcal{F}^*$. This countable subset is then also dense in $C(X)$. \square

The following theorem establishes the result in the opposite direction.

Theorem 1.2.4 *If X is a compact Hausdorff topological space, then the separability of $C(X)$ implies the metrisability of X .*

Proof. Since $C(X)$ is separable, it contains a dense sequence of continuous functions $\{f_n\}_{n \in \mathbf{N}}$. Define $g_n = \frac{f_n}{1 + \|f_n\|}$, this guarantees that $\|g_n\| < 1$.

If $x \neq y$ in X , then by Urysohn's Lemma there exists an $f \in C(X)$ such that $f(x) \neq f(y)$. Therefore, there must exist a g_n such that $g_n(x) \neq g_n(y)$, and so $\{g_n\}_{n \in \mathbf{N}}$ separates the points of X .

Let us define

$$d(x, y) = \sup_{n \in \mathbf{N}} 2^{-n} |g_n(x) - g_n(y)|.$$

Since $\{g_n\}_{n \in \mathbf{N}}$ separates the points of X , $d(x, y) = 0$ implies that $x = y$. Therefore, d is a metric on X .

We shall now examine the open balls of d . Let $x \in X$, $0 < \varepsilon < 1$, and choose $N \in \mathbf{N}$ such that $2^{-N} < \varepsilon$. We write

$$U_n = g_n^{-1}\{z \in \mathbf{C} : |g_n(x) - z| < \varepsilon\},$$

and

$$U = \bigcap_{n=0}^N U_n.$$

Each U_n is open with respect to the original topology on X and, as a result, U is also open. Let $y \in U$. If $n \leq N$, then

$$2^{-n} |g_n(x) - g_n(y)| \leq 2^{-n} \varepsilon \leq \varepsilon.$$

If $n > N$, then, since $\|g_n\| < 1$,

$$2^{-n} |g_n(x) - g_n(y)| < 2^{-n} 2 < \varepsilon.$$

It follows that

$$d(x, y) = \sup_{n \in \mathbf{N}} 2^{-n} |g_n(x) - g_n(y)| < \varepsilon.$$

Thus, $U \subset B_\varepsilon(x)$, and so the open balls of the metric are open with respect to the original topology. Let us now consider the identity map from the compact space X , endowed with its original topology, to X , endowed with the metric topology of d . Clearly, this map is a continuous bijection, thus, since the metric topology of d is Hausdorff, it is a homeomorphism. \square

Thus, a topological space is metrisable if, and only if, its algebra of continuous functions is separable. This gives us the final generalisation of this section:

$$\text{metrisable compact space} \longrightarrow \text{separable unital } C^*\text{-algebra.}$$

Before we finish it should be noted that the transition from topology to algebra is not always as smooth as in the examples above. It quite often happens that there is more than one option for the generalisation of a topological feature (or differential feature), and it may not always be obvious which one is the ‘correct’ choice; we will see an example of this when we come to generalise the de Rham calculus in Chapter 2.

1.3 Vector Bundles and Projective Modules

In this section we shall present the Serre–Swan Theorem following more or less Swan’s original proof in [104] (for a more modern category style proof see [7]). This result will give us a noncommutative generalisation of vector bundles, one of the basic objects in differential geometry. Thus, this section can be seen as our first venture into noncommutative geometry proper.

Much of the material presented in this section will be of use to us later on when we discuss elementary K -theory and when we present the theory of geometric Dirac operators.

1.3.1 Vector Bundles

Definition 1.3.1. A (*complex*) *vector bundle* is a triple (E, π_E, X) , consisting of a topological space E called the *total space*, a topological space X called the *base space*, a continuous surjective map $\pi_E : E \rightarrow X$, called the *projection*, and a complex linear space structure defined on each *fibre* $E_x = \pi_E^{-1}(\{x\})$, such that the following conditions hold:

1. For every point $x \in X$, there is an open neighborhood U of x , a natural number n , and a homeomorphism

$$\varphi_U : \pi_E^{-1}(U) \rightarrow U \times \mathbf{C}^n;$$

such that, for all $x \in X$,

$$\varphi_U(E_x) = \{x\} \times \mathbf{C}^n.$$

2. The map φ_U restricted to E_x is a linear mapping between E_x and $\{x\} \times \mathbf{C}^n$; (where $\{x\} \times \mathbf{C}^n$ is regarded as a linear space in the obvious way).

We could alternatively define a vector bundle to have a real linear space structure on each fibre. Such a vector bundle is called a *real vector bundle*. All of the results presented below would hold equally well in this case. However, since we shall always work with the complex-valued functions on a topological space, it suits us better to work with complex vector bundles.

The canonical examples of a vector bundle are the tangent and cotangent bundles of a manifold.

Whenever possible, we shall denote a vector bundle (E, π_E, X) by E and suppress explicit reference to the projection and the base space. Also, when no confusion arises, we shall use π instead π_E . An open set of the form U is called a *base*

neighbourhood and the corresponding homeomorphism φ_U is called the *associated local trivialisation*. Note that the definition implies that the dimension of the fibres is locally constant.

Let E and F be two vector bundles over the same base space X . A *bundle map* f from E to F is a continuous mapping $f : E \rightarrow F$, such that $\pi \circ f = \pi$, and f_x , the restriction of f to E_x , is a linear mapping from E_x to F_x . If f is also a homeomorphism between E and F , then it is called a *bundle isomorphism*. A vector bundle (E, π, X) is called *trivial* if it is isomorphic to the bundle $(X \times \mathbf{C}^n, \pi, X)$, for some natural number n .

Let U be a base neighbourhood for E and F , and let n and m be the dimensions of E and F respectively. When restricted to $\pi^{-1}(U)$, any bundle map $f : E \rightarrow F$ induces a map $\tilde{f} : U \times \mathbf{C}^n \rightarrow U \times \mathbf{C}^m$, defined by $\tilde{f} = \psi_U \circ f \circ \varphi_U^{-1}$. This in turn induces a map $\hat{f} : U \rightarrow M_{m \times n}(\mathbf{C})$, which is determined by the formula $\tilde{f}(x, v) = (x, \hat{f}(x)v)$, for $(x, v) \in U \times \mathbf{C}^n$. As a little thought will verify, if \hat{f} is continuous for each such neighbourhood U , then f will be continuous. An important use of this fact arises when the bundle map takes each E_x to the corresponding fibre F_x by a linear isomorphism. In this case, $m = n$, and so f^{-1} will determine a mapping $\widehat{f^{-1}}$ from U to $M_{n \times n}(\mathbf{C})$. Now, as a little more reflection will verify, $\widehat{f^{-1}}(x) = (\hat{f}(x))^{-1}$. Thus, f^{-1} is continuous. This implies that f is a homeomorphism, which in turn implies that f is a bundle isomorphism. Thus, we have established the following lemma.

Lemma 1.3.2 *Let E and F be two vector bundles over the same base space X . A continuous bundle map $f : E \rightarrow F$ is an isomorphism if, and only if, it maps E_x to F_x by a linear isomorphism, for all $x \in X$.*

A *subbundle* of (E, π_E, X) is a vector bundle (S, π_S, X) , such that:

1. $S \subseteq E$,
2. π_S is the restriction of π_E to S ,
3. the linear structure on each fibre S_x is the linear structure induced on it by E_x .

Transition Functions

Let E be a vector bundle over X , let $\{U_\alpha\}$ be an open covering of X by base neighbourhoods, and let φ_α denote the corresponding trivialisation maps. Now,

consider the following diagram:

$$\begin{array}{ccc}
 & \pi^{-1}(U_\alpha \cap U_\beta) & \\
 \varphi_\beta \swarrow & & \searrow \varphi_\alpha \\
 (U_\alpha \cap U_\beta) \times \mathbf{C}^n & \xrightarrow{\varphi_\alpha \circ \varphi_\beta^{-1}} & (U_\alpha \cap U_\beta) \times \mathbf{C}^n.
 \end{array}$$

Clearly $\varphi_\alpha \circ \varphi_\beta^{-1}$ is a bundle map from $(U_\alpha \cap U_\beta) \times \mathbf{C}^n$ to itself. As explained above, this means that it determines a continuous map from $U_\alpha \cap U_\beta$ to $M_n(\mathbf{C})$; we denote this map by $g_{\alpha\beta}$. In fact, since $\varphi_\beta \circ \varphi_\alpha^{-1}$ is a bundle map that is inverse to $\varphi_\alpha \circ \varphi_\beta^{-1}$, it must hold that $g_{\alpha\beta}(x) \in \text{GL}(n, \mathbf{C})$, for all $x \in U_\alpha \cap U_\beta$. We call the collection

$$\{g_{\alpha\beta} : U_\alpha \cap U_\beta \neq \emptyset\}$$

the *transition functions* of the vector bundle for the covering $\{U_\alpha\}$.

Three observations about any set of transition functions can be made immediately;

1. $g_{\alpha\alpha} = 1$,
2. $g_{\beta\alpha} \circ g_{\alpha\beta} = 1$,
3. if $U_\alpha \cap U_\beta \cap U_\gamma \neq \emptyset$, then $g_{\alpha\beta} \circ g_{\beta\gamma} \circ g_{\gamma\alpha} = 1$.

The last property is known as the *cocycle condition*.

The following result is of great importance in the theory of vector bundles, and shows that the transition functions for any covering completely determine the bundle.

Proposition 1.3.3 *Given a cover $\{U_\alpha\}$ of X , and a continuous map*

$$g_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow \text{GL}(n, \mathbf{C}),$$

for every non-empty intersection $U_\alpha \cap U_\beta$, such that the conditions 1, 2, 3 listed above hold, then there exists a vector bundle (unique up to bundle isomorphism) for which $\{g_{\alpha\beta}\}$ are the transition functions.

Sections

Let E be a vector bundle over a base space X . A continuous mapping s from X to E is called a *section* if $\pi(s(x)) = x$, for all $x \in X$. The set of sections of E is denoted by $\Gamma(E)$. If $s, t \in \Gamma(E)$ and $a \in C(X)$, we define

$$(s + t)(x) = s(x) + t(x), \quad \text{and} \quad (sa)(x) = a(x)s(x).$$

Here it is understood that the addition and scalar multiplication on the right hand side of each equality takes place in E_x . The mappings $s + t$ and sa are continuous since the composition of either with a local trivialisation is continuous on the corresponding base neighbourhood. Hence, the mappings $s + t$ and sa are sections. With respect to these definitions $\Gamma(E)$ becomes a right- $C(X)$ -module. Let G be another vector bundle over X and let $f : E \rightarrow G$ be a bundle map. We define $\Gamma(f)$ to be the unique module mapping from $\Gamma(E)$ to $\Gamma(G)$ such that

$$[\Gamma(f)(s)](x) = (f \circ s)(x).$$

Let U be a base neighbourhood of $x \in X$ and let φ_U be its associated trivialisation. Consider the set of n continuous mappings on U

$$\widehat{s}_i : U \rightarrow U \times \mathbf{C}^n, \quad x \mapsto (x, e_i),$$

where $\{e_i\}_{i=1}^n$ is the standard basis of \mathbf{C}^n . If we define $s_i = \varphi_U^{-1} \circ \widehat{s}_i$, then the set $\{s_i(y)\}_{i=1}^n$ forms a basis for E_y , for every $y \in U$. If we assume that X is normal, then using the Tietze Extension Theorem, it can be shown that for each s_i , there exists a section $s'_i \in \Gamma(E)$ such that s_i and s'_i agree on some neighbourhood of x . We call the set $\{s'_i(x)\}_{i=1}^n$ a *local base* at x .

Direct sum

Let E and F be two vector bundles and let $\{U_\alpha\}$ be the family of subsets of X that are base neighbourhoods for both bundles. Now, if $U_\alpha \cap U_\beta$ is a non-empty intersection, then we denote the corresponding transition functions for E and F by $g_{\alpha\beta}^E$ and $g_{\alpha\beta}^F$ respectively. Using $g_{\alpha\beta}^E$ and $g_{\alpha\beta}^F$ we can define a matrix-valued function on $U_\alpha \cap U_\beta$ by

$$g_{\alpha\beta} : x \mapsto \begin{pmatrix} g_{\alpha\beta}^E(x) & 0 \\ 0 & g_{\alpha\beta}^F(x) \end{pmatrix}.$$

Clearly, $g_{\alpha\beta}$ satisfies the conditions required by Proposition (1.3.3), for every non-empty intersection $U_\alpha \cap U_\beta$. Hence, there exists a vector bundle for which $\{g_{\alpha\beta}\}$ is the set of transition functions. The nature of the construction of each $g_{\alpha\beta}$ implies that the fibre of the bundle over any point $x \in X$ will be isomorphic to $E_x \oplus F_x$. This prompts us to denote the bundle by $E \oplus F$ and to call it the *direct sum* of E and F .

Using an analogous construction, we can produce bundles whose fibres over any $x \in X$ are equal to E_x^* , $\text{Hom}(E_x)$, or $E_x \otimes F_x$; we denote these bundles by E^* , $\text{Hom}(E)$, and $E \otimes F$ respectively.

Inner Products and Projections

An *inner product* on a vector bundle E is a continuous mapping from $D(E) = \{(v_1, v_2) : v_1, v_2 \in E, \pi(v_1) = \pi(v_2)\}$ to \mathbf{C} such that its restriction to $E_x \times E_x$ is an inner product, for all $x \in X$. Using locally defined inner products and a partition of unity, an inner product can be defined on any vector bundle.

Let S be a subbundle of E , and let $\langle \cdot, \cdot \rangle$ be an inner product on E . Using $\langle \cdot, \cdot \rangle$ we can define an orthogonal projection $P_x : E_x \rightarrow S_x$, for each $x \in X$. This defines a map $P : E \rightarrow S$, which we shall call the *projection* of E onto S . To see that this map is continuous we shall examine the mapping it induces on $U \times \mathbf{C}^n$, for some arbitrary base neighbourhood U . Let $\{t_i\}_{i=1}^n$ be a *local basis* over U , that is, let $\{t_i\}_{i=1}^n$ be a set of sections such that, for each $p \in X$, $\{t_i(p)\}_{i=1}^n$ is a basis of E_x ; then consider the map

$$U \times \mathbf{C}^n \rightarrow D, \quad (x, v) \mapsto (\varphi_U^{-1}(x, v), t_i(x)).$$

It is clearly continuous, for all $i = 1, \dots, n$, and as a result its composition with the inner product is continuous; that is, the map $(x, v) \mapsto \langle \varphi_U^{-1}(x, v), t_i(x) \rangle$ is continuous. Therefore, the induced map

$$U \times \mathbf{C}^n \rightarrow \mathbf{C}, \quad (x, v) \mapsto \sum_{i=1}^n \langle (x, v), t_i(x) \rangle t_i(x)$$

is continuous. The continuity of the projection easily follows.

1.3.2 Standard results

The following four results are standard facts in vector bundle theory. As above, their proofs consist of routine arguments involving local neighbourhoods and local sections. We shall briefly sketch how the results are established, and refer the interested reader to [104], [3], or [52].

Lemma 1.3.4 *Let s_1, \dots, s_k be sections of a vector bundle E such that for some $x \in X$, $s_1(x), \dots, s_k(x)$ are linearly independent in E_x . Then there is a neighbourhood V of x such that $s_1(y), \dots, s_k(y)$ are linearly independent, for all $y \in V$.*

In fact, this lemma is a simple consequence of the continuity of the determinant function. It has the following easy corollary.

Corollary 1.3.5 *Let F be a vector bundle over X and let $f : E \rightarrow F$ be a bundle map. If $\dim(\text{im}(f_x)) = n$, then $\dim(\text{im}(f_y)) \geq n$, for all y in some neighbourhood of x .*

Note that in the following theorem $\ker(f)$, the *kernel* of a bundle map f , is the topological subspace $\bigcup_{x \in X} \ker(f_x)$. We define $\operatorname{im}(f)$ similarly.

Theorem 1.3.6 *Let $f : E \rightarrow F$ be a bundle map. Then the following statements are equivalent:*

1. $\operatorname{im}(f)$ is a subbundle of F ;
2. $\ker(f)$ is a subbundle of E ;
3. the dimensions of the fibres of $\operatorname{im}(f)$ are locally constant;
4. the dimensions of the fibres of $\ker(f)$ are locally constant.

(Note that the linear structure on each $\operatorname{im}(f_x)$ and $\ker(f_x)$ is understood to be that induced by F_x and E_x respectively.)

Clearly statements (3) and (4) are equivalent and are implied by either statement (1) or (2). Thus, the theorem would be proved if (3) could be shown to imply (1) and (4) could be shown to imply (2). Now, if one assumes local constancy of the fibres, then Lemma 1.3.4 can be used to construct local bases for $\operatorname{im}(f)$ and $\ker(f)$ at each point of X . A little thought will verify that the existence of a local base implies local triviality, and the result follows.

Theorem 1.3.7 *If X is a compact Hausdorff space, then for any module homomorphism $G : \Gamma(E) \rightarrow \Gamma(F)$, there is a unique bundle map $g : E \rightarrow F$ such that $G = \Gamma(g)$.*

The first step in establishing this theorem is to show that $\Gamma(E)/I_x \simeq F_x$, for all $x \in X$, where $I_x = \{f \in C(X) : f(x) = 0\}$. Since F clearly induces a map from $\Gamma(E)/I_x$ to $\Gamma(F)/I_x$, it must now induce a map from E to F . This map can then be shown to satisfy the required properties and the result follows.

1.3.3 Finite Projective Modules

Recall that a right A -module \mathcal{E} is said to be *projective* if it is a direct summand of a free module; that is, if there exists a free module \mathcal{F} and a module \mathcal{E}' , such that

$$\mathcal{F} = \mathcal{E} \oplus \mathcal{E}'.$$

Recall also that this is equivalent to the following alternative definition: A right A -module \mathcal{E} is *projective* if, given a surjective homomorphism $\tau : \mathcal{M} \rightarrow \mathcal{N}$ of right A -modules, and a homomorphism $\lambda : \mathcal{E} \rightarrow \mathcal{N}$, there exists a homomorphism

$\tilde{\lambda} : \mathcal{E} \longrightarrow \mathcal{M}$ such that $\tau \circ \tilde{\lambda} = \lambda$, or equivalently, such that the following diagram is commutative

$$\begin{array}{ccc} & & \mathcal{E} \\ & \swarrow \tilde{\lambda} & \downarrow \lambda \\ \mathcal{M} & \xrightarrow{\tau} & \mathcal{N}. \end{array}$$

Suppose now that \mathcal{E} is a projective and finitely-generated module over A . Clearly, there exists a surjective homomorphism $\tau : A^n \rightarrow \mathcal{E}$, for some natural number n . Since $\text{id}_{\mathcal{E}}$ is a homomorphism from \mathcal{E} to \mathcal{E} , the projective properties allow us to find a map $\tilde{\lambda} : \mathcal{E} \rightarrow A^n$ such that $\tau \circ \tilde{\lambda} = \text{id}_{\mathcal{E}}$, or equivalently, such that the following diagram is commutative

$$\begin{array}{ccc} & & \mathcal{E} \\ & \swarrow \tilde{\lambda} & \downarrow \text{id}_{\mathcal{E}} \\ A^n & \xrightarrow{\tau} & \mathcal{E}. \end{array}$$

We then have an idempotent element p of $\text{End}A^n$, given by

$$p = \tilde{\lambda} \circ \tau.$$

We can see that p is idempotent from

$$p^2 = \tilde{\lambda} \circ \tau \circ \tilde{\lambda} \circ \tau = \tilde{\lambda} \circ \tau = p.$$

This allows one to decompose the free module A^n , in the standard manner, as a direct sum of submodules,

$$A^n = \text{im}(p) \oplus \ker(p) = pA^n \oplus (1 - p)A^n.$$

Now, since $\tilde{\lambda} \circ \tau = \text{id}_{pA^n}$ and $\tau \circ \tilde{\lambda} = \text{id}_{\mathcal{E}}$, we have that \mathcal{E} and pA^n are isomorphic as right A -modules. Thus, a module \mathcal{E} over A is finitely-generated and projective if, and only if, there exists an idempotent $p \in \text{End}A^n$ such that $\mathcal{E} = pA^n$.

1.3.4 Serre–Swan Theorem

As a precursor to the Serre–Swan Theorem, we shall show that if X is a compact Hausdorff space, and E is a vector bundle over X , then $\Gamma(E)$ is a finitely-generated projective right- $C(X)$ -module.

Let E be a vector bundle over X , and let S be a subbundle of E . We endow E with an inner product $\langle \cdot, \cdot \rangle$, and we denote the projection of E onto S by P . If S_x^\perp

is the subspace of E_x that is orthogonal to S_x , then $S^\perp = \bigcup_{x \in X} S_x^\perp$ is the kernel of P . Since the image of P is S , Theorem 1.3.6 implies that S^\perp is a subbundle of E . Now, the mapping

$$S \oplus S^\perp \rightarrow E, \quad (v, w) \mapsto v + w,$$

is clearly a continuous mapping. Moreover, it is an isomorphism on each fibre. Thus, Lemma 1.3.2 implies that the two spaces are isomorphic. We now summarise what we have established in the following lemma.

Lemma 1.3.8 *Let E be a vector bundle equipped with an inner product, and let S be a subbundle of E . If S_x^\perp is the subspace of E_x that is orthogonal to S_x and $S^\perp = \bigcup_{x \in X} S_x^\perp$, then S^\perp is a subbundle of E and $E \simeq S^\perp \oplus S$.*

Lemma 1.3.9 *Let X be a compact Hausdorff space and let E be a vector bundle. Then there exists a surjective bundle map f from the trivial bundle $X \times \mathbf{C}^n$ to E , for some positive integer n .*

Proof. For each $x \in X$, choose a set of sections $s_1^x, \dots, s_{k_x}^x \in \Gamma(E)$ that form a local base over U_x , a base neighbourhood of x . A finite number of these neighbourhoods cover X . Therefore, there are a finite number of sections $s_1, \dots, s_n \in \Gamma(E)$ such that $s_1(x), \dots, s_n(x)$ span E_x , for every $x \in X$.

Now, $\Gamma(X \times \mathbf{C}^n)$ is a free module over $C(X)$ generated by sections $t_i(x) = (x, e_i)$, where $\{e_i\}_{i=1}^n$ is the standard basis of \mathbf{C}^n . There is a unique module map from $\Gamma(X \times \mathbf{C}^n)$ to $\Gamma(E)$ that maps each t_i onto s_i . By Theorem 1.3.7 this mapping is induced by a map $f : X \times \mathbf{C}^n \rightarrow E$. Since,

$$f(t_i(x)) = [\Gamma(f)(t_i)](x) = s_i(x), \quad \text{for all } x \in X,$$

it is clear that f is surjective. □

Corollary 1.3.10 *If X is a compact Hausdorff space, then any vector bundle E is a direct summand of a trivial bundle, and $\Gamma(E)$ is a finitely-generated projective right- $C(X)$ -module.*

Proof. Let $f : X \times \mathbf{C}^n \rightarrow E$ be the map defined in the previous lemma. Since $\text{im}(f) = E$, Theorem 1.3.6 implies that $\ker(f)$ is a subbundle of $X \times \mathbf{C}^n$. If we put an inner product on $X \times \mathbf{C}^n$, then by Lemma 1.3.8,

$$\ker(f) \oplus \ker(f)^\perp \simeq X \times \mathbf{C}^n.$$

Restricting f to $\ker(f)^\perp$, we see that it is a linear isomorphism on each fibre, and so $\ker(f)^\perp \simeq E$.

We can identify $\Gamma(\ker(f) \oplus F)$ and $\Gamma(\ker(f)) \oplus \Gamma(F)$, using the module isomorphism

$$\Gamma(\ker(f)) \oplus \Gamma(F) \rightarrow \Gamma(\ker(f) \oplus F), \quad s \oplus t \mapsto (s, t);$$

where $(s, t)(x) = (s(x), t(x))$. Hence, we have that $\Gamma(\ker(f))$ is a direct summand of the finitely generated free $C(X)$ -module $\Gamma(X \times \mathbf{C}^n)$. \square

Finally, we are now in a position to prove the principal result of this section, the Serre–Swan Theorem. It was first published in 1962 [104], and was inspired by a paper of Serre [100] that established an analogous result for algebraic vector bundles over affine varieties.

Theorem 1.3.11 (Serre–Swan) *Let X be a compact Hausdorff space. Then a module \mathcal{E} over $C(X)$ is isomorphic to a module of the form $\Gamma(E)$ if, and only if, \mathcal{E} is finitely generated and projective.*

Proof. If \mathcal{E} is finitely-generated and projective, then, as explained earlier, there exists an idempotent endomorphism $p : C(X)^n \rightarrow C(X)^n$, with $\mathcal{E} \simeq \text{im}(p)$, for some natural number n . Clearly $C(X)^n$ can be associated with the sections of the trivial vector bundle $X \times \mathbf{C}^n$ by mapping (f_1, \dots, f_n) to the section s , defined by $s(x) = (x, f_1(x), \dots, f_n(x))$.

By Theorem 1.3.7, p is the image under Γ of a bundle map $f : X \times \mathbf{C}^n \rightarrow X \times \mathbf{C}^n$. Since $p^2 = p$, and since $p(s) = f \circ s$, we have that

$$f^2 \circ s = p(f \circ s) = p^2(s) = p(s) = f \circ s,$$

for all sections s . Thus, since $\{s_i(x)\}_{i=1}^n$ spans $(X \times \mathbf{C}^n)_x$, for all $x \in X$, it holds that $f^2 = f$.

Let us now define the map

$$(1 - f) : X \times \mathbf{C}^n \rightarrow X \times \mathbf{C}^n, \quad v \mapsto v - f(v),$$

where the addition takes place fibrewise. As usual we denote the restriction of f to the fibre $(X \times \mathbf{C}^n)_x$ by f_x . Since f_x is an idempotent linear map, it holds that $\text{im}(1 - f_x) = \ker(f_x)$. Clearly this implies that $\ker(f) = \text{im}(1 - f)$.

Suppose that $\dim(\text{im}(f_x)) = h$ and that $\dim(\ker(f_x)) = k$, then Lemma 1.3.4 implies that $\dim(\text{im}(f_y)) \geq h$, and $\dim(\ker(f_y)) \geq k$, for all y in some neighbourhood of x . However, since $(X \times \mathbf{C}^n)_x = \text{im}(f_x) \oplus \ker(f_x)$, for all $x \in X$,

$$\dim(\text{im}(f_y)) + \dim(\ker(f_y)) = \dim(X \times \mathbf{C}^n)_y = h + k$$

is a constant. Thus, $\dim(\text{im}(f_y))$ must be locally constant. This implies that $\text{im}(f)$ is a subbundle of $X \times \mathbf{C}^n$.

If we make the observation that

$$\Gamma(\text{im}(f)) = \{f \circ s : s \in \Gamma(X \times \mathbf{C}^n)\},$$

then we can see that

$$\Gamma(\text{im}(f)) = \text{im}(\Gamma(f)) = \text{im}(p) = \mathcal{E}.$$

Thus, \mathcal{E} is indeed isomorphic to the sections of a vector bundle.

The proof in the other direction follows from Corollary 1.3.10. \square

Thus, the Serre–Swan Theorem shows that the vector bundles over compact Hausdorff space X are in one-to-one correspondence with the finite projective modules over $C(X)$. This motivates us to view finitely-generated projective modules over noncommutative C^* -algebras as non-commutative vector bundles.

Smooth Vector Bundles

Let E be a vector bundle over a manifold X . It is not too hard to see that we can use the differential structure of X to canonically endow E with a differential structure. A routine check will establish that the local trivialisations of the bundle then become smooth maps. In general, if (E, π, X) is a vector bundle such that E and X are manifolds and all the local trivialisations are smooth, then we call (E, π, X) a *smooth vector bundle*. *Smooth vector bundle maps* and *smooth vector bundle isomorphisms* are defined in the obvious way. When we speak of the *smooth sections of an ordinary vector bundle* we mean the smooth sections of the bundle endowed with the canonical differential structure discussed above. Clearly, the tangent and cotangent bundles of a manifold are smooth vector bundles. An important point to note is that a direct analogue of Proposition 1.3.3 holds for smooth vector bundles.

When a section of E is also a smooth map between X and E , then we call it a *smooth section*; we denote the set of smooth sections by $\Gamma^\infty(E)$. The canonical example of a smooth section is a smooth vector field over a manifold; it is a smooth section of the tangent bundle. Now, just as we gave $\Gamma(E)$ the structure of a right module over $C(X)$, we can give $\Gamma^\infty(E)$ the structure of a right module over $C^\infty(X)$, where $C^\infty(X)$ is the algebra of smooth complex-valued functions on X .

It can be shown, using an argument quite similar to the one above, that the modules of smooth sections of the smooth vector bundles over X are in one-to-one correspondence with the finitely-generated projective modules over $C^\infty(X)$; for details see [39].

1.4 Von Neumann Algebras

In this section we give a brief presentation of noncommutative measure theory. We shall not venture too far into the details since, with the exception of the material presented on locally convex spaces, we shall not return to this area again. It is introduced here for its heuristic value only. We refer the interested reader to [80] for the details of von Neumann algebras, and to [12] for an in-depth presentation of noncommutative measure theory.

The parallels between noncommutative measure theory and noncommutative topology are obvious. In fact, both areas are part of an overall trend in mathematics towards viewing function algebras as special commutative cases of operator algebras. This area is loosely known as *quantum mathematics*, an obvious reference to the quantum mechanical origins of operator theory. An excellent overview of this trend towards ‘quantization’ can be found in [107].

Locally Convex Topological Vector Spaces

Let \mathcal{P} be a non-empty family of seminorms on a linear space X . If $x \in X$, $\varepsilon \geq 0$, and \mathcal{P}_0 is a finite subfamily of \mathcal{P} , then we define

$$B(x, \mathcal{P}_0, \varepsilon) = \{y \in X : p(x - y) \leq \varepsilon, p \in \mathcal{P}_0\}.$$

It is straightforward to show that

$$\mathcal{B} = \{B(x, \mathcal{P}_0, \varepsilon) : x \in X, \mathcal{P}_0 \text{ a finite subfamily of } \mathcal{P}, \varepsilon \geq 0\}$$

is a base for a topology on X ; it is called the *topology generated by \mathcal{P}* . It is clear that a net $(x_\lambda)_\lambda$ in X converges to x , with respect to this topology, if, and only if, $p(x - x_\lambda) \rightarrow 0$, for all $p \in \mathcal{P}$. If $x_\lambda \rightarrow x$ and $y_\lambda \rightarrow y$, then it is easily seen that $p(x_\lambda + y_\lambda - x - y) \rightarrow 0$, for all $p \in \mathcal{P}$. Thus, addition is continuous with respect to the topology generated by \mathcal{P} . Similarly, scalar multiplication can be shown to be continuous. Hence, when X is endowed with this topology it is a topological vector space. (Recall that a *topological vector space* is a linear space for which the linear space operations of addition and scalar multiplication are continuous.) It can be shown that the topology is Hausdorff if, and only if, for each $x \in X$, there exists a $p \in \mathcal{P}$ such that $p(x) \neq 0$.

We call a topological vector space whose topology is determined by a family of seminorms a *locally convex (topological vector) space* (*locally convex* refers to the fact that \mathcal{B} forms a locally convex base for the topology).

The simplest example of a locally convex space is a normed vector space. The family of seminorms is just the one element set containing the norm, and the topology generated is the norm topology.

Von Neumann Algebras

If H is a Hilbert space then the *strong operator* topology is the topology generated by the family of seminorms $\{\|\cdot\|_x : x \in H\}$, where

$$\|T\|_x = \|Tx\|, \quad T \in B(H).$$

We denote the strong operator topology by τ_S . Thus, $T_\lambda \rightarrow T$ with respect to τ_S if, and only if, $\|(T_\lambda - T)x\| \rightarrow 0$, for all $x \in H$.

Recall that if τ_1, τ_2 are two topologies on a set X such that convergence of a net with respect to τ_1 implies convergence of the net with respect to τ_2 , then $\tau_2 \subseteq \tau_1$. Suppose now that $A_\lambda \rightarrow A$ in $B(H)$ with respect to the norm topology. Since $\|(A_\lambda - A)x\| \leq \|A_\lambda - A\|\|x\|$, for all $x \in X$, we have that $A_\lambda \rightarrow A$ with respect to τ_S . Thus, the strong operator topology is weaker than the norm topology. In fact, it can be strictly weaker. This happens if, and only if, H is infinite-dimensional.

Definition 1.4.1. Let H be a Hilbert space. If \mathcal{A} is a $*$ -subalgebra of $B(H)$ that is closed with respect to τ_S , then we call \mathcal{A} a *von Neumann algebra*.

Since every subset that is closed with respect to the strong operator topology is also closed with respect to the norm topology, every von Neumann algebra is also a C^* -algebra. In fact, it can be shown that every von Neumann algebra is a unital C^* -algebra.

If A is an algebra and C is a subset of A , then we define C' , the *commutant* of C , to be the set of elements of A that commute with all the elements of C . We define C'' to be $(C)'$ and call it the *double commutant* of C . Clearly $C \subseteq C''$. The following famous theorem shows us that there exists a very important relationship between the double commutants and von Neumann algebras; see chapter 4 of [80] for a proof.

Theorem 1.4.2 (von Neumann) *If \mathcal{A} is a $*$ -subalgebra of $B(H)$, for some Hilbert space H , such that $\text{id}_H \in \mathcal{A}$, then \mathcal{A} is a von Neumann algebra if, and only if, $\mathcal{A}'' = \mathcal{A}$.*

1.4.1 Noncommutative Measure Theory

L^∞ as a von Neumann Algebra

Recall that if (M, μ) is a measure space, then $L^\infty(M, \mu)$ is the algebra of all equivalence classes of measurable functions on M that are bounded almost everywhere

(two functions being equivalent if they are equal almost everywhere). If it is equipped with the L^∞ -norm defined by

$$\|f\|_\infty = \inf\{C \geq 0 : |f| \leq C \text{ a.e.}\},$$

then it is a unital Banach algebra. We can define an involution on $L^\infty(M, \mu)$ by $f^* = \overline{f}$, and this clearly gives it the structure of a C^* -algebra.

Recall also that $L^2(M, \mu)$ is the algebra of equivalence classes of measurable functions on M such that if $f \in L^2(M, \mu)$, then $|f|^2$ has finite integral. If it is equipped with the L^2 -norm, defined by setting

$$\|f\|_2 = \left(\int_M |f|^2 d\mu \right)^{\frac{1}{2}},$$

then it is a Banach space. In fact, $L^2(M, \mu)$ is a Hilbert space since the norm is generated by the inner product $\langle f, g \rangle = \int_M f \overline{g} d\mu$.

Let us now assume that M is a compact Hausdorff space and that μ is a finite positive regular Borel measure. For any $f \in L^\infty(M, \mu)$, consider the mapping

$$M_f : L^2(M, \mu) \rightarrow L^2(M, \mu), \quad g \mapsto fg.$$

Clearly each M_f is a linear mapping, and since

$$\|M_f g\|_2^2 = \int_M |fg|^2 d\mu \leq \|f\|_\infty^2 \int_M |g|^2 d\mu = \|f\|_\infty^2 \|g\|_2^2,$$

each M_f is bounded by $\|f\|_\infty$. In fact, using the regularity of the measure, we can show that $\|M_f\| = \|f\|_\infty$. The adjoint of M_f is equal to $M_{\overline{f}}$, and so the mapping

$$M : L^\infty(M, \mu) \rightarrow B(L^2(M, \mu)), \quad f \mapsto M_f$$

is an isometric $*$ -isomorphism between $L^\infty(M, \mu)$ and $\mathcal{L} = M(L^\infty(M, \mu))$. Moreover, it can be shown that if $T \in \mathcal{L}'$, then $T = M_f$, for some $f \in L^\infty(M, \mu)$. Thus, by Theorem 1.4.2, $L^\infty(M, \mu)$ is an abelian von Neumann algebra.

General Abelian von Neumann Algebras

The following theorem shows that all abelian von Neumann are of the form $L^\infty(M, \mu)$, for some measure space (M, μ) . We can consider this result the analogue of the Gelfand–Naimark Theorem for von Neumann algebras; for a proof see [20].

Theorem 1.4.3 *Let \mathcal{A} be an abelian von Neumann algebra on a Hilbert space H . Then there exists a locally compact Hausdorff space M and a positive Borel measure μ on M such that \mathcal{A} is isometrically $*$ -isomorphic to $L^\infty(M, \mu)$.*

It is interesting to note that the space \mathcal{M} is produced as the spectrum of a C^* -subalgebra of \mathcal{A} that is dense in \mathcal{A} with respect to the strong operator topology. Furthermore, if H is separable, then M can be shown to be compact and second-countable.

Noncommutative Measure Theory

In the same spirit as noncommutative topology, we now think of noncommutative von Neumann algebras as ‘noncommutative measure spaces’. A lot of work has been put into finding von Neumann algebra structures that correspond to measure theoretic structures in the commutative case. As would be expected, this area of mathematics is called *noncommutative measure theory*.

Central to most of this work is the notion of a factor. A *factor* is a von Neumann algebra with a trivial centre, that is, a von Neumann algebra \mathcal{A} for which $\mathcal{A}' \cap \mathcal{A} = \mathbf{C}1$. The simplest example is $B(H)$, for any Hilbert space H . Factors have been classified into three types according to the algebraic properties of the projections they contain. The important thing about them is that every von Neumann algebra is isomorphic to a direct integral of factors. (A *direct integral* of linear spaces is a continuous analogue of the direct sum of linear spaces.)

Much of Connes’ original work was in this area. Building on the Tomita–Takesaki Theorem, he established a noncommutative version of the Radon–Nikodym Theorem. This result furnishes a canonical homomorphism from the additive group \mathbf{R} to the group of outer automorphisms of any noncommutative von Neumann algebra. It has no parallel in the commutative case and it inspired Connes to write that ‘noncommutative measure spaces evolve with time’. This work then led on to a classification of all hyperfinite type III factors (type III being one of the three types of factors).

Connes has found applications for his results in the study of the type of singular spaces discussed in the introduction. He has met with particular success in the study of foliations of manifolds.

1.5 Summary

We conclude this chapter by summarising the algebraic generalisations of the elements of topology, differential geometry, and measure theory collected above:

locally compact space	→	C^* -algebra,
compact space	→	unital C^* -algebra,
homeomorphism	→	*-isomorphism,
image of a function	→	spectrum of a element,
positive function	→	positive element,
regular Borel complex measure	→	bounded linear functionals on a C^* -algebra,
one-point compactification of a space	→	unitisation of a C^* -algebra,
closed subset of a compact space	→	quotient of a unital C^* -algebra,
open subset of a compact space	→	ideal of a unital C^* -algebra,
connected compact space	→	projectionless unital C^* -algebra,
metrisable compact space	→	separable unital C^* -algebra,
vector bundle	→	finite projective module
over a locally compact space		over a C^* -algebra,
measure space	→	von Neumann algebra.

Remark 1.5.1. Most of the results of this chapter can be expressed in a very satisfactory manner using the language of category theory. For example, most of the above correspondences can be viewed as functors between the respective categories. For a presentation of this approach see [39].

Chapter 2

Differential Calculi

We are now ready to consider generalised differential structures on quantum spaces. Following on from the last chapter, we shall begin with a compact manifold M , and then attempt to express its structure in terms of the algebra of its continuous functions.

A natural starting point would be to try and establish an algebraic relationship between the smooth and the continuous functions of M . Unfortunately, however, there does not appear to be any simple way of doing this. In fact, there does not seem to be any simple algebraic properties that characterise $C^\infty(M)$ at all. The best that we can do is establish that $C^\infty(M)$ is dense in $C(M)$ using the Stone–Weierstrass Theorem. Consequently, we shall use an arbitrary associative (involutive) algebra A to generalise $C^\infty(M)$. We could assume that A is dense in some C^* -algebra, but there is no major technical advantage in doing so. Neither is there any advantage in assuming that A is unital.

We shall begin this chapter by reviewing differential calculus on a manifold in global algebraic terms. This will naturally lead us to the definition of a differential calculus: this object is a generalisation of the notion of the de Rham calculus, and it is of fundamental importance in noncommutative geometry.

We shall then introduce derivation-based differential calculi, as formulated by M. Dubois–Violette and J. Madore [27]. Their work has been strongly influenced by that of J.L. Koszul who in [58] described a powerful algebraic version of differential geometry in terms of a general commutative associative algebra. We present derivation-based calculi here because we shall refer to them in our discussion of fuzzy physics in Chapter 5, and because they provide a pleasingly straightforward example of a noncommutative differential calculus.

2.1 The de Rham Calculus

Let M be an n -dimensional manifold, and let $C^\infty(M)$ denote the algebra of smooth complex-valued functions on M . Unless otherwise stated, we shall always assume that the manifolds we are dealing with are smooth, real, compact, and without boundary. We define a *smooth vector field* X on M to be a *derivation* on $C^\infty(M)$; that is, a linear mapping on $C^\infty(M)$ such that, for all $f, g \in C^\infty(M)$,

$$X(fg) = X(f)g + fX(g).$$

The set of smooth vector fields on M is denoted by $\mathcal{X}(M)$. We give it the structure of a (left) $C^\infty(M)$ -module in the obvious way. We denote the dual module of $\mathcal{X}(M)$ by $\Omega^1(M)$, and we call it the *module of differential 1-forms* over M . It is easily shown that these definitions are equivalent to defining a smooth vector field to be a smooth section of the tangent bundle of M , and a differential 1-form to be a smooth section of the cotangent bundle of M .

We denote the $C^\infty(M)$ -module of $C^\infty(M)$ -valued, $C^\infty(M)$ -multilinear mappings on

$$\Omega^{\times p}(M) \times \mathcal{X}^{\times q}(M) = \underbrace{\Omega^1(M) \times \dots \times \Omega^1(M)}_{p \text{ times}} \times \underbrace{\mathcal{X}(M) \times \dots \times \mathcal{X}(M)}_{q \text{ times}}$$

by $\mathfrak{T}_q^p(M)$, and we call it the *space of rank- (p, q) smooth tensor fields* on M . Note that $\mathfrak{T}_1^0(M) = \Omega^1(M)$.

An important point about tensor fields is their ‘locality’. Let A_1 and A_2 be two elements of $\Omega^{\times p}(M) \times \mathcal{X}^{\times q}(M)$ such that $A_1(p) = A_2(p)$, for some $p \in M$. If T is a rank- (p, q) tensor field, then, using a simple bump-function argument, it can be shown that $T(A_1)(p) = T(A_2)(p)$.

Let \mathcal{E} be a module over R , and let S be a multilinear mapping on \mathcal{E}^n with values in R . We say that S is *anti-symmetric* if, for every $e_1, \dots, e_n \in \mathcal{E}$,

$$S(e_1, \dots, e_i, \dots, e_j, \dots, e_n) = -S(e_1, \dots, e_j, \dots, e_i, \dots, e_n),$$

for all $1 \leq i < j \leq n$. Consider the submodule of $\mathfrak{T}_p^0(M)$ consisting of the anti-symmetric p -linear maps on $\mathcal{X}^1(M)$. We call it the module of *smooth exterior differential p -forms* on M , or, more simply, the module of *p -forms* on M ; we denote it by $\Omega^p(M)$. Again, it is easily shown that this is equivalent to defining a p -form to be a smooth section of p^{th} -exterior power of the cotangent bundle of M . This equivalent definition implies that $\Omega^p(M) = \{0\}$, for $p > n$.

Let us define

$$\Omega(M) = \bigoplus_{p=0}^{\infty} \Omega^p(M),$$

with $\Omega^0(M) = C^\infty(M)$. Then consider the associative bilinear mapping

$$\Lambda : \Omega(M) \times \Omega(M) \rightarrow \Omega(M), \quad (\omega, \omega') \mapsto \omega \wedge \omega';$$

where if $\omega_p \in \Omega^p(M)$, and $\omega_q \in \Omega^q(M)$, then

$$\begin{aligned} & \omega_p \wedge \omega_q(X_{\pi(1)}, \dots, X_{\pi(k)}) \\ &= \sum_{\pi \in \text{Perm}(p+q)} \text{sgn}(\pi) \omega_p(X_{\pi(1)}, \dots, X_{\pi(k)}) \omega_q(X_{\pi(p+1)}, \dots, X_{\pi(p+q)}) \end{aligned} \quad (2.1)$$

(Note that $\sum_{\pi \in \text{Perm}(p+q)}$ means summation over all permutations of the numbers $1, 2, \dots, p+q$, and $\text{sgn}(\pi)$ is the sign of the permutation π .) When $\Omega(M)$ is endowed with Λ we call it the *algebra of exterior differential forms* on M . Upon examination we see that Λ is *graded commutative*; that is, if $\omega_p \in \Omega^p(M)$ and $\omega_q \in \Omega^q(M)$, then

$$\omega_q \wedge \omega_p = (-1)^{pq} \omega_p \wedge \omega_q.$$

2.1.1 The Exterior Derivative

Consider the canonical mapping from the smooth functions of M to the smooth 1-forms of M defined by

$$d : \Omega^0(M) \rightarrow \Omega^1(M), \quad f \mapsto df, \quad (2.2)$$

where

$$df(X) = X(f).$$

This mapping admits a remarkable extension to a linear operator d on $\Omega(M)$, such that,

$$d(\Omega^p(M)) \subset \Omega^{p+1}(M), \quad \text{for all } p \geq 0.$$

It is called the *exterior differentiation operator*, and it is defined on $\Omega^p(M)$ by

$$\begin{aligned} d\omega(X_0, \dots, X_p) &= \sum_{i=0}^p (-1)^i X_i(\omega(X_0, \dots, \widehat{X}_i, \dots, X_p)) \\ &\quad + \sum_{0 \leq i < j \leq p} (-1)^{i+j} \omega([X_i, X_j], X_0, \dots, \widehat{X}_i, \dots, \widehat{X}_j, \dots, X_p), \end{aligned} \quad (2.3)$$

where \widehat{X}_i means that X_i is omitted, and $[X_i, X_j] = X_i X_j - X_j X_i$. (Note that $[X_i, X_j]$ is indeed a derivation on $C^\infty(M)$; it is called the *Lie bracket* of X_i and X_j .) It is routine to check that $d\omega$ is an anti-symmetric $(p+1)$ -linear mapping on $\mathcal{X}(M)$. Moreover, when $p = 0$, the definition reduces to $df(X) = X(f)$, showing that d does indeed extend the operator defined in (2.2).

The pair $(\Omega(M), d)$ is called the *de Rham calculus* of M . Using a standard partition of unity argument, it is easy to show that every element of $\Omega^k(M)$ is a sum of the elements of the form

$$f_0 df_1 \wedge df_2 \cdots \wedge df_k.$$

Furthermore, it is also easy to show that $d^2 = 0$, and that

$$d(\omega_p \wedge \omega_q) = d\omega_p \wedge \omega_q + (-1)^p \omega_p \wedge d\omega_q.$$

Remark 2.1.1. All of the above constructions and definitions are equally well defined if one uses $C^\infty(M; \mathbf{R})$, the algebra of smooth real-valued functions on M , instead of $C^\infty(M)$. An important point to note is that if $T_p(M)$ denotes the tangent plane to M at p , and if $T_p(M; \mathbf{R})$ denotes the real tangent plane to M at p , then $T_p(M) = T_p(M; \mathbf{R}) \otimes \mathbf{C}$.

2.2 Differential calculi

A *positively-graded algebra* is an algebra of the form $\Omega = \bigoplus_{n=0}^{\infty} \Omega^n$, where each Ω^n is a linear subspace of Ω , and $\Omega^p \Omega^q \subseteq \Omega^{p+q}$, for all $p, q \geq 0$. (Since there will be no risk of confusion, we shall refer to a positively graded algebra simply as a graded algebra.) If $\omega \in \Omega^p$, then we say that ω is of *degree* p . A *homogenous mapping of degree* k from a graded algebra Ω to a graded algebra Λ is a linear mapping $h : \Omega \rightarrow \Lambda$ such that if $\omega \in \Omega^p$, then $h(\omega) \in \Lambda^{p+k}$. A *graded derivation* on a graded algebra Ω is a homogenous mapping of degree 1 such that

$$d(\omega\omega') = d(\omega)\omega' + (-1)^p \omega d\omega',$$

for all $\omega \in \Omega^p$, and $\omega' \in \Omega$. A pair (Ω, d) is a *differential algebra* if Ω is a graded algebra and d is a graded derivation on Ω such that $d^2 = 0$. The operator d is called the *differential* of the algebra. A *differential algebra homomorphism* φ between two differential algebras (Ω, d) and (Λ, δ) is an algebra homomorphism between Ω and Λ , which is also a homogenous mapping of degree 0 that satisfies $\varphi \circ d = \delta \circ \varphi$.

Definition 2.2.1. A *differential calculus* over an algebra A is a differential algebra (Ω, d) , such that $\Omega^0 = A$, and

$$\Omega^n = d(\Omega^{n-1}) \oplus Ad(\Omega^{n-1}), \quad \text{for all } n \geq 1. \quad (2.4)$$

We note that some authors [26] prefer to omit condition (2.4) from the definition of a differential calculus.

Clearly $(\Omega(M), d)$ is a differential calculus over $C^\infty(M)$. However, it should be noted that for a general calculus there is no analogue of the graded commutativity of classical differential forms.

The definition has some immediate consequences. Firstly, if a differential calculus (Ω, d) is unital (as an algebra) then the unit of Ω must belong to $\Omega^0 = A$, and so it must be a unit for A . It then follows that

$$d(1) = d(1.1) = d(1).1 + 1.d(1) = 2.d(1).$$

Thus, if 1 is the unit of a differential calculus, then $d(1) = 0$. If $a_0 da_1 \in \Omega^1$, then

$$d(a_0 da_1) = da_0 da_1 + a_0 d(da_1) = da_0 da_1.$$

In general, an inductive argument will establish that

$$d(a_0 da_1 \dots da_n) = da_0 da_1 \dots da_n.$$

If $a_0 da_1 da_2 \in \Omega^2$, and $b_0 db_1 \in \Omega^1$, then

$$\begin{aligned} (a_0 da_1 da_2)(b_0 db_1) &= a_0 da_1 d(a_2 b_0) db_1 - a_0 (da_1) a_2 db_0 db_1 \\ &= a_0 da_1 d(a_2 b_0) db_1 - a_0 d(a_1 a_2) db_0 db_1 \\ &\quad + a_0 a_1 da_2 db_0 db_1. \end{aligned}$$

Using an inductive argument again, it can be established that, in general,

$$\begin{aligned} (a_0 da_1 \dots da_n)(a_{n+1} da_{n+2} \dots da_{n+k}) &= (-1)^n a_0 a_1 da_2 \dots da_{n+k} \\ &\quad + \sum_{r=1}^n (-1)^{n-r} a_0 da_1 \dots d(a_r a_{r+1}) \dots da_{n+k}. \end{aligned}$$

A *differential ideal* of a differential calculus (Ω, d) is a two-sided ideal $I \subseteq \Omega$, such that $I \cap \Omega^0 = \{0\}$, and $d(I) \subseteq I$. Let π denote the projection from Ω to Ω/I , and let \tilde{d} denote the mapping on Ω/I defined by $\tilde{d}(\pi(\omega)) = \pi(d(\omega))$ (it is well defined since $d(I) \subseteq I$). It is easy to see that, with respect to the natural grading that Ω/I inherits from Ω , $(\Omega/I, \tilde{d})$ is a differential algebra. Furthermore, since condition (2.4) is obviously satisfied, and $\pi(\Omega^0) \cong A$, we have that $(\Omega/I, \tilde{d})$ is a differential calculus over A .

We say that a graded algebra $\Omega = \bigoplus_{n=0}^{\infty} \Omega_n$ is *n-dimensional* if there exists a positive integer n , such that $\Omega^n \neq \{0\}$ and $\Omega^m = \{0\}$, for all $m > n$. If (Ω, d) is *n-dimensional*, for some positive integer n , then we say that it is *finite-dimensional*; otherwise we say that it is *infinite-dimensional*. Given an infinite-dimensional calculus (Ω, d) there is a standard procedure for abstracting from it a new calculus of any finite dimension n . Define the algebra $\Omega' = \bigoplus_{k=0}^{\infty} \Omega'^k$ by setting $\Omega'^k = \Omega^k$, if $k \leq n$, and $\Omega'^k = \{0\}$, if $k > n$. Then define a multiplication \cdot on Ω' as follows: for $\omega \in \Omega^k$, and $\omega' \in \Omega'^l$, define $\omega \cdot \omega' = \omega \omega'$, if $k + l \leq n$, and define $\omega \cdot \omega' = 0$, if $k + l > n$. Define $d'(\omega) = d(\omega)$, if $k < n$, and define $d'(\omega) = 0$, if $k \geq n$. We call (Ω', d') the *n-dimensional calculus obtained from (Ω, d) by truncation*.

Differential *-Calculi

Just as for $C(M)$, we can use complex-conjugation to define an algebra involution on $C^\infty(M)$. This involution extends to a unique involutive conjugate-linear map $*$ on $\Omega(M)$, such that $d(\omega^*) = (d\omega)^*$, for all $\omega \in \Omega(M)$. However, $*$ is not anti-multiplicative; if $\omega_p \in \Omega^p(M)$, and $\omega_q \in \Omega^q(M)$, then

$$(\omega_p \wedge \omega_q)^* = \omega_p^* \wedge \omega_q^* = (-1)^{pq} \omega_q^* \wedge \omega_p^*.$$

Let (Ω, d) be a differential calculus over a $*$ -algebra A . Then, there exists a unique extension of the involution of A to an involutive conjugate-linear map $*$ on Ω , such that $d(\omega^*) = (d\omega)^*$, for all $\omega \in \Omega$. If it holds that

$$(\omega_p \omega_q)^* = (-1)^{pq} \omega_q^* \omega_p^*, \quad \text{for all } \omega_p \in \Omega^p, \omega_q \in \Omega^q,$$

then we say that (Ω, d) is a *differential *-calculus*.

2.2.1 The Universal Calculus

As would be expected from the quite general nature of the definition, there can exist several distinct differential calculi over the same algebra. In fact, as we shall see, there even exist calculi over $C^\infty(M)$ that are different from $(\Omega(M), d)$. Thus, a differential calculus is not a strict generalisation of the notion of the de Rham calculus.

This is an example of a common feature of moving from the commutative to noncommutative setting: it often makes more sense to formulate a noncommutative version of a classical structure that is more (or sometimes less) general than the structure one had originally intended to generalise.

Over any algebra A , we can construct a differential calculus $\Omega_u(A)$ called the universal calculus. Before we present its construction, it should be said that while the universal calculus has a central role to play to the theory of differential calculi, it is in a sense ‘too large’ to be considered as a suitable generalisation of the de Rham calculus. The most important thing about $\Omega_u(A)$ is its ‘universal property’, which we shall discuss later.

Construction of $\Omega_u(\mathcal{A})$

Let A be an algebra and define $\Omega_u^1(A) = \tilde{A} \oplus A$; where \tilde{A} is linear space $A \oplus \mathbf{C}$ endowed with a multiplication as defined in (1.1). Define the structure of an A -bimodule on $\Omega_u^1(A)$ by

$$x((a + \lambda 1) \oplus b) = (xa + \lambda x) \oplus b,$$

and

$$((a + \lambda 1) \oplus b)y = (a + \lambda 1) \oplus by - (ab + \lambda b) \oplus y,$$

for $a, b, x, y \in A$. The mapping

$$d_u : A \rightarrow \Omega_u^1(A), \quad a \mapsto 1 \oplus a$$

is a derivation since

$$d_u(ab) = 1 \oplus ab = (1 \oplus a)b + a \oplus b = d_u(a)b + ad_u(b).$$

In general, if E is an A -bimodule, and δ a derivation from A to E , then we call the pair (E, δ) a *first order differential calculus* over A . We can define a mapping $i_\delta : \Omega_u^1(A) \rightarrow E$, by

$$i_\delta((a + \lambda 1) \oplus b) = a\delta(b) + \lambda\delta(b).$$

Clearly, $i_\delta \circ d_u = \delta$. Furthermore, as a straightforward calculation will verify, i_δ is an A -bimodule homomorphism. The existence of this surjective mapping, for any first order calculus (E, δ) , is known as the *universal property* of $(\Omega_u^1(A), d_u)$.

Let $\Omega_u^n(A)$ denote the n -fold tensor product of $\Omega_u^1(A)$ over A , and let $\Omega_u(A)$ denote the tensor algebra of $\Omega_u^1(A)$ over A . Clearly, $\Omega_u(A)$ is canonically a graded algebra, with $\Omega_u^0 = A$.

We would like to define an operator $d_u : \Omega_u^1(A) \rightarrow \Omega_u^2(A)$ such that, for all $a \in A$, $d_u(d_u a) = 0$. The fact that

$$(a + \lambda 1) \oplus b = a \oplus b + \lambda 1 \oplus b = ad_u b + \lambda d_u b,$$

implies that we must have

$$d_u[(a + \lambda 1) \oplus b] = (d_u a)(d_u b).$$

Let us now progress to define the operator

$$d_u : \Omega_u^2(A) \rightarrow \Omega_u^3(A), \quad \omega_1 \otimes \omega_2 \mapsto (d_u \omega_1) \otimes \omega_2 - \omega_1 \otimes (d_u \omega_2).$$

It is well defined since (as a routine calculation will verify) $d_u(\omega a) = (d_u \omega)a - \omega d_u a$, and $d_u(a\omega) = (d_u a)\omega + ad_u \omega$. In fact, it is the unique operator on $\Omega_u^2(A)$ for which $d_u(d_u \omega) = 0$, for all $\omega \in \Omega_u^1(A)$.

It is not very hard to build upon all of this to define a square-zero graded derivation d_u on $\Omega_u(A)$ that extends each d_u defined above. Neither is it too hard to see that this extension is necessarily unique.

A little reflection will verify that $\Omega_u^n(A)$ is spanned by the elements of the form $a_0 d_u(a_1) \otimes \cdots \otimes d_u(a_n)$ and $d_u(a_1) \otimes \cdots \otimes d_u(a_n)$, for every positive integer n . Hence, $(\Omega_u(A), d_u)$ is a differential calculus over A ; we call it the *universal differential calculus* over A .

From now on, for sake of simplicity, we shall denote $d_u(a_1) \otimes \cdots \otimes d_u(a_n)$ by $d_u a_1 \cdots d_u a_n$.

Some Remarks on $\Omega_u(A)$

Just like $\Omega_u^1(A)$, $\Omega_u(A)$ also has a ‘universal property’. If (Λ, δ) is another differential calculus over A , then, by the universal property of $\Omega_u^1(A)$, there exists a bimodule mapping $i_\delta : \Omega_u^1(A) \rightarrow \Lambda$ such that $i_\delta \circ d_u = \delta$. This extends to a unique surjective differential algebra homomorphism from $\Omega_u(A)$ to Λ , whose restriction to A is the identity. The existence of this mapping for any differential calculus over A is known as the *universal property* of $(\Omega_u(A), d)$. It is easily seen that the universal property defines the universal calculus uniquely up to isomorphism.

For any differential calculus (Λ, δ) the kernel of $i_\delta : \Omega_u \rightarrow \Lambda$ is a differential ideal. Hence, $(\Omega_u(A)/\ker(i_\delta), \tilde{d})$ is well defined as a differential calculus; clearly it is isomorphic to (Λ, δ) . This means that every differential calculus is isomorphic to a quotient of the universal calculus.

If $A = C^\infty(M)$, then it is reasonably clear that the universal calculus of A will not correspond to the de Rham calculus. The most obvious reasons are that it is neither unital nor finite-dimensional. The best we can say is that $(\Omega(M), d)$ is isomorphic to a quotient of $(\Omega_u(C^\infty(M)), d)$.

We remark that there is a natural isomorphism of linear spaces

$$j : \tilde{A} \otimes A^{\otimes n} \rightarrow \Omega_u^n(A),$$

defined by setting

$$j((a_0 + \lambda 1) \otimes a_1 \otimes \cdots \otimes a_n) = a_0 da_1 \cdots da_n + \lambda da_1 \cdots da_n.$$

This fact has an important consequence: for any positive integer n , let T_1 be a linear map from A^{n+1} to a linear space B , and let T_2 be a linear map from A^n to the same space. Then there exists a linear map $T : \Omega_u^n(A) \rightarrow B$ such that $T(a_0 d_u a_1 d_u a_2 \cdots d_u a_n) = T_1(a_0, a_1, \dots, a_n)$, and $T(d_u a_1 d_u a_2 \cdots d_u a_n) = T_2(a_1, a_2, \dots, a_n)$, for all $a_0, a_1, \dots, a_n \in A$.

2.3 Derivations-Based Differential Calculi

We shall now present an example of a differential calculus that is based on derivations. This approach is a direct generalisation of the construction of the de Rham calculus presented earlier. It was originally introduced by M. Dubois–Violette and J. Madore and it is based upon the work of Koszul. Dubois–Violette and Madore originally became interested in derivation based calculi because of their suitability for use in fuzzy physics. In Chapter 5 we shall discuss fuzzy physics, and the role derivation based calculi originally played in it.

Given an algebra A we define the set of *vector fields* over A to be $\text{Der}(A)$, the set of derivations on A . We give it the structure of a complex linear space in the obvious manner. To generalise the fact that $\mathcal{X}(M)$ is a module over $C^\infty(M)$ is a little more problematic because of the noncommutativity of A ; if $X \in \text{Der}(A)$, and $a \in A$, then, in general, it does not follow that aX or Xa are in $\text{Der}(A)$. According to Dubois–Violette, the most natural solution to this problem is to regard $\text{Der}(A)$ as a (left) module over $Z(A)$, the centre of A . In the classical case $Z(C^\infty(M)) = C^\infty(M)$, and so our definition reproduces the original module structure.

Generalising the classical case directly, we define $\underline{\Omega}_{\text{Der}}^n(A)$ to be the $Z(A)$ -module of anti-symmetric $Z(A)$ -multilinear mappings from $\text{Der}(A)^n$ to A . We then define $\underline{\Omega}_{\text{Der}}^0(A) = A$, and $\underline{\Omega}_{\text{Der}}(A) = \bigoplus_{n=0}^{\infty} \underline{\Omega}_{\text{Der}}^n(A)$. We endow $\underline{\Omega}_{\text{Der}}(A)$ with a product that is the direct analogue of the classical product defined in equation (2.1). With respect to this product $\underline{\Omega}_{\text{Der}}(A)$ is canonically a graded algebra. Then we define a differential d on $\underline{\Omega}_{\text{Der}}(A)$ that is the direct analogue of the classical exterior differentiation operator defined in equation (2.3). Unfortunately, the pair $(\underline{\Omega}_{\text{Der}}(A), d)$ is not necessarily a differential calculus over A , since it may happen that condition (2.4) is not satisfied.

If M is a compact manifold, then it is obvious that $(\underline{\Omega}_{\text{Der}}(C^\infty(M)), d)$ will coincide with the de Rham calculus. However, if M is not compact, or more specifically if M is not paracompact, then it may happen that the derivation based construction of the de Rham calculus will no longer coincide with the standard construction based on vector bundles. More explicitly, the module of sections of the p -exterior power of the cotangent bundle of M may be properly contained in the module of anti-symmetric $C^\infty(M)$ -multilinear maps on $\mathcal{X}(M)^n$. But, it has been observed that the calculus constructed using vector bundles will coincide with the smallest differential subalgebra of $(\underline{\Omega}_{\text{Der}}(C^\infty(M)), d)$ that contains $C^\infty(M)$ (see [27] and references therein for details).

This motivates us to consider $(\Omega_{\text{Der}}(A), d)$ the smallest differential subalgebra of $(\underline{\Omega}_{\text{Der}}(A), d)$ that contains A . It can be shown that $\Omega_{\text{Der}}(A)$ is the canonical image of $\Omega_u(A)$ in $\underline{\Omega}_{\text{Der}}(A)$. This means that it consists of finite sums of elements of the form $a_0 da_1 \dots da_n$, and $da_0 \dots da_n$. Hence, $(\Omega_{\text{Der}}(A), d)$ is a differential calculus over A . Dubois–Violette has proposed $(\Omega_{\text{Der}}(C^\infty(M)), d)$ as the most natural noncommutative generalisation of the de Rham calculus.

Using derivation based calculi we can formulate noncommutative versions of many elements of classical differential geometry. We present the following examples:

Noncommutative Lie Derivative

If X is a vector field over an n -dimensional manifold M , and $\omega \in \Omega^n(M)$, then $\mathcal{L}_X(\omega)$, the *Lie derivative* of ω with respect to X , is defined by setting

$$\mathcal{L}_X(\omega) = i_X \circ d(\omega) + d \circ i_X(\omega); \quad (2.5)$$

where i_X is the mapping from $\Omega^n(M)$ to $\Omega^{n-1}(M)$ defined by setting

$$i_X(\omega)(X_1, X_2, \dots, X_{n-1}) = \omega(X, X_1, X_2, \dots, X_{n-1}); \quad (2.6)$$

we call $i_X(\omega)$ the *contraction* of ω by X .

The Lie derivative is usually presented in terms of the pullback of the flow of a vector field. It is intuitively thought of as a ‘generalised directional derivative’ of ω with respect to X .

Returning to the noncommutative world we see that we can effortlessly generalise i_X . This then enables us to define the *Lie derivative* of $\omega \in \underline{\Omega}_{\text{Der}}(A)$ with respect to $X \in \text{Der}(A)$ by setting

$$\mathcal{L}_X(\omega) = i_X \circ d(\omega) + d \circ i_X(\omega).$$

In addition, it can be shown that $\underline{\Omega}_{\text{Der}}(A)$ is invariant under contraction by any element of $\text{Der}(A)$. Thus, we can also define a generalised Lie derivative for $\underline{\Omega}_{\text{Der}}(A)$. Obviously, both these definitions correspond to the classical Lie derivative when $A = C^\infty(M)$.

Symplectic Structures

Symplectic manifolds are objects of central importance in differential geometry and modern physics. The phase space of every manifold is canonically a symplectic manifold, and modern Hamiltonian mechanics is formulated in terms of symplectic manifolds.

If M is a manifold and d is its exterior differentiation operator, then a form $\omega \in \Omega(M)$ is said to be *closed* if $\omega \in \ker(d)$. A 2-form ω on M is said to be *non-degenerate at $p \in M$* if, when $\omega(X, Y)(p) = 0$, for all $Y \in \mathcal{X}(M)$, it necessarily holds that $X(f)(p) = 0$, for all $f \in C(M)$. A *symplectic manifold* is a pair (M, ω) consisting of a manifold M , and a closed 2-form $\omega \in \Omega^2(M)$ that is nondegenerate at each point.

We say that $\omega \in \underline{\Omega}_{\text{Der}}^2(A)$ is *non-degenerate* if, for all $a \in A$, there is a derivation $\text{Ham}(a) \in \text{Der}(A)$ such that $\omega(X, \text{Ham}(a)) = X(a)$, for all $X \in \text{Der}(A)$. If ω is an ordinary 2-form on M , then it is nondegenerate in this sense, if, and only if, it is nondegenerate at each point. This motivates us to define a *symplectic structure* on A to be closed nondegenerate element of $\underline{\Omega}_{\text{Der}}^2(A)$.

Noncommutative connections

Our next generalisation makes sense for any differential calculus, not just the derivation based ones. However, we feel that its inclusion at this point is appropriate.

If (E, π, M) is a smooth vector bundle over a manifold M , then a *connection* for E is a linear mapping

$$\nabla : \Gamma^\infty(E) \rightarrow \Gamma^\infty(E) \otimes_{C^\infty(M)} \Omega^1(M), \quad (2.7)$$

that satisfies the *Leibniz rule*

$$\nabla(sf) = (\nabla s)f + s \otimes df, \quad s \in \Gamma^\infty(E), f \in C^\infty(M).$$

Equivalently, one can define a *connection* to be a bilinear mapping

$$\nabla : \mathcal{X}(M) \times \Gamma^\infty(E) \rightarrow \Gamma^\infty(E), \quad (X, s) \mapsto \nabla_X s, \quad (2.8)$$

such that $\nabla_X(sf) = sX(f) + (\nabla_X s)f$, and $\nabla_{(fX)}s = f\nabla_X s$.

A connection ∇ in the sense of (2.7) corresponds to the connection in the sense of (2.8) given by

$$\nabla(X, s) = i_X(\nabla s);$$

where, in analogy with (2.6), $i_X(s \otimes \omega) = s\omega(X)$, for $s \in \Gamma^\infty(E)$, $\omega \in \Omega(M)$.

If the vector bundle in question is the tangent bundle, then we see that a connection is a generalisation of the directional derivative of one vector field with respect to another. One important application of connections is that they are used to define curvature tensors for manifolds.

We can use the Serre–Swan Theorem to generalise connections to the noncommutative case. If E is a finitely-generated projective right A -module, and (Ω, d) is a differential calculus over A , then a *connection* for E is a linear mapping

$$\nabla : E \rightarrow E \otimes_A \Omega^1,$$

that satisfies the generalised Leibniz rule

$$\nabla(sa) = (\nabla s)a + s \otimes da, \quad s \in E, a \in A.$$

With a noncommutative generalisation of connections in hand, one can progress to define a noncommutative version of curvature. We shall not pursue this path here, instead, we refer the interested reader to [26].

Chapter 3

Cyclic Cohomology and Quantum Groups

In this chapter we shall motivate and introduce cyclic (co)homology. This is a noncommutative generalisation of de Rham (co)homology due to Alain Connes. In the process of doing so we shall also introduce noncommutative generalisations of volume integrals and de Rham currents. Cyclic (co)homology has been one of Connes' most important achievements. As just one example of its usefulness we cite the spectacular applications it has had to one of the central problems in algebraic topology, the Novikov conjecture; for details see [12].

We shall also introduce compact quantum groups. These objects, which generalise compact topological groups, were developed by S. L. Woronowicz independently of Connes. While the relationship between the two theories is still not very well understood, it seems that there are deep connections between them. One important link that has recently emerged is twisted cyclic cohomology. It is a generalisation of cyclic cohomology, and it was developed in Cork by J. Kustermans, G. J. Murphy, and L. Tuset.

3.1 Cyclic Cohomology

3.1.1 The Chern Character

If X is a compact Hausdorff space, then a *characteristic class* for X is a mapping from $V(X)$, the family of vector bundles over X , to $H^*(X)$, the singular cohomology of X . (We note that this definition of a characteristic class is quite general. Usually a mapping from $V(X)$ to $H^*(X)$ is required to satisfy certain 'natural' conditions before it qualifies as a characteristic class. However, we have

no need to concern ourselves with such details here.) Characteristic classes are of great value in the study of vector bundles and can be used in their classification. Very loosely speaking, characteristic classes measure the extent to which a bundle is ‘twisted’, that is, ‘how far’ it is from being a trivial bundle. As an example of their usefulness, we cite the very important role characteristic classes play in the Atiyah–Singer index theorem. For an accessible introduction to the theory of characteristic classes see [76, 48].

A prominent characteristic class is the Chern character. In general, it is defined to be the unique mapping that satisfies a certain distinguished set of conditions. However, if M is a compact manifold, then its Chern character admits a useful explicit description: If E is a vector bundle, then the *Chern character* $\text{ch} : V(M) \rightarrow H^*(M)$ is defined by setting

$$\text{ch}(E) = \overline{\text{tr}}(\exp(\nabla^2/2\pi i)), \quad E \in V(M); \quad (3.1)$$

where ∇ is a connection for E . The value of $\text{ch}(E)$ can be shown to be independent of the choice of connection.

To give the reader a little feeling for the Chern character it is worth our while to take some time to carefully present equation (3.1). For some connection ∇ for E , the extension

$$\nabla : \Gamma^\infty(E) \otimes_{C^\infty(M)} \Omega(M) \rightarrow \Gamma^\infty(E) \otimes_{C^\infty(M)} \Omega(M)$$

is defined by setting

$$\nabla(s \otimes \omega) = \nabla(s) \wedge \omega + s \otimes d\omega;$$

and

$$\exp(\nabla^2) : \Gamma^\infty(E) \rightarrow \Gamma^\infty(E) \otimes_{C^\infty(M)} \Omega(M)$$

is defined by setting

$$\exp(\nabla^2) = 1 + \nabla^2 + \nabla^4/2! + \dots + \nabla^{2n}/n!.$$

Clearly, we can consider the restriction of $\exp(\nabla^2)$ to $\Gamma^\infty(E)$ as an element of $\text{End}(\Gamma^\infty(E)) \otimes_{C^\infty(M)} \Omega(M)$. Thus, if we define $\overline{\text{tr}}$ to be the unique mapping

$$\overline{\text{tr}} : \text{End}(\Gamma^\infty(E)) \otimes_{C^\infty(M)} \Omega^{2n}(M) \rightarrow \Omega^{2n}(M),$$

for which $\overline{\text{tr}}(A \otimes \omega) = \text{tr}(A)\omega$, then equation (3.1) is well defined. It can be shown that $\text{ch}(E)$ is always a closed form. Therefore, because of the de Rham theorem, we can consider $\text{ch}(E)$ as an element of $H^*(M)$.

Let X be a compact Hausdorff space and let $V(X)$ denote the set of isomorphism classes of vector bundles over X . We note that with respect to the direct sum

operation, $V(X)$ is an abelian semigroup. Now, for any abelian semigroup S , there is a standard construction for ‘generating’ a group from S . Let $F(S)$ be the free group generated by the elements of S , and let $E(S)$ be the subgroup of $F(S)$ generated by the elements of the form $s+s'-(s\oplus s')$; where \oplus denotes addition in S , and $+$ denotes the addition of $F(S)$. The quotient group $F(S)/E(S)$ is called the *Grothendieck group* of S , and it is denoted by $G(S)$. (It is instructive to note that $G(\mathbf{N}) = \mathbf{Z}$.) For any compact Hausdorff space X , we define $K_0(X) = G(V(X))$. The study of this group is known as *topological K-theory*. It is a very important tool in topology, and was first used by Grothendieck and Atiyah. For examples of its uses see [3]. If X is a compact smooth manifold, then an important point for us to note is that the above construction works equally well for smooth vector bundles. Pleasingly, the group produced turns out to be isomorphic to $K_0(X)$.

It is not too hard to show that the Chern character can be extended to a homomorphism from $K_0(M)$ to $H_{dR}^*(M)$ (see Section 3.1.3 for a definition of $H_{dR}^*(M)$). This extension plays a central role in topological K -theory.

There also exists an algebraic version of K -theory. Let A be an arbitrary associative algebra and define $P[A]$ to be the semigroup of finitely generated projective left A -modules, where the semigroup addition is the module direct sum. Then define $K_0(A)$ to be the Grothendieck group of $P[A]$. (It is more usual to give a definition of K_0 in terms of projections in $M_n(A)$. However, the above definition is equivalent, and it is more suited to our needs.)

Recalling the Serre–Swan Theorem, we see that if M is a compact manifold, then $K_0(C(M)) = K_0(C^\infty(M)) = K_0(M)$. The important thing about the algebraic formulation, however, is that it is well defined for any associative algebra A . The study of the K_0 -groups of algebras is called *algebraic K-theory*. It has become a very important tool in the study of noncommutative C^* -algebras and it has been used to establish a number of important results; see [80] for details.

There also exists a theory that is, in a certain sense, dual to K -theory; it is called K -homology. While K -theory classifies the vector bundles over a space, K -homology classifies the Fredholm modules over a space. (A Fredholm module is an operator theoretic structure, introduced by Atiyah, that axiomatizes the important properties of a special type of differential operator on a manifold called an elliptic operator; for details see [3].) Just like K -theory, K -homology has a straightforward noncommutative generalisation. It has also become a very important tool in the study of noncommutative C^* -algebras. (In recent years, G. Kasparov unified K -theory and K -homology into a single theory called KK -theory.)

The existence of a noncommutative version of K -theory prompts us to ask the following question: Could a noncommutative generalisation of de Rham cohomology

be defined; and if so, could equation (3.1) be generalised to the noncommutative setting? Connes showed that the answer to both these questions is yes. In the late 1970's he was investigating foliated manifolds. (Informally speaking, a foliation is a kind of 'clothing' worn on a manifold, cut from a 'stripy fabric'. On each sufficiently small piece of the manifold, these stripes give the manifold a local product structure.) Canonically associated to a foliated manifold is a noncommutative algebra called the foliation algebra of the manifold. While studying the K -homology of this algebra, Connes happened upon a new cohomology theory, and a means of associating one of its cocycles to each Fredholm module. After investigating this cohomology theory, which he named cyclic cohomology, Connes concluded that it was a noncommutative generalisation of de Rham homology. He also concluded that its associated homology theory was a noncommutative generalisation of de Rham cohomology. As we shall discuss later, Connes then proceeded to define noncommutative generalisations of the Chern mapping.

In this section we shall show how cyclic (co)homology is constructed, we shall examine its relationship with classical de Rham (co)homology, and we shall briefly discuss Connes' noncommutative Chern characters.

3.1.2 Traces and Cycles

While investigating the K -homology of the foliation algebra of foliated manifolds, Connes developed a method for constructing noncommutative differential calculi from Fredholm modules. (Connes' method for constructing calculi from Fredholm modules is much the same as his method for constructing calculi from spectral triples; see Chapter 4). Using operator traces he then constructed a canonical linear functional on the n -forms of this calculus. Connes' functional had properties similar to those of the volume integral of a manifold. By axiomizing these properties he came up with the notion of a graded trace. This new structure is considered to be a noncommutative generalisation of the volume integral. We shall begin this section by introducing graded traces. Then we shall show how they easily propose cyclic cohomology theory.

Graded Traces

If M is a manifold and ω is a form on M , then we say that ω is *non-vanishing* at $p \in M$, if there exists an $f \in C^\infty(M)$ such that $\omega(f)(p) \neq 0$. If M is n -dimensional, then we say that it is *orientable* if there exists an n -form $\omega \in \Omega^n(M)$ that is non-vanishing at each point. When M is orientable it is well known that one can define a complex-valued linear mapping \int on $\Omega^n(M)$ that, in a certain sense, generalises the ordinary n -dimensional volume integral; we call \int the *volume integral* of M .

An important result about volume integrals is Stokes' Theorem; it says that if M is a manifold, then \int vanishes on $d(\Omega^{n-1}(M))$; that is, if $\omega \in \Omega^{n-1}(M)$, then $\int d\omega = 0$.

We shall now generalise volume integrals to the noncommutative setting. Let A be an algebra and let (Ω, d) be an n -dimensional differential calculus over A . If n is a positive integer and \int is a linear functional on Ω^n , then we say that \int is *closed* if $\int d\omega = 0$, for all $\omega \in \Omega^{n-1}$. Since $da_1 da_2 \cdots da_n = d(a_1 da_2 \cdots da_n)$, it must hold that $\int da_1 da_2 \cdots da_n = 0$, for all closed functionals \int . Moreover, if $\omega_p \in \Omega^p$, $\omega_q \in \Omega^q$, and $p + q = n - 1$, then, since $d(\omega_p \omega_q) = d(\omega_p) \omega_q + (-1)^p \omega_p d\omega_q$, it must hold that

$$\int d(\omega_p) \omega_q = (-1)^{p+1} \int \omega_p d(\omega_q). \quad (3.2)$$

An n -dimensional graded trace \int on Ω is a linear functional on Ω^n such that, whenever $p + q = n$, we have that

$$\int \omega_p \omega_q = (-1)^{pq} \int \omega_q \omega_p,$$

for all $\omega_p \in \Omega^p(A)$, $\omega_q \in \Omega^q(A)$. Clearly, closed graded traces generalise volume integrals. A point worth noting is that the generalisation is not strict, that is, for a manifold M , there can exist closed graded traces on $C^\infty(M)$ that are not equal to the volume integral. Another point worth noting is that the definition of a closed graded trace reintroduces a form of graded commutativity.

If Ω is an n -dimensional calculus over A and \int is an closed n -dimensional graded trace on Ω , then the triple (A, Ω, \int) is called an n -dimensional cycle over A . Cycles can be thought of as noncommutative generalisations of orientable manifolds.

Cyclic Cocycles

Let A and (Ω, d) be as above, and let \int be a closed, n -dimensional graded trace on Ω . We can define a complex-valued linear functional on A^{n+1} by

$$\varphi(a_0, a_1, \dots, a_n) = \int a_0 da_1 \cdots da_n. \quad (3.3)$$

From the graded commutativity of \int we see that

$$\int a_0 da_1 da_2 \cdots da_n = (-1)^{n-1} \int da_n a_0 da_1 \cdots da_{n-1}.$$

Equation (3.2) then implies that

$$\int a_0 da_1 da_2 \cdots da_n = (-1)^n \int a_n da_0 da_1 \cdots da_{n-1}.$$

This means that

$$\varphi(a_0, a_1, \dots, a_n) = (-1)^n \varphi(a_n, a_0, \dots, a_{n-1}). \quad (3.4)$$

For every positive integer n , let $C^n(A)$ denote the space of complex-valued multilinear mappings on A^{n+1} . As we shall see later, it is common to consider the sequence of mappings

$$b_n : C^n(A) \rightarrow C^{n+1}(A), \quad \psi \mapsto b(\psi); \quad (3.5)$$

where

$$\begin{aligned} b_n(\psi)(a_0, \dots, a_{n+1}) &= \sum_{i=0}^n (-1)^i \psi(a_0, \dots, a_{i-1}, a_i a_{i+1}, a_{i+2}, \dots, a_{n+1}) \\ &\quad + (-1)^{n+1} \psi(a_{n+1} a_0, a_1, \dots, a_n). \end{aligned} \quad (3.6)$$

We call this sequence, the sequence of *Hochschild coboundary operators*.

If φ is defined as in equation (3.3), then it turns out that $b_n(\varphi) = 0$. To show this we shall need to use the following identity:

$$\sum_{i=1}^n (-1)^i da_1 \cdots d(a_i a_{i+1}) \cdots da_{n+1} = -a_1 da_2 \cdots da_{n+1} + (-1)^n da_1 \cdots da_n(a_{n+1}).$$

(It is not too hard to convince oneself of the validity of this identity. For example, if we take the instructive case of $n = 3$, then we see that

$$\begin{aligned} -d(a_1 a_2) da_3 + da_1 d(a_2 a_3) &= -d(a_1) a_2 da_3 - a_1 da_2 da_3 + da_1 d(a_2) a_3 + d(a_1) a_2 da_3 \\ &= -a_1 da_2 da_3 + d(a_1) a_2 da_3. \end{aligned}$$

With this identity in hand we see that

$$\begin{aligned} b_n(\varphi)(a_0, \dots, a_{n+1}) &= \int a_0 a_1 da_2 \cdots da_{n+1} + \sum_{i=1}^n (-1)^i \int a_0 da_1 \cdots d(a_i a_{i+1}) \cdots da_{n+1} \\ &\quad + (-1)^{n+1} \int a_{n+1} a_0 da_1 \cdots da_n \\ &= \int a_0 a_1 da_2 \cdots da_{n+1} - \int a_0 a_1 da_2 \cdots da_{n+1} \\ &\quad + (-1)^n \int a_0 da_1 \cdots da_n(a_{n+1}) + (-1)^{n+1} \int a_{n+1} a_0 da_1 \cdots da_n. \\ &= 0. \end{aligned}$$

In general, if A is an algebra, and ψ is an $(n + 1)$ -multilinear mapping on A such that ψ satisfies equation (3.4), and $b_n(\psi) = 0$, then we call it an n -dimensional cyclic cocycle. Interestingly, it turns out that every cyclic cocycle arises from a closed graded trace; that is, if ψ is an n -dimensional cyclic cocycle, then there exists an n -dimensional cycle (A, Ω, f) such that

$$\psi(a_0, a_1, \dots, a_n) = \int a_0 da_1 \dots da_n,$$

for all $a_0, a_1, \dots, a_n \in \Omega$. To see this take $\Omega_u(A)$, the universal differential calculus over A , and consider the linear functional

$$\int : \Omega_u^n(A) \rightarrow \mathbf{C}, \quad \omega \mapsto \int \omega;$$

where, if $\omega = da_1 \dots da_n$, then $\int \omega = 0$, and if $\omega = a_0 da_1 \dots da_n$, then $\int a_0 da_1 \dots da_n = \psi(a_0, \dots, a_n)$. By definition, \int is closed. Using the fact that $b_n(\psi) = 0$, and that ψ satisfies equation (3.4), it is not too hard to show that \int is also a graded trace. Hence, if we denote by $\Omega'_u(A)$ the truncation of $\Omega_u(A)$ to an n -dimensional calculus, then $(A, \Omega'_u(A), \int)$ is an n -dimensional cycle from which ψ arises.

It appears that cyclic cocycles are of great importance in noncommutative geometry. Moreover, all the pieces are now in place to define a cohomology theory that is based on them. In the following sections we shall carefully present this theory. We shall begin with an exposition of Hochschild (co)homology, and then progress to cyclic (co)homology. We shall also show why cyclic (co)homology is considered a noncommutative generalisation of de Rham homology.

3.1.3 (Co)Chain Complex (Co)Homology

Before we begin our presentation of Hochschild and cyclic (co)homology it seems wise to recall some basic facts about (co)chain complex (co)homology in general.

A *chain complex* (C_*, d) is a pair consisting of a sequence of modules $C_* = \{C_n\}_{n=0}^\infty$ (all over the same ring), and a sequence of module homomorphisms $d = \{d_n : C_n \rightarrow C_{n-1}\}_{n=1}^\infty$, satisfying $d_n d_{n+1} = 0$. A chain complex is usually represented by a diagram of the form

$$\dots \xrightarrow{d_4} C_3 \xrightarrow{d_3} C_2 \xrightarrow{d_2} C_1 \xrightarrow{d_1} C_0.$$

Each d_n is called a *differential operator*; for ease of notation we shall usually omit the subscript and write d for all differentials. The elements of each C_n are called

n-chains. Those chains that are elements of $Z_n(C_*) = \ker(d_n)$ are called *n-cycles*, and those that are elements of $B_n(C_*) = \text{im}(d_{n+1})$ are called *n-boundaries*. Since the composition of two successive differentials is 0, it holds that $B_n(C_*) \subseteq Z_n(C_*)$. For any positive integer n , we define the *nth-homology group* of a complex (C_*, d) to be the quotient

$$H_n(C_*) = Z_n(C_*)/B_n(C_*),$$

We define the 0th-homology group of C_* to be $C_0/B_0(C_*)$. (Note that the homology groups are actually modules, the use of the term group is traditional.) If $Z_n(C_*) = B_n(C_*)$, for all n , then it is clear that the homology groups of the complex are trivial; we say that such a complex is *exact*.

A cochain complex is a structure that is, in a certain sense, dual to the structure of a chain complex. A *cochain complex* (C^*, d) is a pair consisting of a sequence of modules $C^* = \{C^n\}_{n=0}^\infty$, and a sequence of module homomorphisms $d = \{d^n : C^n \rightarrow C^{n+1}\}_{n=0}^\infty$, satisfying $d^{n+1}d^n = 0$. A chain complex is usually represented by a diagram of the form

$$\dots \xleftarrow{d^3} C^3 \xleftarrow{d^2} C^2 \xleftarrow{d^1} C^1 \xleftarrow{d^0} C^0.$$

Each d^n is called a *boundary operator*; as with differentials, we shall usually omit the superscript. The definitions of *n-cochains*, *n-cocycles*, and *co-boundaries* are analogous to the chain complex definitions. For $n > 0$, the *nth-cohomology groups* are defined in parallel with the homological case; the 0th-group is defined to be $Z^0(C^*)$.

Given a chain complex (C_*, d) , we can canonically associate a cochain complex to it: For all non-negative integers n , let C^n be the linear space of linear functionals on C_n , and let $d^n : C^n \rightarrow C^{n+1}$ be the mapping defined by $d^n(\varphi)(a) = \varphi(d_n(a))$, $\varphi \in C^n$, $a \in C_{n+1}$. Clearly, the sequence of modules and the sequence of homomorphisms form a cochain complex. We denote this complex by (C^n, d) , and we say that it is the cochain complex *dual* to (C_n, d) . This approach is easily amended to produce a chain complex dual to a cochain complex.

If (C_*, d) and (D_*, d') are two chain complexes, then a *chain map* $f : (C_*, d) \rightarrow (D_*, d')$ is a sequence of maps $f_n : C_n \rightarrow D_n$, such that, for all positive integers n , the following diagram commutes

$$\begin{array}{ccc} C_n & \xrightarrow{d} & C_{n-1} \\ f_n \downarrow & & \downarrow f_{n-1} \\ D_n & \xrightarrow{d'} & D_{n-1}. \end{array}$$

It is clear that a chain map brings cycles to cycles and boundaries to boundaries. This implies that it induces a map from $H_n(C_*)$ to $H_n(D_*)$; we denote this map

by $H_n(f)$. A *chain homotopy* between two chain maps $f, g : C \rightarrow D$ is a sequence of homomorphisms $h_n : C_n \rightarrow D_{n+1}$ satisfying

$$d'_{n+1}h_n + h_{n-1}d_n = f_n - g_n;$$

(this is usually written more compactly as $d'h + hd = f - g$). If there exists a chain homotopy between two chain maps, then we say that the two maps are *homotopic*. Note that, for any $a \in Z_n(C)$, $d'_{n+1}(h_n(a)) + h_{n-1}(d_n(a)) \in B_n(D)$. Thus, for any two homotopic maps f and g , $H(f)$ and $H(g)$ are equal. If the identity map on a chain complex (C_*, d) is homotopic to the zero map, then the chain is called *contractible*; the map h is called a *contracting homotopy*. A complex that admits a contracting homotopy is obviously exact.

Clearly, directly analogous results hold for cochain complexes.

Examples

The standard example of a cochain complex is the *de Rham complex* of a manifold M :

$$\cdots \xleftarrow{d} \{0\} \xleftarrow{d} \Omega^n(M) \xleftarrow{d} \cdots \xleftarrow{d} \Omega^1(M) \xleftarrow{d} \Omega^0(M).$$

The cohomology groups of this complex are called the *de Rham cohomology groups* of M ; we denote them by $H^n_{dR}(M)$. They are very important because, as we shall now see, they provide an intimate link between the differential structure of M and its underlying topology.

The typical example of a chain complex is the *simplicial complex* of a manifold M : an n -*simplex* in M is a diffeomorphic image of the standard n -simplex in \mathbf{R}^n ; and the n -cochains of the simplicial complex are the formal sums of n -simplices with complex coefficients. It is easy to use the canonical boundary operator on the standard n -simplices to define an operator from the n -simplices to the $(n - 1)$ -simplices. This operator can then be extended by linearity to a boundary operator on the n -cochains. The homology groups of the simplicial complex are called the *singular homology groups* and are denoted by $H_n(M)$. The *singular cohomology groups* $H^n(M)$ are constructed by duality, as explained above.

In one of the fundamental theorems of differential geometry, de Rham showed that the singular cohomology groups and the de Rham cohomology groups coincide, that is, $H^n_{dR}(M) \cong H^n(M)$, for all $n \geq 0$. This is a remarkable result since singular cohomology is a purely topological object while de Rham cohomology is derived from a differential structure. For example, $H^0_{dR}(M) = \bigoplus_{i=0}^k \mathbf{C}$, where k is the number of connected components of M ; we note that the value of n is a purely topological invariant.

Another fundamental result about de Rham cohomology is Poincaré duality: It states that if M is an n -dimensional oriented manifold, then the k^{th} -cohomology

group of M is isomorphic to the $(n - k)^{\text{th}}$ -homology group of M , for all $k = 0, 1, \dots, n$.

Another example of a cochain complex is

$$\dots \xleftarrow{d_u} \Omega_u^2(A) \xleftarrow{d_u} \Omega_u^1(A) \xleftarrow{d_u} A.$$

This complex is exact. To see this consider the unique sequence of maps $h = \{h_n : \Omega_u^n \rightarrow \Omega_u^{n+1}\}$ for which

$$h(a_0 da_1 \cdots da_n) = 0, \quad \text{and} \quad h(da_0 da_1 \cdots da_n) = a_0 da_1 \cdots da_n.$$

It is easily seen that $dh + hd = 1$. Hence, h is a contracting homotopy for the complex.

Let $[\Omega_u(A), \Omega_u(A)]$ denote the smallest subspace of $\Omega_u(A)$ containing the elements of the form

$$[\omega_p, \omega_q] = \omega_p \omega_q - (-1)^{pq} \omega_q \omega_p,$$

where $\omega_p \in \Omega_u^p(A), \omega_q \in \Omega_u^q(A)$. It was shown in [60] that $[\Omega_u(A), \Omega_u(A)]$ is invariant under the action of d_u . Thus, d_u reduces to an operator on the linear quotient $\Omega_u(A)/[\Omega_u(A), \Omega_u(A)]$. We shall denote the image of $\Omega_u^n(A)$, under the projection onto this quotient, by $\Omega_{dR}^n(A)$. In general, the cohomology of the complex

$$\dots \xleftarrow{d_u} \Omega_u^2(A) \xleftarrow{d_u} \Omega_u^1(A) \xleftarrow{d_u} A.$$

is not trivial, and its cohomology groups have some very useful properties.

3.1.4 Hochschild (Co)Homology

The Hochschild cohomology of associative algebras was introduced by G. Hochschild in [51]. One of his original motivations was to formulate a cohomological criteria for the separability of algebras.

Let A be an algebra. For every non-negative integer n , define $C_n(A)$, the space of *Hochschild n -chains* of A , to be the $(n + 1)$ -fold tensor product of A with itself; that is, define

$$C_n(A) = A^{\otimes(n+1)}. \quad (3.7)$$

We denote the sequence $\{C_n(A)\}_{n=0}^{\infty}$ by $C_*(A)$. For all positive integers n , let b' be the unique sequence of linear maps $b' = \{b' : C_n(A) \rightarrow C_{n-1}(A)\}_{n=0}^{\infty}$ for which

$$b'(a_0 \otimes \dots \otimes a_n) = \sum_{i=0}^{n-1} (-1)^i (a_0 \otimes \dots \otimes a_i a_{i+1} \otimes \dots \otimes a_n). \quad (3.8)$$

Due to a simple cancellation of terms it holds that $b'^2 = b' \circ b' = 0$. Hence, $(C_*(A), b')$ is well defined as a complex; it is known as the *bar complex* of A . If A is unital, then the bar complex is exact. To see this consider the unique sequence of linear maps $s = \{s : C_n(A) \rightarrow C_{n+1}(A)\}_{n=0}^\infty$, for which

$$s(a_0 \otimes \dots \otimes a_n) = 1 \otimes a_0 \otimes \dots \otimes a_n.$$

It is not too hard to show that $b's + sb' = 1$. Hence, s is a contracting homotopy for the complex. It has been noted by Wodzicki [108] that non-unital algebras whose bar complex is exact have useful properties. Algebras with this property are now called *H-unital* (homologically unital).

Define b to be the unique sequence of linear maps $b = \{b : C_n(A) \rightarrow C_{n-1}(A)\}_{n=0}^\infty$ for which

$$b(a_0 \otimes \dots \otimes a_{n-1} \otimes a_n) = b'(a_0 \otimes \dots \otimes a_{n-1} \otimes a_n) + (-1)^n (a_n a_0 \otimes \dots \otimes a_{n-1}). \quad (3.9)$$

We call b the *Hochschild boundary operator*. It is easy to conclude from the fact that $b'^2 = 0$, that $b^2 = 0$. Hence, $(C_*(A), b)$ is a complex; we call it the *Hochschild chain complex* of A . The n^{th} -homology group of the Hochschild chain complex is called the n^{th} -*Hochschild homology group* of A ; it is denoted by $HH_n(A)$. (There also exists a more general formulation of the Hochschild chain complex in terms of a general A -bimodule M ; when $M = A^*$ the two formulations coincide.)

Note that the image of b in A is $[A, A]$, the *commutator subalgebra* of A ; that is, the smallest subalgebra containing all the elements of the form $[a_0, a_1] = a_0 a_1 - a_1 a_0$. Thus, $HH_0(A) = A/[A, A]$. If A is commutative, then $HH_0(A) = A$.

To gain some familiarity with the definition, we shall calculate the Hochschild cohomology groups of \mathbf{C} . Firstly, we note that $\mathbf{C}^{\otimes n+1} \simeq \mathbf{C}$. Thus, b reduces to 0 if n is odd, and to the identity if n is even. This means that $HH_0(\mathbf{C}) = \mathbf{C}$, (confirming that $HH_0(A) = A$ when A is commutative), and $HH_n(\mathbf{C}) = 0$, if $n \geq 1$.

As would be expected, the cochain complex dual to $(C_*(A), b)$ is called the *Hochschild cochain complex*; it is denoted by $(C^*(A), b)$. The n^{th} -cohomology group of $(C^*(A), b)$ is called the n^{th} -Hochschild cohomology group; it is denoted by $HH^n(A)$. Note that b , the sequence of Hochschild coboundary operators, is the same as the sequence defined by equation (3.6).

Just like $HH_0(A)$, $HH^0(A)$ has a simple formulation. Recall that, by definition, $HH^0(A)$ is the set of 0-cocycles of the complex. Now, $\tau \in Z^0(C^*(A))$ if, and only if, $\tau(a_0 a_1 - a_1 a_0) = 0$, for all $a_0, a_1 \in A$. Thus, $HH^0(A)$ is the space of traces on A .

Just as in the homological case, the 0^{th} -cohomology group of the Hochschild complex of \mathbf{C} is \mathbf{C} , while all the others groups are trivial.

Projective Resolutions

In general, direct calculations of Hochschild homology groups can be quite complicated. Fortunately, however, there exists an easier approach involving projective resolutions. If \mathcal{E} is an R -module, then a *projective resolution* of \mathcal{E} is an exact complex

$$\cdots \xrightarrow{b_3} P_2 \xrightarrow{b_2} P_1 \xrightarrow{b_1} P_0 \xrightarrow{b_0} \mathcal{E},$$

where each P_i is a projective R -module. Let A be a unital algebra. Define A^{op} , the *opposite algebra* of A , to be the algebra that is isomorphic to A as a linear space by some isomorphism $a \mapsto a^o$, and whose multiplication is given by $a^o b^o = (ba)^o$. Let us denote $B = A \otimes A^{\text{op}}$, and let us give A the structure of a B -module by setting $(a \otimes b^o)c = acb$. If (P_*, b) is a projective resolution of A over B , then

$$\cdots \xrightarrow{1 \otimes b} A \otimes_B P_3 \xrightarrow{1 \otimes b} A \otimes_B P_2 \xrightarrow{1 \otimes b} A \otimes_B P_1 \xrightarrow{1 \otimes b} A \otimes_B P_0$$

is clearly a complex. Using chain homotopies it can be shown that the homology of this complex is equal to the Hochschild homology of A . (Those readers familiar with the theory of derived functors will see that this result can be more succinctly expressed as $HH_n(A) = \text{Tor}_n^B(A, A)$.) Thus, by making a judicious choice of projective resolution of A , the job of calculating its Hochschild homology groups can be considerably simplified. A similar result holds for Hochschild cohomology.

Hochschild–Konstant–Rosenberg Theorem

The following theorem is a very important result about Hochschild homology that comes from algebraic geometry. It inspired Connes to prove Theorem 3.1.2, which can be regarded as the analogue of this result for manifolds. (We inform the reader unacquainted with algebraic geometry that any introductory book on the subject will explain any of the terms below that are unfamiliar.)

Theorem 3.1.1 (Hochschild–Konstant–Rosenberg) *Let Y be a smooth complex algebraic variety, let $\mathcal{O}[Y]$ be the ring of regular functions on Y , and let $\Omega^q(Y)$ be the space of algebraic q -forms on Y . Then there exists a map χ , known as the Hochschild–Konstant–Rosenberg map, that induces an isomorphism*

$$\chi : HH_q(\mathcal{O}[Y]) \rightarrow \Omega^q(Y), \quad \text{for all } q \geq 0.$$

Continuous Hochschild (Co)Homology

The Hochschild cohomology of an algebra is a purely algebraic object, that is, it does not depend on any topology that could possibly be defined on the al-

gebra. However, there exists a version of Hochschild cohomology, called continuous Hochschild cohomology, that does take topology into account. For applications to noncommutative geometry it is often crucial that we consider topological algebras. For example, while the Hochschild cohomology groups of the algebra of smooth functions on a manifold are unknown, their continuous version have been calculated by Connes, as we shall see below.

The continuous Hochschild homology is defined for complete locally convex algebras and not for general topological algebras: A *locally convex algebra* A is an algebra endowed with a locally convex Hausdorff topology for which the multiplication $A \times A \rightarrow A$ is continuous. (The use here of the word *complete* requires some clarification. A net $\{x_\lambda\}_{\lambda \in D}$ in a locally convex vector space with a Hausdorff topology is called a *Cauchy net* if, for every open set U containing the origin, there exists a $\Lambda \in D$, such that, for $\kappa, \lambda \geq \Lambda$, $x_\kappa - x_\lambda \in U$. The space is *complete* if every Cauchy net is convergent.)

To define the continuous Hochschild homology we shall need to make use of Grothendieck's projective tensor product of two locally convex vector spaces [41] which is denoted by \otimes_π . We shall not worry too much about the details of the product, but we shall make two important points about it: Firstly, the projective tensor product of two locally convex vector spaces (with Hausdorff topologies) is a complete locally convex vector space (with a Hausdorff topology); and secondly, the projective tensor product satisfies a universal property, for jointly continuous bilinear maps, that is analogous to the universal property that the algebraic tensor product satisfies for bilinear maps.

If A is a locally convex algebra then we define $C_n^{\text{cts}}(A)$, the space of *continuous Hochschild n -chains*, to be the completion of the $(n+1)$ -fold projective tensor product of A with itself. Because of the universal property of \otimes_π , and the joint continuity of multiplication, each Hochschild boundary operator b has a unique extension to a continuous linear map on the continuous Hochschild n -chains. The n^{th} -homology group of the complex (HH_n^{cts}, b) is called the *continuous Hochschild n^{th} -homology group* of A ; it is denoted by $HH_n^{\text{cts}}(A)$.

We define $C_{\text{cts}}^n(A)$, the space of *continuous Hochschild n -cochains* of A , to be the continuous dual of $C_n^{\text{cts}}(A)$. Note that the Hochschild differentials canonically induce a unique sequence of linear operators on the continuous cochains. The n^{th} -cohomology group of the resulting complex is called the *continuous Hochschild n^{th} -cohomology group* of A ; it is denoted by $HH_{\text{cts}}^n(A)$.

We saw that the Hochschild (co)homology groups of an algebra could be computed using projective resolutions. A similar result holds true for continuous Hochschild (co)homology. However, in the continuous version one must use topological projective resolutions, for details see [12] and references therein.

Continuous Hochschild Cohomology of $C^\infty(M)$

The algebra of smooth functions on a manifold M can be canonically endowed with a complete, locally convex topological algebra structure (the seminorms are defined using local partial derivatives). With a view to constructing its continuous Hochschild cohomology groups, Connes constructed a topological projective resolution of $C^\infty(M)$ over $C^\infty(M) \otimes_\pi C^\infty(M)$ [11]. Each of the modules of the resolution arose as the module of sections of a vector bundle over $M \times M$. (In fact, each of the modules arose as the module of sections of the vector bundle pullback of $\bigwedge^n(M)$, for some positive integer n , through the map $M \times M \rightarrow M$, $(p, q) \mapsto q$.) Connes then used the aforementioned result relating Hochschild cohomology and projective resolutions (or more correctly its continuous version) to conclude the following theorem. It identifies the continuous Hochschild n^{th} -homology group of $C^\infty(M)$ with \mathcal{D}_n , the continuous dual of $H_{dR}^n(M)$. The elements of \mathcal{D}_n are called the n -dimensional de Rham currents.

Theorem 3.1.2 *Let M be a smooth compact manifold. Then the map*

$$HH_{cts}^n(C^\infty(M)) \rightarrow \mathcal{D}_n(M), \quad \varphi \mapsto C_\varphi,$$

where

$$C_\varphi(f_0 df_1 \cdots df_n) = \frac{1}{k!} \sum_{\pi \in \text{Perm}(p+q)} (-1)^{\text{sgn}(\pi)} \varphi(f_0, f_1, \dots, f_n),$$

is an isomorphism.

As a corollary to this result it can be shown that

$$HH_n^{cts}(C^\infty(M)) \simeq \Omega^n(M).$$

Thus, Hochschild homology groups generalise differential forms.

3.1.5 Cyclic (Co)Homology

Let A be an algebra, and let $C^n(A)$ and b be as above. Define $\lambda = \{\lambda : C^n(A) \rightarrow C^n(A)\}_{n=0}^\infty$ to be the unique sequence of linear operators for which

$$\lambda(\varphi)(a_0 \otimes \cdots \otimes a_{n-1} \otimes a_n) = (-1)^n \varphi(a_n \otimes a_0 \otimes \cdots \otimes a_{n-1}), \quad \varphi \in C^n(A).$$

If $\varphi \in C^n(A)$, and $\lambda(\varphi) = \varphi$, then we call φ a *cyclic n -cochain*. We denote the subspace of cyclic n -cochains by $C_\lambda^n(A)$.

Let b' be the operator defined in equation (3.8). A direct calculation will show that $(1 - \lambda)b = b'(1 - \lambda)$. Thus, since $C_\lambda^n(A) = \ker(1 - \lambda)$, we must have that $b(C_\lambda^n(A)) \subseteq C_\lambda^{n+1}(A)$. Thus, if we denote $C_\lambda^*(A) = \{C_\lambda^n(A)\}_{n=0}^\infty$, then the pair $(C_\lambda^*(A), b)$ is a subcomplex of the Hochschild cochain complex; we call it the *cyclic complex* of A . The n^{th} -cohomology group of $(C_\lambda^*(A), b)$ is called the n^{th} -*cyclic cohomology group* of A ; we denote it by $HC^n(A)$ (HC stands for homologie cyclique). Note that each cocycle of the cyclic complex is a cyclic cocycle, as defined in Section 3.1.2.

To get a little feeling for these definitions, we shall calculate the cyclic cohomology groups of \mathbf{C} : If f is a non-zero Hochschild n -cochain of \mathbf{C} , then it is clear that $\lambda(f) = f$ if, and only if, n is even. Thus, $C_\lambda^{2n+1}(\mathbf{C}) = \{0\}$, and $C_\lambda^{2n}(\mathbf{C}) = \mathbf{C}$. This means that the cyclic complex of \mathbf{C} is

$$\cdots \xleftarrow{0} \mathbf{C} \xleftarrow{\text{id}} 0 \xleftarrow{0} \mathbf{C} \xleftarrow{\text{id}} 0,$$

and

$$HC^{2n}(\mathbf{C}) = \mathbf{C}, \quad \text{and} \quad HC^{2n+1}(\mathbf{C}) = 0.$$

Cyclic Homology

As we stated earlier, there is also a cyclic homology theory. Define $\lambda = \{\lambda : C_n(A) \rightarrow C_n(A)\}_{n=0}^\infty$ to be the unique sequence of linear maps for which

$$\lambda(a_0 \otimes \cdots \otimes a_n) = (-1)^n(a_n \otimes a_0 \otimes \cdots \otimes a_{n-1});$$

and define the *space of cyclic n -chains* to be

$$C_n^\lambda(A) = C_n(A)/\text{im}(1 - \lambda).$$

Using an argument analogous to the cohomological case, it can be shown that $b(C_n^\lambda(A)) \subseteq C_{n-1}^\lambda(A)$. Thus, if we denote $C_*^\lambda(A) = \{C_n^\lambda(A)\}_{n=0}^\infty$, then the pair $(C_*^\lambda(A), b)$ is a subcomplex of the Hochschild chain complex; we call it the *cyclic complex* of A . The n^{th} -homology group of this complex is called the n^{th} -*cyclic homology group* of A ; it is denoted by $HC_n(A)$.

The theory of cyclic homology was developed after the cohomological version. In [59], Loday and Quillen related it to the Lie algebra homology of matrices over a ring.

Connes' Long Exact Sequence

If $I : C_\lambda^n(A) \rightarrow C^n(A)$ is the inclusion map of the space of cyclic n -cochains into the space of Hochschild n -cochains, and $\pi : C^n(A) \rightarrow C^n(A)/C_\lambda^n(A)$ denotes the

projection map, then the sequence

$$0 \longrightarrow C_\lambda^m(A) \xrightarrow{I} C^m(A) \xrightarrow{\pi} C^m(A)/C_\lambda^m(A) \longrightarrow 0$$

is exact. By a standard result of homological algebra, this sequence induces a long exact sequence of cohomology groups

$$\dots HC^m(A) \xrightarrow{I} HH^n(A) \xrightarrow{B^n} HC^{n-1}(A) \xrightarrow{S^n} HC^{n+1}(A) \xrightarrow{I} \dots$$

The operators S_n are called the *periodicity maps*, the operators B_n are called the *connecting homomorphisms*, and the sequence is known as *Connes' long exact sequence*, or the *SBI-sequence*. The *SBI* sequence is of great importance in the calculation of cyclic cohomology groups because it ties cyclic cohomology up with the tools of homological algebra available to calculate Hochschild cohomology from projective resolutions.

There also exists a more efficient approach to calculating cyclic cohomology groups that is based on the cyclic category (this is an abelian category introduced by Connes). Connes used this new structure to show that cyclic cohomology can be realised as a derived functor. For details see [12].

The periodicity maps $S_n : HC^{2n}(A) \rightarrow HC^{n+2}(A)$ define two directed systems of abelian groups. Their inductive limits

$$HP^0(A) = \varinjlim HC^{2n}(A), \quad HP^1(A) = \varinjlim HC^{2k+1}(A),$$

are called the *periodic cyclic cohomology groups* of A . (See [80] for details on directed systems and inductive limits.)

For cyclic homology there also exists a version of Connes' long exact sequence,

$$\dots HC_n(A) \xleftarrow{I} HH_n(A) \xleftarrow{B_n} HC_{n-1}(A) \xleftarrow{S_n} HC_{n+1}(A) \xleftarrow{I} \dots$$

Moreover, there exists two *periodic cyclic homology* groups $HP_0(A)$ and $HP_1(A)$, whose definitions are, in a certain sense, analogous to the cohomological case. (We pass over the exact nature of these definitions since their presentation would require an excessive digression. Details can be found in [12].)

Continuous Cyclic (Co)Homology

Beginning with the definition of continuous Hochschild cohomology, it is relatively straightforward to formulate continuous versions of cyclic, and periodic cyclic (co)homology for locally convex algebras. Building on the proof of Theorem 3.1.2, Connes established the following result [12].

Theorem 3.1.3 *If M is an n -dimensional manifold, then*

$$HP_0^{\text{cts}}(C^\infty(M)) \oplus HP_1^{\text{cts}}(C^\infty(M)) \simeq \bigoplus_{i=0}^n H_{dR}^i,$$

and

$$HP_{\text{cts}}^0(C^\infty(M)) \oplus HP_{\text{cts}}^1(C^\infty(M)) \simeq \bigoplus_{i=0}^n H_i,$$

This result justifies thinking of periodic cyclic homology as a noncommutative generalisation of de Rham cohomology, and thinking of cyclic cohomology as a noncommutative generalisation of singular homology.

3.1.6 The Chern–Connes Character Maps

As we discussed earlier, Connes discovered a method for constructing differential calculi from Fredholm modules while he was studying foliated manifolds. Moreover, he also discovered a way to associate a cyclic cocycle to each Fredholm module. In fact, this cocycle was none other than the cocycle corresponding to the module’s calculus and Connes’ trace functional. For any algebra A , this association induces a mapping

$$Ch^0 : K^0(A) \rightarrow HP^0(A);$$

it is called the *Chern–Connes character mapping* for K^0 . We should consider it as a ‘dual’ noncommutative Chern character.

Let us now return to the question we posed much earlier: Does there exist a noncommutative Chern character that maps $K_0(A)$ to $HP_0(A)$? By generalising equation (3.1) directly to the noncommutative case Connes produced just such a mapping

$$Ch_0 : K_0(A) \rightarrow HP_0(A);$$

it is called the *Chern–Connes character map* for K_0 . The existence of Ch_0 is the second principle reason why periodic cyclic homology is considered to be a noncommutative generalisation of de Rham cohomology. (We inform readers familiar with the K_1 and K^1 groups that, for any algebra A , there also exist Chern–Connes maps $Ch_1 : K_1(A) \rightarrow HP_1(A)$ and $Ch^1 : K^1(A) \rightarrow HP^1(A)$.)

Connes used these powerful new tools to great effect in study of foliated manifolds. He then went on to apply them to a number of other prominent problems in mathematics. The Chern maps are tools of fundamental importance in Connes’ work. Some mathematicians have even described them as the “backbone of noncommutative geometry”. For an in-depth discussion of the Connes–Chern maps see [12] and references therein; for a more accessible presentation see [7].

3.2 Compact Quantum Groups

A *character* on a locally compact group G is a continuous group homomorphism from G to the circle group T . If G is abelian, then set of all characters on G is itself canonically a locally compact abelian group; we call it the *dual group* of G and denote it by \widehat{G} . The following result was established in 1934 by Pontryagin; for a proof see [97].

Theorem 3.2.1 (Pontryagin duality) *If G is a locally compact abelian group, then the dual of \widehat{G} is isomorphic to G .*

Unfortunately, the dual of a non-abelian group is not itself a group. Thus, in order to extend Pontryagin's result one would have to work in a larger category that included both groups and their duals. The theory of quantum groups has its origin in this idea. (The term *quantum group* is only loosely defined. It is usually taken to mean some type of noncommutative generalisation of a locally compact topological group.) G. I. Kac, M. Takesaki, M. Enock and J.-M. Schwartz did pioneering work in this direction, see [31] and references therein. One of the more important structures to emerge during this period was Kac algebras, a theory based on von Neumann algebras. However, it became apparent in the 1980s that the then current theories did not encompass new and important examples such as S. L. Woronowicz's $SU_q(2)$ [110] and others found in V. G. Drinfeld's work [29]. In response a number of new formulations emerged. Guided by the example of $SU_q(2)$, Woronowicz developed a new theory based on C^* -algebras that he called compact quantum group theory. It is Woronowicz's approach that we shall follow here. Woronowicz later went on to define a more general structure called an algebraic quantum group. In this setting an extended version of Pontryagin duality was established, for further details see [106]. Recently, another approach called locally compact quantum group theory has also emerged. It generalises a large number of previous formulations including Kac algebras, compact quantum groups, and algebraic quantum groups, for details see [63].

3.2.1 Compact Quantum Groups

Recall that a *topological group* G is a group endowed with a Hausdorff topology with respect to which the group multiplication $(x, y) \mapsto xy$ and the inverse map $x \mapsto x^{-1}$ are continuous. If G is compact as a topological space, then we say that G is a *compact group*. *Topological*, and *compact semigroups* are defined similarly. Theorem 1.1.4 tells us that we can recover G , as a topological space, from $C(G)$. However, it is clear that the semigroup structure of G cannot be recovered; two compact semigroups can be homeomorphic as topological spaces

without being homomorphic as semigroups. With a view to constructing a non-commutative generalisation of compact groups, we shall attempt below to express the group structure of G in terms of $C(G)$. However, before we do this we shall need to present some facts about C^* -algebra tensor products.

If H and K are two Hilbert spaces, then it is well known that there is a unique inner product $\langle \cdot, \cdot \rangle$ on their algebraic tensor product $H \otimes K$ such that

$$\langle x \otimes y, x' \otimes y' \rangle = \langle x, x' \rangle \langle y, y' \rangle, \quad x, x' \in H, \quad y, y' \in K.$$

The Hilbert space completion of $H \otimes K$ with respect to the induced norm is called the *Hilbert space tensor product* of H and K ; we shall denote it by $H \overline{\otimes} K$. If $S \in B(H)$ and $T \in B(K)$, then it can be shown that there exists a unique operator $S \overline{\otimes} T \in B(H \overline{\otimes} K)$ such that

$$(S \overline{\otimes} T)(x \otimes y) = S(x) \otimes T(y), \quad x \in H, y \in K.$$

Moreover, it can also be shown that

$$\|S \overline{\otimes} T\| = \|S\| \|T\|. \quad (3.10)$$

If A and B are two $*$ -algebras and $A \otimes B$ is their algebra tensor product, then it is well known that one can define an algebra involution $*$ on $A \otimes B$, such that, for $a \otimes b \in A \otimes B$, $(a \otimes b)^* = a^* \otimes b^*$. This involution makes $A \otimes B$ into a $*$ -algebra called the *$*$ -algebra tensor product* of A and B . If \mathcal{A} and \mathcal{B} are two C^* -algebras and (U, H) and (V, K) are their respective universal representations, then it can be shown that there exists a unique injective $*$ -algebra homomorphism $W : \mathcal{A} \otimes \mathcal{B} \rightarrow B(H \overline{\otimes} K)$, such that $W(a \otimes b) = U(a) \overline{\otimes} V(b)$, for $a \in \mathcal{A}$, $b \in \mathcal{B}$. We can use W to define a norm $\|\cdot\|_*$ on $\mathcal{A} \otimes \mathcal{B}$ by setting

$$\|c\|_* = \|W(c)\|, \quad c \in \mathcal{A} \otimes \mathcal{B}.$$

We call it the *spatial C^* -norm*. Clearly, $\|cc^*\|_* = \|c\|_*^2$. Hence, the C^* -algebra completion of $\mathcal{A} \otimes \mathcal{B}$ with respect to $\|\cdot\|_*$ exists. We call this C^* -algebra the *spatial tensor product* of \mathcal{A} and \mathcal{B} , and we denote it by $\widehat{\mathcal{A} \otimes \mathcal{B}}$. If we recall that U and V are isometries, then it is easy to see that equation (3.10) implies that

$$\|a \otimes b\|_* = \|a\| \|b\|, \quad a \in \mathcal{A}, \quad b \in \mathcal{B}. \quad (3.11)$$

Now, if $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$ are C^* -algebras, and $\varphi : \mathcal{A} \rightarrow \mathcal{C}$ and $\psi : \mathcal{B} \rightarrow \mathcal{D}$ are two $*$ -algebra homomorphisms, then a careful reading of Chapter 6 of [80] will verify that $\varphi \otimes \psi : \mathcal{A} \otimes \mathcal{B} \rightarrow \mathcal{C} \otimes \mathcal{D}$ is continuous with respect to $\|\cdot\|_*$. Hence, it has a unique extension to a continuous linear mapping from $\widehat{\mathcal{A} \otimes \mathcal{B}}$ to $\widehat{\mathcal{C} \otimes \mathcal{D}}$; we denote this mapping by $\widehat{\varphi \otimes \psi}$.

Let G be a compact semigroup, and consider the injective $*$ -algebra homomorphism $\pi : C(G) \otimes C(G) \rightarrow C(G \times G)$ determined by $\pi(f \otimes g)(s, t) = f(s)g(t)$, for $f, g \in C(G)$, $s, t \in G$. Since equation (3.11) implies that $\|f \otimes g\|_* = \|\pi(f \otimes g)\|_\infty$, it is easy to use the Stone–Weierstrass theorem to show that π has a unique extension to a $*$ -isomorphism from $C(G) \widehat{\otimes} C(G) \rightarrow C(G \times G)$. We shall, from here on, tacitly identify these two C^* -algebras.

Consider the $*$ -algebra homomorphism

$$\Delta : C(G) \rightarrow C(G \times G), \quad f \mapsto \Delta(f),$$

where

$$\Delta(f)(r, s) = f(rs), \quad r, s \in G;$$

we call it the *composition with the multiplication* of G . As a straightforward examination will verify,

$$[(\text{id} \widehat{\otimes} \Delta)\Delta(f)](r, s, t) = f(r(st)), \quad \text{and} \quad [(\Delta \widehat{\otimes} \text{id})\Delta(f)](r, s, t) = f((rs)t),$$

for all $r, s, t \in G$. Thus, since $C(G)$ separates the points of G , the associativity of its multiplication is equivalent to the equation

$$(\text{id} \widehat{\otimes} \Delta)\Delta = (\Delta \widehat{\otimes} \text{id})\Delta. \tag{3.12}$$

This motivates the following definition: A pair (\mathcal{A}, Δ) is a *compact quantum semigroup* if \mathcal{A} is a unital C^* -algebra and Δ is a *comultiplication* on \mathcal{A} , that is, if Δ is a unital $*$ -algebra homomorphism from \mathcal{A} to $\mathcal{A} \widehat{\otimes} \mathcal{A}$ that satisfies equation (3.12). We say that $(C(G), \Delta)$ is the *classical* compact quantum semigroup associated to G .

Let (\mathcal{A}, Δ) be an *abelian* compact quantum semigroup, that is, a compact quantum semigroup whose C^* -algebra is abelian. By the Gelfand–Naimark Theorem $\mathcal{A} = C(G)$, where G is the compact Hausdorff space $\Omega(\mathcal{A})$. Thus, we can regard Δ as a mapping from $C(G)$ to $C(G \times G)$. By modifying the argument of the second generalisation in Section 1.2.1, one can show that there exists a unique continuous mapping $m : G \times G \rightarrow G$, such that $\Delta(f) = f \circ m$, for $f \in C(G)$. The fact the Δ is a comultiplication implies that m is associative. Thus, G is a semigroup and (\mathcal{A}, Δ) is the compact quantum semigroup associated to it. All this means that the abelian compact quantum semigroups are in one-to-one correspondence with the compact semigroups.

It is natural to ask whether or not we can identify a ‘natural’ subfamily of the family of compact quantum semigroups whose abelian members are in one-to-one

correspondence with the compact groups. Pleasingly, it turns out that we can. Let \mathcal{A} be a C^* -algebra and consider the two subsets of $\widehat{\mathcal{A}} \widehat{\otimes} \mathcal{A}$ given by

$$\Delta(\mathcal{A})(1 \widehat{\otimes} \mathcal{A}) = \{\Delta(a)(1 \widehat{\otimes} b) : a, b \in \mathcal{A}\}, \quad (3.13)$$

and

$$\Delta(\mathcal{A})(\mathcal{A} \widehat{\otimes} 1) = \{\Delta(a)(b \widehat{\otimes} 1) : a, b \in \mathcal{A}\}. \quad (3.14)$$

When $\mathcal{A} = C(G)$, for some compact semigroup G , then the linear spans of (3.13) and (3.14) are both dense in $C(G) \widehat{\otimes} C(G)$ if, and only if, G is a group. With a view to showing this consider T the unique automorphism of $C(G) \widehat{\otimes} C(G)$ for which $T(f)(s, t) = f(st, t)$, for $f \in C(G \times G)$, $s, t \in G$. If we assume that G is a group, then this map is invertible; its inverse T^{-1} is the unique endomorphism for which $T^{-1}(f)(s, t) = f(st^{-1}, t)$. Since T is continuous, $T(C(G) \otimes C(G))$ is dense in $C(G) \widehat{\otimes} C(G)$. If we now note that

$$T(g \otimes h)(s, t) = g \otimes h(st, t) = g(st)h(t) = [\Delta(g)(1 \otimes h)](s, t),$$

for all $g, h \in C(G)$, then we see that the linear span of $\Delta(C(G))(1 \widehat{\otimes} C(G))$ is indeed dense in $C(G) \widehat{\otimes} C(G)$. A similar argument will establish that $\Delta(C(G))(C(G) \widehat{\otimes} 1)$ is dense in $C(G) \widehat{\otimes} C(G)$. Conversely, suppose that the linear spans of (3.13) and (3.14) are each dense in $C(G) \widehat{\otimes} C(G)$, and let $s_1, s_2, t \in G$ such that $s_1 t = s_2 t$. As a little thought will verify, the points (s_1, t) and (s_2, t) are not separated by the elements of the linear span of $\Delta(C(G))(1 \widehat{\otimes} C(G))$. Thus, the density of the latter in $C(G) \widehat{\otimes} C(G)$ implies that $s_1 = s_2$. Similarly, the density of $\Delta(C(G))(C(G) \widehat{\otimes} 1)$ in $C(G) \widehat{\otimes} C(G)$ implies that if $st_1 = st_2$, then $t_1 = t_2$. Now, it is well known that a compact semigroup that satisfies the left- and right-cancellation laws is a compact group. Consequently, G is a compact group. This motivated Woronowicz [111] to make the following definition.

Definition 3.2.2. A compact quantum semigroup (\mathcal{A}, Δ) is said to be a *compact quantum group* if the linear spans of $(1 \widehat{\otimes} \mathcal{A})\Delta(\mathcal{A})$ and $(\mathcal{A} \widehat{\otimes} 1)\Delta(\mathcal{A})$ are each dense in $\widehat{\mathcal{A}} \widehat{\otimes} \mathcal{A}$.

Motivated by the classical case, the density conditions on the linear spans of $(1 \widehat{\otimes} \mathcal{A})\Delta(\mathcal{A})$ and $(\mathcal{A} \widehat{\otimes} 1)\Delta(\mathcal{A})$ are sometimes called *Woronowicz's left- and right-cancellation laws* respectively.

The Haar State

Let G be a locally compact topological group, and let μ be a non-zero regular Borel measure on G . We call μ a *left Haar measure* if it is *invariant under left translation*, that is, if $\mu(gB) = \mu(B)$, for all $g \in G$, and for all Borel subsets B . A

right Haar measure is defined similarly. It is well known that every locally compact topological group G admits a left and a right Haar measure that are unique up to positive scalar multiples. If the left and right Haar measures coincide, then G is said to be *unimodular*. It is a standard result that every compact group is unimodular. Thus, since we shall only consider compact groups here, we shall speak of *the* Haar measure. Furthermore, the regularity of μ ensures that $\mu(G)$ is finite. This allows us to work with the *normalised* Haar measure, that is, the measure μ for which $\mu(G) = 1$. We define the *Haar integral* to be the integral over G with respect to μ . It is easily seen that the left- and right-invariance of the measure imply *left-* and *right-invariance* of the integral, that is,

$$\int_G f(hg)dg = \int_G f(gh)dg = \int_G f(g)dg, \quad (3.15)$$

for all $h \in G$, and for all integrable functions f . (Note that since there is no risk of confusion, we have suppressed explicit reference to μ .)

Considered as a linear mapping on $C(G)$, \int is easily seen to be a positive linear mapping of norm one. In general, if φ is a positive linear mapping of norm one on a C^* -algebra, then we call it a *state*. Now, as direct calculation will verify,

$$(\text{id} \widehat{\otimes} \int_G) \Delta(f)(h) = \int f(hg)dg,$$

and

$$\left(\int_G \widehat{\otimes} \text{id} \right) \Delta(f)(h) = \int f(gh)dg.$$

Thus, the left- and right-invariance of the integral is equivalent to the equation

$$(\text{id} \widehat{\otimes} \int_G) \Delta(f) = \left(\int_G \widehat{\otimes} \text{id} \right) \Delta(f) = \int_G f d\mu.$$

This motivates the following definition: Let (\mathcal{A}, Δ) be a compact quantum group and let h be a state on \mathcal{A} . We call h a *Haar state* on \mathcal{A} if

$$(\text{id} \widehat{\otimes} h) \Delta(a) = (h \widehat{\otimes} \text{id}) \Delta(a) = h(a),$$

for all $a \in \mathcal{A}$.

The following result is of central importance in the theory of compact quantum groups. It was first established by Woronowicz under the assumption that \mathcal{A} was separable [112], and it was later proved in the general case by Van Daele [105].

Theorem 3.2.3 *If (\mathcal{A}, Δ) is a compact quantum group, then there exists a unique Haar state on \mathcal{A} .*

A state φ on a C^* -algebra \mathcal{A} is called *faithful* if $\ker(\varphi) = \{0\}$. Obviously, the Haar state of any classical compact quantum group is faithful. However, this does not carry over to the noncommutative setting. It is a highly desirable property that h be faithful. Necessary and sufficient conditions for this to happen are given in [82].

Hopf Algebras

In the definition of a compact quantum group we find no mention of a generalised identity nor any generalisation of the inverse of an element. This might cause one to suspect that identities and inverses have no important quantum analogues. However, this is certainly not the case.

Consider the *classical co-unit* $\varepsilon : C(G) \rightarrow \mathbf{C}$, defined by $\varepsilon(f) = f(e)$, where e is the identity of G . Direct calculation will verify that

$$(\text{id} \widehat{\otimes} \varepsilon)\Delta(f) = (\varepsilon \widehat{\otimes} \text{id})\Delta(f) = f. \quad (3.16)$$

Consider also the *classical anti-pode* $S : C(G) \rightarrow C(G)$ defined by setting $(Sf)(t) = f(t^{-1})$, for $t \in G$. If m is the linearisation of the multiplication of $C(G)$ on $C(G) \otimes C(G)$, then

$$m(S \widehat{\otimes} \text{id})\Delta(f) = m(\text{id} \widehat{\otimes} S)\Delta(f) = \varepsilon(f)1. \quad (3.17)$$

In this regard $C(G)$ bears comparison to a standard structure in mathematics.

Definition 3.2.4. Let A be a unital $*$ -algebra, let $A \otimes A$ be the $*$ -algebra tensor product of A with itself, and let $\Delta : A \rightarrow A \otimes A$ be a unital $*$ -algebra homomorphism such that $(\Delta \otimes \text{id})\Delta = (\text{id} \otimes \Delta)\Delta$. The pair (A, Δ) is called a *Hopf $*$ -algebra* if there exist linear maps $\varepsilon : A \rightarrow \mathbf{C}$ and $S : A \rightarrow A$ such that

$$(\varepsilon \otimes \text{id})\Delta(a) = (\text{id} \otimes \varepsilon)\Delta(a) = a, \quad (3.18)$$

$$m(S \otimes \text{id})\Delta(f) = m(\text{id} \otimes S)\Delta(f) = \varepsilon(f)1,$$

for all $a \in A$, where m is the linearisation of the multiplication of A on $A \otimes A$.

When no confusion arises, we shall, for sake of simplicity, usually denote a Hopf algebra (A, Δ) by A . We call ε and S the *co-unit* and *anti-pode* of the Hopf algebra. It can be shown that the co-unit and anti-pode of any Hopf algebra are unique. It can also be shown that the co-unit is multiplicative and that the anti-pode is anti-multiplicative. A standard reference for Hopf algebras is [2].

If $f \in C(G)$, then in general $\Delta(f)$ will not be contained in $C(G) \otimes C(G)$. Thus, $(C(G), \Delta)$ is not a Hopf $*$ -algebra. However, there exists a distinguished unital $*$ -subalgebra of $C(G)$ for which this problem does not arise. Let

$$U : G \rightarrow M_n(\mathbf{C}), \quad g \mapsto U_g,$$

be a unitary representation of G , and let $U_{ij}(g)$ denote the ij^{th} -matrix element of U_g . Clearly, $u_{ij} : g \rightarrow U_{ij}(g)$ is a continuous function on G . We call any such function arising from a finite-dimensional unitary representation a *polynomial function* on G . We denote by $\text{Pol}(G)$ the smallest $*$ -subalgebra of $C(G)$ that contains all the polynomial functions. Since $U_{gh} = U_g U_h$,

$$\Delta(u_{ij})(g, h) = \sum_{k=1}^n u_{ik}(g)u_{kj}(h) = \sum_{k=1}^n u_{ik} \otimes u_{kj}(g, h).$$

Thus, $\Delta(u_{ij}) \in \text{Pol}(G) \otimes \text{Pol}(G)$. Moreover, since U is unitary, $S(u_{ij}) = \overline{u_{ji}} \in \text{Pol}(G)$. Hence, the pair $(\text{Pol}(G), \Delta)$ is a Hopf $*$ -algebra with respect to the restrictions of S and ε to $\text{Pol}(G)$. Using the Stone–Weierstrass Theorem it can be shown that $\text{Pol}(G)$ is dense in $C(G)$.

Using quantum group corepresentations, the natural quantum analogue of group representations, Woronowicz established the following very important generalisation of $(\text{Pol}(G), \Delta)$ to the quantum setting.

Theorem 3.2.5 *Let (\mathcal{A}, Δ) be a compact quantum group. Then there exists a unique Hopf $*$ -algebra (A, Φ) such that A is a dense unital $*$ -subalgebra of \mathcal{A} and Φ is the restriction of Δ to A .*

We call (A, Φ) the Hopf $*$ -algebra *underlying* G (or simply the Hopf algebra underlying G when we have no need to consider its $*$ -structure.)

The fact that the anti-pode and co-unit are only defined on a dense subset of the C^* -algebra of a compact quantum group might seem a little unnatural. In fact, earlier formulations of the compact quantum group definition included generalisations of the classical co-unit and anti-pode that were defined on all of \mathcal{A} . However, examples were later to emerge that would not fit into this restrictive framework, and the greater generality of the present definition was required to include them. Probably the most important such example is $SU_q(2)$.

Quantum $SU(2)$

Recall that the *special unitary group* of order 2 is the group

$$SU(2) = \{A \in M_2(\mathbf{C}) : A^* = A^{-1}, \det(A) = 1\}.$$

Recall also that

$$SU(2) = \left\{ \begin{pmatrix} z & -\bar{w} \\ w & \bar{z} \end{pmatrix} : z, w \in \mathbf{C}, |z|^2 + |w|^2 = 1 \right\}.$$

Consider α' and γ' two continuous functions on $SU(2)$ defined by

$$\alpha' \begin{pmatrix} z & -\bar{w} \\ w & \bar{z} \end{pmatrix} = z, \quad \text{and} \quad \gamma' \begin{pmatrix} z & -\bar{w} \\ w & \bar{z} \end{pmatrix} = w.$$

It can be shown that the smallest $*$ -subalgebra of $C(SU(2))$ containing α' and γ' is $\text{Pol}(SU(2))$. Clearly, $\text{Pol}(SU(2))$ is a commutative algebra.

We shall now construct a family of not necessarily commutative algebras $\{A_q\}_{q \in I}$, $I = [-1, 1] \setminus \{0\}$, such that when $q = 1$, the corresponding algebra is $\text{Pol}(SU(2))$. Define A_q to be the universal unital $*$ -algebra generated by two elements α and γ satisfying the relations

$$\begin{aligned} \alpha^* \alpha + \gamma^* \gamma &= 1, & \alpha \alpha^* + q^2 \gamma \gamma^* &= 1, \\ \gamma \gamma^* &= \gamma^* \gamma, & q \gamma \alpha &= \alpha \gamma, & q \gamma^* \alpha &= \alpha \gamma^*. \end{aligned}$$

It is straightforward to show that A_1 is commutative, and that the relations satisfied by its generators are also satisfied by $\alpha', \gamma' \in \text{Pol}(SU(2))$. This means that there exists a unique surjective $*$ -algebra homomorphism θ from A_1 to $\text{Pol}(SU(2))$ such that $\theta(\alpha) = \alpha'$ and $\theta(\gamma) = \gamma'$. Now, if λ is a character on A_1 , then

$$[\lambda] = \begin{pmatrix} \lambda(\alpha) & -\lambda(\gamma^*) \\ \lambda(\gamma) & \lambda(\alpha^*) \end{pmatrix} \in SU(2).$$

Clearly $\theta(x)[\lambda] = \lambda(x)$ if $x = \alpha$ or γ . This immediately implies that $\theta(x)[\lambda] = \lambda(x)$, for all $x \in A_1$. Thus, if the characters separate the points of A_1 , then θ is injective. In [110] Woronowicz showed that A_1 can be embedded into a commutative C^* -algebra. If we recall that we saw in Chapter 1 that the characters of a commutative C^* -algebra always separate its elements, then we can see that Woronowicz's result implies that the characters of A_1 do indeed separate the points of A_1 . Hence, $\text{Pol}(SU(2))$ is isomorphic to A_1 .

Now, we would like to give each A_q the structure of a Hopf algebra. As some routine calculations will verify, there exist unital $*$ -algebra homomorphisms $\Delta : A_q \rightarrow A_q \otimes A_q$, $\varepsilon : A_q \rightarrow \mathbf{C}$, and a unital algebra anti-homomorphism $S : A_q \rightarrow A_q$ such that:

$$\Delta(\alpha) = \alpha \otimes \alpha - q \gamma^* \otimes \gamma, \quad \Delta(\gamma) = \gamma \otimes \alpha + \alpha^* \otimes \gamma,$$

$$\begin{aligned}
S(\alpha) &= \alpha^*, & S(\alpha^*) &= \alpha, & S(\gamma) &= -q\gamma, & S(\gamma^*) &= -q^{-1}\gamma^*, \\
\varepsilon(\alpha) &= 1, & & & \varepsilon(\gamma) &= 0. & &
\end{aligned}$$

As some more routine calculations will verify, the pair (A_q, Δ) is a Hopf $*$ -algebra with ε and S as counit and antipode respectively. If we denote the composition with the multiplication of $SU(2)$ by Φ , then it is not too hard to show that when $q = 1$, Δ and Φ coincide on $\text{Pol}(SU(2))$.

Let us now define a norm on each A_q by

$$\|a\|_u = \sup_{(U, H)} \|U(a)\|, \quad a \in A_q;$$

where the supremum is taken over all pairs (U, H) with H a Hilbert space and U a unital $*$ -representation of A_q on H . Let us denote the completion of A_q with respect to $\|\cdot\|_u$ by \mathcal{A}_q . It can be shown that Δ has a unique extension to a continuous mapping $\Delta_u : \mathcal{A}_q \rightarrow \mathcal{A}_q \widehat{\otimes} \mathcal{A}_q$ such that the pair $(\mathcal{A}_q, \Delta_u)$ is a compact quantum group. It is called *quantum $SU(2)$* and it is denoted by $SU_q(2)$.

It is easy to show that the mapping

$$\Omega(\mathcal{A}_1) \rightarrow SU(2), \quad \tau \mapsto \begin{pmatrix} \tau(\alpha) & -\tau(\gamma) \\ \tau(\gamma) & \tau(\alpha^*) \end{pmatrix}$$

is a homeomorphism. Thus, since \mathcal{A}_1 is clearly commutative, the Gelfand–Naimark theorem implies that $\mathcal{A}_1 = C(SU(2))$. If we recall that Δ and Φ coincide on the dense subset $\text{Pol}(SU(2)) \subseteq C(SU(2))$, then we can see that (\mathcal{A}_1, Δ) is the classical compact quantum group associated to $SU(2)$.

When $q \neq 1$ what we have is a purely quantum object. Each such structure is a prototypical example of a compact quantum group and is of central importance in the theory of compact quantum groups.

3.2.2 Differential Calculi over Quantum Groups

Earlier, we stated that Woronowicz introduced the concept of a compact quantum group in the context of a general mathematical movement to extend Pontrgagin duality. This is not the whole truth: Woronowicz was also heavily influenced by physical considerations. In the early 1980s examples of quantum groups arose in the work of certain Leningrad based physicists studying the inverse scattering problem [89]. Woronowicz (as well as Drinfeld and many others) took considerable inspiration from these structures. Furthermore, the main reason Woronowicz was interested in quantum groups in the first place was because he felt that they might have applications to theoretical physics.

All the topological groups generalised by these physicists were Lie groups. (Recall that a *Lie group* G is a topological group endowed with a differential structure, with respect to which the group multiplication $(x, y) \mapsto xy$ and the inverse map $x \mapsto x^{-1}$ are smooth.) However, none of the quantum groups that appeared in the physics literature had a generalised differential structure associated to them. Woronowicz felt that the introduction of differential calculi into the study of quantum groups would be of significant benefit to the theory. Consequently, he formulated the theory of covariant differential calculi: Let (A, Δ) be the Hopf algebra underlying a compact quantum group G , let ε be the counit of (A, Δ) , and let (Ω, d) be a differential calculus over A . If there exists an algebra homomorphism $\Delta_\Omega : \Omega \rightarrow A \otimes \Omega$ such that

$$(\text{id}_A \otimes d)\Delta_\Omega = \Delta_\Omega d; \quad \Delta_\Omega(a) = \Delta(a), \quad \text{for all } a \in A; \quad \text{and } d1 = 0,$$

then we call the triple $(\Omega, d, \Delta_\Omega)$ a *left-covariant differential calculus* over G . The fact that Ω is generated as an algebra by the elements a and da , for $a \in A$, means that there can only exist one such left action Δ_Ω making $(\Omega, d, \Delta_\Omega)$ a left-covariant calculus. (For this reason we shall usually write (Ω, d) for $(\Omega, d, \Delta_\Omega)$.) Differential calculi that are left-covariant are of prime importance in theory of differential calculi over quantum groups. (There also exists an analogous definition of a right-covariant calculus. Calculi that are both left- and right-covariant are called *bicovariant*.)

Let G be a compact Lie group and let $(\Omega_{\text{Pol}}(G), d)$ denote the de Rham calculus over the algebra of polynomial functions of G ; the construction of this calculus is the same as the construction of the ordinary de Rham calculus except that one uses $\text{Pol}(G)$ instead of $C^\infty(M)$. It is not too difficult to show that $(\Omega_{\text{Pol}}(G), d)$ can be endowed with the structure of a left-covariant calculus over G . Woronowicz gave the first example of a left-covariant differential calculus over a non-classical compact quantum group in his seminal paper [110]. It is a 3-dimensional calculus over $SU_q(2)$, and it is the prototypical example for the theory.

Let (Ω, d) be a left-covariant differential calculus over the Hopf algebra A underlying a compact quantum group; and let \int be a linear functional on Ω^n . If

$$(\text{id} \otimes \int)\Delta_\Omega(\omega) = (\int \omega)1, \quad \text{for all } \omega \in \Omega^n,$$

then we say that \int is *left-invariant*. If whenever $\omega \in \Omega$ is such that $\int \omega \omega' = 0$, for all $\omega' \in \Omega$, we necessarily have $\omega = 0$, then we say that \int is *left-faithful*. A little thought will verify that both these definitions are classically motivated. The natural linear functionals to study on differential calculi over compact quantum groups are the left-invariant, left-faithful, linear functionals.

Recall that earlier in this chapter we studied graded traces on differential calculi as generalisations of volume integrals. It would be quite pleasing if ‘most’ of the natural examples of closed linear functionals (that is, all the closed, left-invariant, left-faithful, linear functionals) on left-covariant differential calculi were of this form. However, this is not the case. When Woronowicz constructed his differential calculus over $SU_q(2)$, he also constructed a canonical 3-dimensional, closed, left-invariant, left-faithful, linear functional on it. This linear functional was not a graded trace but a twisted graded trace.

Definition 3.2.6. Let (Ω, d) be a differential calculus over an algebra A and let \int be a linear functional on Ω^n . We say that \int is an *n-dimensional twisted graded trace* if there exists a differential algebra automorphism $\sigma : \Omega \rightarrow \Omega$ of degree 0 such that, whenever $p + q = n$,

$$\int \omega_p \omega_q = (-1)^{pq} \int \sigma(\omega_q) \omega_p,$$

for all $\omega_p \in \Omega^p$, $\omega_q \in \Omega^q$. We say that σ is a *twist automorphism* associated to \int .

Clearly, if \int is a twisted graded trace with id_A as an associated twist automorphism, then \int is a graded trace. There may exist more than one twist automorphism for a twisted graded trace. However, if the twisted graded trace is left-faithful, then its twist automorphism is unique (this a sufficient but not a necessary condition).

The definition of a twisted graded trace is analogous to the definition of a KMS state on a C^* -algebra: A *KMS state* h on a C^* -algebra \mathcal{A} is a state for which there exists an algebra automorphism σ , defined on a dense $*$ -subalgebra of \mathcal{A} , such that $h(ab) = h(\sigma(b)a)$, for all elements a and b in the $*$ -subalgebra.

It turns out that Woronowicz’s linear functional is not an isolated case, that is, there exist many other examples of closed, left-invariant, left-faithful, linear functionals that are twisted graded traces, but not graded traces. The following very pleasing result, due to Kustermans, Murphy, and Tuset [65], shows why twisted graded traces are so important in the theory of differential calculi over compact quantum groups.

Theorem 3.2.7 *Let (Ω, d) be an n-dimensional left-covariant differential calculus over a compact quantum group; and let $\int : \Omega^n \rightarrow \mathbf{C}$ be an n-dimensional linear functional. If \int is closed, left-invariant, and left-faithful, then it is necessarily a twisted graded trace.*

We note that the notion of a left-covariant differential calculus is well defined over any Hopf algebra, not just those associated to compact quantum groups. However,

for Theorem 3.2.7 to hold we must assume the existence of a Haar integral on the algebra. A unital linear functional h on a Hopf algebra is said to be a *Haar integral* if

$$(\text{id}_A \otimes h)\Delta(a) = (h \otimes \text{id}_A)\Delta(a) = h(a)1, \quad \text{for all } a \in A.$$

Before we leave this section it is interesting to note that Kustermans, Murphy, and Tuset's work on twisted graded traces led them to a new method for constructing left-covariant differential calculi over Hopf algebras. In their approach one essentially starts with a twisted graded trace and then constructs a differential calculus. (Woronowicz's construction ran in the other direction.) Their method seems to be a more natural approach than others, and in [65] they used it to reconstruct Woronowicz's 3-dimensional calculus in an entirely different manner.

3.2.3 Twisted Cyclic Cohomology

The fact that the natural linear functionals to study on differential calculi over quantum groups are twisted graded traces, and not graded traces, poses a natural question: Can one construct a cohomology theory from twisted graded traces in the same way that we constructed cyclic cohomology from graded traces? It turns out that one can.

Let A be a unital algebra and let σ be an algebra automorphism of A . As before, for any positive integer n , we let $C^n(A)$ denote the linear space of complex-valued multilinear maps on A^{n+1} , and we define $C^*(A) = \{C^n(A)\}_{n=0}^\infty$. Let us introduce the unique sequence of maps $b_\sigma = \{b_\sigma : C^n(A) \rightarrow C^{n+1}(A)\}_{n=0}^\infty$, called the *twisted Hochschild coboundary operators*, such that for $\varphi \in C^n(A)$ and $a_0, \dots, a_n \in A$,

$$(b_\sigma\varphi)(a_0, \dots, a_{n+1}) = \sum_{i=0}^n (-1)^i \varphi(a_0, \dots, a_{i-1}, a_i a_{i+1}, a_{i+2}, \dots, a_{n+1}) \\ + (-1)^{n+1} \varphi(\sigma(a_{n+1})a_0, a_1, \dots, a_n).$$

Just as for the ordinary Hochschild coboundary operators, a straightforward calculation will show that $b_\sigma^2 = 0$. Let us also introduce the unique sequence of maps $\lambda_\sigma = \{\lambda_\sigma : C^n(A) \rightarrow C^n(A)\}_{n=0}^\infty$, called the *twisted permutation operators*, such that for $\varphi \in C^n(A)$ and $a_0, \dots, a_n \in A$,

$$\lambda_\sigma(\varphi)(a_0, a_1, \dots, a_n) = (-1)^n \varphi(\sigma(a_n), a_0, a_1, \dots, a_{n-1}).$$

Define $C^n(A, \sigma) = \{\varphi \in C^n(A) : \lambda_\sigma^{n+1}(\varphi) = \varphi\}$ (it is instructive to note that $\lambda_\sigma^{n+1}(\varphi)(a_0, \dots, a_n) = \varphi(\sigma(a_0), \dots, \sigma(a_n))$). It can be shown that $b_\sigma(C^n(A, \sigma)) \subseteq C^{n+1}(A, \sigma)$. Thus, if we denote $C^*(A, \sigma) = \{C^n(A, \sigma)\}_{n=0}^\infty$, then the pair $(C^*(A, \sigma), b_\sigma)$ is a cochain complex. We denote its n^{th} -cohomology group

by $HH^n(A, \sigma)$, and we call it the n^{th} -twisted Hochschild cohomology group of (A, σ) . Clearly, the twisted Hochschild cochain complex of (A, id_A) is equal to the Hochschild cochain complex of A .

Let us now define $C_\lambda^n(A, \sigma) = \{\varphi : \varphi \in C^n(A, \sigma), \lambda_\sigma(\varphi) = \varphi\}$, and $C_\lambda^*(A) = \{C_\lambda^n(A, \sigma)\}_{n=0}^\infty$. It can be shown that $b_\sigma(C_\lambda^n(A, \sigma)) \subseteq C_\lambda^{n+1}(A, \sigma)$, and so, the pair $(C_\lambda^*(A), \sigma)$ is a subcomplex of the twisted Hochschild cochain complex. We call the n^{th} -cohomology group of this complex the n^{th} -twisted cyclic cohomology group of (A, σ) , and we denote it by $HC^n(A, \sigma)$. Furthermore, we denote the set of n -cocycles of the complex by $Z_\lambda^n(A, \sigma)$; we call its elements *twisted cyclic n -cocycles*. Clearly, when $\sigma = \text{id}_A$, the twisted cyclic cochain complex and the cyclic cochain complex coincide.

Recall that if \int is an n -dimensional closed graded trace on an n -dimensional differential calculus over A , then the multilinear map $\varphi : A^{n+1} \rightarrow \mathbf{C}$ defined by setting

$$\varphi(a_0, a_1, \dots, a_n) = \int a_0 da_1 \cdots da_n,$$

is a cyclic n -cocycle. In the twisted setting we have the following result.

Theorem 3.2.8 *Let (Ω, d) be an n -dimensional differential calculus over a unital algebra A and suppose that \int is an n -dimensional, closed, twisted graded trace on Ω . Define the function, $\varphi : A^{n+1} \rightarrow \mathbf{C}$, by setting*

$$\varphi(a_0, \dots, a_n) = \int a_0 da_1 \cdots da_n.$$

Let σ be an automorphism of A for which $\int \sigma(a)\omega = \int \omega a$, for all $a \in A$, $\omega \in \Omega^n$. Then it holds that $\varphi \in Z_\lambda^n(A, \sigma)$. We call φ the twisted cyclic n -cocycle associated to (Ω, d) and \int .

Recall also that if φ is a cyclic n -cocycle of a unital algebra A , then there exists an n -dimensional cycle (Ω, d, \int) such that

$$\varphi(a_0, \dots, a_n) = \int a_0 da_1 \cdots da_n,$$

for all a_0, a_1, \dots, a_n . This result generalises to the twisted case.

Theorem 3.2.9 *Let σ be an automorphism of a unital algebra A and let $\varphi \in Z_\lambda^n(A, \sigma)$, for some integer $n \geq 0$. Then there exists an n -dimensional differential calculus (Ω, d) over A and an n -dimensional, closed, twisted graded trace \int on Ω such that φ is the twisted cyclic n -cocycle associated to (Ω, d) and \int .*

The proofs of Theorem 3.2.8 and Theorem 3.2.9 amount to suitably modified versions of the proofs in the non-twisted case, for details see [65].

It is very interesting to note that Connes' *S-B-I* sequence generalises directly to the twisted setting. It relates the twisted Hochschild, and twisted cyclic cohomologies; for details see [65]. Moreover, it is possible to take the cyclic category technique for calculating cyclic cohomology, and adapt it for use in the twisted setting; for details see [45].

All this is a very pleasing generalisation of Connes' work. The straightforwardness with which everything carries over to the twisted setting suggests that a twisted version of the Connes–Chern maps could also be constructed. This area is the subject of active research.

In [17] Connes and Moscovici introduced a version of cyclic cohomology theory for Hopf algebras. For a discussion of the relationship between Hopf cyclic cohomology and twisted cyclic cohomology see [47].

3.2.4 Twisted Hochschild Homology and Dimension Drop

Recall that in our construction of $SU_q(2)$ we produced a family of Hopf algebras dependant upon a parameter q ; when $q = 1$, the corresponding Hopf algebra was equal to $\text{Pol}(SU(2))$. There exists many other examples of q -parameterised families of Hopf algebras that give the polynomial algebra of a Lie group when $q = 1$. Such a parameterised family is called a *quantisation* of the Lie group. The Hopf algebra corresponding to a particular value of q is known as its *q -deformed polynomial algebra*, or simply a *quantum group*. A wealth of examples can be found in [61]; most of these appeared for the first time in the physics literature.

Motivated by the Hochschild–Kostant–Rosenberg Theorem, and Theorem 3.1.2, we make the following definition: For any algebra A , we define its *Hochschild dimension* to be

$$\sup\{n : HH_n(A) \neq 0\}.$$

For a quantisation of a Lie group it can happen that the Hochschild dimension of the polynomial algebra of the group is greater than the Hochschild dimension of the deformed algebras. This occurrence is known as *Hochschild dimension drop*. For example, take $SL_q(2)$ the standard q -deformed algebra of $SL(2) = \{A \in M_n(2) : \det(A) = 1\}$; it is a Hopf algebra $(\text{Pol}(SL_q(2)), \Delta)$, where $\text{Pol}(SL_q(2))$ is the algebra generated by the symbols a, b, c, d with relations

$$ab = qba, \quad ac = qca, \quad bd = qdb, \quad cd = qdc, \quad bc = cb,$$

$$ad - qbc = 1, \quad da - q^{-1}bc = 1, \quad (3.19)$$

for q not a root of unity. It can easily be shown that it admits a Haar integral. The Hochschild homology of $SL_q(2)$ was calculated in [78]: Its Hochschild dimension was found to be 1, whereas the Hochschild dimension of $\text{Pol}(SL(2))$ is 3. Other examples of Hochschild dimension drop can be found in the work of Feng and Tsygan [33]. This loss of homological information has led many to believe that the Hochschild and cyclic theories are ill-suited to the study of quantum groups, and that a generalisation of Hochschild homology should be introduced to overcome it.

Now, there also exists a twisted version of Hochschild homology. Let A be an algebra and let $\sigma : A \rightarrow A$ be an algebra automorphism. Denote by $C^n(A, \sigma)$ the quotient $C^n(A)/\text{im}(1 - \sigma)$; where $\text{im}(1 - \sigma)$ denotes the image of the map $(1 - \sigma) : C_n(A) \rightarrow C_n(A)$ defined by setting

$$(1 - \sigma)(a_0 \otimes a_1 \otimes \cdots \otimes a_n) = a_0 \otimes a_1 \otimes \cdots \otimes a_n - \sigma(a_0) \otimes \sigma(a_1) \otimes \cdots \otimes \sigma(a_n).$$

A routine calculation will show that the pair $(C^*(A, \sigma), b)$ is a subcomplex of $(C^*(A), b)$. We call the n^{th} -homological group of this subcomplex the *twisted Hochschild n^{th} -homological group* of (A, σ) , and we denote it by $HH_n(A, \sigma)$. (It is interesting to note that it is not very hard to build upon these definitions and define a twisted version of cyclic homology.)

In [45] Tom Hadfield, a former postdoctoral assistant to Gerard Murphy in Cork, and Ulrich Krähmer, calculated the twisted Hochschild and cyclic homology of $SL_q(2)$. Following Feng and Tsygan, they carried out these calculations using noncommutative Koszul resolutions. They showed that for certain automorphisms, the corresponding twisted Hochschild dimension was 3, equal to the Hochschild dimension of $\text{Pol}(SL(2))$. In fact, the simplest of these automorphisms arose in a natural way from the Haar integral on $SL_q(2)$. The two authors would later generalise this result to one that holds for all $SL_q(N)$ [46]. Hadfield [44] went on to verify that a similar situation holds for all Podleś quantum spheres, and Andrzej Sitarz established analogous findings for quantum hyperplanes [102]. Moreover, quite recently, Brown and Zhang [8] have produced some very interesting new results in this area.

Chapter 4

Dirac Operators

The Schrödinger equation for a free particle,

$$i\hbar \frac{\partial}{\partial t} \psi = -\frac{\hbar^2}{2m} \nabla^2 \psi,$$

is based on $E = \frac{p^2}{2m}$, the non-relativistic relation between momentum and kinetic energy, and not on the relativistic one,

$$E = c(m^2 c^2 + p^2)^{\frac{1}{2}}. \quad (4.1)$$

Furthermore, the fact that the time derivative is of first order and the spatial derivatives are of second order implies that the equation is not Lorentz invariant. One of the first attempts to construct a quantum mechanical wave equation that was in accord with special relativity was the Klein–Gordan equation,

$$-\hbar^2 \frac{\partial^2}{\partial t^2} \psi = c^2(m^2 c^2 - \hbar^2 \nabla^2) \psi.$$

It is obtained by canonically quantizing the square of both sides of equation (4.1). While the equation is Lorentz invariant, it does have some problems: it allows solutions with negative energy; and that which one would wish to interpret as a probability distribution turns out not to be positive definite. Dirac hoped to overcome these shortcomings by reformulating the equation so that the time derivative would be of first order. He did this by rewriting it as

$$\frac{i\hbar}{c} \frac{\partial}{\partial t} \psi = (m^2 c^2 - \hbar^2 \nabla^2)^{\frac{1}{2}} \psi.$$

The obvious problem with above expression is that the right hand side is ill-defined. Dirac assumed that it corresponded to a first order linear operator of the form

$$(mcA_0 + \hbar \sum_{i=1}^3 A_i \frac{\partial}{\partial x_i}); \quad (4.2)$$

where each A_i is a matrix, such that

$$A_0^2 = 1, \quad A_i A_0 = -A_0 A_i, \quad i = 1, 2, 3, \quad (4.3)$$

and

$$A_i A_j + A_j A_i = -2\delta_{ij}, \quad i = 1, 2, 3. \quad (4.4)$$

He then found examples of such matrices in $M_4(\mathbf{C})$, namely;

$$A_0 = \begin{pmatrix} 1_2 & 0 \\ 0 & -1_2 \end{pmatrix}; \quad A_i = \begin{pmatrix} 0 & \sigma_i \\ -\sigma_i & 0 \end{pmatrix}, \quad i = 1, 2, 3;$$

where 1_2 is the identity of $M_2(\mathbf{C})$, and σ_i are the *Pauli spin matrices*:

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

(It can be shown that $n = 4$ is the lowest value for which solutions to equations (4.3) and (4.4) can be found in $M_n(\mathbf{C})$.) For all this to make sense it must be assumed that ψ in equation (4.2) takes values in \mathbf{C}^4 .

While Dirac's reformulation of the Klein–Gordon equation is only suitable for describing electrons, or more correctly spin- $\frac{1}{2}$ particles, it was still a great success. Firstly, the allowed solutions no longer have badly behaved probability densities. Also, a natural implication of the equation is the existence of electron spin. (This quantity had previously required a separate postulate.) But, despite Dirac's efforts, the equation does allow apparently 'unphysical' negative energy solutions. Dirac initially considered this a 'great blemish' on his theory. However, after closer examination of these solutions he proposed that they might actually correspond to previously unobserved 'antielectron particles'. According to Dirac these particles would have positive electrical charge, mass equal to that of the electron, and when an electron and an antielectron came into contact they would annihilate each other with the emission of energy according to Einstein's equation $E = mc^2$. Experimental evidence would later verify the existence of such particles; we now call them positrons. Thus was introduced the notion of antimatter. We shall return to this topic in our discussion of quantum field theory in Chapter 5.

An important point to note is that the operator

$$\sum_{i=1}^3 A_i \frac{\partial}{\partial x_i} \quad (4.5)$$

is in fact a square root of the Laplacian

$$-\nabla^2 : C^\infty(\mathbf{R}^4, \mathbf{C}^4) \rightarrow C^\infty(\mathbf{R}^4, \mathbf{C}^4).$$

Moreover, it is easy to produce many similar examples of Laplacian square roots. As we shall see below, one can use Clifford algebras to unite all these examples into a formal method for constructing square roots for any Laplacian; we shall call these ‘square root’ operators, Dirac operators.

The definition and construction of a geometric, or non-Euclidean, version of Dirac operators is more involved. In fact, it was not until the 1960s that examples began to appear. The first was the Kähler–Dirac operator in 1961, and the second was the Atiyah–Singer–Dirac operator in 1962. We shall present both of these operators in this chapter.

At this stage one may well ask what connection there is between Dirac operators and noncommutative geometry: Dirac operators are important in noncommutative geometry because of Connes’ discovery that all the structure of a compact Riemannian (spin) manifold M can be re-expressed in terms of an algebraic structure based on a Dirac operator. This structure, which is known as a spectral triple, admits a straightforward noncommutative generalisation that can then be considered as a ‘noncommutative Riemannian manifold’.

In the first section of this chapter we shall present the basic theory of Dirac operators, and in the second section we shall provide an overview of Connes’ theory of spectral triples.

Recall that in the previous chapter we saw that the borderline between compact quantum groups and cyclic (co)homology is a very active area of research. The same is true of the borderline between compact quantum groups and spectral triples. In the final section of this chapter we shall present some of the work that was done in Cork to construct generalised Dirac operators on quantum groups, and we shall discuss how it relates to Connes’ theory.

4.1 Euclidean and Geometric Dirac Operators

As we stated above, the formal method for constructing Euclidean Dirac operators involves a special type of algebra called a Clifford algebra. These algebras generalise the properties of Dirac’s matrices given in equation (4.4).

4.1.1 Clifford Algebras

Let V be a linear space, over \mathbf{K} ($\mathbf{K} = \mathbf{R}$ or \mathbf{C}), endowed with a symmetric bilinear form B . We shall denote by J the smallest two-sided ideal of $\mathcal{T}(V)$, the tensor algebra of V , that contains all elements of the form $v \otimes v + B(v, v)1$, for $v \in V$.

The algebra $Cl(V) = \mathcal{T}(V)/J(V)$ is called the *Clifford algebra* of V . For sake of convenience we shall denote the coset $(v_1 \otimes \cdots \otimes v_n + J)$ by $v_1 \cdots v_n$.

If we denote the canonical injection of V into $Cl(V)$ by j , then it must hold that $j(v)^2 = -B(v, v)1$; this is a very important property of $Cl(V)$. If A is another unital algebra over \mathbf{K} for which there exists a linear map $j' : V \rightarrow A$ such that

$$j'(v)^2 = -B(v, v)1, \quad (4.6)$$

then it is not very difficult to show that there exists a unique linear homomorphism $h : Cl(V) \rightarrow A$ such that $j' = h \circ j$; or equivalently, such that the following diagram commutes:

$$\begin{array}{ccc} V & \xrightarrow{j} & Cl(V) \\ & \searrow j' & \downarrow h \\ & & A. \end{array}$$

This fact is known as the *universal property* of $Cl(V)$, and it defines $Cl(V)$ uniquely. From now on, for sake of convenience, we shall suppress any reference to j and not distinguish notationally between V and its image in $Cl(V)$.

Let us assume that V is finite-dimensional and that $\{e_i\}_{i=1}^n$ is an orthonormal basis for the space. It easily follows from the definition of $Cl(V)$ that

$$e_i^2 = -1, \quad \text{and} \quad e_i e_j = -e_j e_i, \quad (4.7)$$

for all $i, j = 1, \dots, n$, $i \neq j$. Thus, Clifford algebras generalise the properties of Dirac's matrices given in equation (4.4). The relations in (4.7) imply that the set

$$S = \{1, e_{i_1} \cdots e_{i_k} : 1 \leq i_1 < i_2 < \cdots < i_k \leq n, 1 \leq k \leq n\}$$

spans $Cl(V)$. Moreover, its elements can be shown to be linearly independent. Hence, it forms a basis for $Cl(V)$ of dimension 2^n . An important consequence of this fact is that $Cl(V)$ is linearly isomorphic to the exterior algebra of V . The unique mapping that sends

$$v_1 \wedge \cdots \wedge v_k \mapsto \sum_{\pi \in \text{Perm}(k)} \text{sgn}(\pi) v_{\pi(1)} \cdots v_{\pi(k)}, \quad (4.8)$$

for $k = 1, 2, \dots, k$, is a canonical isomorphism between the two spaces.

The linear map

$$V \rightarrow Cl(V), \quad v \mapsto -v,$$

satisfies equation (4.6). Hence, by the universal property of Clifford algebras, it extends to an algebra automorphism $\chi : Cl(V) \rightarrow Cl(V)$. Clearly, χ is an involution operator, that is, $\chi^2 = 1$. This means that its eigenvalues are ± 1 , and that one can decompose $Cl(V)$ into positive and negative eigenspaces $Cl^+(V)$ and $Cl^-(V)$. As a moment's thought will verify, $Cl^+(V)$ is spanned by products of even numbers of elements of V , and $Cl^-(V)$ is spanned by products of odd numbers of elements of V .

Examples and Representations

If we take $V = \mathbf{R}$, with multiplication as the bilinear form, then $Cl(\mathbf{R})$ has $\{1, e_1\}$ as a basis, where 1 denotes the identity of the Clifford algebra, and e_1 denotes the identity of \mathbf{R} . Now $e_1^2 = -1$, therefore $Cl(\mathbf{R})$ is isomorphic to \mathbf{C} .

Take $V = \mathbf{R}^2$ with the Euclidean inner product as the bilinear form. If $\{e_1, e_2\}$ is the standard basis of \mathbf{R}^2 , then the set $\{1, e_1, e_2, e_1e_2\}$ forms a basis for $Cl(\mathbf{R}^2)$. If we denote

$$i = e_1, \quad j = e_2, \quad k = e_1e_2,$$

then the following relations are satisfied:

$$ij = k, \quad jk = i, \quad ki = j, \quad i^2 = j^2 = k^2 = -1.$$

Thus, we see that $Cl(\mathbf{R}^2) \simeq \mathbf{H}$, the algebra of quaternions.

Now let $\{e_i\}_{i=1}^n$ be the standard basis of \mathbf{C}^n , and let B be the unique symmetric bilinear form on \mathbf{C}^n for which $B(e_i, e_j) = \delta_{ij}$. It can be shown that for any natural number k , then

$$Cl(\mathbf{C}^{2k}) \simeq M_{2^k}(\mathbf{C}); \tag{4.9}$$

and

$$Cl(\mathbf{C}^{2k+1}) \simeq M_{2^k}(\mathbf{C}) \oplus M_{2^k}(\mathbf{C}). \tag{4.10}$$

This algebra representation of $Cl(\mathbf{C}^n)$ is called the *spin representation*. Elementary representation theory now implies that all irreducible representations of $Cl(\mathbf{C}^{2k})$, and $Cl(\mathbf{C}^{2k+1})$, are of dimension 2^k .

Unfortunately, the representation theory of $Cl(\mathbf{R}^n)$, the Clifford algebra of \mathbf{R}^n endowed with the Euclidean inner product, is not as straightforward. However, it does follow similar lines.

4.1.2 Euclidean Dirac Operators

We are now ready to define a generalised version of Dirac's operator. (Note that in this definition we consider \mathbf{R}^n as equipped with its usual inner product.)

Definition 4.1.1. Let $c : Cl(\mathbf{R}^n) \rightarrow \text{End}(\mathbf{C}^m)$ be an algebra representation of $Cl(\mathbf{R}^n)$ on \mathbf{C}^m , and let $\{e_i\}$ be the standard basis of \mathbf{R}^n . The operator

$$D : C^\infty(\mathbf{R}^n, \mathbf{C}^m) \rightarrow C^\infty(\mathbf{R}^n, \mathbf{C}^m), \quad f \mapsto \sum_{i=1}^n c(e_i) \frac{\partial f}{\partial x_i},$$

is called the *Dirac operator associated to c* .

The relations in (4.7) easily imply that $D^2 = -\nabla^2$.

As an example, let us construct a Dirac operator for the Laplacian on $C^\infty(\mathbf{R}^2, \mathbf{C}^2)$. We saw above that $Cl(\mathbf{C}^2) \simeq \mathbf{H}$. Now a routine calculation will show that there exists a unique homomorphism of real algebras that maps $i \mapsto i\sigma_1$, $j \mapsto i\sigma_2$, and $k \mapsto -i\sigma_3$. The Dirac operator associated to this representation is

$$D = \begin{pmatrix} 0 & i\partial_1 + \partial_2 \\ i\partial_1 - \partial_2 & 0 \end{pmatrix},$$

where $\partial_i = \frac{\partial}{\partial x_i}$. As a straightforward calculation will verify, the square of D is indeed the Laplacian. Alternatively, by mapping $i \mapsto i\sigma_1$, $j \mapsto i\sigma_3$, and $k \mapsto i\sigma_2$, we get the Dirac operator

$$D = i \begin{pmatrix} \partial_2 & \partial_1 \\ \partial_1 & -\partial_2 \end{pmatrix}.$$

Again, a straightforward calculation will verify that D squares to give the Laplacian.

Although the details are a little more involved, it is not too hard to show that the operator of equation (4.5) also fits into our framework.

4.1.3 Geometric Dirac Operators

The Dirac operators constructed above can be viewed as operating on the smooth sections of the trivial bundle $\mathbf{R}^n \times \mathbf{C}^m$. This makes it natural to consider the idea of constructing generalised Dirac operators that would operate on the sections of vector bundles over manifolds. We call such operators *geometric Dirac operators*.

Riemannian Manifolds

The construction of a geometric Dirac operator for a manifold M requires a choice of Riemannian metric tensor for M . Thus, we shall need to recall some details about Riemannian manifolds.

If g is a real, or complex, rank-(0, 2) tensor field on an n -dimensional manifold M , then by our comments in Chapter 2 on the locality of tensor fields, g will

induce a bilinear form on each real, or complex, tangent plane of M respectively. A *Riemannian metric tensor* is a real rank- $(0, 2)$ tensor field $g \in \mathfrak{T}_2^0(M; \mathbf{R})$ such that, for each $p \in M$, the induced bilinear form $g_p : T_p(M; \mathbf{R}) \times T_p(M; \mathbf{R}) \rightarrow \mathbf{R}$ is an inner product. Clearly, g_p has a unique extension to a complex-valued symmetric bilinear form on the complex tangent plane $T_p(M)$, which we shall also denote by g_p . It is routine to show that any inner product on $T(M; \mathbf{R})$, in the sense of Section 1.3, induces a Riemannian metric tensor on M . This means that one can find a Riemannian metric tensor for any manifold M . A pair (M, g) consisting of a manifold M and a Riemannian metric g is called a *Riemannian manifold*.

Using g we can define two mutually inverse $C^\infty(M)$ -module isomorphisms between $\Omega^1(M)$ and $\mathcal{X}(M)$. The maps, known as the *flat*, and *sharp, musical isomorphisms* respectively, are

$$\flat : \mathcal{X}(M) \rightarrow \Omega^1(M), \quad X \mapsto X^\flat;$$

where $X^\flat(Y) = g(X, Y)$, for all $Y \in \mathcal{X}(M)$; and

$$\sharp : \Omega^1(M) \rightarrow \mathcal{X}(M), \quad \omega \mapsto \omega^\sharp;$$

where $g(\omega^\sharp, Y) = \omega(Y)$, for all $Y \in \mathcal{X}(M)$. (Note that a simple local argument will show that \sharp is well defined.) We can use the sharp musical isomorphism to define a rank- $(2, 0)$ tensor field g^{-1} on M by

$$g^{-1} : \Omega^1(M) \times \Omega^1(M), \quad (\omega_1, \omega_2) \mapsto g(\omega_1^\sharp, \omega_2^\sharp).$$

This can then be extended to a unique symmetric bilinear mapping

$$g : \Omega^p(M) \times \Omega^p(M) \rightarrow C^\infty(M),$$

by setting

$$g(\omega_1 \wedge \cdots \wedge \omega_p, \omega'_1 \wedge \cdots \wedge \omega'_p) = \det[g^{-1}(\omega_i, \omega'_j)]_{ij}.$$

Clifford Bundles and Clifford Modules

The analogue of the Clifford algebra of \mathbf{R}^n is a smooth algebra bundle over M called the Clifford bundle of M ; the definition of an *algebra bundle* is essentially the same as that of a vector bundle except that the fibres are no longer linear spaces but algebras, and all linear mappings are replaced by algebra mappings. Let us denote the set $\bigcup_{p \in M} Cl(T_p(M))$ by $Cl(M)$ (where $Cl(T_p(M))$ is the Clifford algebra of the tangent plane of M at p , with g_p as the bilinear form) and define a projection $\pi : Cl(M) \rightarrow M$ in the obvious way. Recalling the isomorphism induced by the mappings in (4.8), we see that we can endow $Cl(M)$ with a unique topology that makes it a smooth algebra bundle that is isomorphic, as a smooth vector bundle, to the exterior bundle of M . We call $Cl(M)$ the *Clifford bundle* of M .

Equivalently, one can define the Clifford bundle of a manifold using a transition function argument. Let $\{U_\alpha\}$ be an open covering of M by base neighbourhoods for $T(M)$, and let $\{g_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow \text{GL}(2^n, \mathbf{C})\}_{\alpha\beta}$ be the corresponding set of transition functions. For each $p \in U_\alpha \cap U_\beta$, we can use the Riemannian metric of M to endow the domain and codomain of each $g_{\alpha\beta}(p)$ with canonical symmetric bilinear forms. Obviously, each $g_{\alpha\beta}(p)$ is isometric with respect to these bilinear forms. Using the universal property of $Cl(V)$, it is not too difficult to show that there exists a unique algebra homomorphism $cl(g_{\alpha\beta}(p)) : Cl(\mathbf{R}^n) \rightarrow Cl(\mathbf{R}^n)$ such that the following diagram is commutative:

$$\begin{array}{ccc} \mathbf{R}^n & \xrightarrow{g_{\alpha\beta}(p)} & \mathbf{R}^n \\ j \downarrow & & \downarrow j' \\ Cl(\mathbf{R}^n) & \xrightarrow{cl(g_{\alpha\beta}(p))} & Cl(\mathbf{R}^n). \end{array}$$

We can use this fact to define a set of functions

$$\widehat{g}_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow \text{GL}(2^n, \mathbf{C}), \quad p \rightarrow cl(g_{\alpha\beta}(p)), \quad U_\alpha \cap U_\beta \neq \emptyset.$$

This set can easily be shown to satisfy the conditions of the smooth analogue of Proposition 1.3.3. The associated smooth vector bundle, endowed with the obvious smooth algebra bundle structure, is isomorphic to $Cl(M)$.

There also exists a more formal construction of the Clifford bundle in terms of principle bundles. For details on this approach see [69].

Finally, let us introduce the analogue of a Clifford algebra representation: a *Clifford module* for M is a pair (E, c) where E is a smooth vector bundle over M , called the *spinor bundle*, and c is a module homomorphism from $\Gamma(Cl(M))$ to $\text{End}(\Gamma(E))$. Smooth sections of a spinor bundle are called *spinors*. Using Theorem 1.3.7, it is not hard to show that if $b \in \Gamma^\infty(Cl(M))$, then, for every spinor s , $c(b)s$ is also a spinor.

We note that since $\mathcal{X}(M)$ is canonically a subset of $\Gamma(Cl(M))$, any Clifford module (E, c) induces a module homomorphism $c : \mathcal{X}(M) \rightarrow \text{End}(\Gamma(E))$ by restriction. In turn, this homomorphism induces a homomorphism $c : \Omega^1(M) \rightarrow \text{End}(\Gamma(E))$ defined by setting $c(\omega) = c(\omega^\sharp)$, for $\omega \in \Omega^1(M)$.

Geometric Dirac Operators

The pieces are now in place to define a generalised Dirac operator.

Definition 4.1.2. Let (E, c) be a Clifford module for a manifold M , and let ∇ be a connection for E . The associated *Dirac operator* D is

$$D = \widehat{c} \circ \nabla : \Gamma^\infty(E) \rightarrow \Gamma^\infty(E);$$

where $\widehat{c} : \Gamma^\infty(E) \otimes \Omega^1(M) \rightarrow \Gamma^\infty(E)$ is the unique module homomorphism for which

$$\widehat{c}(s \otimes \omega) = c(\omega)s.$$

A local orthonormal basis $\{E_i\}_{i=1}^n$ of $T(M)$ over a neighbourhood U is a local basis over U , such that $g(E_i, E_j)|_U = \delta_{ij}1|_U$, for $i = 1, 2, \dots, n$. It is clear that a local orthonormal basis can be constructed from any local basis. Now if $\{E_i\}_{i=1}^n$ is a local orthonormal basis of $T(M)$ over a neighbourhood U , then for any $X \in \mathcal{X}(M)$, $X|_U = \sum_{k=1}^n g(E_k, X)E_k|_U$. Therefore, if (E, c) is a Clifford module and $s \in \Gamma^\infty(E)$, then

$$\nabla_X s|_U = \sum_{i=1}^n g(E_i, X) \nabla_{E_i} s|_U.$$

Consequently,

$$\nabla s|_U = \sum_{i=1}^n \nabla_{E_i} s \otimes E_i^\flat|_U.$$

This means that if D is the Dirac operator associated to (E, c) and ∇ , then

$$D(s)|_U = \widehat{c} \left(\sum_{i=1}^n \nabla_{E_i} s \otimes E_i^\flat s \right) |_U = \sum_{i=1}^n c((E_i^\flat)^\sharp) \nabla_{E_i} s|_U \quad (4.11)$$

$$= \sum_{i=1}^n c(E_i) \nabla_{E_i} s|_U. \quad (4.12)$$

An immediate consequence of this is that if $M = \mathbf{R}^n$, $E = \mathbf{R}^n \times \mathbf{C}^m$, and ∇ is the canonical connection for E , then the Dirac operators of Definition 4.1.1 and Definition 4.1.2 coincide.

Hodge Theory and the Kähler–Dirac Operator

An interesting example of a geometric Dirac operator comes from noting that $Cl(M)$ is itself canonically a spinor bundle. Moreover, since the exterior bundle of M is isomorphic to $Cl(M)$ as a vector bundle, $\Lambda(M)$ is also canonically a spinor bundle.

Now for every Riemannian manifold M , there is a unique connection ∇^g for the tangent bundle that is *compatible with the metric*, that is,

$$X(g(Y, Z)) = g(\nabla_X^g Y, Z) + g(Y, \nabla_X^g Z), \quad X, Y, Z \in \mathcal{X}(M);$$

and *torsion-free*, that is,

$$\nabla_X^g Y - \nabla_Y^g X = [X, Y].$$

We call ∇^g the *Levi-Civita connection*. It induces a connection for $\Omega^1(M)$, also denoted by ∇^g , that is defined by

$$[\nabla_X^g \omega](Y) = X(\omega(Y)) - \omega(\nabla_X^g Y), \quad \omega \in \Omega(M).$$

(A simple local argument will show that ∇^g is well defined). This connection can then be extended to a unique connection $\bar{\nabla}$ for $\Omega(M)$ such that

$$\bar{\nabla}(\omega_1 \wedge \cdots \wedge \omega_n) = \sum_{i=1}^n \omega_1 \wedge \cdots \wedge \bar{\nabla} \omega_i \wedge \cdots \wedge \omega_n,$$

for $\omega_i \in \Omega^1(M)$. We call $\bar{\nabla}$ the *Levi-Civita connection* for $\Omega(M)$.

The Dirac operator associated to $\Omega(M)$ and $\bar{\nabla}$ is called the *Kähler-Dirac operator*. It has a pleasing representation in terms of the Hodge codifferential, which we shall now introduce.

As is well known, there exists a unique invertible linear mapping $*$: $\Omega(M) \rightarrow \Omega(M)$ called the *Hodge operator* such that if $\omega_1, \omega_2 \in \Omega^p(M)$ then

$$\omega_1 \wedge * \omega_2 = g(\omega_1, \omega_2) d\mu;$$

where $d\mu$ is the Riemannian volume form, (see Section (4.2.1) for details). Clearly, $*$ must map p -forms to $(n - p)$ -forms.

Using the Hodge operator, we can endow $\Omega(M)$ with an inner product by defining

$$\langle \omega, \omega' \rangle = \int \bar{w}' \wedge * \omega, \tag{4.13}$$

if ω and ω' have the same degree, and

$$\langle \omega, \omega' \rangle = 0 \tag{4.14}$$

otherwise. If $\omega \in \Omega^{p-1}(M)$ and $\omega' \in \Omega^p(M)$, then

$$d(\bar{\omega} \wedge * \omega') = (d\bar{\omega}) \wedge * \omega' + (-1)^{p-1} \bar{\omega} \wedge * (*^{-1} d* \omega').$$

Stokes' Theorem now implies that

$$\int d\bar{\omega} \wedge *\omega' = (-1)^p \int \bar{\omega} \wedge (*^{-1}d*)\omega'.$$

If we denote

$$d^* = (-1)^p(*^{-1}d*), \quad (4.15)$$

then $\langle \omega', d\omega \rangle = \langle d^*\omega', \omega \rangle$. Hence, d^* is the adjoint of d . We call d^* the *Hodge codifferential*.

It can be shown that the Kähler–Dirac operator is equal to $d + d^*$, see [69, 39] for details. Since $d^2 = (d^*)^2 = 0$, it holds that $(d + d^*)^2 = dd^* + d^*d$. We call $\nabla = dd^* + d^*d$ the *Hodge–Laplacian* operator. If $\nabla(\omega) = 0$, then we call ω a *harmonic form*. The space of harmonic forms is denoted by $\Omega_{\nabla}(M)$. An important result involving the differential, codifferential, and Laplacian is the *Hodge Decomposition Theorem*:

$$\Omega(M) = \Omega_{\nabla}(M) \oplus d(\Omega(M)) \oplus d^*(\Omega(M)).$$

4.1.4 Spin Manifolds and Dirac Operators

Let M be an n -dimensional Riemannian manifold and let (S, c) be a spinor module over M . Theorem 1.3.7 implies that, for any $p \in M$, c induces an action of $Cl(M)_p$ on S_p . A spinor module for which this action is irreducible, for all p , is called an *irreducible spinor bundle*. If the dimension of M is $2k$, or $2k + 1$, then equations (4.9) and (4.10) imply that any irreducible spinor bundle will have rank 2^k .

Irreducible spinor bundles do not exist for every Riemannian manifold. However, there does exist a distinguished type of Riemannian manifold for which one always does: the Riemannian spin manifolds. Usually, a *spin manifold* is defined to be an n -dimensional orientable Riemannian manifold whose $SO(n)$ -principle bundle of oriented orthonormal frames can be ‘lifted’ to a $\text{spin}(n)$ -principle bundle. (The group $\text{spin}(n)$ is the universal covering group of $SO(n)$ and it arises as a subspace of $Cl(\mathbf{C}^n)$.) The restriction of the spin representation, given in equations (4.9) and (4.10), to $\text{spin}(n)$ gives a representation of $\text{spin}(n)$. The vector bundle associated to this representation is an irreducible spinor bundle.

The ability to ‘lift’ the bundle of oriented orthonormal frames to a $\text{spin}(n)$ -bundle can be shown to be equivalent to the vanishing of the second Stiefel–Whitney cohomological class of $T(M)$. This means that a spin manifold can be alternatively defined as Riemannian manifold whose second Stiefel–Whitney class vanishes.

A very thorough introduction to the principle bundle approach to spin manifolds can be found in [69].

When M is compact, there also exists an operator theoretic formulation of the definition. While the idea for this approach originally came from Connes, it was Roger Plymen [86] who first published a written account of it. As one would expect, it is the operator theoretic approach that we shall follow here. We shall make the assumption that M is of even dimension. The odd dimensional case follows along very similar lines but it has some added technical difficulties, for details see [86] or [39].

Morita Equivalence

Let \mathcal{B} be a C^* -algebra, and let \mathcal{E} be a right B -module. We call \mathcal{E} a *pre-Hilbert \mathcal{B} -module* if there exists a map $(\cdot, \cdot) : \mathcal{E} \times \mathcal{E} \rightarrow \mathcal{B}$, called the *Hilbert module inner product*, such that, for all $x, y, z \in \mathcal{E}$, $b \in \mathcal{B}$, $\lambda \in \mathbf{C}$,

1. $(x, y + z) = (x, y) + (x, z)$;
2. $(x, yb) = (x, y)b$;
3. $(x, y) = (y, x)^*$;
4. $(x, x) \geq 0$; and if $(x, x) = 0$, then $x = 0$.

We can define a norm on \mathcal{E} by setting $\|x\| = \sqrt{\|(x, x)\|}$, for $x \in \mathcal{E}$. If \mathcal{E} is complete with respect to this norm, then we say that \mathcal{E} is a *Hilbert \mathcal{B} -module*. Clearly, every Hilbert space (with a right linear inner product) is a Hilbert \mathbf{C} -module. Every C^* -algebra \mathcal{A} can be given a Hilbert \mathcal{A} -module structure by defining $(a, b) = a^*b$, $a, b \in \mathcal{A}$. (It is not too hard to see that Hilbert modules generalise the notion of a Hilbert bundle; for details see [66].)

A Hilbert \mathcal{B} -module is called *full* if the closure of the linear span of $\{(x, y) : x, y \in \mathcal{E}\}$ is \mathcal{B} .

Let \mathcal{E} be a Hilbert module and let T be a module mapping from \mathcal{E} to \mathcal{E} . We say that T is *adjointable* if there exists a module mapping $T^* : \mathcal{E} \rightarrow \mathcal{E}$ satisfying $(Tx, y) = (x, T^*y)$, for all $x, y \in \mathcal{E}$. We call T^* the *adjoint* of T . It is routine to verify that if an operator is adjointable, then its adjoint is necessarily unique, and that $(T^*)^* = T$. Furthermore, a straightforward application of the closed graph theorem will show that any adjointable operator is bounded. However, unlike the special case of Hilbert space operators, not all bounded module maps are adjointable. We denote the space of adjointable module maps on \mathcal{E} by $L(\mathcal{E})$.

For $x, y \in \mathcal{E}$, consider the mapping

$$\theta_{x,y} : \mathcal{E} \rightarrow \mathcal{E}, \quad z \mapsto x(y, z).$$

It is easy to see that $\theta_{x,y}$ is adjointable, for all $x, y \in \mathcal{E}$, with $\theta_{x,y}^* = \theta_{y,x}$. Let $K(\mathcal{E})$ denote the closure in $L(\mathcal{E})$ of the linear span of $\{\theta_{x,y} : x, y \in \mathcal{E}\}$. We call an element of $K(\mathcal{E})$ a *compact operator*.

Let \mathcal{A}, \mathcal{B} be two C^* -algebras. If there exists a full Hilbert \mathcal{B} -module such that $\mathcal{A} \simeq K(\mathcal{E})$, or a full Hilbert \mathcal{A} -module such that $\mathcal{B} \simeq K(\mathcal{E})$, then we say that \mathcal{A} and \mathcal{B} are *Morita equivalent*. We call \mathcal{E} an *\mathcal{A} - \mathcal{B} -equivalence bimodule*, or a *\mathcal{B} - \mathcal{A} -equivalence bimodule*, depending on which case holds. (An interesting fact is that if two algebras are Morita equivalent, then their Hochschild and cyclic (co)homology groups are the same.)

The Morita Equivalence of $C(M)$ and $\Gamma(Cl(M))$

We can endow $\Gamma(Cl(M))$ with a norm by defining $\|b\|_\infty = \sup\{\|b(x)\|, x \in M\}$, for $b \in \Gamma(Cl(M))$; where $\|\cdot\|$ is the unique C^* -norm on $Cl(M)_p \simeq M_{2^n}(\mathbf{C})$. It is straightforward to show that $\Gamma(Cl(M))$ is a C^* -algebra with respect to this norm. A *spin^c structure* for M is a pair (ω, Σ) consisting of an orientation ω , and a $\Gamma(Cl(M))$ - $C(M)$ -equivalence bimodule Σ . A manifold endowed with a spin^c structure is called a *spin^c manifold*.

It is by no means guaranteed that an arbitrary manifold can be endowed with a spin^c structure. Those manifolds that can be so endowed can be categorised in cohomological terms. Specifically, an oriented Riemannian manifold can be equipped with a spin^c structure if, and only if, the Dixmier–Douady class of its Clifford bundle is zero. (The Dixmier–Douady class of $Cl(M)$ is a characteristic class of $Cl(M)$ that takes values in $H^3(M, \mathbf{Z})$; it is equal to the third integral Stiefel–Whitney class of $T(M)$; for details see [86]. This subtle interaction of spin^c structures with the underlying topology of the manifold is one of the reasons why they are so interesting.)

Now if M is a spin^c manifold and if (ω, Σ) is its spin^c structure, then it can be shown that Σ is projective and finitely-generated. Thus, by the Serre–Swan Theorem, there exists a vector bundle S over M such that $\Sigma = \Gamma(S)$. (It is not too hard to see that the Hilbert module inner product of Σ can be induced by a vector bundle inner product on S .) Using Theorem 1.3.7, it is easy to show that S is a spinor bundle. Moreover, using Dixmier–Douady theory, S can be shown to be an irreducible spinor bundle [39, 86].

Spin Manifolds and the Atiyah–Singer–Dirac Operator

While it is certainly possible to construct Dirac operators that act on the smooth sections of the spinor bundle associated to a spin^c manifold, our interest lies in a more specific structure. Before we introduce this structure, however, we shall

need to define a new operator: Let $\chi : \Gamma(Cl(M)) \rightarrow \Gamma(Cl(M))$ be the unique linear mapping such that, for $b \in \Gamma(Cl(M))$,

$$\chi(b) : p \mapsto \chi(b(p))$$

(where on the right hand side χ is the operator defined in Section 4.1.1).

Definition 4.1.3. A *spin structure* on an orientable Riemannian manifold M is a triple $(\omega, \Gamma(S), C)$, where $(\omega, \Gamma(S))$ is a spin^c structure on M , and C is a bijective module endomorphism of $\Gamma(S)$ such that

$$1. C(bsf) = \chi(\bar{b})(Cs)\bar{f}, \quad f \in C(M), b \in \Gamma(Cl(M)), s \in \Gamma^\infty(S);$$

and, if (\cdot, \cdot) is the Hilbert module inner product of S , then

$$2. (Cs, Cs') = (s', s), \quad s, s' \in \Gamma(S).$$

A *spin manifold* is an orientable compact Riemannian manifold endowed with a spin structure.

For sake of clarity, it is worthwhile to show exactly what is meant by \bar{b} , the ‘complex conjugate’ of $b \in \Gamma(Cl(M))$. Since $T_p(M) \simeq T_p(M; \mathbf{R}) \oplus iT_p(M; \mathbf{R})$, there exists a canonical complex conjugation operator on $T_p(M)$, defined by $(v_p, \lambda v_p) = (v_p, \bar{\lambda} v_p)$, for $v_p \in Cl(M; \mathbf{R})$, $\lambda \in \mathbf{C}$. This has a unique extension to an antilinear mapping on $\mathcal{T}(T_p(M))$, the tensor algebra of $T_p(M)$, which in turn descends to an antilinear mapping on $Cl(T_p(M))$. This ‘complex conjugation’ on $Cl(T_p(M))$ then induces a complex conjugation on $\Gamma(Cl(M))$ in an obvious manner; we denote the image of $b \in \Gamma(Cl(M))$ under this mapping by \bar{b} .

One may well ask why we are interested in the existence, or not, of a spin structure for a manifold. Unfortunately, it is a little difficult to give an intuitive way of looking at the operator C without engaging in an excessive digression. We shall, instead, attempt to justify its introduction as follows. Classically, there are two principal motivations: firstly, we have the important properties of the (Atiyah–Singer–)Dirac operator that is canonically associated to each spin structure. In the general spin^c case such a well-behaved Dirac operator is not guaranteed to exist. Secondly, we have the formulation of spin structures in terms of principle bundles. In this setting the condition corresponding to the existence of the operator C is much more natural; for details see [39]. Finally, from a noncommutative point of view, we are interested in spin structures because of Rennie’s Spin Manifold Theorem, as discussed below. A version of Rennie’s Theorem is not known to hold in the general spin^c case.

Let $(\omega, \Gamma(S), C)$ be a spin structure for a Riemannian manifold M . It can be shown that there exists a unique connection ∇^S for $\Gamma^\infty(S)$ called the *spin connection*, such that

1. ∇^S commutes with C ;
2. $\nabla_X^S(c(\omega)s) = c(\nabla_X^g \omega)s + c(\omega)\nabla_X^S s$, $\omega \in \Omega^1(M)$, $s \in \Gamma^\infty(S)$;

and, if (\cdot, \cdot) is the Hilbert module inner product of S , then

$$3 \ (\nabla_X^S s_1, s_2) + (s_1, \nabla_X^S s_2) = X(s_1, s_2), \quad s_1, s_2 \in \Gamma^\infty(S).$$

The Dirac operator associated to S and ∇^S is called the *Atiyah–Singer–Dirac operator* and it is denoted by \mathcal{D} . It was introduced by Atiyah and Singer while they were working on their famous index theorem.

The Atiyah–Singer–Dirac operator has many important applications in modern mathematics and physics. In recent years, for example, it has been used in the study of 4-dimensional manifolds through the Seiberg–Witten invariants. Details on the mathematical applications of Dirac operators can be found in [69, 36]; and details on some of its physical applications can be found in [32, 83].

4.1.5 Properties of the Atiyah–Singer–Dirac Operator

Let S be the irreducible spinor bundle associated to a spin manifold M and let (\cdot, \cdot) be the Hilbert module inner product of S . If $d\mu$ is the Riemannian measure on M , then we can define an inner product on $\Gamma^\infty(S)$ by

$$\langle s_1, s_2 \rangle = \int (s_2, s_1) d\mu, \quad s_1, s_2 \in \Gamma^\infty(S).$$

We denote by $L^2(S)$ the Hilbert space completion of $\Gamma^\infty(S)$ with respect to the norm induced by this inner product. We call $L^2(S)$ the *Hilbert space of square-integral spinors* on M . In this context \mathcal{D} becomes a linear operator defined on a dense subspace of \mathcal{H} . It can be shown that \mathcal{D} is always unbounded.

If A is an operator on a Hilbert space H , then the *graph* of A is the set

$$G(A) = \{(x, Ax) : x \in \text{dom}(A)\} \subseteq H \oplus H.$$

If $\langle \cdot, \cdot \rangle$ is the inner product of H , then one can define an inner product on $H \oplus H$ by setting

$$\langle (x, y), (u, v) \rangle = \langle x, u \rangle + \langle y, v \rangle.$$

If the closure of $G(A)$ with respect to the norm induced by $\langle \cdot, \cdot \rangle$ is also the graph of an operator B , then we say that A is closable, and we call B the *closure* of A . Obviously, the closure of A will extend A .

It can be shown that \mathcal{D} is a closable operator [39]. From now on, we shall always use \mathcal{D} to denote the closure of the Dirac operator, and we shall refer to the closure of the Dirac operator simply as the Dirac operator.

The theory of unbounded operators is notoriously problematic, to the extent that substantial results about general unbounded operators are rare. It is only when an unbounded operator is closed (that is, equal to its closure) that it becomes somewhat ‘manageable’.

Given a densely defined linear operator A on a Hilbert space H , its *adjoint* A^* is defined as follows: the domain of A^* consists of all vectors $x \in H$ such that the linear map

$$\text{dom}(A) \rightarrow \mathbf{C}, \quad y \mapsto \langle x, Ay \rangle,$$

is a bounded linear functional. Since the $\text{dom}(A)$ is dense in H , each such functional will extend to a unique bounded linear functional defined on all H . Now if x is in the domain of A^* , then the Riesz representation theorem implies that there is a unique vector $z \in H$ such that

$$\langle x, Ay \rangle = \langle z, y \rangle, \quad \text{for all } y \in \text{dom}(A).$$

It is routine to show that the dependence of z on x is linear. We define A^* to be the unique linear operator for which $A^*x = z$. An operator A is said to be *self-adjoint* if $G(A) = G(A^*)$. Self-adjoint operators have many very useful properties that are not necessarily possessed by non-self-adjoint operators. For example, it can easily be shown that every self-adjoint operator is closed.

In 1973 Wolf [109] showed that \mathcal{D} is a self-adjoint operator.

We can regard any $f \in C^\infty(M)$ as a linear operator on $\Gamma^\infty(S) \subseteq L^2(S)$ that acts by multiplication. This representation of $C^\infty(M)$ is obviously faithful. Furthermore, the definition of the inner product on $\Gamma^\infty(S)$ implies that f is bounded with norm $\|f\|_\infty$. Hence, it extends to a unique bounded linear operator on all of $L^2(S)$ with norm $\|f\|_\infty$. We shall not distinguish notationally between f and this operator.

Now

$$[\mathcal{D}, f]s = \hat{c}(\nabla^S(sf)) - f\hat{c}(\nabla^S s) = \hat{c}(\nabla^S(sf) - (\nabla^S s)f) = \hat{c}(s \otimes df) = c(df)s,$$

for all $f \in C^\infty(M)$, $s \in \Gamma^\infty(S)$. Thus, we have that

$$[\mathcal{D}, f] = c(df), \tag{4.16}$$

for all $f \in C^\infty(M)$. As we shall see, this formula allows one to recover the differential structure of M from \mathcal{D} . Moreover, it can easily be used to show that, for all $f \in C^\infty(M)$, the operator norm of the densely defined operator $[\mathcal{D}, f]$ is equal to $\|(df)^\sharp\|_\infty$; where $\|\mathcal{X}\|_\infty = \sup_{p \in M} g_p(X(p), X(p))$, for $X \in \mathcal{X}(M)$. Consequently, each such operator is bounded and has a unique extension to a bounded linear operator defined on all of $L^2(S)$. We shall not distinguish notationally between $[\mathcal{D}, f]$ and this operator.

Finally, we shall list one more important property of \mathcal{D} : the Dirac operator has *compact resolvent*; that is, $(\mathcal{D} - \lambda)^{-1}$ is a compact operator, for all $\lambda \in \rho(\mathcal{D}) = \mathbf{C}/\sigma(\mathcal{D})$.

4.2 Spectral Triples

As we stated earlier, and as we shall see below, much of the structure of a compact Riemannian spin manifold can be expressed in terms of its associated Dirac operator. This fact motivated Connes to try and construct a noncommutative generalisation of spin manifolds based on a type of generalised Dirac operator. In Connes' work the Dirac operator is no longer an object associated to a manifold but rather one of the data defining it.

Definition 4.2.1. A *spectral triple* (A, H, D) consists of a $*$ -algebra A faithfully represented on a Hilbert space H , together with a (possibly unbounded) self-adjoint operator D on H such that:

1. $\text{dom}(D) \subseteq H$ is a dense subset of H , and $a \text{ dom}(D) \subseteq \text{dom}(D)$, for all $a \in A$;
2. the operator $[D, a]$ is bounded on $\text{dom}(D)$, for all $a \in A$, and so it extends to a unique bounded operator on H ;
3. $(D - \lambda)^{-1}$ is a *compact operator*, for all $\lambda \notin \sigma(D)$.

(Since no confusion will arise, we shall not distinguish notationally between an element of A and its image in $B(H)$, nor between $[D, a]$ and its unique extension. Moreover, when we speak of \bar{A} , the closure of A , we mean the closure of its image in $B(H)$ with respect to the operator norm.)

If M is a compact Riemannian spin manifold, then, from our comments above, $(C^\infty(M), L^2(S), \mathcal{D})$ is a spectral triple; we call it the *canonical triple* associated to M .

The definition of a spectral triple is partly motivated by the notion of a Fredholm module. We shall not enter into a discussion of such involved topics here, instead, we refer the interested reader to [12, 39].

Associated to every spectral triple (A, H, D) is a distinguished differential calculus. Let $\Omega_u(A)$ denote the universal calculus of A , and let π denote the unique algebra homomorphism from $\Omega_u(A)$ to $B(H)$ for which

$$\pi(a_0 da_1 \cdots da_n) = a_0 [D, a_1] \cdots [D, a_n].$$

(The fact that π is a homomorphism easily follows from the fact that d and $[D, \cdot]$ are derivations.) Now $\pi(\Omega_u(A)) \simeq \Omega_u(A)/\ker(\pi)$ is canonically a graded algebra, and it would be natural to define a differential d on it by $d(\pi(\omega)) = \pi(d\omega)$. However, this is not always a well-defined mapping since there can exist forms $\omega \in \Omega_u(A)$ such that $\omega \in \ker(\pi)$ and $d\omega \notin \ker(\pi)$.

With a view to overcoming this problem, let us consider the subalgebra $J = \ker(\pi) + d\ker(\pi)$. If $a_0 + da_1 \in J \cap \Omega_u^k(A)$ and $b \in \Omega_u(A)$, then

$$(a_0 + da_1)b = a_0b + (da_1)b = a_0b + d(a_1b) - (-1)^{k-1}a_1db.$$

Since a_0b , a_1b , and a_1db are elements of $\ker(\pi)$, we have that $(a_0 + da_1)b \in J$. Similarly, $b(a_0 + da_1) \in J$. Hence, J is an ideal of $\Omega_u(A)$. Moreover, since $dJ = d(\ker(\pi)) \subseteq J$, J is a differential ideal. This means that we can canonically give the quotient algebra

$$\Omega_D(A) = \pi(\Omega_u(A))/\pi(J) = \pi(\Omega_u(A))/\pi(d(J))$$

the structure of a differential algebra; we call it the *differential algebra of D-forms*. Furthermore, since $\Omega_D^0(A) = \pi(\Omega_u^0(A))/\pi(J_0)$, and $J_0 = \{0\}$, we have that $\Omega_D^0(A) = A$. Hence, $\Omega_D(A)$ is a differential calculus over A .

Now if M is a Riemannian spin manifold and $f_0 d_u f_1 \cdots d_u f_n \in \Omega_u^n(C^\infty(M))$, then, since $[\mathcal{D}, f] = c(df)$, for all $f \in C^\infty(M)$,

$$\begin{aligned} \pi(f_0 d_u f_1 \cdots d_u f_n) &= f_0 [\mathcal{D}, f_1] \cdots [\mathcal{D}, f_n] \\ &= f_0 c(df_1) \cdots c(df_n). \end{aligned}$$

One can build upon this fact to show that $\Omega_{\mathcal{D}}^p(C^\infty(M)) \simeq \Omega^p(M)$, for all $p \geq 0$. Moreover, one can show that $(\Omega_{\mathcal{D}}^p(C^\infty(M)), d)$ and $(\Omega(M), d)$ are isomorphic as differential algebras. Thus, the differential calculus of a Riemannian spin manifold is entirely encoded in its canonical spectral triple.

Rennie's Spin Manifold Theorem

A natural question to ask is which spectral triples arise from spin manifolds and which do not. With a view to answering this question Connes introduced a refinement of the notion of a spectral triple called a noncommutative geometry. (The definition of a noncommutative geometry is rather lengthy, and we shall not go into the details here. It is, however, interesting to note that it contains noncommutative generalisations of orientability and Poincaré duality.) Connes showed that every spectral triple arising from a compact spin manifold is a noncommutative geometry. He went on to claim that all commutative noncommutative geometries (that is, all noncommutative geometries such that the $*$ -algebras of their underlying spectral triples are commutative) arise from compact spin manifolds [13]; Rennie and Varilly [88] would later prove this. At present, there is no extension of this work to the locally compact case.

The Connes–Moscovici Index Theorem

From what we have presented above, one might get the impression that spectral triples and cyclic cohomology have little in common. However, the two areas are intimately linked. In fact, in [16] Connes and Moscovici used spectral triples to construct new versions of the Chern–Connes maps. Their reformulation is extremely attractive for practical calculations. In a related development, they also formulated a very important noncommutative generalisation of the Atiyah–Singer index theorem; it is called the Connes–Moscovici index theorem.

4.2.1 The Noncommutative Riemannian Integral

Let M be an n -dimensional manifold and let (U, φ) be one of its coordinate neighbourhoods. For $i = 1, 2, \dots, n$, let the function x_i be some global extension of $\pi_i \circ \varphi^{-1}$, where π_i is the canonical projection onto the i^{th} coordinate of \mathbf{R}^n . We can define a matrix-valued function on U by $g_U(p) = [g_p(dx_i, dx_j)(p)]_{ij}$ (we shall assume that $\sqrt{\det(g_U(p))} \geq 0$, for all $p \in U$, and for all coordinate neighbourhoods U). Now, it is straightforward to show that there exists a unique $d\mu \in \Omega^n(M)$ such that, for any coordinate neighbourhood U ,

$$d\mu(p) = \sqrt{\det(g_U(p))} dx_1 \wedge dx_2 \wedge \cdots \wedge dx_n, \quad \text{for all } p \in U.$$

We call $d\mu$ the *Riemannian volume form*. For any $f \in C^\infty(M)$, we call $\int f d\mu$ the *Riemannian integral* of f over M .

A major achievement of Connes has been to re-express the Riemannian integral of a spin manifold purely in terms of its Dirac operator. More explicitly, Connes

established the following equation:

$$\frac{1}{c_n} \int f d\mu = \text{Tr}_\omega(f|D|^{-n}), \quad f \in C^\infty(M);$$

where c_n is a constant, and Tr_ω denotes the Dixmier trace. (A precise exposition of this equation is well outside the scope of this presentation. Instead, we refer the interested reader to [12, 39, 1].) Using this formula Connes was able to introduce a generalised noncommutative Riemannian integral for certain suitable types of spectral triples.

One of Connes' principal uses for his noncommutative integral has been to construct noncommutative action functionals. Using one of these, he reconstructed the standard model of particle physics in terms of a distinguished non-classical spectral triple. While Connes' reformulation does not address any questions of renormalisation (see Chapter 5 for a discussion of renormalisation), it does have far greater simplicity than the usual presentation. For further details see [12, 15]. The noncommutative integral has also found use in the study of fractals, see [42].

4.2.2 Connes' State Space Metric

Let $c : [a, b] \rightarrow M$ be a smooth curve in a connected Riemannian manifold M , and let $\dot{c}(t)$ denote the tangent vector to c at $t \in [a, b]$. We can define a nonnegative continuous function $s_c : [a, b] \rightarrow \mathbf{R}$ by

$$s_c(t) = \sqrt{g_{c(t)}(\dot{c}(t), \dot{c}(t))}, \quad t \in [a, b].$$

The nonnegative real number

$$L(c) = \int_a^b s_c dt$$

is called the *arc-length* of c . If c is a piecewise smooth curve, then its *arc length* is defined to be the sum of the arc lengths of its components.

For any two points $p, q \in M$, let us use the symbol $\Omega(p, q)$ to denote the set of all piecewise smooth curves in M from p to q . It is easily shown that $\Omega(p, q)$ is non-empty. The non-negative real number

$$d(p, q) = \inf\{L(c) : c \in \Omega(p, q)\},$$

is called the *Riemannian distance* from p to q . It can be shown that the mapping

$$d : M \times M \rightarrow \mathbf{R}, \quad (p, q) \rightarrow d(p, q),$$

is a metric on M ; it is called the *Riemannian metric*. An important result of Riemannian geometry is that the topology determined by this metric coincides with the original topology of M .

If M is a spin manifold and \mathcal{D} is its associated Dirac operator, then Connes [12] showed that

$$d(p, q) = \sup\{|f(p) - f(q)| : f \in C(M), \|[\mathcal{D}, f]\| \leq 1\}. \quad (4.17)$$

Thus, the Riemannian distance can be recovered from the canonical spectral triple of M .

Let M be a Riemannian spin manifold and let $(C^\infty(M), L^2(S), \mathcal{D})$ be its canonical spectral triple. We can define a metric on $\overline{\Omega(C^\infty(M))}$ by

$$d(\varphi, \psi) = \sup\{|\varphi(f) - \psi(f)| : f \in C(M), \|[\mathcal{D}, f]\| \leq 1\}. \quad (4.18)$$

equation (4.17) and Theorem 1.1.4 imply that M and $\overline{\Omega(C^\infty(M))}$ are identical as metric spaces. (We note that, since the representation of $C^\infty(M)$ on $L^2(S)$ is isometric, the closure of $C^\infty(M)$ with respect to the supremum norm is isometrically isomorphic to the closure of the image of $C^\infty(M)$ in $B(L^2(S))$ with respect to the operator norm. Thus, no ambiguity arises when we speak of the closure of $C^\infty(M)$.)

If (A, H, D) is a spectral triple, then the metric in (4.18) can easily be generalised to a metric on $\overline{\Omega(A)}$. As we noted in Chapter 1, however, $\overline{\Omega(A)}$ is not guaranteed to be non-empty when A is noncommutative. Thus, this metric is of limited interest.

Recall that a positive linear functional of norm 1 on a C^* -algebra \mathcal{A} is called a state. We denote the set of all states on \mathcal{A} by $S(\mathcal{A})$, and we call it the *state space* of \mathcal{A} . The state space of a C^* -algebra is always non-empty. Now since every positive element of a C^* -algebra \mathcal{A} is of the form aa^* , for some $a \in \mathcal{A}$, and every character on \mathcal{A} is Hermitian and multiplicative, it holds that every character is positive. If we also recall that every character is of norm 1, then we see that $\Omega(\mathcal{A}) \subseteq S(\mathcal{A})$. State spaces are of great importance in the study of C^* -algebras and in quantum mechanics.

Now the metric defined above in (4.18) can easily be modified to define a metric on $S(\overline{A})$; it is called *Connes' state space metric*. While Connes did not explore this metric very much for the noncommutative case, Marc Rieffel used it as his starting point when he was developing the theory of compact quantum metric spaces. We shall present the theory of compact quantum metric spaces in detail in Chapter 6.

4.3 Dirac Operators and Quantum Groups

As we saw in Chapter 3, the boundary between cyclic (co)homology and quantum groups is a very active area of research. The same is true of the boundary between spectral triples and quantum groups. At present there is no comprehensive theory linking Connes' calculi to covariant differential calculi. While examples of spectral triples have been constructed over the coordinate algebra of $SU_q(2)$ [9, 10], it is known that the basic examples of calculi over $SU_q(2)$ can not be realized by spectral triples, see [98].

With a view to better understand the relationship between the two theories, a lot of effort has been put into constructing generalised Dirac operators on quantum groups. One approach is to introduce quantum analogues of Clifford and spinor bundles and to define a Dirac operator in terms of them, see [62, 84] for example. Here, however, we shall present the approach pursued by Kustermans, Murphy, and Tuset in [64]. They generalised the Kähler–Dirac operator $d + d^*$ by introducing a generalised codifferential. This generalisation has led to some very interesting results, most notably a generalised version of the Hodge decomposition. We shall briefly discuss why their work led them to call for a generalisation of spectral triples.

4.3.1 The Dirac and Laplace Operators

Let G be a compact quantum group and let (A, Δ) be the Hopf $*$ -algebra underlying it. A left-covariant $*$ -differential calculus over G is a $*$ -differential calculus Ω over A such that there exists an algebra homomorphism $\Delta_\Omega : A \rightarrow A \otimes \Omega$ that satisfies $\Delta_\Omega(a) = \Delta(a)$, for all $a \in A$; and $(\text{id}_A \otimes d)\Delta_\Omega = \Delta_\Omega d$. Now the multiplication of $A \otimes \Omega$ respects the grading given by $A \otimes \Omega = \bigoplus_{n=0}^{\infty} A \otimes \Omega^n$. Hence, $A \otimes \Omega$ can be considered as a graded algebra. Moreover, since

$$\Delta_\Omega(a_0 da_1 \cdots da_n) = \Delta(a_0)[(1 \otimes d)\Delta(a_1)] \cdots [(1 \otimes d)\Delta(a_n)] \in A \otimes \Omega^n,$$

we have that Δ_Ω is a mapping of degree zero with respect to this grading.

We say that $\omega \in \Omega$ is *left-invariant* if $\Delta(\omega) = 1 \otimes \omega$. Clearly, the set of all left-invariant elements forms a subalgebra of Ω ; we denote it by Ω_{inv} . It is obvious that $\Omega_{\text{inv}}^k = \Omega^k \cap \Omega_{\text{inv}}$ is also a subalgebra. We say that a differential calculus over a compact quantum group is *strongly finite-dimensional* if Ω_{inv} is finite-dimensional as a linear space. Using results from general Hopf algebra theory, it can be shown that every strongly finite-dimensional calculus is finite-dimensional as a graded algebra. It is easy to show that in the classical case left-invariant elements correspond to left-invariant differential forms, as defined in the next chapter. As is

well known and easily seen, the algebra of differential forms of every Lie group is strongly finite-dimensional.

Now if $\omega = \sum_{k=1}^n \omega_k$ is a left-invariant element of a strongly finite-dimensional calculus Ω , and each $\omega_k \in \Omega^k$, then we have that $\sum_{k=1}^n 1 \otimes \omega_k = \sum_{k=1}^n \Delta(\omega_k)$. Since Δ_Ω is a mapping of degree 0, it follows that $\Delta(\omega_k) = 1 \otimes \omega_k$, for $k = 1, \dots, n$. Therefore, $\Omega_{\text{inv}} = \bigoplus_{k=0}^{\infty} \Omega_{\text{inv}}^k$. Furthermore, the fact that $(\text{id}_A \otimes d)\Delta_\Omega = \Delta_\Omega d$, implies that Ω_{inv} is invariant under the action of d . Hence, $(\Omega_{\text{inv}}, \partial)$ is a graded differential algebra, where ∂ denotes the restriction of d to Ω_{inv} .

Using results from general Hopf algebra theory again, it can be shown that $A \otimes \Omega_{\text{inv}}$ and Ω are isomorphic as left A -modules. An isomorphism is provided by the unique mapping that sends $a \otimes \omega$ to $a\omega$. Clearly, we also have that $A \otimes \Omega_{\text{inv}}^k \simeq \Omega^k$.

Inner products on Ω will play an important part in our work. We call an inner product on Ω *graded* if the subspaces Ω_k are orthogonal with respect to it. We call an inner product $\langle \cdot, \cdot \rangle$ *left-invariant* if $\langle a\omega_1, b\omega_2 \rangle = h(ba^*) \langle \omega_1, \omega_2 \rangle$, for all $a, b \in A$, and $\omega_1, \omega_2 \in \Omega$; where h denotes the restriction of the Haar integral of G to A . We shall only consider graded left-invariant inner products here. Clearly, the inner product defined in equations (4.13) and (4.14) is graded and left invariant.

We shall now introduce some useful linear operators. If χ is a linear functional on the Hopf algebra A , and $a \in A$, then we write $\chi * a$ for $(\text{id} \otimes \chi)\Delta(a)$; and we denote by E_χ the linear operator defined by setting $E_\chi(a) = \chi * a$. We define χ^* , the *adjoint* of χ , by setting $\chi^*(a) = \overline{\chi(S(a)^*)}$; where S is the antipode of A . And, if $\omega \in \Omega_{\text{inv}}$, then we denote by M_ω the linear operator on Ω_{inv} defined by setting $M_\omega(\eta) = \omega\eta$, $\eta \in \Omega_{\text{inv}}$.

Now if Ω is endowed with an inner product, then it can be shown (see [64] and references therein) that E_χ is adjointable with respect to this inner product and that $E_\chi^* = E_{\chi^*}$. It can also be shown that if $\omega_1, \dots, \omega_m$ is an orthonormal basis for Ω_{inv}^1 , then there exist unique linear functionals χ_1, \dots, χ_m on A such that

$$da = \sum_{r=1}^m E_{\chi_r}(a)\omega_r, \quad (4.19)$$

for all $a \in A$, (see [61] for details).

Theorem 4.3.1 *Let (Ω, d) be a strongly finite-dimensional $*$ -differential calculus over a compact quantum group G , and let (A, Δ) be the Hopf $*$ -algebra underlying G . If Ω is endowed with an inner product, then d is adjointable, and d^* is of degree -1 . Indeed, if $\omega_1, \dots, \omega_m$ and χ_1, \dots, χ_m are as in equation (4.19), then*

$$d = \text{id}_A \otimes \partial + \sum_{j=1}^m E_{\chi_j} \otimes M_{\omega_j} \quad (4.20)$$

and

$$d^* = \text{id}_A \otimes \partial^* + \sum_{j=1}^m E_{\chi_j^*} \otimes M_{\omega_j}^*. \quad (4.21)$$

Proof. Recall that we can identify Ω and $A \otimes \Omega_{\text{inv}}$ by identifying $a\omega$ with $a \otimes \omega$, for $a \in A$, $\omega \in \Omega_{\text{inv}}$. If $a \in A$, and $\omega \in \Omega_{\text{inv}}$, then since

$$d(a\omega) = (da)\omega + ad\omega = \sum_{j=1}^m E_{\chi_j}(a)\omega_j\omega + ad\omega,$$

we have that $d = \text{id}_A \otimes \partial + \sum_{j=1}^m E_{\chi_j} \otimes M_{\omega_j}$. As remarked above, the operators E_{χ_j} are adjointable; since Ω_{inv} is finite-dimensional, M_{ω_j} and ∂ are also adjointable. The adjointability of d and the formula for d^* in the statement of the theorem follow immediately. Since d is of degree 1 and the inner product is graded, it is easily verified that $d^*(a) = 0$, for all $a \in A$, and if $\omega_k \in \Omega^k$ and $k \geq 1$, then $d^*(\omega_k) \in \Omega^{k-1}$.

The operator d^* is called the *codifferential* of d ; the sum $D = d + d^*$ is called the *Dirac operator*; and the square $\nabla = (d + d^*)^2$ is called the *Laplacian*. Since $d^2 = 0$ implies that $d^{*2} = 0$, we have that $\nabla = dd^* + d^*d$. We call $\omega \in \Omega$ a *harmonic form* if $\nabla(\omega) = 0$; we denote the linear space of harmonic forms by Ω_{∇} . Clearly, these operators generalise the classical Kähler–Dirac and Hodge–Laplacian operators (or, more correctly, it generalises its restriction to $\Omega_{\text{Pol}}(M)$).

4.3.2 The Hodge Decomposition

Now that we have defined generalised Dirac and Laplacian operators, we are ready to present the following important theorem. It contains a generalisation of the Hodge decomposition, which is surprising, since such a decomposition is known not to exist in the setting of general noncommutative geometry. (Recall that an operator T on a linear space V is *diagonalisable* if V admits a Hamel basis consisting of eigenvectors of T .)

Theorem 4.3.2 (Hodge decomposition [64]) *Let (Ω, d) be a strongly finite-dimensional $*$ -differential calculus over a compact quantum group G and suppose that we have a graded, left invariant, inner product on Ω . Then*

1. *The Dirac operator D and the Laplacian ∇ are diagonalisable;*
2. *The space Ω admits the orthogonal decomposition*

$$\Omega = \Omega_{\nabla} \oplus d(\Omega) \oplus d^*(\Omega).$$

4.3.3 The Hodge Operator

Let (Ω, d) be a strongly finite-dimensional differential calculus over a compact quantum group G , and assume that (Ω, d) is endowed with a graded, left-invariant, inner product. As we remarked earlier, the fact that (Ω, d) is strongly finite-dimensional implies that (Ω, d) is finite-dimensional; let us assume that it is of dimension n . It is straightforward to show that Ω_{inv}^n contains a self-adjoint element θ of norm 1. If $\dim(\Omega_{\text{inv}}^n) = 1$, and we shall always that it is, then this element is unique up to a change of sign. (The assumption that $\dim(\Omega_{\text{inv}}^n) = 1$ is motivated by the fact it holds in the classical case.) Indeed, since $A\Omega_{\text{inv}}^n = \Omega^n$, it holds that $\{a\theta : a \in A\} = \Omega^n$. Let us define a linear functional \int on Ω^n by setting $\int \omega = h(a)$, when $\omega = a\theta$, for some $a \in A$. We call \int the *integral* associated to θ . Now we say that a calculus (Ω', d') is *non-degenerate* if, whenever $k = 0, \dots, n$ and $\omega \in \Omega'_k$, and $\omega'\omega = 0$, for all $\omega' \in \Omega'_{n-k}$, then necessarily $\omega = 0$. If we further assume that (Ω, d) is non-degenerate and that $d(\Omega_{\text{inv}}^{n-1}) = \{0\}$, then it can be shown that \int is a closed twisted graded trace, for details see [65]. (Again, these assumptions are motivated by the classical case.)

A proof of the following important theorem can be found in [64].

Theorem 4.3.3 *Suppose that (Ω, d) is a non-degenerate. Then there exists a unique left A -linear operator L on Ω such that $L(\Omega^k) = \Omega^{n-k}$, for $k = 0, \dots, n$, and such that $\int \omega^* L(\omega') = \langle \omega', \omega \rangle$, for all $\omega, \omega' \in \Omega$. Moreover, L is bijective.*

The operator L is called the *Hodge operator* on Ω ; we see that it generalises the classical Hodge operator defined in Section 4.1.3. Moreover, since $L(\Omega_{\text{inv}}^k) = \Omega_{\text{inv}}^{n-k}$, for $k = 0, \dots, n$, we have $\dim(\Omega_{\text{inv}}^k) = \dim(\Omega_{\text{inv}}^{n-k})$.

The following pleasing result for the codifferential generalises equation (4.15) to the quantum setting.

Corollary 4.3.4 *Suppose (Ω, d) is non-degenerate and $d(\Omega_{n-1}^{\text{inv}}) = \{0\}$. Then, if $k = 0, \dots, n$ and $\omega \in \Omega_k$, we have*

$$d^*\omega = (-1)^k L^{-1} dL(\omega).$$

Proof. If $k = 0$, then clearly $d^*\omega = 0$. Also, $dL(\omega) = 0$, since $L(\omega)$ is in Ω_n and $d(\Omega_n) = 0$. Hence, $d^*\omega = (-1)^k L^{-1} dL(\omega)$ in this case.

Suppose now that $k > 0$. Clearly, $L^{-1} dL(\omega) \in \Omega_{k-1}$. Hence, if $\omega' \in \Omega_{k-1}$, then

$$\begin{aligned} \langle (-1)^k L^{-1} dL(\omega), \omega' \rangle &= (-1)^k \int \omega'^* dL(\omega) = \int d(\omega')^* L(\omega) \\ &= \langle \omega, d(\omega') \rangle = \langle d^*\omega, \omega' \rangle; \end{aligned}$$

(note that we have used the fact that $d(\Omega_{n-1}^{\text{inv}}) = 0$ implies that \int is closed). Thus, $d^*\omega = (-1)^k L^{-1} dL(\omega)$ in this case also. \square

4.3.4 A Dirac Operator on Woronowicz's Calculus

As we mentioned in Chapter 3, the first ever example of a noncommutative differential calculus was Woronowicz's 3-dimensional calculus over $SU_q(2)$. It is strongly finite-dimensional, and in [64] Kustermans, Murphy, and Tuset calculated the eigenvalues of the Dirac and Laplace operators arising from a canonical choice of inner product for this calculus. From their investigations they feel that it is very unlikely that this Dirac operator fits into the framework of Connes' spectral triples; although they do not have a proof of this fact. This provides an indication that generalisations of spectral triples will have to be studied. Their Dirac operator does, however, have properties that are close to those required to enable it to fit into A. Jaffe's [53] extension of Connes' theory. It may be that Jaffe's theory can, and should, be further developed to cover this example.

A very interesting recent development in this area is Connes and Moscovici's introduction of a twisted version of the spectral triple definition [18]. Let (A, H, D) be a triple consisting of an algebra A , a Hilbert space H upon which A is represented, and a self-adjoint operator D with compact resolvent. If σ is an algebra automorphism of A such that $Da - \sigma(a)D$ is bounded, for all $a \in A$, then we say that (A, H, D) is a σ -spectral triple. Clearly, if $\sigma = \text{id}_A$, then we recover the ordinary spectral triple definition. Connes and Moscovici believe that, since the domain of quantum groups is an arena where twisting frequently occurs, σ -spectral triples could be useful in the study of quantum groups. It is interesting to note that they cite Hadfield and Krähmer's paper [46]. The obvious question to ask is whether or not Kustermann, Murphy, and Tuset's Dirac operator fits into this framework. At present no work has been done in this area.

Chapter 5

Fuzzy Physics

In this chapter we shall give a brief account of some of the work being done by the Dublin node of the European Union Operator Algebras Network: the Dublin Institute for Advanced Studies (DIAS). The members of the DIAS are pursuing research in an area called fuzzy physics. Essentially, fuzzy physics is an application of noncommutative geometry to the problem of UV -divergences in quantum field theory. We shall begin this chapter by introducing quantum field theory, we shall then discuss the role noncommutative geometry plays, and finally we shall discuss some of the DIAS work.

Unfortunately, because of the physical nature of the material, our presentation in this chapter will sometimes be a little vague. Moreover, we shall also assume some prior knowledge of physics.

5.1 Quantum Field Theory

As we saw in Chapter 4, Dirac's relativistic equation of the electron introduces into physics the notion of antimatter. The key property of antimatter is that a particle and its antiparticle can come together and annihilate one another, their combined mass being converted into energy in accordance with Einstein's equation: $E = mc^2$. Conversely, if sufficient energy is introduced into a system, localised in a suitably small region, then there arises the strong possibility that this energy might serve to create some particle together with its antiparticle. This does not violate any conservation laws since the conserved quantum numbers of a particle and its antiparticle have opposite signs. For example, for an electron-positron pair, the electron has an electric charge of -1 and the positron has an electric charge of 1 , thus, the addition of the pair has no effect on the total charge of the system. All this means that the number of particles in a system is always variable.

However, ordinary non-relativistic quantum mechanics does not allow us to describe systems in which the number of particles is variable. In order to cope with the introduction of anti-matter a new framework was required. The theory that emerged was called *quantum field theory*. There exist a number of different approaches to the subject. Any reader interested in the details of quantum field theory can find a good, well referenced, non-technical introduction in [85].

The laws of classical mechanics were formulated by Isaac Newton in 1680's. In the centuries that followed two important reformulations emerged: one due to Joseph Louis Lagrange, and another due to William Rowan Hamilton. The standard approach to dynamics in quantum mechanics is based on Hamilton's reformulation. In the 1940's the brilliant American physicist Richard Feynman introduced an approach that is based on Lagrange's reformulation. It is based on the idea that if a particle travels between two points, then, in a certain 'quantum mechanical sense', it must travel every path between those two points. The basic structure in Feynman's approach is the Feynman path integral. Feynman used his ideas to formulate an approach to quantum field theory. At present, it is arguably the most widely used framework.

Soon after its introduction, however, it became clear that Feynman's approach suffers from serious conceptual and technical difficulties. From a mathematical point of view the whole construction is horribly ill-defined. From a practical point of view physical processes occurring at arbitrarily small distances cause the theory to give infinite answers to questions that should have a finite answers; these are the so called *UV-divergences*. In the late 1940's, the ground breaking work of Feynman, Tomonaga, Dyson, Swinger and others on these divergence problems produced successful, if somewhat ad hoc, methods for extracting information from the theory. Their approach is generally known as *renormalisation*. When renormalisation was applied to systems involving the electromagnetic field, the result was quantum electrodynamics (QED). Astonishingly, the predictions of QED match experimental data to a level of accuracy never seen before. QED has been described as the 'most successful scientific theory ever'. Later, a quantum field theory that united the electromagnetic, strong nuclear, and weak nuclear fields would emerge; it is called the standard model. It is regarded as a high point in twentieth century physics. However, despite all efforts, gravity (the last remaining fundamental field) has not been incorporated into this framework.

5.2 Noncommutative Regularization

Despite the experimental successes of quantum field theory, it is still wholly unsatisfactory from a mathematical point of view. Moreover, renormalization (which

is in essence a perturbation theory) is not a ‘fundamental’ process (even Feynman himself was of this opinion). The fact that physics has failed to incorporate gravity into the framework of the standard model seems to confirm that fundamental changes need to be made. Study in this area is known as *nonperturbative regularization*.

At present, the conventional approach to nonperturbative regularization is *lattice field theory*: it works by replacing a continuous manifold with a discrete lattice. In this context Feynman’s formulation can be well defined mathematically, and the divergences that arise from processes occurring at arbitrarily small distances disappear. There is, however, one feature that must be criticized: lattices do not retain the symmetries of the original theory, except in some rough sense. This is a very serious matter because questions of symmetry lie at the heart of modern physics. A related feature is that the topology and differential geometry of the underlying manifold are only treated indirectly; the lattice points are generally manipulated like a trivial topological set. There do exist radical attempts to overcome these limitations using partially ordered sets [35], but their potentials have yet to be fully realised.

5.2.1 The Fuzzy Sphere

In the early 1990’s John Madore began to propose that noncommutative geometry be used in nonperturbative regularisation. (Madore was not the first to consider such an idea; as early as the 1940’s Synder made a suggestion that space-time coordinates should be noncommutative.) Madore showed that to every compact coadjoint orbit (a distinguished type of manifold) one could associate a canonical sequence of noncommutative algebras, each of which retains the symmetries of the original space. As n goes to infinity, the algebras approach the algebra of continuous functions of the space (in a certain loose sense). Since each of these noncommutative algebras is finite-dimensional, it was hoped that a well-defined version of quantum field theory could be expressed in terms of them. Moreover, it was also hoped that the resulting theory would not suffer from UV -divergences.

Madore’s construction is not an entirely mathematical exercise. His approach is also physically motivated. We quote Hawkins [50]:

‘The observation of structures at very small distances requires radiation of very short wavelength and correspondingly large energy. Attempting to observe a sufficiently small structure would thus require such a high concentration of energy that a black hole would be formed and no observation could be made. If this is so, then distances below about the Planck scale are unobservable and thus operationally meaningless. If short distances are meaningless, then perhaps precise locations are as well. This suggests the possibility of uncertainty relations between position and

position, analogous to the standard ones between position and momentum. An uncertainty relation between, say, x -position and y -position, would mean that the x and y coordinates do not commute (see [24]).'

We shall begin by outlining Madore's 'fuzzification' of the 2-sphere. Let us recall that S^2 , the 2-sphere, is the submanifold of \mathbf{R}^3 consisting of all points (x_1, x_2, x_3) for which

$$x_1^2 + x_2^2 + x_3^2 = 1. \quad (5.1)$$

It is felt that if a nonperturbative regularized version of quantum field theory could be constructed on S^2 , then a regularized version of quantum field theory on ordinary space-time would follow.

Let us define three functions \hat{x}_i , $i = 1, 2, 3$, on S^2 by setting $\hat{x}_i(x_1, x_2, x_3) = x_i$, for $(x_1, x_2, x_3) \in S^2$. We shall denote by $\text{Cor}(S^2)$ the smallest subalgebra of $C(S^2)$ containing the functions \hat{x}_i . Clearly, $\text{Cor}(S^2)$ is dense in $C(M)$, and each element of $\text{Cor}(S^2)$ is of the form

$$z_0 + \sum_{i=1}^3 z_i \hat{x}_i + \sum_{1 \leq i < j \leq 3} z_{ij} \hat{x}_i \hat{x}_j + \sum_{1 \leq i < j < k \leq 3} z_{ijk} \hat{x}_i \hat{x}_j \hat{x}_k + \dots$$

Madore's method uses truncations of the above sum to construct a sequence of noncommutative algebras approximating $\text{Cor}(S^2)$.

If we discard all functions with nonconstant terms, then $\text{Cor}(S^2)$ reduces to \mathbf{C} ; this is our first approximation and we denote it by S_1^2 .

If we add to S_1^2 all terms linear in \hat{x}_i , then we get a four dimensional linear space; we shall denote it by S_2^2 . Consider the unique isomorphism $\phi_2 : S_2^2 \rightarrow M_2(\mathbf{C})$ for which $\phi_2(1) = 1$ and

$$\phi_2(\hat{x}_i) = \frac{1}{\sqrt{3}} \sigma_i;$$

where σ_i are the Pauli spin matrices. We can use ϕ_2 to give S_2^2 a noncommutative algebra structure in an obvious way. Note that equation (5.1) holds in this context.

We shall now construct S_3^2 . To begin with, we include all the elements of S_2^2 , and all terms of the form $z_{ij} \hat{x}_i \hat{x}_j$, for $i < j$. Since we no longer intend to use the multiplication of $C(S^2)$, we shall also include terms of the form $z_{ij} \hat{x}_i \hat{x}_j$, for $j < i$, where $\hat{x}_i \hat{x}_j \neq \hat{x}_j \hat{x}_i$. It will also be assumed that equation (5.1) holds. We can give S_3^2 the structure of a noncommutative algebra (that is consistent with our assumptions) in much the same way as we did for S_2^2 ; however, we shall need to present some facts about the representation theory of $\mathfrak{su}(2)$ first.

Let $\mathfrak{su}(2)$ be the Lie algebra of $SU(2)$ ($\mathfrak{su}(2)$ is a three dimensional algebra canonically associated to $SU(2)$, in a sense that we shall explain below). As is well known, for every positive integer n , there is a unique n -dimensional representation U_n of $\mathfrak{su}(2)$. As is also well known, it is possible to choose a self-adjoint basis $\{J_i\}_{i=1}^3$ of $U_n(\mathfrak{su}(2))$ such that $J_i^2 = 1$, and

$$[J_j, J_k] = 2i\epsilon_{jkl}J_l, \quad 1 \leq j, k, l \leq 3; \quad (5.2)$$

where ϵ is the Levi-Civita permutation symbol. (For $n = 2$, the set of Pauli spin matrices is such a basis.) It can be shown that every element of $M_n(\mathbf{C})$ is expressible as a polynomial in J_1, J_2, J_3 .

Let $\phi_3 : S_3^2 \rightarrow M_3(\mathbf{C})$ be the unique linear space isomorphism for which $\phi_3(1) = 1$ and

$$\phi_3(\hat{x}_i) = \frac{1}{\sqrt{8}}J_i.$$

Obviously, we can use ϕ_3 to give S_3^2 the structure of a noncommutative algebra. Note this new structure is consistent with our earlier assumptions about the linear structure of S_3^2 .

For $n > 3$, it is now clear how to define the linear space S_n^2 . As before, we assume that $\hat{x}_i\hat{x}_j \neq \hat{x}_j\hat{x}_i$, and that equation (5.1) holds. We define a noncommutative multiplication for S_n^2 using the unique isomorphism $\phi_n : S_n^2 \rightarrow M_n(\mathbf{C})$ for which $\phi_n(1) = 1$ and

$$\phi_n(\hat{x}_i) = \frac{1}{\sqrt{n^2 - 1}}J_i.$$

Note that this new structure is again consistent with the linear space structure of S_n^2 . We call S_n^2 the *n-fuzzy sphere*.

The fuzzy spheres are intuitively thought of as ‘noncommutative lattice approximations to S^2 ’. Unlike the lattice approximations, however, each $M_n(\mathbf{C})$ retains the symmetries of the sphere: The symmetry group of S^2 is clearly $SU(2)$. As is well known, there exists an irreducible representation of $SU(2)$ on \mathbf{C}^n , for each $n > 0$. Thus, the primary shortcoming of the lattice approach is rectified.

The dimension of the algebra S_n^2 is thought of as the number of points in the noncommutative lattice approximation. (Note that in the commutative case, the dimension of the function algebra of the lattice is equal to the number of points in the lattice.) Thus, as $n \rightarrow \infty$, we think of the noncommutative lattice as ‘approaching’ the continuum. Furthermore, it is easy to see that equation (5.2) implies that

$$[\hat{x}_i, \hat{x}_j] = \frac{2i}{\sqrt{n^2 - 1}}\epsilon_{ijk}\hat{x}_k, \quad \hat{x}_i, \hat{x}_j \in S_n.$$

This motivates the intuitive statement: ‘In the limit $n \rightarrow \infty$, the algebra S_n^2 becomes commutative’. For these reasons, physicists write that ‘ S_n^2 converges to $\text{Cor}(S^2)$ ’.

It is worth our while to comment briefly on the mathematical form that the fuzzy sphere takes. Its presentation owes a lot to its physical origins. In the quantum theory of rotations, $SU(2)$ and $\mathfrak{su}(2)$ play central roles. As a result, physicists are usually well versed in their properties.

5.2.2 Fuzzy Coadjoint Orbits

As we have presented it above, Madore’s construction of the fuzzy sphere seems somewhat ad hoc. However, it is in fact a special case of a much more general procedure known as Berezin–Toeplitz quantization. This is a method used in quantum mechanics for ‘quantising’ a phase space M that is a Kähler manifold (a Kähler manifold is a special type of symplectic manifold). By *quantising* we mean producing a sequence of finite-dimensional Hilbert spaces H_n , and a sequence of maps $T_n : C(M) \rightarrow B(H_n)$. The maps T_n are called the *Toeplitz maps* for the quantization. The dimension of the Hilbert spaces are governed by the Riemann–Roch formula; see [50] for details. Madore’s insight was that he could use Berezin–Toeplitz quantisation to construct fuzzy versions of spaces.

Coadjoint Orbits

If the Kähler manifold in question is a coadjoint orbit, then Berezin–Toeplitz quantisation takes a simpler form. Let us recall some of the basic theory of coadjoint orbits. Let G be a compact Lie group (we shall assume that all the Lie groups we work with are connected and semi-simple). The map

$$\lambda_g : G \rightarrow G, \quad h \mapsto gh,$$

is clearly a diffeomorphism, for all $g \in G$. Moreover, λ_g induces a map $\lambda_g^* : T_h(G) \rightarrow T_{gh}(G)$ that is defined by setting

$$\lambda_g^*(v_h) = v_h \circ \lambda_g, \quad h \in G. \tag{5.3}$$

If X is a vector field such that $[\lambda_g^*(X)](h) = X(gh)$, for all $g, h \in G$, then we say that X is *left-invariant*. The left-invariant vector fields on a Lie group form an algebra that is closed under the Lie bracket; we call it the *Lie algebra* of G and we denote it by \mathfrak{g} . As a little thought will verify, \mathfrak{g} is canonically isomorphic to the tangent plane at any point of G . It is common practice to equate it with the

tangent plane at e , where e is the identity of G . Thus, \mathfrak{g} is a complex linear space whose dimension is equal to the manifold dimension of G . (It is straightforward to build on the definition of a left-invariant field and define the notion of a left-invariant differential form, as referred to in the previous chapter.)

Let us consider the *conjugation map*

$$I_g : G \rightarrow G, \quad h \mapsto ghg^{-1}, \quad g \in G.$$

In direct analogy with equation (5.3), I_g induces a map $I_g^* : T_h(G) \rightarrow T_{I_g(h)}(G)$, for all $h \in G$. Since $I_g(e) = e$, I_g^* is a linear mapping on $T_e(G)$. The fact that we equated \mathfrak{g} and $T_e(G)$ means that we can also regard I_g^* as a linear map on \mathfrak{g} . A routine calculation will now verify that

$$\text{Ad} : G \rightarrow \text{End}(\mathfrak{g}), \quad g \mapsto I_g^*,$$

is a representation of G ; we call it the *adjoint representation* of G . The *coadjoint representation* of G is the mapping

$$\text{Ad}^* : G \rightarrow \text{End}(\mathfrak{g}^*), \quad g \mapsto \text{Ad}_g^*;$$

where Ad_g^* is the linear dual of Ad_g , that is, it is the unique mapping for which

$$[\text{Ad}_g^*(\mu)](\zeta) = \mu(\text{Ad}_g(\zeta)), \quad \zeta \in \mathfrak{g}, \mu \in \mathfrak{g}^*.$$

We define the *coadjoint orbit* of $\mu \in \mathfrak{g}^*$ to be the subset

$$\mathcal{O}_\mu = \{\text{Ad}_g^*(\mu) : g \in G\}.$$

Consider the subgroup $H = \{h \in G : \text{Ad}^*(\mu) = \mu\}$. It can be shown that \mathcal{O}_μ is homeomorphic to G/H , and that G/H is a homogeneous space for G (recall that a *homogeneous space* is a topological space with a transitive group action).

Quantized Coadjoint Orbits

Each coadjoint orbit $\mathcal{O}_\mu \simeq G/H$ comes naturally equipped with a symplectic form. One can use this form to endow \mathcal{O}_μ with the structure of a Kähler manifold. When Berezin–Toeplitz quantization is applied to \mathcal{O}_μ , each Hilbert space H_n produced is the representation space of a unitary irreducible representation of G . (Strictly speaking, we should only be considering *integral* coadjoint orbits, but we shall not need to concern ourselves with such details here. See [49] for further details.) Thus, each *fuzzy coadjoint orbit* $B(H_n)$ retains the homogeneous space symmetries of \mathcal{O}_μ . As n goes to infinity, the fuzzy coadjoint orbits $B(H_n)$ converge to $C(\mathcal{O}_\mu)$ in a loose intuitive sense that is analogous to the fuzzy sphere case.

The 2-sphere is a coadjoint orbit of $SU(2)$. When the Berezin–Toeplitz quantisation is applied to it, one gets the sequence of n -fuzzy spheres described earlier. The Toeplitz maps T_n produced correspond to the unique set of maps $F_n : \text{Cor}(S^2) \rightarrow S_n^2$, for which $F_n(\hat{x}_i) \rightarrow \hat{x}_i$.

Besides the fuzzy sphere, many other examples of fuzzy coadjoint orbits have been explicitly described. For example, fuzzy complex projective spaces [4], fuzzy complex Grassmannian spaces [21], and fuzzy orbifolds [77] have appeared in the papers of DIAS members.

5.2.3 Fuzzy Physics

Now that we have introduced the notion of a fuzzy space, we are ready to (briefly) discuss fuzzy physics. In essence, the subject consists of noncommutative generalisations of physical theories expressed in terms of fuzzy space matrix algebras. Its long term goal is to construct a fuzzy version of the standard model. Towards this end, a lot of work has been put into the construction of fuzzy vector bundles, fuzzy Lagrangians, fuzzy Dirac operators, and fuzzy gauge theories. While a good deal of progress has been made, it seems that a fuzzy version of the full standard model is still some way off. A large portion of this work has focused on the fuzzy physics of the fuzzy sphere; in particular the fuzzification of scalar field theories on S^2 .

An obvious question to ask is if these new theories give finite answers where ordinary quantum field theory fails to do so. Unfortunately, the answer in general is no. At present, it is not clear whether these divergences are due to ‘incorrect’ noncommutative generalisations; or whether there are fundamental problems with the present form of fuzzy physics. It should be noted, however, that there are a large number of examples that *are* naturally regularised; a notable example is [87].

An interesting (albeit troublesome) feature of fuzzy physics is UV – IR mixing. In physics it is usually possible to organize physical phenomena according to the energy scale or distance scale. The short-distance, ultraviolet (UV) physics does not directly affect qualitative features of the long-distance, infrared physics (IR), and vice versa. However, in fuzzy physics interrelations between UV and IR physics start to emerge. This occurrence is known as UV – IR -mixing. Unfortunately, this mixing leads to divergences. In [22] members of the DIAS showed how to overcome this problem in the case of the fuzzy sphere.

In recent years DIAS members Denjoe O’ Connor and Xavier Martin have been using computer based numerical simulations to study the fuzzy physics of the fuzzy sphere S_n^2 . Their work studies the behaviour of certain fuzzy models as $n \rightarrow \infty$. For details see [79, 34].

Recently, an application of the fuzzy sphere to the study of black hole entropy has emerged. It has been proposed that the event horizon of a black hole should be modeled by a fuzzy sphere. A paper on the subject [23] has been written by DIAS member Brian Dolan.

Before we finish our discussion of fuzzy physics, we should note that there exist other ways of applying noncommutative geometry to quantum field theory. As one notable example, we cite Seiberg and Witten's paper [99].

Chapter 6

Compact Quantum Metric Spaces

As we saw in the previous chapter, a number of physicists are now working with noncommutative geometry in the hope that it will provide a means to solve the problems of quantum field theory. In practice, their work involves the construction of ‘field theories’ on algebras of $n \times n$ matrices. As n goes to infinity, the matrix algebras converge (in some intuitive sense) to the algebra of continuous functions on a coadjoint orbit; and the field theories converge to some classical field theory. The prototypical example of this process is the fuzzification of the two sphere.

After reading the fuzzy physics literature, Marc Rieffel suspected that the matrix algebra convergence referred to therein, involved some kind of continuous C^* -field structure. (A C^* -field is a type of bundle construction where each fibre is a C^* -algebra.) However, he later changed his mind. In [93] he wrote:

‘There is much more in play than just the continuous-field aspect. Almost always there are various lengths involved, and the writers are often careful in their bookkeeping with these lengths as n grows. This suggested to me that one is dealing here with metric spaces in some quantum sense, and with the convergence of quantum metric spaces’.

Rieffel set himself the formidable task of formulating a mathematical framework in which he could precisely express these ideas. Previous to this the only major example of metric considerations in the literature of quantum mathematics was Connes’ state space metric, as discussed in Chapter 4. Rieffel used this as his initial point of reference. However, for reasons that we shall explain later, Connes’ metric proves inadequate for the task at hand. So Rieffel decided to base his work on that of Kantorovich instead.

In this chapter we shall begin by presenting compact quantum metric spaces, quantum Gromov–Hausdorff convergence, and some important examples of both. We shall then go on to establish what is arguably the most important achievement

of the theory to date, a formalisation of the convergence of the fuzzy spheres to the sphere. Finally, we shall present some recently proposed modifications to Rieffel's definitions.

The theory of compact quantum metric spaces is still a very young area. Connes' state space metric, which can be considered its starting point, was only introduced in 1989. Rieffel's first papers on the subject emerged in the late nineties, and the subject only began to take definite shape after the year 2000. In fact, some of the more recent developments discussed in this chapter occurred in 2005 and 2006.

With regard to future work, it seems to be Rieffel's intention to expand his theory so that it will be able to define a distance between some appropriately defined noncommutative version of vector bundles. This would then enable one to discuss the convergence of field theories on fuzzy spaces as well as the convergence of the spaces themselves. Two recent papers of Rieffel [95, 96] indicate that work in this area is well under way.

We should also mention that Nik Weaver has formulated his own noncommutative generalisation of metric spaces based on von Neumann algebras and w^* derivations. Weaver's book [107] gives a very good presentation of this theory. While the two subjects have a very different feel to them, there does seem to be some common ground. For a brief discussion of their relationship see [92].

6.1 Compact Quantum Metric Spaces

6.1.1 Noncommutative Metrics and the State Space

Recall that if X is a compact Hausdorff space, then by Theorem 1.1.4 it is homeomorphic to $\Omega(C(X))$ endowed with the weak* topology. Thus, metrics that metrize the topology of X are in one-to-one correspondence with metrics that induce the weak* topology on $\Omega(C(X))$. As a result a natural candidate for a noncommutative metric space would be a pair (\mathcal{A}, ρ) , where \mathcal{A} is a C^* -algebra and ρ is a metric on $\Omega(\mathcal{A})$. However, as we showed earlier, the character space of a noncommutative C^* -algebra is not guaranteed to be non-empty. Thus, any generalisation based upon it is unlikely to be very fruitful. A better proposal might be to examine metrics that induce the weak* topology on some non-empty set of functionals that is equal to $\Omega(\mathcal{A})$ in the commutative case and properly contains it in the noncommutative case.

Recall that a state μ on a C^* -algebra is a positive linear functional of norm 1, and that $\Omega(\mathcal{A}) \subseteq S(\mathcal{A})$. Unlike $\Omega(\mathcal{A})$, $S(\mathcal{A})$ is always non-empty; it is well known that for every normal element $a \in \mathcal{A}$, there exists a state μ such that $|\mu(a)| = \|a\|$.

Using the Banach–Alaoglu Theorem $S(\mathcal{A})$ can easily be shown to be compact with respect to the weak* topology. Furthermore, it is quite easy to show that $S(\mathcal{A})$ is a convex subset of \mathcal{A}^* , and that when \mathcal{A} is commutative, $\Omega(\mathcal{A})$ is equal to the set of extreme points of $S(\mathcal{A})$; or, as they are better known, the *pure-states* of \mathcal{A} . The Krein–Milman Theorem implies that it is also non-empty in the noncommutative case. Thus, the set of pure-states of a C^* -algebra seems to be the set of functionals we are looking for. Unfortunately, however, this set is quite badly-behaved with respect to our needs; and, even though the generalisation is no longer strict, it turns out to be much more profitable to use the entire state space of a C^* -algebra instead. As we shall see, the compactness of $S(\mathcal{A})$ makes it well suited to our needs.

6.1.2 Lipschitz Seminorms

To restate, if \mathcal{A} is the C^* -algebra of continuous functions on some compact Hausdorff space X , then there is a canonical correspondence between metrics that metrize the topology of X and metrics that induce the weak* topology on $\Omega(\mathcal{A})$. We have no such correspondence for metrics on $S(\mathcal{A})$. It seems reasonable to assume that a study of all metrics that induce the weak* topology on $S(\mathcal{A})$ would be too broad. However, it is not clear what subfamily of metrics we should restrict our attention to.

Recall that for a spectral triple (A, H, D) Connes defined a metric on $S(\overline{A})$ by

$$\rho_D(\mu, \nu) = \sup\{|\mu(a) - \nu(a)| : \|[D, a]\| \leq 1\}.$$

In general, it proves quite difficult to say when Connes’ metric induces the weak* topology on $S(\overline{A})$. This makes ρ_D unsuitable for our later work; not only because of Theorem 1.1.4, but also because the weak* topology is needed for Rieffel’s definition of quantum Gromov–Hausdorff distance to work smoothly.

Connes’ work on noncommutative metrics was in part based upon Kantorovich’s work [55] with Lipschitz seminorms. Thus, a different generalisation of Kantorovich’s work might identify a ‘suitable’ subfamily of state space metrics.

Kantorovich’s Work

For a compact metric space (X, ρ) , the *Lipschitz seminorm* is a commonly used seminorm on $\mathcal{A} = C(X)$ defined by

$$L_\rho(f) = \sup\{|f(x) - f(y)|/\rho(x, y) : x \neq y\}.$$

Finiteness of the seminorm is not guaranteed, but the subset $\{f : L_\rho(f) < \infty\}$ is always dense in $C(X)$. (Note that in this chapter we shall allow seminorms and

metrics to take infinite values.) Just as we recovered the geodesic distance using ρ_D , the well known result

$$\rho(x, y) = \sup\{|f(x) - f(y)|/\rho(x, y) : L_\rho(f) \leq 1\}, \quad (6.1)$$

allows us to recover the metric using L_ρ . We now have a reformulation of the metric space data in terms of the commutative C^* -algebra $C(X)$ and the seminorm L_ρ . (If $(C^\infty(M), L^2(M, S), \mathcal{D})$ is the canonical spectral triple on a compact Riemannian spin manifold, then it holds that $\|[\mathcal{D}, f]\| = L_\rho(f)$, where ρ is the geodesic distance of M . Thus, in this case equation (6.1) and equation (4.17) coincide. However, we should note that equation (6.1) is true for any metric space X , and so, it is a much more general result.)

Anticipating Connes, Kantorovich defined a metric on the state space of $C(X)$ by

$$\rho_L(\mu, \nu) = \sup\{|\mu(f) - \nu(f)| : L_\rho(f) \leq 1\}. \quad (6.2)$$

He showed that, amongst other properties, the topology induced by ρ_L on $S(C(X))$ always coincides with the weak* topology.

All of this suggested to Rieffel that a noncommutative metric space should consist of a noncommutative unital C^* -algebra \mathcal{A} endowed with a suitable seminorm. A metric ρ_L could then be defined on the state space by the obvious generalisation of equation (6.2) to the noncommutative case.

6.1.3 Order-unit Spaces

One might now assume that we were ready to give a definition of a noncommutative metric space. However, our proposed formulation must undergo a simplification.

Let \mathcal{A} be a C^* -algebra and let L be a seminorm on \mathcal{A} . If L is to be considered as a generalised Lipschitz seminorm, then it seems reasonable to assume that, just as in the classical case, $L(a^*) = L(a)$, for all $a \in \mathcal{A}$. Let us define a possibly infinite-valued metric on $S(\mathcal{A})$ using the obvious generalisation of equation (6.2), and let $\mu, \nu \in S(\mathcal{A})$, and $\delta > 0$ be given. Then there exists an $a \in \mathcal{A}$ such that $L(a) \leq 1$ and

$$\rho_L(\mu, \nu) - \delta \leq |\mu(a) - \nu(a)| = \mu(a) - \nu(a).$$

Define an element of \mathcal{A}_{sa} by $b = (a + a^*)/2$ and note that since $L(a^*) = L(a)$, $L(b) \leq 1$. Now it is well known that all states are Hermitian (see [80] for a proof). Therefore

$$|\mu(b) - \nu(b)| = \Re(\mu(a) - \nu(a)) = \mu(a) - \nu(a) \geq \rho_L(\mu, \nu) - \delta.$$

This, means that when we are calculating the value of ρ_L it suffices to take the supremum over the self-adjoint elements of \mathcal{A} . This fact seems to suggest that self-adjoint elements have a distinguished role to play in any formulation of non-commutative metric spaces. In fact, it prompted Rieffel to suggest that order-unit spaces should play a part.

Definition 6.1.1. An *order-unit space* is a pair (A, e) , where A is a real partially-ordered linear space, and $e \in A$ is a distinguished element called the *order-unit*, such that;

1. For each $a \in A$, there is an $r \in \mathbf{R}$ such that $a \leq re$ (*order-unit property*).
2. If $a \in A$ and if $a \leq re$, for all $r \in \mathbf{R}$ with $r > 0$, then $a \leq 0$ (*Archimedean property*).

For any C^* -algebra \mathcal{A} , \mathcal{A}_{sa} is an order-unit space with respect to the partial order defined by setting $a \geq b$ if $(a - b) \geq 0$, $a, b \in \mathcal{A}$. We shall take some care to establish this. For any $a \in \mathcal{A}_{sa}$, note that $C^*(a, 1)$, the smallest C^* -subalgebra of \mathcal{A} containing a and 1, is a commutative unital C^* -algebra. By the Gelfand–Naimark Theorem $C^*(a, 1) \cong C(X)$, for some compact Hausdorff space X . Using this fact, it can easily be shown that \mathcal{A}_{sa} satisfies properties 1 and 2. Hence, it is indeed an order-unit space. It is interesting to note that just as any C^* -algebra can be concretely realised as a norm closed $*$ -subalgebra of $B(H)$, for some Hilbert space H , any order-unit space can be concretely realised as a real linear subspace of the self-adjoint operators on a Hilbert space.

The *standard norm* on an order-unit space is defined by

$$\|a\| = \inf\{r \in \mathbf{R} : -re \leq a \leq re\}.$$

By returning to the fact that $C^*(a, 1) \cong C(X)$, for any $a \in \mathcal{A}_{sa}$, we can establish that the order-unit norm on \mathcal{A}_{sa} coincides with the restriction of the C^* -norm to \mathcal{A}_{sa} .

A *state* on an order-unit space (A, e) is a bounded linear functional μ on A such that $\mu(e) = 1 = \|\mu\|$ (where $\|\mu\| = \sup\{\mu(a) : \|a\| \leq 1, a \in A\}$). It can be shown that all states are automatically positive. We denote the set of all states on A by $S(A)$ and call it the *state-space* of A . Just as for C^* -algebras, we can show that $S(A)$ is a weak* closed subset of the unit ball of A^* . Thus, by the Banach–Alaoglu theorem, it is compact with respect to the weak* topology on A^* . Again, just as in the C^* -algebra case, $S(A)$ is a convex subset of A^* .

It is well known [80] that a bounded linear functional μ on a unital C^* -algebra is positive if, and only if, $\mu(1) = \|\mu\|$. Thus, if \mathcal{A} is unital and $\mu \in$

\mathcal{A}^* , then μ is a state if, and only if, $1 = \mu(1) = \|\mu\|$. As a little thought will verify, this implies that the restriction of a state on \mathcal{A} , to \mathcal{A}_{sa} , will be a state in the order-unit sense, and conversely, that every state on \mathcal{A}_{sa} has a unique extension to a C^* -algebra state on \mathcal{A} . Thus, there is a one-to-one correspondence between the elements of $S(\mathcal{A})$ and $S(\mathcal{A}_{sa})$. If we endow both spaces with their respective weak* topologies, then it is a simple exercise to show that they are homeomorphic. As we saw above any seminorm L on \mathcal{A} satisfying $L(a^*) = L(a)$ will induce the same metric on $S(\mathcal{A})$ as its restriction will induce on $S(\mathcal{A}_{sa})$. Thus, two C^* -algebras whose collection of self-adjoint elements are isomorphic as order-unit spaces will, loosely speaking, produce the same ‘metric data’. When we also take into account the greater technical flexibility afforded by working with order-unit spaces, it seems that they are the natural structure upon which to base a noncommutative metric theory.

6.1.4 Compact Quantum Metric Spaces

Now that we have settled on the category that we shall be working in, we are ready to formulate a generalisation of the Lipschitz seminorm.

Definition 6.1.2. Let A be an order-unit space and let L be a seminorm on A taking finite values on a dense order-unit subspace of A . We say that L is a *Lip-norm* if the following conditions hold:

1. $L(a) = 0$ if, and only if, $a \in \mathbf{Re}$.
2. The topology on $S(A)$ induced by the metric

$$\rho_L(\mu, \nu) = \sup\{|\mu(a) - \nu(a)| : L(a) \leq 1\}, \quad \mu, \nu \in S(\mathcal{A}),$$

coincides with the weak* topology.

The first requirement on L is a direct generalisation of the fact that $L_\rho(f) = 0$, if, and only if, f is a constant function, while the second requirement has already been motivated. As might be expected, Lip-norms do not directly generalise Lipschitz seminorms; that is, if X is a compact Hausdorff space, then there can exist Lip-norms on $C(X)$ that are not the Lipschitz seminorm for any metric on X .

We now are ready to define our noncommutative generalisation of compact metric spaces.

Definition 6.1.3. A *compact quantum metric space* is a pair (A, L) where A is an order-unit space and L is a Lip-norm on A .

Rieffel uses the term quantum metric space instead of noncommutative metric space because there is no multiplication in an order-unit space, and so, there is no noncommutativity to speak of. He also cites the central role that states play in quantum mechanics as a motivation.

The metric ρ_L only takes finite values. To see this assume that $\rho_L(\mu_0, \nu_0) = +\infty$, for some $\mu_0, \nu_0 \in S(A)$. The proper subset $\{\mu : \rho_L(\mu, \nu_0) < \infty\}$ is both open and closed, which is impossible since the convexity of $S(A)$ implies that it is connected with respect to the weak* topology. Since $S(A)$ is compact, ρ_L is bounded, and so we can speak of the radius of the metric space $S(A)$. We define the *radius* of a compact quantum metric space to be the radius of its state space.

In naturally occurring examples it can often be quite difficult to verify directly whether or not a metric induces the weak* topology. Fortunately, however, Rieffel [91] managed to reformulate this property in simpler terms.

Theorem 6.1.4 *Let L be a seminorm on an order-unit space A , such that $L(a) = 0$ if, and only if, $a \in \mathbf{Re}$, and let $B_1 = \{a \in A : L(a) \leq 1\}$. Then ρ_L induces the weak* topology on $S(A)$ exactly if*

1. (A, L) has finite radius, and
2. B_1 is totally bounded in A .

An analogous reformulation is not known to exist for Connes' state space metric.

It is interesting to note that recently Frederic Latrémolière [68], a former doctoral student of Rieffel, has shown how to extend the definition of a compact quantum metric space to one that generalises locally compact metric spaces.

6.1.5 Examples

Let (X, ρ) be a compact metric space and let L_ρ be the Lipschitz seminorm on $C(X)$. Let us also use L_ρ to denote the restriction of L_ρ to $C(X)_{sa}$, the space of real-valued functions on X . The pair $(C(X)_{sa}, L_\rho)$ is clearly a compact quantum metric space; we call it the *classical compact quantum metric space associated to (X, ρ)* . It is not too hard to show that given $(C(X)_{sa}, L_\rho)$ one can reproduce (X, ρ) , and so no information is lost by focusing on the order-unit spaces.

What we would now like to see are some purely quantum examples. Most of these are constructed using the actions of groups on C^* -algebras. Let \mathcal{A} be a C^* -algebra and let G be a compact group. An *action of G on \mathcal{A} by automorphisms*

is a homomorphism $\alpha : G \rightarrow \text{Aut}(A)$, $g \mapsto \alpha_g$, where $\text{Aut}(A)$ is the space of automorphisms of \mathcal{A} . An action is called *strongly continuous* if, for each $a \in \mathcal{A}$, the mapping

$$G \rightarrow \mathcal{A}, \quad g \mapsto \alpha_g(a)$$

is continuous. A *length-function* ℓ on G is a function that takes values in \mathbf{R}^+ such that $\ell(xy) \leq \ell(x) + \ell(y)$, $\ell(x^{-1}) = \ell(x)$, and $\ell(x) = 0$ if, and only if, $x = e$. (We note that every length function ℓ on a group gives a metric ρ that is defined by $\rho(x, y) = \ell(xy^{-1})$. Moreover, since

$$\rho(xz, yz) = \ell(xz(yz)^{-1}) = \ell(xzz^{-1}y^{-1}) = \rho(x, y),$$

the metric is right invariant. Conversely, given a right-invariant metric on G we can define a length function by $\ell(x) = \rho(x, e)$.) Define a seminorm L on \mathcal{A} by

$$L(a) = \sup\{\|\alpha_x(a) - a\|/\ell(x) : x \neq e\}. \quad (6.3)$$

(Note that $L(a^*) = L(a)$.) The restriction of L to \mathcal{A}_{sa} seems like a reasonable candidate for a Lip-norm. If we are to have that $L(a) = 0$ only when $a \in \mathbf{C}1$, then it is clear that α must be ergodic on \mathcal{A} ; an action α of a group on a unital algebra A is called *ergodic* if $\alpha(a) = a$ only when $a \in \mathbf{C}1$. In fact, in [91] Rieffel showed that ergodicity is all we need.

Theorem 6.1.5 *If G is a compact group endowed with a continuous length function and α is a strongly continuous ergodic action of G on \mathcal{A} , then L , as defined in equation (6.3), restricted to \mathcal{A}_{sa} , is a Lip-norm.*

This result is established by verifying the criteria of Theorem 6.1.4. (It easily fails if G is not compact; as yet no effort has been made to construct a noncompact version.) This result provides us with a large stock of good examples and motivates us to make the following definition: Let \mathcal{A} be a unital C^* -algebra and let L be a seminorm on \mathcal{A} satisfying $L(a^*) = L(a)$, for all $a \in \mathcal{A}$. If the restriction of L to \mathcal{A}_{sa} is a Lip-norm, then we call the pair (\mathcal{A}, L) a *Lip-normed C^* -algebra* and we call (\mathcal{A}_{sa}, L) its associated compact quantum metric space. Note that if X is a compact metric space, then $(C(X), L_\rho)$ is a Lip-normed C^* -algebra.

Our search for compact quantum metric spaces now turns into a search for ergodic actions of compact groups on unital C^* -algebras. If G is a compact group endowed with a continuous length function ℓ , and if U is an irreducible unitary representation of G on a Hilbert space H , then we can define a strongly continuous group action $\alpha : G \rightarrow B(H)$ by setting $\alpha_g(B) = U_g B U_g^*$. We shall examine this example in greater detail in Section 6.3. The corresponding compact quantum metric space will be of great importance to us.

Quantum Tori

One of the most important families of examples of spaces in noncommutative geometry is the family of quantum tori. For $\hbar \in \mathbf{R}$, the quantum torus $C_\hbar(\mathbf{T}^2)$ is a C^* -subalgebra constructed as follows: let H be the Hilbert space $L^2(\mathbf{T}^2)$, where $\mathbf{T}^2 = \mathbf{R}^2/2\pi\mathbf{Z}^2$; and let U and V be the two bounded linear operators defined on H by setting

$$\begin{aligned} U f(x_1, x_2) &= e^{ix_1} f(x_1, x_2 - \frac{1}{2}\hbar), \\ V f(x_1, x_2) &= e^{ix_2} f(x_1 + \frac{1}{2}\hbar, x_2), \end{aligned}$$

for $f \in H, x_i \in \mathbf{T}^2$. These are unitary operators and they obey the commutation relation $UV = e^{i\hbar}VU$. We define the *quantum torus*, for \hbar , to be the closed span in $B(L^2(\mathbf{T}^2))$ of the operators $U^m V^n$, for $m, n \in \mathbf{Z}$.

It can be shown that when $\hbar = 0$, $C_\hbar(\mathbf{T}^2) \simeq C(T^2)$. This is the motivation for the name quantum torus.

Quantum tori have a number of important applications, the most notable being in Connes' study of the quantum Hall effect [12]. They are also a central example in cyclic cohomology theory.

In [94] Rieffel defined a canonical strongly continuous ergodic action of \mathbf{T}^2 on $C_\hbar(\mathbf{T}^2)$ (where \mathbf{T}^2 is considered as a compact group in the obvious way). Thus, by choosing a continuous length function on \mathbf{T}^2 , which it is always possible to do, one can give $C_\hbar(\mathbf{T}^2)_{sa}$ the structure of a compact quantum metric space.

Other examples of compact quantum metric spaces have been produced from Connes and Landi's θ -deformed spheres [14]. Lip-normed AF-algebras have been produced using Bratteli's non-commutative spheres [5]. Quite recently Li [72] has used Podlès definition of an action of a compact quantum group on a C^* -algebra as a means to generate compact quantum metric spaces. He has used this structure to good effect in studying the types of convergence in quantum field theory that motivated Rieffel.

6.1.6 Spectral Triples

At this stage it might be interesting for us to reflect on what connection, if any, exists between compact quantum metric spaces and spectral triples. We recall that for a compact Riemannian spin manifold M , the Lipschitz seminorm L_ρ and the Dirac operator \mathcal{D} are related by the equation $L_\rho(f) = \|[\mathcal{D}, f]\|$, for all $f \in C^\infty(M)$. Rieffel has made some progress towards establishing a similar relation in the non-commutative case. For an arbitrary compact quantum metric space (A, L) he

constructed a faithful representation of A on a Hilbert space H that preserves the order-unit structure, and a self-adjoint operator D on H such that $L(a) = \|[D, a]\|$, for all $a \in A$. A major shortcoming of his construction is that in general D does not have compact resolvent. If (\mathcal{A}, L) is a Lip-normed C^* -algebra, then the representation of \mathcal{A}_{sa} can be extended to a linear representation of \mathcal{A} on H such that $L(a) = \|[D, a]\|$, for all $a \in \mathcal{A}$. However, the representation is not always a $*$ -algebra homomorphism. While there exist examples for which these problems does not arise, it is an open question as to what additional conditions a compact quantum metric space would have to satisfy in order to ensure that they did not arise in general.

6.2 Quantum Gromov–Hausdorff distance

As explained in the introduction, Rieffel introduced compact quantum metric spaces in the hope that they could be used to formalise statements about matrix algebras converging to the sphere. Now, that we have presented compact quantum metric spaces we shall move onto defining what it means for them to converge.

6.2.1 Gromov–Hausdorff Distance

When speaking of convergence of ordinary metric spaces the most frequently used formulism is that of Gromov–Hausdorff distance. It is a generalisation of Hausdorff distance and it is most commonly used in the study of compact Riemannian manifolds. When the manifold is a spin manifold, the associated Dirac operator plays a prominent role [73]. This hinted to Rieffel that he might be able to discuss convergence of compact quantum metric spaces in terms of a suitable ‘quantum version’ of Gromov–Hausdorff distance.

We shall begin by recalling the definition of Hausdorff distance. Let (X, ρ) be a compact metric space and let Y be a subset of X . For any positive real number r , define $\mathcal{N}_r(Y)$, the *open r -neighborhood* of Y , by

$$\mathcal{N}_r(Y) = \{x \in X : \rho(x, y) < r, \text{ for some } y \in Y\}.$$

We define $\text{dist}_H^\rho(Y, Z)$, the *Hausdorff distance* between two closed subsets Y and Z of X , by setting

$$\text{dist}_H^\rho(X, Y) = \inf\{r : Y \subseteq \mathcal{N}_r(Z) \text{ and } Z \subseteq \mathcal{N}_r(Y)\}.$$

The Hausdorff distance defines a metric on the family of all closed subsets of X . The resulting metric space can be shown to be compact, and it is complete if X is complete.

Gromov generalised this metric to one that defines a distance between any two compact metric spaces. Let (X, ρ_X) and (Y, ρ_Y) be two compact metric spaces and let $X \dot{\cup} Y$ denote the disjoint union of X and Y . Let $\mathcal{M}(\rho_X, \rho_Y)$ denote the set of all metrics on $X \dot{\cup} Y$ that induce its topology, and whose restrictions to X and Y are ρ_X and ρ_Y respectively. We call the elements of $\mathcal{M}(\rho_X, \rho_Y)$ *admissible metrics*. An element $\rho \in \mathcal{M}(\rho_X, \rho_Y)$ can be produced as follows: if $x, x' \in X$, define $\rho(x, x') = \rho_X(x, x')$; if $y, y' \in Y$, then define $\rho(y, y') = \rho_Y(y, y')$; if $x \in X$, $y \in Y$, then, for some fixed $x_0 \in X$, some fixed $y_0 \in Y$, and some fixed positive number L , define $\rho(x, y) = \rho_X(x, x_0) + L + \rho_Y(y, y_0)$. It is routine to show that this defines a metric that induces the topology of $X \dot{\cup} Y$.

Now, for each metric ρ in $\mathcal{M}(\rho_X, \rho_Y)$, it is clear that $X \dot{\cup} Y$ is compact, and that X and Y are closed subsets of $X \dot{\cup} Y$. Thus, the Hausdorff distance between them is well defined. Their *Gromov–Hausdorff distance* $\text{dist}_{GH}(X, Y)$ is defined by setting

$$\text{dist}_{GH}(X, Y) = \inf\{\text{dist}_H^\rho(X, Y) : \rho \in \mathcal{M}(\rho_X, \rho_Y)\}.$$

Gromov showed that if $\text{dist}_{GH}(X, Y) = 0$, then X and Y are isometric as metric spaces. He went on to establish that if \mathcal{CM} denotes the family of all isometry classes of compact metric spaces, then the pair $(\mathcal{CM}, \text{dist}_{GH})$ is a complete metric space. He also established necessary and sufficient conditions for a subset of \mathcal{CM} to be totally bounded.

It is interesting to consider the relationship between Hausdorff and Gromov–Hausdorff distance. As a little thought will verify, if (X, ρ) is a compact Hausdorff space, and Y and Z are closed subsets of X , then

$$\text{dist}_{GH}(Y, Z) \leq \text{dist}_H^\rho(Y, Z).$$

6.2.2 Quantum Gromov–Hausdorff

We shall now construct a generalised version of Gromov–Hausdorff distance that will define a distance between two compact quantum metric spaces (A, L_A) and (B, L_B) . An obvious, if somewhat crude, way to do this would be to take the Gromov–Hausdorff distance between $S(A)$ and $S(B)$. However, a distance that involved the Lip-norms of (A, L_A) and (B, L_B) more directly would be more natural. Let us look to the classical case for some intuition on how to do this. The space of continuous real-valued functions on $X \dot{\cup} Y$ can be identified with the order-unit space $C(X; \mathbf{R}) \oplus C(Y; \mathbf{R})$. Thus, for any metric ρ on $X \dot{\cup} Y$, we have a corresponding Lipschitz seminorm L_ρ on $C(X; \mathbf{R}) \oplus C(Y; \mathbf{R})$. This prompts us to generalise metrics on the disjoint union of two compact spaces by Lip-norms on direct sum of two order-unit spaces. (The *direct sum* of two order-unit spaces A and B is defined in the obvious way: take $A \oplus B$, the direct sum of A and B as linear spaces

and define (e_A, e_B) to be the order-unit, then define a partial order by setting $(a, b) \leq (c, d)$ if $a \leq c$ and $b \leq d$. The standard norm on the direct sum is easily seen to satisfy $\|(a, b)\| = \max\{\|a\|, \|b\|\}$.

We now need to generalise to the quantum case the notion of an admissible metric. Let (X, ρ) be an arbitrary compact metric space, let Y be a closed subset of X , and let f be an element of $C(X)$. Denote the restriction of ρ to Y by ρ_Y and denote the restriction of f to Y by $\pi(f)$. If $g \in C(Y; \mathbf{R})$ and $f \in C(X; \mathbf{R})$ such that $\pi(f) = g$, then it is clear that

$$L_{\rho_Y}(g) \leq L_\rho(f).$$

Consider the function

$$h(x) = \inf_{y \in Y} (g(y) + L_{\rho_Y}(g)\rho(x, y)) \in C(X; \mathbf{R}).$$

A short computation will verify that $\pi(h) = g$ and that $L_\rho(h) = L_{\rho_Y}(g)$. Thus,

$$L_{\rho_Y}(g) = \inf\{L_\rho(f) : \pi(f) = g\}.$$

(In fact, it is not always possible to find an analogue of h for complex-valued functions. This provides another important reason for our emphasis on real-valued functions.)

This motivates us to make the following definition: let (A, L_A) and (B, L_B) be two compact quantum metric spaces. We call a Lip-norm L on $A \oplus B$ *admissible* if

$$L_A^q(a) = \inf\{L(a, b) : b \in B\}$$

for all $a \in A$, and

$$L_B^q(b) = \inf\{L(a, b) : a \in A\}$$

for all $b \in B$. We denote the set of all admissible Lip-norms by $\mathcal{M}(L_A, L_B)$.

In a short while we shall use admissible Lip-norms to define a distance between the state spaces of A and B . Firstly, however, for sake of clarity, we shall spell out some details about the relationship between $S(A)$, $S(B)$, and $S(A \oplus B)$. Denote the canonical injection of $S(A)$ into $S(A \oplus B)$ by i ; that is, if $\varphi \in S(A)$, then $i(\varphi)(a \oplus b) = \varphi(a)$. It is easily seen that i is injective, and that it is continuous with respect to the weak* topology. Thus, since $S(A)$ and $S(A \oplus B)$ are both compact Hausdorff spaces, $S(A)$ is homeomorphic to $i(S(A))$. In this sense we shall consider $S(A)$ to be a closed subset of $S(A \oplus B)$. Similarly, we shall consider $S(B)$ to be a closed subset of $S(A \oplus B)$.

It can be shown [94] that if L is an admissible Lip norm on $A \oplus B$, then the restriction of ρ_L to $S(A)$ is equal to ρ_{L_A} , and the restriction of ρ_L to $S(B)$ is equal

to ρ_{L_B} . (The proof relies upon the fact that the metric on each space induces the weak* topology.) This pleasing result allows us to define a quantum version of Gromov–Hausdorff distance.

Definition 6.2.1. Let (A, L_A) and (B, L_B) be two compact quantum metric spaces. The *quantum Gromov–Hausdorff distance* between them is

$$\text{dist}_q(A, B) = \inf\{\text{dist}_H^{\rho_L}(S(A), S(B)) : L \in \mathcal{M}(L_A, L_B)\}.$$

Quantum Gromov–Hausdorff distance is clearly symmetric, that is, for two compact quantum metric spaces (A, L_A) and (B, L_B) , $d_q(A, B) = d_q(B, A)$. In [94] Rieffel showed that if (A, L_A) , (B, L_B) , and (C, L_C) are compact quantum metric spaces, then

$$\text{dist}_q(A, C) \leq \text{dist}_q(A, B) + \text{dist}_q(B, C).$$

He also showed that if $\text{dist}_q(A, B) = 0$, then, with respect to an appropriately defined notion of isometry based on Lip-norms, (A, L_A) is isometric to (B, L_B) . Thus, if we denote the family of isometry classes of compact quantum metric spaces by \mathcal{CQM} , then the pair $(\mathcal{CQM}, \text{dist}_q)$ is a metric space. In fact, Rieffel went on to show that it is a complete metric space, and that analogues of Gromov’s results on the total boundedness of subsets of \mathcal{CM} also hold.

When the definition of quantum Gromov–Hausdorff distance is applied to compact metric spaces it does not in general agree with Gromov–Hausdorff distance. The basic reason why the two definitions fail to agree is not too difficult to understand. For ordinary Gromov–Hausdorff distance one is looking, loosely speaking, at the distance between the pure states of $C(X)$ and the pure states of $C(Y)$. In the case of quantum Gromov–Hausdorff distance one is looking at the distance between the states of $C(X)$ and the states of $C(Y)$. It turns out that the quantum Gromov–Hausdorff distance between two compact metric spaces is always less than the Gromov–Hausdorff distance. Loosely speaking, this is because it is ‘more difficult’ to find a pure state that is close to a pure state than it is to find a state that is close to a state.

The set of pure states is badly behaved in the noncommutative case and it is not clear how one would develop a useful theory that would define a distance between the pure states of two C^* -algebras. In fact, Rieffel is unsure as to whether the non-equivalence of the two definitions should be viewed as a defect or as a ‘quantum feature’. For a more detailed discussion of the relationship between Gromov–Hausdorff distance and its quantum version see [94].

If (\mathcal{A}, L_A) and (\mathcal{B}, L_B) are two Lip-normed C^* -algebras, then it is unfortunate, but true, that the quantum Gromov–Hausdorff distance between (\mathcal{A}_{sa}, L_A) and

$(\mathcal{B}_{sa}, L_{\mathcal{B}})$ can be zero even when \mathcal{A} and \mathcal{B} are not isomorphic as C^* -algebras. (As we noted earlier, this cannot happen in the commutative case.) Both David Kerr and Hanfeng Li have worked towards addressing this shortcoming of the theory and we shall review their work in the last section of this chapter.

6.2.3 Examples

In general it proves quite difficult to find the quantum Gromov–Hausdorff distance between two compact quantum metric spaces. Usually the best one can do is to establish an upper bound for it (lower bounds are also quite difficult to find). An analogous situation holds for classical Gromov–Hausdorff distance.

The first major example of quantum Gromov–Hausdorff convergence that Rieffel established involved a sequence of quantum tori [94]. (The quantum tori being considered as compact quantum metric spaces in the sense explained earlier.) It was shown that if $\{\hbar(n)\}_n$ is a sequence of real numbers converging to a real number \hbar , then the corresponding quantum tori $C_{\hbar(n)}(T^2)$ converge to $C_{\hbar}(T^2)$ with respect to quantum Gromov–Hausdorff distance. In other words, the mapping from \mathbf{R} to \mathcal{CQM} given by $\hbar \mapsto C_{\hbar}(T^2)$ is continuous with respect to the canonical topology of \mathbf{R} and the topology induced on \mathcal{CQM} by quantum Gromov–Hausdorff distance.

Another interesting example, that builds on Rieffel’s work, comes from Latrémolière [67]. He has recently shown that with respect to quantum Gromov–Hausdorff distance any quantum torus can be approximated by a sequence of finite-dimensional C^* -algebras. He loosely terms these finite-dimensional C^* -algebras *fuzzy tori*. His motivation for establishing such a result came again from various statements in quantum field theory. According to Rieffel there is a wealth of other examples in the physics literature that could be given formal description using the language of compact quantum metric spaces.

In the next section we shall present what is arguably the most famous example of Gromov–Hausdorff convergence. It involves a sequence of canonically constructed compact quantum metric spaces converging to a classical compact quantum metric space associated to the sphere. It is interesting because it gives rigorous expression to our earlier discussion of fuzzy spheres converging to S^2 and it demonstrates very well the interplay between theoretical physics and mathematics that is so prevalent in noncommutative geometry.

6.3 Matrix Algebras Converging to the Sphere

As we saw in the previous chapter, the two sphere is a coadjoint orbit of $SU(2)$. For sake of convenience and generality, most of the discussion in this section will

be in terms of a general coadjoint orbit \mathcal{O}_μ . (Strictly speaking we should only be considering *integral* coadjoint orbits. However, just as in the previous chapter, we are going to be a little careless about this.) It is only as we near the end of the exposition that we shall return to the special case of the two sphere.

To show that \mathcal{O}_μ is the limit of a sequence of matrix algebras, we must first find an ‘appropriate’ way to give it the structure of a compact quantum metric space. Recall that all coadjoint orbits of a Lie group G are of the form G/H , for some subgroup H . We shall use this fact to endow \mathcal{O}_μ with a compact quantum metric space structure. Let ℓ be a continuous length function on G . As we noted before, ℓ induces a metric on G that is defined by $\rho(g, h) = \ell(gh^{-1})$. If we assume that

$$\ell(xgx^{-1}) = \ell(x), \quad (6.4)$$

then

$$\rho(xg, xh) = \rho(g, h), \quad \text{for all } x, g, h \in G, \quad (6.5)$$

then ρ in turn induces a metric ρ_π on G/H that is defined by setting

$$\rho_\pi([x], [y]) = \inf\{\rho(x, y) : x \in [x], y \in [y]\}.$$

We shall use $L_{\mathcal{A}}$ to denote the Lipschitz seminorm that ρ_π induces on $\mathcal{A}_{sa} \subseteq \mathcal{A} = C(G/H)$.

We note that it is always possible to define a continuous length function on a compact Lie group G that satisfies equation (6.4). In fact, ‘most’ canonical metrics on coadjoint orbits arise in this way. For example, the usual round metric on the 2-sphere is of this for

In its standard form $L_{\mathcal{A}}$ is somewhat awkward to work with. Fortunately, however, there exists a more convenient formulation. Let λ be an action of G on \mathcal{A} defined by setting

$$(\lambda_h f)[g] = f([h^{-1}g]), \quad f \in \mathcal{A}, \quad g, h \in G. \quad (6.6)$$

A series of straightforward calculations will show that

$$L_{\mathcal{A}}(f) = \sup_{g \neq e} \{\|\lambda_g(f) - f\|_\infty / \ell(g)\}.$$

(This means that $L_{\mathcal{A}}$ is the Lip norm on \mathcal{A} arising from the ergodic action λ .)

Compact Quantum Metric Spaces from Group Representations

We now need a suitable way to endow the fuzzy space matrix algebras with a compact quantum metric space structure. Earlier in the chapter, we saw that one

could endow the algebra of operators on a Hilbert space with a compact quantum metric space structure using the action of a compact group. Since every fuzzy coadjoint orbit comes naturally endowed with a compact group action, this seems like a very suitable formulation. Let us present what is involved in detail: let α be a strongly continuous ergodic action of a compact group G on a unital C^* -algebra \mathcal{A} , and let ℓ be a continuous length function defined on G . Then Theorem 6.1.5 states that the seminorm L_α , defined by setting

$$L_\alpha(a) = \sup\{\|\alpha_g(a) - a\|/\ell(g) : g \neq e_G\}, \quad (6.7)$$

is a Lip-norm on \mathcal{A}_{sa} . Now, let U be an irreducible unitary representation of G on a Hilbert space H . (We note that every irreducible representation of a compact group is finite-dimensional, as is well known.) We can define an action α of G on $\mathcal{B} = B(H)$ by setting

$$\alpha_g(T) = U_g T U_g^*, \quad g \in G, T \in \mathcal{B}.$$

Let us show that this action is strongly continuous and ergodic: if $\alpha_g(T) = T$, for all $g \in G$, then $U_g T = T U_g$, for all $g \in G$. Thus, if λ is an eigenvalue of T , and v is an element of the corresponding eigenspace E_λ , then $T U_g v = U_g T v = \lambda U_g v$. This means that $U_g v \in E_\lambda$, and so E_λ is invariant under U . To avoid a contradiction we conclude that α is ergodic. To see that α is strongly continuous take a net g_λ in G that converges to g and note that since $U_{g_\lambda} \rightarrow U_g$, and

$$\|U_g T U_g^* - U_{g_\lambda} T U_{g_\lambda}^*\| \leq \|U_g T\| \|U_g^* - U_{g_\lambda}^*\| + \|U_g - U_{g_\lambda}\| \|T U_{g_\lambda}^*\|,$$

we must have that $\alpha_{g_\lambda}(T) \rightarrow \alpha_g(T)$. Hence, the seminorm defined by equation (6.7) is a Lip-norm, and $(\mathcal{B}_{sa}, L_\alpha)$ is a compact quantum metric space.

6.3.1 The Berezin Covariant Transform

Let $\mathcal{O}_\mu \simeq G/H$ and $(\mathcal{A}_{sa}, L_{\mathcal{A}})$ be as above, and let $(\mathcal{B}_{sa}, L_{\mathcal{B}})$ be the compact quantum metric space associated to H_n , the n -fuzzification of \mathcal{O}_μ , for some $n > 0$. Most of the rest of this section will be spent trying to find an upper bound for the quantum Gromov–Hausdorff distance between $(\mathcal{A}_{sa}, L_{\mathcal{A}})$ and $(\mathcal{B}_{sa}, L_{\mathcal{B}})$. If we calculated $\text{dist}_H^{\rho_L}(S(\mathcal{A}_{sa}), S(\mathcal{B}_{sa}))$, for some admissible Lip-norm L on $\mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$, then this would give us such an upper bound.

Before we try to do this, however, we need to introduce an important reformulation of \mathcal{O}_μ . Let U_n be the representation of \mathcal{O}_μ on H_n , and let ξ be a *highest-weight vector* of U_n (see [101] for details on highest weight vectors). We define P , the

projection operator *corresponding* to ξ , by setting $Px = \langle x, \xi \rangle \xi$, for $x \in H_n$. We define R , the *stabilizer* of P , by setting

$$R = \{g \in G \mid \alpha_g(P) = P\}.$$

It can be shown [93] that R is equal to H , and so $\mathcal{O}_\mu \simeq G/R$.

Inspired by previous work on Gromov–Hausdorff distance in [94] Rieffel made the following guess at a Lip norm L on $\mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$:

$$L(f, T) = L_{\mathcal{A}}(f) \vee L_{\mathcal{B}}(T) \vee N(f, T), \quad \gamma \in \mathbf{C};$$

where $a \vee b$ denotes the maximum of a and b , and N is a seminorm on $\mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$ that satisfies $N(1_{\mathcal{A}}, 1_{\mathcal{B}}) = 0$, among a number of other natural conditions; see [94] for details. By verifying the criteria of Theorem (6.1.4), Rieffel [94] showed that L induces the weak* topology on $S(\mathcal{A}_{sa} \oplus \mathcal{B}_{sa})$. Thus, since it is clear that $L(1_{\mathcal{A}}, 1_{\mathcal{B}}) = 0$, L must be a Lip norm on $\mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$. For L to be of use to us, however, it must be admissible; that is, its quotient seminorms on \mathcal{A}_{sa} and \mathcal{B}_{sa} must be $L_{\mathcal{A}}$ and $L_{\mathcal{B}}$ respectively. We shall begin by establishing that the quotient of L on \mathcal{B}_{sa} is equal to $L_{\mathcal{B}}$. This will require us to construct a specific form for N using the Berezin covariant transform.

The Berezin Covariant Transform

For $T \in \mathcal{B}$, and τ the trace, the *Berezin covariant symbol* of T , with respect to P , is the continuous mapping

$$\sigma_T : G \rightarrow \mathbf{C}, \quad g \rightarrow \tau(T\alpha_g(P)).$$

The mapping

$$\sigma : \mathcal{B} \rightarrow C(G), \quad T \mapsto \sigma_T,$$

is called the *Berezin covariant transform*. We can easily see that $\sigma_T(gh) = \sigma_T(g)$, for all $h \in R$. Thus, the function

$$\sigma_T : G/R \rightarrow \mathbf{C}, \quad [g] \mapsto \tau(T\alpha_g(P))$$

is a well-defined element of $\mathcal{A} = C(G/R)$; and $\sigma : \mathcal{B} \rightarrow \mathcal{A}$, $T \mapsto \sigma_T$ is a well defined mapping. It has a number of useful properties. Firstly, $\sigma_1 = 1$, as can be seen from

$$\sigma_1([g]) = \tau(1\alpha_g(P)) = \tau(U_g P U_g^*) = \tau(P) = 1.$$

If T is a positive element of \mathcal{B} , then, since $\alpha_g(P)$ is clearly positive, $\tau(T\alpha_g(P)) \geq 0$. Hence, σ is a positive operator. If $T \in \mathcal{B}_{sa}$, then $\sigma_T \in \mathcal{A}_{sa}$. If λ is the action of G

on G/R , as defined in equation (6.6), then it is easily seen that σ is λ - α -equivariant, that is, $\lambda_g \sigma_T = \sigma_{\alpha_g(T)}$, for all $g \in G$, $T \in \mathcal{B}$. Finally, we have that $\|\sigma_T\|_\infty \leq \|T\|$, for all $T \in \mathcal{B}$. To see why this is so, note that since $\alpha_g(P)$ is a rank-one projection, it is of the form $\alpha_g(P)x = \langle x, e_0 \rangle e_0$, where e_0 is some norm-one element in image of $\alpha_g(P)$. Now, if $\{e_i\}_{i=0}^n$ is an orthonormal basis of H containing e_0 , then, for all $g \in G$,

$$|\sigma_T[g]| = \left| \sum_{i=0}^n \langle T \alpha_g(P) e_i, e_i \rangle \right| = \langle T e_0, e_0 \rangle \leq \|T\| \|e_0\|^2 = \|T\|,$$

and the desired result follows.

The Quotient of L on \mathcal{B}_{sa}

We now define $N(a, b) = \gamma^{-1} \|f - \sigma_T\|_\infty$, for some constant γ . It is clear from the definition of L that $L_B^q(T) = \inf\{L(f, T) : f \in \mathcal{A}_{sa}\} \leq L_B(T)$. Thus, to establish equality between L_B^q and L_B , it would suffice to show that, for every $T \in \mathcal{B}_{sa}$, there exists an $f_T \in \mathcal{A}_{sa}$, such that $L(f_T, T) = L_B(T)$. This is where the Berezin covariant transform comes into play. For any given $T \in \mathcal{B}_{sa}$, try $f_T = \sigma_T$. Since σ is λ - α -equivariant, we have that

$$\begin{aligned} L_A(\sigma_T) &= \sup_{g \neq e} \{\|\lambda_g(\sigma_T) - \sigma_T\|_\infty / \ell(g)\} = \sup_{g \neq e} \{\|\sigma_{(\alpha_g(T)-T)}\|_\infty / \ell(g)\} \\ &\leq \sup_{g \neq e} \{\|\alpha_g(T) - T\| / \ell(g)\} = L_B(T), \end{aligned}$$

Thus,

$$L(\sigma_T, T) = L_A(\sigma_T) \vee L_B(T) \vee \gamma^{-1} \|\sigma_T - \sigma_T\|_\infty = L_B(T);$$

and so $L_B^q(T) = L_B$, for all choices of γ .

The quotient of L on \mathcal{A}_{sa} is not as easy to calculate. In fact, we shall only be able to prove that it is equal to L_A for a certain adequately large values of γ . In order to calculate this value we shall need to introduce a suitably defined adjoint of σ called the Berezin contravariant transform.

6.3.2 The Berezin Contravariant Transform

In this section we shall make extensive use of the notion of averaging an operator over a compact group. Therefore, before we begin any presentation of the Berezin contravariant transform, it would be wise to recall what it means to ‘average an operator over a group’.

Averaging Operators over Compact groups

Let G, U, H and α be as above. The compactness of G implies that the continuous mapping

$$g \mapsto \langle y, \alpha_g(T)x \rangle \quad (6.8)$$

is Haar integrable, for all $x, y \in H$. (Note that as usual we shall only consider the normalised Haar measure.) Thus, the mapping

$$y \mapsto \int_G \langle y, \alpha_g(T)x \rangle dg,$$

is a well defined element of the continuous dual of H . Linearity of the functional is obvious, and it is easily seen to be bounded. By the Riesz representation theorem there exists a unique $z \in H$ such that

$$\langle y, z \rangle = \int_G \langle y, \alpha_g(T)x \rangle dg. \quad (6.9)$$

We shall denote z by $\int_G \alpha_g(T)x dg$. Consider the operator

$$\int_G \alpha_g(T)dg : H \rightarrow H, \quad x \mapsto \int_G \alpha_g(T)x dg.$$

It is easy to establish that it is linear and bounded, with norm less than or equal to $\|T\|$. We call $\int_G \alpha_g(T)dg$ the *average* of T over G , and we denote it by \tilde{T} . An important point, that is easily verified, is that if $T \geq 0$, then $\tilde{T} = 0$ if, and only if, $T = 0$.

A little thought will verify that the map defined in (6.8) can be replaced by the map $g \mapsto \langle y, A\alpha_g(T)x \rangle$, for any $A \in B(H)$, and that a well defined meaning can then be ascribed to $\int_G A\alpha_g(T)dg$ as an element of $B(H)$. Similarly, a well defined meaning can be ascribed to $\int_G \alpha_g(T)Adg$. Let us note that since

$$\begin{aligned} \left\langle \int_G A\alpha_g(T)x dg, y \right\rangle &= \int_G \langle \alpha_g(T)x, A^*y \rangle dg = \left\langle \int_G \alpha_g(T)x dg, A^*y \right\rangle \\ &= \left\langle A \int_G \alpha_g(T)x dg, y \right\rangle, \end{aligned}$$

for all $x, y \in H$, it holds that $\int A\alpha_g(T)dg = A \int \alpha_g(T)dg$. It is also easily seen that $\int \alpha_g(T)Adg = \int \alpha_g(T)dgA$.

An important consequence of these two results is that $U_h \tilde{T} = \tilde{T} U_h$, for all $h \in G$. This can be seen from

$$U_h \tilde{T} = \int U_h U_g T U_g^* dg = \int U_{hg} T U_{hg}^* U_{h^{-1}} dg = \tilde{T} U_h.$$

Thus, if T is non-zero, then the ergodicity of α implies that $\tilde{T} = \lambda 1$, for some $\lambda \in \mathbf{C}$. With a view to finding a value for λ , consider the trace of $\int_G S_g dg$, where $S_g = A \alpha_g(T) B$, for some $A, B \in B(H)$. If $\{e_i\}_{i=1}^n$ is an orthonormal basis of H , then

$$\begin{aligned} \tau\left(\int_G S_g dg\right) &= \sum_{i=1}^n \left\langle \int_G S_g e_i dg, e_i \right\rangle = \sum_{i=1}^n \int_G \langle S_g e_i, e_i \rangle dg \\ &= \int_G \sum_{i=1}^n \langle S_g e_i, e_i \rangle dg = \int_G \tau(S_g) dg. \end{aligned}$$

This implies that

$$\tau(\tilde{T}) = \tau\left(\int_G U_g T U_g^* dg\right) = \int_G \tau(U_g T U_g^*) dg = \int_G \tau(T) dg = \tau(T).$$

Since $\tilde{T} = \lambda 1$, we also have that $\tau(\tilde{T}) = \tau(\lambda 1) = \lambda n$, where n is the dimension of H . Thus,

$$\tilde{T} = \frac{\tau(T)}{n} 1. \quad (6.10)$$

The Berezin Contravariant Transform

Endow \mathcal{B} with the *Hilbert-Schmidt* inner product, which is defined by setting

$$\langle T, S \rangle_{HS} = \frac{1}{n} \tau(T S^*).$$

Let μ be the Haar measure on G , and let π be the canonical projection from G to G/R . We denote by $L^2(G/R)$ the linear space of equivalence classes of Borel measurable functions on G/R that are square integrable with respect to the measure $\tilde{\mu} = \mu \circ \pi^{-1}$. We endow $L^2(G/R)$ with its standard inner product, as defined in Section 1.4.1. Since G/R is compact, $C(G/R) \subseteq L^2(G/R)$. Thus, σ can be viewed as a linear mapping from the Hilbert space \mathcal{B} to the Hilbert space $L^2(G/R)$. This means that there exists an operator $\check{\sigma} : L^2(G/R) \rightarrow \mathcal{B}$ such that

$$\langle \sigma_T, f \rangle_{L^2} = \langle T, \check{\sigma} f \rangle_{HS},$$

for all $f \in L^2(G/R)$, $T \in B(H)$. We call $\check{\sigma}$ the *Berezin contravariant mapping*, and we call $\check{\sigma}_f$ the *Berezin contravariant symbol* of f . We shall only consider the restriction of $\check{\sigma}$ to \mathcal{A} , which we denote by the same symbol. This mapping is often viewed as a ‘quantization’ operator since it brings functions to operators. It is related to the Toeplitz maps discussed in the previous chapter.

Using the results that we established above for the average of an operator over a group, we shall find a more explicit formulation of $\check{\sigma}$. To begin with, we note that we can regard $C(G/R)$ as a subset of $C(G)$ (in the sense that there exists a canonical embedding of $C(G/R)$ into $C(G)$; namely the mapping $f \mapsto \tilde{f} = f \circ \pi$). It proves profitable to do so since

$$\int_G \tilde{f} d\mu = \int_{G/R} f d\tilde{\mu};$$

as can be verified by a routine investigation. In what follows we shall tacitly assume this observation, and we shall not distinguish notationally between f and \tilde{f} .

For any $f \in \mathcal{A}$, $T \in \mathcal{B}$, we have that

$$\begin{aligned} \frac{1}{n} \tau(\check{\sigma}_f T^*) &= \langle \check{\sigma}_f, T \rangle_{HS} = \langle f, \sigma_T \rangle_{L^2} = \int_G f(g) \overline{(\sigma_T(g))} dg \\ &= \int_G f(g) \tau(\alpha_g(P) T^*) dg = \tau \left(\int_G f(g) \alpha_g(P) dg T^* \right). \end{aligned}$$

Since this is true for all T , it must hold that

$$\check{\sigma}_f = n \int_G f(g) \alpha_g(P) dg. \quad (6.11)$$

Thus, since

$$\begin{aligned} \alpha_h(\check{\sigma}_f) &= n U_h \int_G f(g) \alpha_g(P) dg U_h^* = n \int_G f(g) \alpha_{hg}(P) dg \\ &= n \int_G f(h^{-1}g) \alpha_g(P) dg = \check{\sigma}_{(\lambda_h f)}, \end{aligned}$$

it also holds that $\check{\sigma}$ is α - λ -equivariant. Following similar lines of argument we can also show that $\check{\sigma}$ is norm-decreasing.

The Quotient of L on \mathcal{A}_{sa}

Now, that we have constructed the Berezin contravariant mapping, we are ready to approach the question of the quotient of L on \mathcal{A}_{sa} . For convenience sake we shall recall here that

$$L(f, T) = L_{\mathcal{A}}(f) \vee L_{\mathcal{B}}(T) \vee \gamma^{-1} \|f - \sigma_T\|_{\infty}.$$

An immediate consequence of the definition of L is that

$$L_{\mathcal{A}}(f) \leq L_{\mathcal{A}}^q(f) = \inf\{L(f, T) : T \in \mathcal{B}\}.$$

Thus, to establish equality between $L_{\mathcal{A}}^q$ and $L_{\mathcal{A}}$, it would suffice to show that, for each $f \in \mathcal{A}$, there exists a $T_f \in \mathcal{B}(H)$ such that $L(f, T_f) \leq L_{\mathcal{A}}(f)$. Recalling the use we made of the Berezin covariant symbol when we were examining the quotient of L on B , it seems reasonable to try $T_f = \check{\sigma}_f$. Since $\check{\sigma}$ is norm-decreasing and α - λ -equivariant, it holds that

$$\begin{aligned} L_{\mathcal{B}}(\check{\sigma}_f) &= \sup_{g \neq e} \|\alpha_g(\check{\sigma}_f) - \check{\sigma}_f\| / \ell(g) = \sup_{g \neq e} \|\check{\sigma}_{(\lambda_g(f)-f)}\| / \ell(g) \\ &\leq \sup_{g \neq e} \|\lambda_g(f) - f\| / \ell(g) = L_{\mathcal{A}}(f). \end{aligned}$$

Thus, $L_{\mathcal{B}}(\check{\sigma}_f) \leq L_{\mathcal{A}}(f)$, for all values of γ . This means that if we could find a value for γ such that

$$L_{\mathcal{A}}(f) \geq \gamma^{-1} \|f - \sigma(\check{\sigma}_f)\|_{\infty}, \quad (6.12)$$

then for the corresponding L , it would hold that $L_{\mathcal{A}}^q = L_{\mathcal{A}}$. We shall spend the remainder of this section trying to find such a value.

The map $f \mapsto \sigma(\check{\sigma}_f)$ is called the *Berezin transform*. We can derive a more explicit formulation of it as follows:

$$\begin{aligned} (\sigma(\check{\sigma}_f))[h] &= \tau(\check{\sigma}_f \alpha_h(P)) = \tau \left(n \int_G f(g) \alpha_g(P) dg \alpha_h(P) \right) \\ &= n \int_G f(g) \tau(\alpha_g(P) \alpha_h(P)) dg = n \int_G f(g) \tau(P \alpha_{g^{-1}h}(P)) dg. \end{aligned}$$

For any rank-one projection P on H , we shall find it useful to introduce a function $k_P \in C(G/R)$ defined by

$$k_P[g] = n \tau(P \alpha_g(P)).$$

(Note that $k_P[g] = n \sigma_P[g]$; we use a distinct symbol for k_P for sake of presentation.) Our formula for the Berezin transform now becomes

$$(\sigma(\check{\sigma}_f))[h] = \int_G f(g) k_P(g^{-1}h) dg.$$

The function k_P has some pleasing properties: for any norm-one vector e_0 contained in the image of P ,

$$k_P([g]) = n | \langle U_g e_0, e_0 \rangle |^2 \geq 0. \quad (6.13)$$

(This is easily established by choosing a specific orthonormal basis for H that contains e_0 .) Thus, k_P is a positive function. Using equation (6.13) we can easily show that $k_P([g^{-1}]) = k_P([g])$. Finally, we also have that

$$\int_G k_P(g) dg = \int_G n \tau(P \alpha_g(P)) dg = \tau \left(P n \int_G \alpha_g(P) dg \right) = \tau(P \check{\sigma}_1) = 1.$$

We shall tacitly make use of these observations below.

We are now in a position to find a value for γ for which equation (6.12) will be satisfied. To begin with, let us note that

$$\begin{aligned} |f([h]) - \sigma(\check{\sigma}_f)([h])| &= \left| \int_G (f([h]) - f(g)) k_P(g^{-1}h) dg \right| \\ &\leq \int_G |f([h]) - f(g)| k_P(g^{-1}h) dg. \end{aligned}$$

Now, if f is an element of the dense order-unit subspace of \mathcal{A}_{sa} on which the Lipschitz seminorm takes finite values, then $|f[h] - f[g]| \leq L_{\mathcal{A}}(f) \rho_{\pi}([h], [g])$, for all $g, h \in G$. Therefore,

$$\begin{aligned} |f([h]) - \sigma(\check{\sigma}_f)([h])| &\leq L_{\mathcal{A}}(f) \int_G \rho_{\pi}([h], [g]) k_P(g^{-1}h) dg \\ &= L_{\mathcal{A}}(f) \int_G \rho_{\pi}([h], [g]) k_P(h^{-1}g) dg. \end{aligned}$$

Since we required the length function ℓ on G to satisfy $\ell(xg, xh) = \ell(g, h)$, for all $x, g, h \in G$, we have that $\rho_{\pi}([xg], [xh]) = \rho_{\pi}([g], [h])$. Consequently,

$$\begin{aligned} |f([h]) - \sigma(\check{\sigma}_f)([h])| &\leq L_{\mathcal{A}}(f) \int_G \rho_{\pi}([h], [hg]) k_P(g) dg \\ &= L_{\mathcal{A}}(f) \int_G \rho_{\pi}([e], [g]) k_P(g) dg. \end{aligned}$$

Thus, if we choose

$$\gamma = n \int_G \rho_{\pi}([e], [g]) \sigma_P(g) dg,$$

then $\|f - \sigma(\check{\sigma}_f)\|_{\infty} \leq \gamma L_{\mathcal{A}}(f)$, for all $f \in \mathcal{A}_{sa}$. This gives us the following proposition.

Proposition 6.3.1 *For γ chosen as above, the seminorm on $\mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$ defined by setting*

$$L(f, T) = L_{\mathcal{A}}(f) \vee L_{\mathcal{B}}(T) \vee \gamma^{-1} \|f - \sigma_T\|_{\infty},$$

has $L_{\mathcal{A}}$ as its quotient norm on \mathcal{A}_{sa} .

6.3.3 Estimating the QGH Distance

Now, that we have shown that L is an admissible Lip-norm on $\mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$, we shall try and estimate the quantum Gromov–Hausdorff distance between $(\mathcal{A}_{sa}, L_{\mathcal{A}})$ and $(\mathcal{B}_{sa}, L_{\mathcal{B}})$.

Proposition 6.3.2 *If γ is a constant chosen such that the quotient of L on \mathcal{A}_{sa} is $L_{\mathcal{A}}$, then $S(\mathcal{A}_{sa})$ is in the γ -neighborhood of $S(\mathcal{B}_{sa})$ for ρ_L .*

Proof. For each $\mu \in S(\mathcal{A}_{sa})$ we must produce a $\nu \in S(\mathcal{B}_{sa})$, such that $\rho_L(\mu, \nu) \leq \gamma$. Let us try $\nu = \mu \circ \sigma$. (Note that σ is unital and positive, and therefore ν is indeed contained in $S(\mathcal{B}_{sa})$). Recall that

$$\rho_L(\mu, \nu) = \sup\{|\mu(f, T) - \nu(f, T)| : (f, T) \in \mathcal{A}_{sa} \oplus \mathcal{B}_{sa}, L(f, T) \leq 1\}.$$

For $(f, T) \in \mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$,

$$\begin{aligned} |\mu(f, T) - \nu(f, T)| &= |\mu(f) - \nu(T)| = |\mu(f) - \mu(\sigma_T)| \\ &= |\mu(f - \sigma_T)| \leq \|\mu\| \|f - \sigma_T\|_{\infty} \\ &= \|f - \sigma_T\|_{\infty}. \end{aligned}$$

Since $L(f, t) \leq 1$ and $L(f, t) = L(f) \vee L(T) \vee \gamma^{-1} \|f - \sigma_T\|_{\infty}$, we have $\|f - \sigma_T\|_{\infty} \leq \gamma$. It follows that $\rho_L(\mu, \nu) \leq \gamma$, and so $S(\mathcal{A}_{sa})$ is contained in the γ -neighbourhood of $S(\mathcal{B}_{sa})$. \square

Consequently, to put a suitably small upper bound on the quantum Gromov–Hausdorff distance between $(\mathcal{A}_{sa}, L_{\mathcal{A}})$ and $(\mathcal{B}_{sa}, L_{\mathcal{B}})$, it only remains to show that $S(\mathcal{B}_{sa})$ is contained in a suitably small neighbourhood of $S(\mathcal{A}_{sa})$. That is, for each $\nu \in S(\mathcal{B}_{sa})$, we must find a $\mu \in S(\mathcal{A}_{sa})$ such that $\text{dist}_q(\nu, \mu)$ is suitably small. Mimicking the proof of Proposition 6.3.2, we propose $\mu = \nu \circ \check{\sigma}$. If $(f, T) \in \mathcal{A}_{sa} \oplus \mathcal{B}_{sa}$, and $L(f, T) \leq 1$, then $L_{\mathcal{B}}(T) \leq 1$ and $\|f - \sigma_T\| \leq \gamma$. Thus,

$$\begin{aligned} |\mu(f, T) - \nu(f, T)| &= |\nu(\check{\sigma}_f) - \nu(T)| \leq \|\nu\| \|\check{\sigma}_f - T\| \\ &\leq \|\check{\sigma}_f - \check{\sigma}(\sigma_T)\| + \|\check{\sigma}(\sigma_T) - T\| \\ &\leq \|f - \sigma_T\|_{\infty} + \|\check{\sigma}(\sigma_T) - T\| \\ &\leq \gamma + \|\check{\sigma}(\sigma_T) - T\|. \end{aligned}$$

This means that any bound that we can obtain on $\|\check{\sigma}(\sigma_T) - T\|$, for $L_{\mathcal{B}}(T) \leq 1$, will give us a bound on the quantum Gromov–Hausdorff distance between $(\mathcal{A}_{sa}, L_{\mathcal{A}})$ and $(\mathcal{B}_{sa}, L_{\mathcal{B}})$.

6.3.4 Matrix Algebras Converging to the Sphere

Let us summarise what we have established: If $(\mathcal{A}_{\text{sa}}, L_{\mathcal{A}})$ is the compact quantum metric space associated to \mathcal{O}_{μ} , and $(\mathcal{B}_{\text{sa}}, L_{\mathcal{B}})$ is the compact quantum metric space associated to the n -fuzzification of \mathcal{O}_{μ} , then the quantum Gromov–Hausdorff distance between them is less than or equal to $\gamma + \|\check{\sigma}(\sigma_T) - T\|$; where $T \in \mathcal{B}_{\text{sa}}$ such that $L_{\mathcal{B}}(T) \leq 1$, and $\gamma = n \int_G \rho_{\pi}([e], g) \sigma_P(g) dg$.

Rieffel went on to show that γ and $\|\check{\sigma}(\sigma_T) - T\|$ are dependent on n , and that as $n \rightarrow \infty$, γ and $\|\check{\sigma}(\sigma_T) - T\|$ become arbitrarily small. Thus, with respect to quantum Gromov–Hausdorff distance, the sequence of fuzzy coadjoint orbits converges to the \mathcal{O}_{μ} . (We shall not outline the proof because it quite lengthy and would require the introduction of an excessive amount of Lie group theory; for details see [93].) A precise meaning has now been given to statements involving the convergence of fuzzy spaces to a coadjoint orbit. Moreover, since S^2 is a coadjoint orbit of $SU(2)$, a precise meaning has also been given to statements involving the convergence of fuzzy spaces to the 2-sphere.

6.4 Matricial Gromov–Hausdorff Distance

As we discussed earlier, a shortcoming of quantum Gromov–Hausdorff distance is that two Lip-normed C^* -algebras can have distance zero yet their C^* -algebras may not be isomorphic. Following a suggestion of Rieffel, David Kerr [56] began to investigate the possibility of defining a modified version of quantum Gromov–Hausdorff using ‘matrix-valued states’. In this context the notion of positivity gives way to the notion of complete positivity. We shall now review this concept.

Complete Positivity and Operator Systems

Let \mathcal{A} be a C^* -algebra, and let $M_n(\mathcal{A})$ denote the algebra of all $n \times n$ matrices with entries in \mathcal{A} . We can define an involution on $M_n(\mathcal{A})$ by setting $[a_{ij}]^* = [a_{ji}^*]$. If φ is a mapping from \mathcal{A} to another C^* -algebra \mathcal{B} , then we define its n -inflation to be the mapping

$$\varphi_n : M_n(\mathcal{A}) \rightarrow M_n(\mathcal{B}), \quad [a_{ij}] \rightarrow [\varphi(a_{ij})].$$

Note that if φ is a $*$ -algebra homomorphism, then its n -inflation is also a $*$ -algebra homomorphism, for all n .

Let H be a Hilbert space, and consider the mapping

$$\psi : M_n(B(H)) \rightarrow B(H^n), \quad u \mapsto \psi(u);$$

where H^n is the orthogonal n -sum of H , and $\psi(u)$ is the operator defined by setting

$$\psi(u)(x_1, x_2, \dots, x_n) = \left(\sum_{j=1}^n u_{1j}(x_j), \dots, \sum_{j=1}^n u_{nj}(x_j) \right).$$

It is straightforward to show that ψ is a $*$ -algebra isomorphism. This means that we can define a norm $\|\cdot\|$ on $M_n(B(H))$ that makes it a C^* -algebra, by setting $\|u\| = \|\psi(u)\|$. The following useful inequalities are easily established:

$$\|u_{ij}\| \leq \|u\|, \quad i, j = 1, \dots, n. \quad (6.14)$$

Let π be a faithful representation of \mathcal{A} in $B(H)$, for some Hilbert space H . (As discussed in Chapter 1, such a π and H can be always be produced using the GNS construction.) Let π_p be the p -inflation of π , and let $u(n)$ be a sequence in $\pi_p(M_p(\mathcal{A}))$ that converges to some $u \in M_p(B(H))$. By (6.14), we have that $u_{ij}(n)$ converges to u_{ij} , for each $i, j = 1, \dots, n$. Since $\pi(\mathcal{A})$ is complete, each u_{ij} is contained in $\pi(\mathcal{A})$. Thus, $\pi(M_p(\mathcal{A}))$ is closed in $M_p(B(H))$, implying that $\pi(M_p(\mathcal{A}))$ is a C^* -algebra. This enables us to define a norm on $M_p(\mathcal{A})$, that makes it a C^* -algebra, by setting

$$\|v\| = \|\pi(v)\|, \quad v \in M_p(\mathcal{A}).$$

As discussed in Chapter 1, it is the unique norm on $M_p(\mathcal{A})$ that does so.

Let φ be a map from \mathcal{A} to another C^* -algebra \mathcal{B} . Then φ is called *completely-positive* if its n -inflation $\varphi_n : M_n(\mathcal{A}) \rightarrow M_n(\mathcal{B})$ is positive, for all n . Not all positive mappings are completely-positive, the standard example of a positive, non-completely-positive, mapping is

$$\varphi : M_2(\mathbf{C}) \rightarrow M_2(\mathbf{C}), \quad A \rightarrow A^T.$$

Complete-positivity is usually studied in the more general setting of operator systems. A (*concrete*) *operator system* is a unital self-adjoint closed linear subspace of a unital C^* -algebra. (We use the adjective *concrete* here because there exists a more general definition of an operator system [30]; each such structure can, however, be represented as a self-adjoint linear subspace of a C^* -algebra. Kerr works with the concrete definition, and so it is the formulation that we shall use here.) Note that for any operator system X , its subset of self-adjoint elements, which we denote by X_{sa} , is an order-unit space.

If \mathcal{A} is a C^* -algebra, and $X \subseteq \mathcal{A}$ is an operator system, then we denote by $M_n(X)$ the subset of $M_n(\mathcal{A})$ whose elements are the matrices with entries in X . It is clear

that $M_n(X)$ is a unital self-adjoint linear subspace of $M_n(\mathcal{A})$, and so it is also an operator system. We define the set of *positive* elements of X to be $X \cap \mathcal{A}_+$, and we define the set of n -positive and completely-positive maps between two operator systems in exactly the same way as for C^* -algebras. Finally, we denote by $UCP_n(X)$, the set of unital completely-positive maps from X to $M_n(\mathbf{C})$; or, more explicitly, the set of maps $\varphi : X \rightarrow M_n(\mathbf{C})$ whose inflation

$$\varphi_p : M_p(X) \rightarrow M_p(M_n(\mathbf{C})) \simeq M_{p \times n}(\mathbf{C})$$

is positive and unital, for all p . Since each $\varphi \in UCP_1(X)$ is clearly a state, the following lemma tells us that $UCP_1(X)$ is equal to the state space of X ; for a proof see [19].

Lemma 6.4.1 *If φ is a state on X , then φ is a unital completely-positive map.*

Lip-Normed Operator Systems

We shall now define the analogue for operator systems of compact quantum metric spaces: let (X, L) be a pair consisting of an operator system X , and a Lip-norm L defined on X_{sa} ; if $D_1(L) = \{x \in D(L) : L(x) \leq 1\}$ is closed in X_{sa} , then we say that (X, L) is a *Lip-normed operator system*. (The technical requirement that $D_1(L)$ be closed will not be of great importance to us here; it is included for the sake of accuracy.) Clearly, every Lip-normed unital C^* -algebra (\mathcal{A}, L) (for which $D_1(L)$ is closed) is a Lip-normed operator system.

Mimicking the manner in which we defined a metric on the state space of an order-unit space using a Lip-norm, we define a metric $\rho_{L,n}$ on $UCP_n(X)$, for each n , by setting

$$\rho_{L,n}(\varphi, \psi) = \sup\{\|\varphi(a) - \psi(a)\| : L(a) \leq 1\},$$

for $\varphi, \psi \in UCP_n(X)$. (Note that by $\|\cdot\|$ we mean the unique norm on $M_n(\mathbf{C})$ that makes it a C^* -algebra.)

In the order-unit case we required that the metric ρ_L induce the weak* topology on $S(A)$. For $UCP_n(X)$ the natural analogue of the weak* topology is the *point-norm topology*; it is defined to be the weakest topology with respect to which the family of functions $\{\hat{x} : x \in X\}$ is continuous, where $\hat{x}(\varphi) = \varphi(x)$. Just as the state space is compact with respect to the weak* topology, each $UCP_n(X)$ is point-norm compact. It would be natural to require that each $\rho_{L,n}$ induce the point-norm topology on $UCP_n(X)$. However, Kerr established that this is a consequence of the fact that L is a Lip-norm on X_{sa} . Therefore, there is no need to impose such a condition.

Complete Gromov–Hausdorff Distance

The definition of quantum Gromov–Hausdorff distance involves the direct sum of two order-unit spaces, and the embedding of their state spaces into the state space of their direct sum. We shall now translate this process to the operator system setting. Let (X, L_X) and (Y, L_Y) be Lip-normed operator systems, and let \mathcal{A} and \mathcal{B} be two C^* -algebras containing X and Y respectively. The *direct sum* of \mathcal{A} and \mathcal{B} is defined to be their direct sum as normed algebras endowed with the pointwise-defined addition and involution; it is denoted by $\mathcal{A} \oplus \mathcal{B}$. Clearly, $\mathcal{A} \oplus \mathcal{B}$ is a C^* -algebra. The direct sum of X and Y as normed linear spaces, endowed with the pointwise-defined multiplication and involution, is clearly a unital self-adjoint linear subspace of $\mathcal{A} \oplus \mathcal{B}$. Hence, it is an operator space. We shall denote it by $X \oplus Y$.

Now, $UCP_n(X)$ can be embedded into $UPC_n(X \oplus Y)$ in an obvious manner, and it is easily seen that this embedding is continuous when we put the point-norm topology on both spaces. Since $UCP_n(X)$ and $UPC_n(X \oplus Y)$ are both compact Hausdorff spaces, the image of $UCP_n(X)$ in $UPC_n(X \oplus Y)$, which we shall equate with $UCP_n(X)$, is closed. Obviously, an entirely analogous situation holds for $UCP_n(Y)$.

We can speak of admissible Lip-norms on $X \oplus Y$, since it is just the sum of two order-unit subspaces. We denote the set of admissible Lip-norms on $X \oplus Y$ by $\mathcal{M}(L_X, L_Y)$. Kerr showed that if $L \in \mathcal{M}(L_X, L_Y)$, then the restriction of $\rho_{L,n}$ to $UCP_n(X)$ is equal to $\rho_{L_X,n}$, and the restriction of $\rho_{L,n}$ to $UCP_n(Y)$ is equal to $\rho_{L_Y,n}$. This is a direct and pleasant generalisation of what happens in the order-unit case. We can now imitate the definition of quantum Gromov–Hausdorff distance. For each natural number n , we define $\text{dist}^n(X, Y)$ the *n-distance* between (X, L_X) and (Y, L_Y) by setting

$$\text{dist}^n(X, Y) = \inf\{\rho_H^{\rho_{L,n}}(UCP_n(X), UCP_n(Y)) : L \in \mathcal{M}(L_X, L_Y)\}.$$

We then define $\text{dist}_c(X, Y)$ the *complete quantum Gromov–Hausdorff distance* by setting

$$\text{dist}_c(X, Y) = \sup_{n \in \mathbf{N}}\{\text{dist}^n(X, Y)\}.$$

If we bear in mind that $UCP_1(X) = S(X)$, then a little careful reflection will verify that the quantum Gromov–Hausdorff distance between two Lip-normed C^* -algebras is equal to their 1-distance. Thus, the complete distance is always greater than or equal to the quantum Gromov–Hausdorff distance.

Clearly, the complete distance is symmetric in its arguments. Kerr showed that it also satisfies the triangle inequality and is positive definite on the family of appropriately defined equivalence classes of Lip-normed operator systems. Hence,

it is well defined as a metric. A direct consequence of the proof of positive definiteness is that two C^* -algebras have complete distance zero if, and only if, they are $*$ -isomorphic. Thus, Kerr's definition overcomes the shortcoming of Rieffel's definition.

Kerr also showed that the continuity of quantum Gromov–Hausdorff distance for non-commutative tori, as described earlier, carries over to the complete quantum Gromov–Hausdorff distance case; as does the convergence of matrix algebras to coadjoint orbits described in Section 6.3. In [57] it was shown that the family of equivalences classes of Lip-normed operator systems endowed with dist_c is a complete metric space.

Operator Gromov–Hausdorff Distance

Hangfeng Li, a doctoral student of Rieffel, devised another strategy for quantizing Gromov–Hausdorff distance that operates entirely at the algebraic level. It also overcomes the shortcoming of Rieffel's distance addressed above. His versatile approach was implemented in both the order-unit and C^* -algebraic contexts under the terminology *order-unit*, and *C^* -algebraic quantum Gromov–Hausdorff distance* respectively [70, 71]. It affords many technical advantages.

In a recent paper [57] Kerr, working jointly with Li, established an analogue for Lip-normed operator systems of Li's distance. The pair then proved that this new distance is in fact equal to the complete quantum Gromov–Hausdorff distance. This consolidation of complete Gromov–Hausdorff distance has motivated Kerr and Li to propose that it be renamed *operator Gromov–Hausdorff distance*.

Completeness and Lip-Ultraproducts

Since it is primarily C^* -algebras, as opposed to operator systems, that we are interested in, it would be pleasing if the subfamily of Lip-normed C^* -algebras were closed in the family of Lip-normed operator systems. In [57] Kerr and Li produced sufficient conditions for a sequence of C^* -algebras to converge to a C^* -algebra. However, in a recent paper Daniele Guido and Tommaso Isola [43] (both members of the European Union Operator Algebras Network) constructed a Cauchy sequence of C^* -algebras that converges, with respect to quantum Gromov–Hausdorff distance, to an operator system that is not a C^* -algebra. Hence, the space of Lip-normed C^* -algebras is not complete with respect to dist_c . Guido and Isola's work is based upon their newly defined notion of a 'Lip-ultraproduct'. For natural reasons they propose that it be viewed as the quantum analogue of the ultralimit of a sequence of compact metric spaces (for details on ultralimits see [6]). Their work has lead them to define a new metric on the space of Lip-normed C^* -algebras with

respect to which it is complete. Consequently, they propose it as a more natural way to define distance.

Bibliography

- [1] P. ALBERTI AND R. MATTHES, Connes Trace Formula and Dirac Realization of Maxwell and Yang-Mills Action, Lecture delivered at “*The Standard Model of Elementary Particle Physics from a Mathematical-Geometrical Viewpoint*”, Hesselberg, March 1999, arXiv:math-ph/9910011 v2.
- [2] E. ABE, Hopf Algebras, *Cambridge University Press*, Cambridge, (1997).
- [3] M. F. ATIYAH, *K-Theory*, *W. A. Benjamin, Inc.* New York–Amsterdam, (1967).
- [4] A. P. BALACHANDRAN, B. P. DOLAN, J. LEE, X. MARTIN, AND D. O’CONNOR, Fuzzy Complex Projective Spaces and their Star-Products, arXiv:hep-th/0107099.
- [5] O. BRATTELI, G. A. ELLIOTT, D. E. EVANS, AND A. KISHIMOTO, Non-commutative spheres I, *Internat. Journ. Math.* **2**, Pages 139-166, (1991).
- [6] M. R. BRIDSON AND A. HAEFLIGER, Metric spaces of non-positive curvature, *Springer-Verlag*, Berlin–NewYork, (1999).
- [7] J. BRODZKI, An Introduction to *K*-theory and Cyclic Cohomology, *Polish Scientific Publishers PWN*, Warsaw, (1998), arXiv:funct-an/9606001.
- [8] K. BROWN AND J. ZHANG. Dualising complexes and twisted Hochschild (co)homology for noetherian Hopf algebras, (2006), arXiv:Math.RA:/0603732.
- [9] P. S. CHAKRABORTY AND A. PAL, On equivariant Dirac operators for $SU_q(2)$, arXiv:math.QA/0501019.
- [10] P. S. CHAKRABORTY AND A. PAL, Equivariant spectral triples for $SU_q(l+1)$ and the odd dimensional quantum spheres, arXiv:math.QA/0503689.

- [11] A. CONNES, Noncommutative differential geometry. Part 1: The Chern character in K -homology, *Preprint I.H.E.S. (M/82/53)*, (1982). Part 2: de Rham homology and noncommutative algebra, *Preprint I.H.E.S. (M/83/19)*, (1983).
- [12] A. CONNES, Noncommutative Geometry, *Academic Press*, San Diego–New York, (1994).
- [13] A. CONNES, Gravity Coupled with Matter and the Foundations of Noncommutative Geometry, *Comm. Math. Phys.* **182**, Pages 155-176, (1996).
- [14] A. CONNES AND M. DUBOIS-VIOLETTE, Noncommutative finite-dimensional manifolds I. spherical manifolds and related examples, *Comm. Math. Phys.* **230**, Number 3, Pages 539-579, (2002), arXiv:math.QA/0107070.
- [15] A. CONNES AND J. LOTT, Particle Models and Noncommutative Geometry, *Nuclear Physics B*, **18B**, Pages 29-47, (1990).
- [16] A. CONNES, H. MOSCOVICI, The local index formula in noncommutative geometry, *Geom. Funct. Anal.* **5**, Pages 174-243, (1995), <http://www.alainconnes.org/downloads.html>.
- [17] A. CONNES, H. MOSCOVICI, Hopf algebras, cyclic cohomology and the transverse index theorem, *Commun. Math. Phys.* **198**, Pages 199-246, (1998).
- [18] A. CONNES, H. MOSCOVICI, Type III and spectral triples, arXiv:math.OA/0609703.
- [19] J. B. CONWAY, A Course in Operator Theory, *American Mathematical Society*, Providence, Rhode Island, (2000).
- [20] J. DIXMIER, Von Neumann Algebras, *North-Holland*, Amsterdam, (1981).
- [21] B. P. DOLAN AND O. JAHN, Fuzzy Complex Grassmannian Spaces and their Star Products, arXiv:hep-th/0111020.
- [22] B. P. DOLAN, D. O'CONNOR AND P. PRESNAJDER, Matrix φ^4 models on the fuzzy sphere and their continuum limits, *J. High Energy Phys.* **03**, (2002).
- [23] B. P. DOLAN, Quantum Black Holes: the Event Horizon as a Fuzzy Sphere, *J. High Energy Phys.* (2005), arXiv:hep-th/0409299.

- [24] S. DOPLICHER, K. FREDENHAGEN, J. ROBERTS: The quantum structure of space-time at the Planck scale and quantum fields, *Commun. Math. Phys.*, **172**, 187, (1995).
- [25] M. DUBOIS–VIOLETTE, Lectures on graded differential algebras and non-commutative geometry, Noncommutative Differential Geometry and Its Applications to Physics, Proceedings of the Workshop at Shonan, Japan, June 1999, *Kluwer Academic Publishers*, Pages 245-306, 2001,
- [26] M. DUBOIS–VIOLETTE, Lectures on graded differential algebras and non-commutative geometry, *Geometry and Nature, Contemporary Mathematics* **203**, Pages 145-157, (1997),
- [27] M. DUBOIS–VIOLETTE, R. KERNER, J. MADORE, Non-commutative differential geometry of matrix algebras, *J. Math. Phys.* **31**, Number 316, Pages 316-322, (1990).
- [28] M. DUBOIS–VIOLETTE, R. KERNER, J. MADORE, Noncommutative differential geometry and new models of gauge theory, *J. Math. Phys.* **31**, Number 316, Pages 323-330, (1990).
- [29] V. G. DRINFELD, Quantum groups, Proceedings of the International Congress of Mathematicians, Berkeley, Pages 793-820, (1986).
- [30] E. G. EFFROS AND Z.-J. RUAN, Operator Spaces, *London Mathematical Society Monographs, New Series* 23, Clarendon, New York, (2000).
- [31] M. ENOCK AND J.-M. SCHWARTZ, Kac Algebras and Duality of Locally Compact Groups, *Springer*, Berlin-Heidelberg, 1992.
- [32] G. ESPOSITO, Dirac Operators and Spectral Geometry, *Cambridge University Press*, Cambridge, 1998.
- [33] P. FENG AND B. TSYGAN, Hochschild and cyclic homology of quantum groups, *Comm. Math. Phys.* **140**, Number 3, 481-521 (1991).
- [34] F. G. FLORES, X. MARTIN, AND D. O'CONNOR, Simulating the scalar field on the fuzzy sphere, *Proceedings of Science*, arXiv:hep-lat/0601012, (2006).
- [35] P. H. FRAMPTON, Gauge Field Theories, *The Benjamin Cummings Publishing Company*, Menlo Park CA, (1987).
- [36] T. FRIEDRICH, Dirac Operators in Riemannian Geometry, *American Mathematical Society*, Providence, Rhode Island, (2000).

- [37] I. GELFAND AND M. NAIMARK, On the embedding of normed rings into the ring of operators in Hilbert space, *Mat. Sbornik* 12 (1943), Pages 197 – 213.
- [38] J. M. GRACIS–BONDÍA, J.C. VARILLY, Connes’ noncommutative differential geometry and the Standard Model, *J. of Geom. and Physics*, 12, Number 4, (1993).
- [39] J. M. GRACIS–BONDÍA, J.C. VARILLY, AND H. FIGUEROA, Elements of Noncommutative Geometry, *Birkhäuser*, Boston–Basel–Berlin, (2001).
- [40] M. GROMOV, Metric Structures for Riemannian and non-Riemannian Spaces, *Birkhäuser Boston Inc.* Boston, MA, 1999.
- [41] A. GROTHENDIECK, Produits tensoriels topologiques et espace nucléaires, *Memoirs of the American Mathematical Society*, **16**, (1955).
- [42] D. GUIDO AND T. ISOLA, Fractals in Noncommutative Geometry, *Proceedings of “Mathematical Physics in Mathematics and Physics”*, Siena 2000.
- [43] D. GUIDO AND T. ISOLA, The problem of completeness for Gromov–Hausdorff metrics on C^* -algebras, *Journal of Functional Analysis*, **233**, Pages 173-205, (2006), arXiv:math.OA/0502013 v2.
- [44] T. HADFIELD, Twisted cyclic homology of all Podles quantum spheres, To appear in *The Jour. of Geom. and Physics*, arXiv:math.QA/0405243 v5.
- [45] T. HADFIELD AND U. KRAEHMER, Twisted homology of quantum $SL(2)$, **34**, Number 4, Pages 327-360, (2005) , *K-theory*, arXiv:math.QA/0405249.
- [46] T. HADFIELD AND U. KRAEHMER, On the Hochschild homology of quantum $SL(N)$, math.QA/0509254.
- [47] P. M. HAJAC, M. KHALKHALI, B. RANGIPOUR, Y. SOMMERHUSER, *C. R. Acad. Sci.* Hopf-cyclic homology and cohomology with coefficients, Paris, Ser. I 338, (2004), <http://www.fuw.edu.pl/~pmh/>.
- [48] A. HATCHER, Vector Bundles and K -theory, <http://www.math.cornell.edu/~hatcher/VBKT/VBpage.html>.
- [49] E. HAWKINS, Quantization of Equivariant Vector Bundles, *Commun. Math. Phys*, 202, Pages 517-546, (1999), arXiv:QA/9708030.
- [50] E. HAWKINS, Noncommutative Regularization for the Practical Man, arXiv:hep-th/9908052.

- [51] G. HOCHSCHILD, On the cohomology groups of an associative algebra, *Ann. of Math.* (2) **46**, Pages 58-67, (1945).
- [52] D. HUSEMOLLER, *Fibre Bundles*, Springer, Berlin-Heidelberg, (1993).
- [53] A. JAFFE, Quantum harmonic analysis and geometric invariants, *Advances in Math.* **143**, Pages 1-110, (1999).
- [54] R. V. KADISON, A representation theory for commutative topological algebra, *Mem. Amer. Math. Soc.* Number 7, Page 39, (1951).
- [55] L. V. KANTOROVICH, On the translocation of masses, *C. R. (Doklady) Acad. Sci. URSS (N.S.)* **37**, Pages 199-201, (1942).
- [56] D. KERR, Matricial quantum Gromov–Hausdorff distance, *Journal of Functional Analysis* **205**, Pages 132-167, (2003), arXiv:math.OA/0207282.
- [57] D. KERR, H. LI, On Gromov–Hausdorff convergence for operator metric spaces, to appear in *J. Operator Theory*, arXiv:math.OA/0411157.
- [58] J.L. KOSZUL, *Fibre bundles and differential geometry*, Tata Institute of Fundamental Research, Bombay, (1960).
- [59] J. L. LODAY AND D. QUILLEN, Cyclic homology and the Lie algebra homology of matrices, *Comment. Math. Helv.* **59**, Pages 569-591. (1984).
- [60] M. KAROUBI, Homologie cyclique des groupes et algèbres, *C.R. Acad. Sci. Paris* **297**, Série I, Pages 381-384, (1983).
- [61] A. U. KLIMYK AND K. SCHMUDGEN, *Quantum Groups and Their Representations*, Springer, Berlin-Heidelberg, (1997).
- [62] U. KRÄHMER, Dirac Operators on Quantum Flag Manifolds, *Lett. Math. Phys.* **67**, (2004), Pages 49-59 arXiv:math.QA/0305071.
- [63] J. KUSTERMANN, AND S. VAES, Locally compact quantum groups, *Ann. Scient. Ec. Norm. Sup.* **33**, Pages 837-934, (2000).
- [64] J. KUSTERMANN, G. J. MURPHY, AND L. TUSET, Quantum groups, differential calculi and the eigenvalues of the Laplacian, *Trans. Amer. Math. Soc.* (2005), arXiv:math.math.QA/0111042 v1.
- [65] J. KUSTERMANN, G. J. MURPHY, AND L. TUSET, Differential calculi over quantum groups and twisted cyclic cocycles, *J. Geom. Phys.* **44**, Number 4, Pages 570-594, (2003).

- [66] N. P. LANDSMANN, Strict quantization of coadjoint orbits, *J. of Math. Physics*, 39, Number 12, (1998).
- [67] F. LATRÉMOLIÉRE, Approximation of Quantum Tori by Finite Quantum Tori for the Quantum Gromov–Hausdorff Distance, *Journal of Functional Analysis*, **223**, Pages 365-395, (2005), arXiv:math.OA/0310214 v4.
- [68] F. LATRÉMOLIÉRE, Bounded-Lipschitz distances on the state space of C^* -algebras, To appear in *Taiwainese Journal of Mathematics*, **11**, Number 2, (2007), math.OA/0510340.
- [69] H. B. LAWSON AND M. L. MICHELSON, Spin Geometry, *Princeton University Press*, Princeton, (1989).
- [70] H. LI, Order-unit quantum Gromov–Hausdorff distance, *Journal of Functional Analysis*, **231**, Number 2, Pages 312-360, (2006), arXiv:math.OA/0312001.
- [71] H. LI, C^* -algebraic quantum Gromov–Hausdorff distance, math.OA/0312003.
- [72] H. LI, Compact quantum metric spaces and ergodic actions of compact quantum groups, arXiv.math.OA/0411178.
- [73] J. LOTT, Collapsing and Dirac-type operators, *Geom. Dedicata* 91, Pages 175-196, 2002, arXiv:math.DG/0005009.
- [74] J. MADORE, The Fuzzy Sphere, *Class. Quant. Grav.* 9, Number 69, (1992).
- [75] J. MADORE, An Introduction to Noncommutative Geometry and its Physical Applications, *Cambridge University Press*, Cambridge, (1999).
- [76] I. MADSEN AND J. TORNEHAVE, From Calculus to Cohomology, *Cambridge University Press*, Cambridge, (1997).
- [77] X. MARTIN, Fuzzy Orbifolds, arXiv:hep-th/0405060.
- [78] T. MASUDA, Y. NAKAGAMI, AND J. WATANABE, Noncommutative differential geometry on the quantum $SU(2)$, I, *K-theory* 4, Pages 157-180, (1990).
- [79] J. MEDINA, W. BIETENHOLZ, F. HOFHEINZ, AND D.O' CONNOR, Field theory simulations on a fuzzy sphere an alternative to the lattice, *Proceedings of Science*, arXiv:hep-lat/0509162, (2005).

- [80] G. J. MURPHY, *C*-Algebras and Operator Theory*, *Academic Press*, San Diego–New York, (1990).
- [81] G. J. MURPHY, Differential calculi over quantum groups, *Noncommutative Geometry and Quantum Groups*, Banach Centre Publ. **61**, 157-170, (2003).
- [82] G. J. MURPHY AND L. TUSET, Aspects of Compact Quantum Group Theory, *Proceedings of the American Mathematical Society* **132**, Number 10, Pages 3055-3067, (2004).
- [83] M. NAKAHARA, *Geometry, Topology and Physics*, *Adam Hilger*, Bristol, (1990).
- [84] M. PASCHKE AND A. SITARZ, On Spin Structures and Dirac Operators on the Noncommutative Torus, arXiv:math.QA/0605191, To appear in *Lett. Math. Phys.*
- [85] R. PENROSE, *The Road to Reality: A Complete Guide to the Laws of the Universe*, *Johnathan Cape*, London, 2004.
- [86] R. PLYMEN, Strong Morita Equivalence, Spinors and Symplectic Spinors, *Journal of Operator Theory*, **16**, Number 2, Pages 305 - 324, (1986).
- [87] H. GROSSE, C. KLIMCIK, AND P. PRESNAJDER , Towards finite quantum field theory in noncommutative geometry, *Int. J. Theor. Phys.* **35**, Pages 231-24, arXiv:hep-th/9505175,(1996).
- [88] A. RENNIE AND J. C. VARILLY, Reconstruction of Manifolds in Noncommutative Geometry, arXiv:math/0610418v3.
- [89] N. Y. RESHETIKHIN, L. A. TAKHTAJAN, AND L. D. FADDEEV, Quantization of Lie groups and Lie algebras, *Leningrad Math. J.* **1**, Pages 193-225, (1990).
- [90] M. A. RIEFFEL *C*-algebras and quantization; C*-algebras: 1943-1953 A Fifty Year Celebration*, *American Mathematical Society*, Providence, Rhode Island, (1993).
- [91] M. A. RIEFFEL, Metrics on states from actions of compact groups, *Doc. Math.* Number 3, Pages 215-229, (1998), arXiv:math.OA/9807084.
- [92] M. A. RIEFFEL, Compact Quantum Metric Spaces, *AMS Special Session Operator Algebras, Quantization, and Noncommutative Geometry: A Centennial Celebration in Honor of J. V. Neumann and M. H. Stone, II*, editors R. S. Doran and R. V. Kadison, math.OA/0308207.

- [93] M. A. RIEFFEL, Matrix algebras converge to the sphere for quantum Gromov–Hausdorff distance, *Mem. Amer. Math. Soc.* 168 , Number 796, Pages 67-91, (2004).
- [94] M. A. RIEFFEL, Gromov–Hausdorff distance for quantum metric spaces, *Mem. Amer. Math. Soc.* 168, Number 796, Pages 1-65, (2004).
- [95] M. A. RIEFFEL, Lipschitz extension constants equal projection constants, arXiv:math.FA/0508097 v2.
- [96] M. A. RIEFFEL, Vector bundles and Gromov-Hausdorff distance, arXiv:math.MG/0608266.
- [97] W. RUDIN, Fourier Analysis on Groups, *Interscience*, New York, (1962).
- [98] K. SCHMÜDGEN, Commutator Representations of Differential Calculi on the Quantum Group $SU_q(2)$, *J. Geom. Phys.* **31**, Pages 241-264, (1999).
- [99] N. SEIBERG AND E. WITTEN, String Theory and Noncommutative Geometry, arXiv:hep-th/9908142.
- [100] J.-P. SERRE, Faisceaux algébriques cohérents, *Ann. of Math.* (2) **6**, Pages 197-278, (1955).
- [101] B. SIMON, Representations of Finite and Compact Groups, *American Mathematical Society*, Providence, RI, (1996).
- [102] A. SITARZ, Twisted Hochschild Homology of Quantum Hyperplanes, *K-theory*, **35**, (1-2), Pages 187 - 198, (2005), arXiv: math.QA/0405240.
- [103] J. SMIT, Introduction to Quantum Fields on a Lattice, *Cambridge University Press*, Cambridge, (2002).
- [104] R. G. SWAN, Vector bundles and projective modules, *Trans. Amer. Math. Soc.* **105** (1962), 264–277.
- [105] A. VAN DEALE, The Haar measure on a compact quantum group, *Proc. Amer. Math. Soc.* **123** Pages 3125–3128, (1995)
- [106] A. VAN DAELE, An algebraic framework for group duality, *Adv. Math.* **140**, Number 2, Pages 323-366, (1998).
- [107] N. WEAVER, Mathematical Quantization, *Chapman and Hall/CRC Press*, Boca Raton, Fla. (2001).

- [108] M. WODZICKI, Excision in cyclic homology and in rational algebraic K -theory, *Annals of Mathematics*, **129**, Pages 591-640, (1989).
- [109] J. A. WOLF, Essential selfadjointness for the Dirac operator and its square, *Indiana Univ. Math. J.* **22**, Pages 611-640, (1973).
- [110] S. L. WORONOWICZ, Twisted $SU(2)$ group, An example of a non-commutative differential calculus, *Publ. RIMS, Kyoto University*, **23**, Number 1, Pages 117–181, (1987).
- [111] S. L. WORONOWICZ, From multiplicative unitaries to quantum groups, *Int. J. Math.* **7**, Pages 127–149, 1996.
- [112] S. L. WORONOWICZ, Compact quantum groups, in ‘Symétries Quantique’ (Les Houches, 1995), North Holland, Amsterdam, Pages 845–884, (1998).
- [113] S. L. WORONOWICZ, ‘Compact and noncompact quantum groups’ Lectures of S. L. Woronowicz, written by L. Dabrowski and P. Norowski, SISSA preprint, <http://info.fuw.edu.pl/~nurowski/prace/wor.ps>.
- [114] K. ZHU, An Introduction to Operator Algebras, *Chapman and Hall/CRC Press*, Boca Raton, Fla. (1993).