

Noncommutative Geometry

Ralf Meyer

CONTENTS

1. Coordinates and manifolds	1
2. Recovering a manifold from its algebra of smooth functions	7
3. Algebraic varieties	12
4. Representations and simple modules	17
5. Endomorphism algebras and finite-dimensional algebras	21
6. Group algebras	27
7. Category algebras and quiver algebras	33
8. The group algebra of the dihedral group	36
9. Involutive algebras	41
10. Crossed products	46
11. Morita equivalence	50
12. Derivations as a noncommutative analogue of differentiation	58
13. More on derivations, automorphisms and Lie algebra structure	63
14. Representations and crossed products for Lie algebras	67
15. The universal derivation — Morita invariance	73
16. From deformation quantisations to Hochschild cohomology	78
17. Computing Hochschild cohomology with projective resolutions	85
18. Hochschild cohomology of algebras of polynomials	89
19. Hochschild cohomology for algebras of smooth functions	94
20. Quasi-free algebras and their Hochschild cohomology	98
21. Hochschild cohomology of the Weyl algebra — De Rham cohomology	102
22. Noncommutative differential forms	107
23. Towards periodic cyclic homology	111
24. Periodic cyclic homology	114
References	120

1. Coordinates and manifolds

The goal of this course is to do geometry with noncommutative algebras. Before we can do this, we must first rewrite geometric objects in terms of commutative algebras. This is usually called algebraic geometry. Having done this, we may then try to encode our favourite geometric concepts in algebraic terms and apply these to noncommutative algebras.

Since I am more fond of differential geometry than of algebraic geometry, I choose to work with smooth manifolds instead of algebraic varieties. I will show that we may recover a smooth manifold from the commutative algebra of smooth

functions on it. In the first lecture, I will recall the concept of a smooth manifold. And I will motivate it by talking about coordinates. This is also a preparation for the algebraisation of smooth manifolds. Roughly speaking, the algebra of smooth functions is a way of choosing all possible coordinate functions on a manifold simultaneously.

A basic result about smooth manifolds is that they may be embedded into \mathbb{R}^n . This result is important for our intuition: it tells us that our definition of a smooth manifold is not too general. And it is crucial for some proofs later on. These are the three topics of the first lecture: coordinates, smooth manifolds, and their embeddings into \mathbb{R}^n .

1.1. Coordinates. René Descartes was, perhaps, the first to use coordinates systematically in order to solve geometric problems. His coordinate systems are also needed to apply mathematical analysis to the physical world. The idea of Descartes is to describe a point in the plane by two numbers. More generally, a point in d -dimensional space is encoded through an element of \mathbb{R}^d , consisting of its coordinates. We may calculate with these coordinates. Thus we may solve geometric problems by algebraic computations. More precisely, coordinates allow to describe a geometric figure through equations or inequalities in the coefficients of a vector. For the figures that appear in classical geometry such as lines, circles or cone sections, these equations and inequalities are at most quadratic functions of the coordinates. Thus we may solve geometric problems by solving linear and quadratic equations.

EXAMPLE 1.1. The *sphere* in 3-space \mathbb{R}^3 of radius R centred at the point (x_0, y_0, z_0) is described by the equation

$$(1.2) \quad (x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2 = R^2;$$

the corresponding inequality $(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2 < R^2$ describes the interior, and $(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2 > R^2$ the exterior of the sphere.

Suppose we want to find the points where three given spheres intersect. In coordinates, this becomes a system of three quadratic equations in three unknowns. For instance, let us intersect the sphere around the origin $(0, 0, 0)$ of radius 3, the sphere around the point $(0, 0, 1)$ of radius $\sqrt{6}$, and the sphere around the point $(0, 1, 1)$ of radius $\sqrt{3}$. This is equivalent to solving the system of equations

$$\begin{aligned} x^2 + y^2 + z^2 &= 9, \\ x^2 + y^2 + (z - 1)^2 &= 6, \\ x^2 + (y - 1)^2 + (z - 1)^2 &= 3. \end{aligned}$$

A short computation shows that the solution set of this system of equations is $\{(-1, 2, 2), (1, 2, 2)\}$.

EXERCISE 1.3. Let P, Q, R be three points on a circle in the plane; assume that P and R are antipodes. Show that the angle $\angle PQR$ at the third point is a right angle. Give two arguments, one algebraic and one geometric. Which argument do you prefer?

Descartes used coordinates to do geometry. Even more importantly, however, coordinates allow to apply *analysis*. Geometric quantities like the length of a curve, the area of a surface, tangent vectors, and so on, may be computed using differentiation and integration of suitable functions. The definition of the *derivative* of a function uses not just limits – a purely topological concept – but also differences and quotients. Without coordinates or a suitable substitute, we cannot apply differential calculus in geometry.

The analytic concepts above are crucial for Newtonian mechanics. The laws of mechanics say that the acceleration of a physical system is described through the forces that act on it. The acceleration is the second derivative. Without coordinates, this would make no sense. So we need something like coordinates in order to even formulate the laws of mechanics.

1.2. Manifolds. The original sense of the word geometry is to measure the earth. Actually, that subject is nowadays called geodesy. One reason for this may be that classical Euclidean geometry does not apply to our earth – it is not flat. The earth is not exactly round either, but it is a reasonable first approximation to replace the earth by the surface of a sphere in \mathbb{R}^3 .

We may describe points on the sphere by the three coordinates from \mathbb{R}^3 . This is highly inefficient, however, because the sphere has only dimension 2. We usually describe a point on the earth by only two coordinates, namely, latitude and longitude. This idea was already proposed by Eratosthenes back in 300 BC.

The coordinates of Eratosthenes do not work well at the two poles, however. There the longitude is ill defined. For an object at one of the two poles, the longitude–latitude coordinates do not give a meaningful concept of velocity. Since the sphere is round, we may use two such coordinate systems centred on different axes. One of the two coordinate systems will work well at any given point on the sphere. When we get near to one of the axes, we switch to the other coordinate system. This is a simple example of a local coordinate system. This way of navigating on the earth is also familiar to those who still use a printed road atlas. Each of its pages shows a certain area. When you get to the boundary of the page, you must go to another page. Hopefully, there is some overlap among the pages, so that you do not get lost near the boundary of a page.

DEFINITION 1.4. A *d-dimensional topological manifold* is a paracompact, Hausdorff topological space (these technical conditions will be explained later) with the property that each point has an open neighbourhood that is homeomorphic to \mathbb{R}^d (or to an open subset of \mathbb{R}^d – this yields an equivalent definition). Such a homeomorphism is called a *chart*, and its domain is a *chart neighbourhood*.

EXAMPLE 1.5. The unit sphere \mathbb{S}^d in \mathbb{R}^{d+1} with its subspace topology is a *d-dimensional topological manifold*. To see this, we must describe a chart near each point $\vec{x} = (x_0, \dots, x_d) \in \mathbb{S}^d$. Since $x_0^2 + \dots + x_d^2 = 1$, we may choose i with $x_i \neq 0$. Then the map $\mathbb{S}^d \rightarrow \mathbb{R}^d$ that forgets the i th coordinate is a chart in the open neighbourhood of \vec{x} defined by the equation $x_i \neq 0$.

These projection charts distort both distances and angles on the sphere. It is, in fact, impossible to find charts on the sphere that preserve distances because the sphere has non-zero curvature. But *stereographic projection* provides another system of charts that at least preserve angles.

Let N be a point in \mathbb{S}^d , let S be its antipode, and let $T \cong \mathbb{R}^d$ be the tangent hyperplane to \mathbb{S}^d through S . The stereographic projection from N maps each point $P \in \mathbb{S}^d \setminus \{N\}$ to the intersection P' of the line NP with the hyperplane T (see Figure 1). It is easy to check that this map is a homeomorphism $\varphi_N: \mathbb{S}^d \setminus \{N\} \rightarrow T$.

Important concepts such as differentiability of functions depend on the choice of charts. They cannot be defined on topological manifolds because these do not carry distinguished system of local coordinates. The concept of smooth manifold remedies this by choosing an “atlas”, which is a system of compatible local charts.

Let $\varphi: U \rightarrow \mathbb{R}^d$ and $\psi: V \rightarrow \mathbb{R}^d$ be charts on overlapping chart neighbourhoods. The *change of coordinate map*

$$(1.6) \quad \psi \circ \varphi^{-1}: \varphi(U \cap V) \rightarrow \psi(U \cap V)$$




FIGURE 1. The stereographic projection. For fixed antipodal points N and S , it maps P to the point P' where the tangent plane T to the sphere at S intersects the line through N and P .

transforms coordinates between the two local coordinate systems on the overlap $U \cap V$. The definition of a topological manifold guarantees that this is a homeomorphism between two open subsets of \mathbb{R}^d . This is the only constraint: any such local homeomorphism is possible. Since homeomorphisms need not be differentiable, there is no relation between the differentiability of $f \circ \varphi^{-1}$ and $f \circ \psi^{-1}$ for a function f . We refine our definition of a manifold to make differentiability well defined. Recall that a map between open subsets of \mathbb{R}^d is called *smooth* if it has partial derivatives of arbitrarily high order.

DEFINITION 1.7. An *atlas* on a d -dimensional topological manifold M is a family of charts $\mathfrak{U} = \{\varphi: U_\varphi \rightarrow \mathbb{R}^d\}$ such that the domains U_φ for $\varphi \in \mathfrak{U}$ cover M and the change of coordinate maps in (1.6) for $\varphi, \psi \in \mathfrak{U}$ are smooth. A *smooth manifold* is a topological manifold together with such an atlas.

Since we will mainly work with smooth manifolds, we agree that, from now on, “manifold” means “smooth manifold.”

REMARK 1.8. It is common and useful to require the atlas in the definition of a smooth manifold to be “maximal”. This means that for any local chart not already in the atlas, the coordinate change function to some local chart in the atlas will fail to be smooth. This maximality assumption makes the atlas more unique. In practice, however, we never specify such a maximal atlas. We always give a rather small family of local charts.

Whereas being a topological manifold is a *property* of topological spaces, being a smooth manifold is an additional *structure* – the atlas. Nevertheless, we usually drop the atlas from our notation for simplicity.

DEFINITION 1.9. Let X with the atlas $\mathfrak{U} = \{\varphi: U_\varphi \rightarrow \mathbb{R}^d\}$ be a smooth manifold and let $k \in \mathbb{N}$. A function $f: X \rightarrow \mathbb{R}^k$ is *smooth* if $f \circ \varphi^{-1}: \mathbb{R}^d \rightarrow \mathbb{R}^k$ is smooth for all $\varphi \in \mathfrak{U}$.

DEFINITION 1.10. Let X and Y be smooth manifolds of dimensions d_X and d_Y , respectively. A map $f: X \rightarrow Y$ is called *smooth* if it is continuous and if, for all $x \in X$ and smooth charts $\varphi: U_x \rightarrow \mathbb{R}^{d_X}$, $\psi: U_{f(x)} \rightarrow \mathbb{R}^{d_Y}$ on neighbourhoods U_x and $U_{f(x)}$ of x and $f(x)$, the map $\psi \circ f \circ \varphi^{-1}$ from $\varphi(f^{-1}(U_{f(x)})) \subseteq \mathbb{R}^{d_X}$ to \mathbb{R}^{d_Y} is smooth.

The smooth manifolds with smooth maps form a category, the category of smooth manifolds. The isomorphisms in this category are the diffeomorphisms:

DEFINITION 1.11. A *diffeomorphism* between two smooth manifolds is a smooth bijection whose inverse is also smooth. Two smooth manifolds are *diffeomorphic* if there is a diffeomorphism between them.

REMARK 1.12. There are topological manifolds without a smooth structure. More surprisingly, a topological manifold may have many smooth structures that are essentially different: the resulting smooth manifolds are not diffeomorphic. For instance, this happens for high-dimensional spheres. Topologists have developed homological tools to classify different smooth structures on a given topological manifold, but we do not discuss this here.

REMARK 1.13. The concept of a smooth manifold allows many variations because we may impose other regularity conditions on the change of coordinate maps. We may require them to be piecewise linear, k times continuously differentiable for some $k \in \mathbb{N}$, Lipschitz, or real-analytic; this leads to other kinds of manifolds: PL-manifolds, C^k -manifolds, Lipschitz manifolds, and real-analytic manifolds. Some of these concepts are closely related; for instance, any C^1 -manifold carries an essentially unique real-analytic structure. Replacing \mathbb{R}^k by \mathbb{C}^k , we may require the change of coordinate maps to be holomorphic or given by rational functions. This leads to complex manifolds and complex algebraic varieties, respectively.

The definition of a topological manifold contains two global assumptions – Hausdorffness and paracompactness. The Hausdorffness assumption means that any two points may be separated by open neighbourhoods. We will see some interesting non-Hausdorff manifolds later and learn how to describe them in noncommutative geometry. Paracompactness is a condition on open coverings of the space that makes partitions of unity work. These are the main tools to globalise local constructions in a manifold.

DEFINITION 1.14. Let \mathcal{U} be an open covering of a topological space X , that is, a set of open subsets with $\bigcup_{U \in \mathcal{U}} U = X$. A *partition of unity subordinate to \mathcal{U}* is a family of continuous functions $\varphi_U: X \rightarrow [0, 1]$ for $U \in \mathcal{U}$ with the following properties:

- the support of φ_U is contained in U ; recall that the *support* of φ_U is the closure of the set $\{x \in X : \varphi_U(x) \neq 0\}$;
- each $x \in X$ has a neighbourhood V such that $\varphi_U|_V = 0$ for all but finitely many φ_U ;
- $\sum_{U \in \mathcal{U}} \varphi_U(x) = 1$.

The partition of unity is called *smooth* if X is a smooth manifold and all the functions φ_U are smooth.

THEOREM 1.15. A Hausdorff topological space is paracompact if any open covering of it has a subordinate partition of unity.

Any open covering of a smooth manifold has a subordinate smooth partition of unity.

We take the first part of this theorem as our definition of paracompactness. The second part is proved using the existence of smooth functions with compact support on \mathbb{R} . These may be constructed from the function $\exp(-1/x^2): (0, \infty) \rightarrow \mathbb{R}$, which remains smooth when we extend it by 0 on $(-\infty, 0]$. Since there are no real-analytic or holomorphic functions with compact support, there are no real-analytic or holomorphic partitions of unity.

We need one more condition, which is not part of the standard definition of a manifold, but which we are going to assume anyway:

DEFINITION 1.16. A Hausdorff space X is σ -compact if there is a sequence of compact subsets $L_n \subseteq X$ with $\bigcup_{n \in \mathbb{N}} L_n = X$.

EXAMPLE 1.17. An uncountable set with the discrete topology is a 0-dimensional manifold that is not σ -compact.

1.3. Embedding smooth manifolds. Our next theorem asserts that any smooth manifold may be identified with a closed submanifold of \mathbb{R}^N for some N . This means that our abstract concept of manifold is equivalent to the more concrete concept of a submanifold of \mathbb{R}^n , showing that our definition is not more general than necessary. Before we can identify abstract manifolds with submanifolds, we have to discuss the concept of “embedding.”

DEFINITION 1.18. A map between two locally compact topological spaces is called *proper* if pre-images of compact subsets are again compact.

DEFINITION 1.19. An *embedding* $X \hookrightarrow Y$ for two smooth manifolds X and Y is an injective, smooth map $f: X \rightarrow Y$ that is a homeomorphism onto its image and whose first derivatives – computed in local coordinates – are injective at each point.

The definition of the derivative of a smooth map requires the tangent bundle. Nevertheless, the *injectivity* of the derivative is already defined using only local coordinates. It is coordinate-independent because of the chain rule.

The range $f(X)$ of an embedding $f: X \rightarrow Y$ is a submanifold of Y . The embedding f is proper if and only if this submanifold $f(X)$ is closed.

The *Tietze Extension Theorem* asserts that a continuous function on a closed subset of a “normal” topological space extends to a continuous function on the whole space. Proper embeddings are important because of a smooth analogue of this theorem:

THEOREM 1.20. *Let X and Y be smooth manifolds and let $f: X \rightarrow Y$ be a proper embedding. Then any smooth function $h: X \rightarrow \mathbb{R}$ extends to a smooth function $\bar{h}: Y \rightarrow \mathbb{R}$, that is, $\bar{h}(f(x)) = h(x)$ for all $x \in X$. In addition, if $U \subseteq Y$ is a neighbourhood of $f(X)$, then we can assume that \bar{h} vanishes outside U .*

PROOF. This follows easily from the much stronger Tubular Neighbourhood Theorem (see [8, p. 109ff]). The latter yields a smooth local retraction $r: U \rightarrow f(X) \cong X$ in a neighbourhood U of $f(X)$; here the various requirements on a proper embedding are needed, compare the following exercise. In addition, let $\varphi_U + \varphi_{Y \setminus f(X)} = 1$ be a partition of unity on Y subordinate to the open covering $\{U, Y \setminus f(X)\}$. Then $\bar{h} := \varphi_U \cdot (h \circ r)$ does the job. \square

EXERCISE 1.21. For each of the three conditions for a proper embedding – proper, injective, injective derivative – find a smooth map $\mathbb{R} \rightarrow \mathbb{R}^2$ that lacks this property but has the other two properties. Check that the Extension Theorem becomes false in each of these examples.

THEOREM 1.22 (see [3, Theorem 10.8]). *Let X be a smooth d -dimensional manifold. Assume that X is also σ -compact. There is a proper embedding $X \rightarrow \mathbb{R}^N$ for some $N \in \mathbb{N}$. Even more, we can choose $N = 2d + 1$.*

PROOF. We only sketch the construction of an embedding $X \rightarrow \mathbb{R}^N$ for some $N \in \mathbb{N}$ if X is compact. So we do not treat the non-compact case, and we do not reduce the dimension N of the target space. Compactness ensures that X can be covered by finitely many chart neighbourhoods. Let $\varphi_j: U_j \rightarrow \mathbb{R}^d$, $j = 1, \dots, k$, be the corresponding charts and let $\psi_j: X \rightarrow \mathbb{R}$ be a subordinate smooth partition of unity (see Theorem 1.15). Then $(\psi_j \cdot \varphi_j, \psi_j)$ extends by 0 to a smooth function $X \rightarrow \mathbb{R}^{d+1}$. Let

$$f := (\psi_j \cdot \varphi_j, \psi_j)_{1 \leq j \leq k}: X \rightarrow \mathbb{R}^{(d+1)k}.$$

It is routine to check that f is an embedding. Since X is compact, any continuous map $X \rightarrow \mathbb{R}^N$ is proper. \square

This proof depends on partitions of unity. In fact, the theorem becomes false if we drop the paracompactness assumption in the definition of a manifold; the “long line” provides a counterexample.

REMARK 1.23. Any closed subset of \mathbb{R}^N is paracompact, σ -compact and Hausdorff. So the Embedding Theorem requires all the technical extra conditions. A hidden assumption in Theorem 1.22 is that the dimensions of the connected components of a manifold should be bounded above. For instance, the disjoint union $\bigsqcup_{n=1}^{\infty} \mathbb{R}^n$ is a smooth manifold that does not embed into any \mathbb{R}^N . In the following, when we speak of manifolds, we tacitly assume them to be σ -compact and to have bounded above dimension. Then they admit a proper embedding into \mathbb{R}^N for some $N \in \mathbb{N}$.

2. Recovering a manifold from its algebra of smooth functions

Although local coordinates are used to define manifolds, they are only an auxiliary notion. Important definitions should be independent of the local coordinates. How can we specify a smooth manifold without singling out one atlas? We may avoid this by restricting to a maximal atlas, which consists of *all* possible local coordinate systems – unlike a real-life atlas. The algebra of smooth functions we are about to define solves the problem in another, more algebraic way. We will see that it uniquely determines the underlying manifold up to diffeomorphism. We also describe smooth maps between manifolds in terms of their algebras of smooth functions.

First we need some concepts from algebra. Then we define the algebra $C^\infty(X)$ of smooth functions on a smooth manifold. The central concept in this section are characters. We show that the points of the manifold are in bijection with characters $C^\infty(X) \rightarrow \mathbb{C}$. We describe the topology on the manifold as the Zariski topology on the space character. We show that smooth maps $X \rightarrow Y$ between manifolds are in bijection with algebra homomorphisms $C^\infty(Y) \rightarrow C^\infty(X)$.

As a result, up to diffeomorphism a smooth manifold X may be recovered from the algebra $C^\infty(X)$ of smooth functions on X . This is a key result: it says that we may replace a geometric object by an algebraic one without losing any information. Noncommutative geometry goes a step further and considers a larger class of algebraic objects as describing some generalised geometric objects. One open problem here is that we cannot characterise nicely when a commutative algebra corresponds to a smooth manifold. (A partial result in this direction is [13, Lemma 2.4].)

A technically difficult aspect in our results and proofs is the continuity of characters. We treat arbitrary characters and prove along the way that they are automatically continuous. This uses an embedding of X into \mathbb{R}^n and some results about the local structure of smooth functions. It would have been easier to prove a result only for the continuous characters. This result would, however, use more structure, namely, the standard topology on the algebra $C^\infty(X)$. I prefer to prove a simpler result with a slightly more difficult proof.

DEFINITION 2.1. Let K be a field. A K -algebra – briefly algebra – is a K -vector space A with a map $m: A \times A \rightarrow A$, $(x, y) \mapsto x \cdot y$, called *multiplication*, which is bilinear and *associative*:

$$(2.2) \quad (a \cdot b) \cdot c = a \cdot (b \cdot c) \quad \text{for all } a, b, c \in A.$$

An algebra A is *commutative* if

$$(2.3) \quad a \cdot b = b \cdot a \quad \text{for all } a, b \in A.$$

A *unit element* is an element $1_A \in A$ that satisfies

$$(2.4) \quad 1_A \cdot a = a = a \cdot 1_A \quad \text{for all } a \in A.$$

An algebra with a unit element is called *unital*.

LEMMA 2.5. *A unit element is unique if it exists.*

PROOF. Copy the argument for groups. \square

EXAMPLES 2.6. The ground field K with the usual multiplication is the simplest example of a unital K -algebra.

The polynomials in one variable with coefficients in K form a commutative, unital K -algebra $K[X]$. So do the polynomials in n commuting variables X_1, \dots, X_n .

The $n \times n$ -matrices with entries in K form a unital K -algebra $M_n K$. It is commutative only for $n = 1$.

We will meet many more examples of algebras, mainly over the fields \mathbb{R} or \mathbb{C} . We often write \mathbb{K} to denote either \mathbb{R} or \mathbb{C} .

DEFINITION 2.7. Let \mathbb{K} be either \mathbb{R} or \mathbb{C} , and let X be a smooth manifold. Let $C^\infty(X, \mathbb{K})$ be the set of all smooth functions $X \rightarrow \mathbb{K}$, equipped with the pointwise addition, scalar multiplication, and multiplication:

$$(f_1 + f_2)(x) := f_1(x) + f_2(x), \quad (\lambda \cdot f)(x) := \lambda \cdot f(x), \quad (f_1 \cdot f_2)(x) := f_1(x) \cdot f_2(x)$$

for $f_1, f_2, f \in C^\infty(X, \mathbb{K})$, $\lambda \in \mathbb{K}$.

It is straightforward to verify that $C^\infty(X, \mathbb{K})$ with this additional structure is a commutative, unital \mathbb{K} -algebra; the unit element is the constant function 1. For most of the following, it makes no difference whether we work over \mathbb{R} or \mathbb{C} . Then $C^\infty(X)$ denotes either $C^\infty(X, \mathbb{R})$ or $C^\infty(X, \mathbb{C})$.

Our next goal is to recover the manifold X from $C^\infty(X)$.

DEFINITION 2.8. Let K be a field and let A and B be K -algebras. An *algebra homomorphism* from A to B is a map $f: A \rightarrow B$ that is K -linear and satisfies $f(a \cdot b) = f(a) \cdot f(b)$ for all $a, b \in A$. If A and B are unital, we call f *unital* if $f(1_A) = 1_B$.

EXERCISE 2.9. Define $f^*(x) := \overline{f(x)}$ for all $x \in X$, $f \in C^\infty(X, \mathbb{C})$. Does this define an algebra homomorphism $*$: $C^\infty(X, \mathbb{C}) \rightarrow C^\infty(X, \mathbb{C})$?

DEFINITION 2.10. Let K be a field. A (K -valued) *character* on a K -algebra A is a non-zero algebra homomorphism $A \rightarrow K$.

LEMMA 2.11. *Let A be a unital K -algebra. A K -algebra homomorphism $A \rightarrow K$ is unital if and only if it is non-zero.*

PROOF. If $f: A \rightarrow K$ is unital, then $f \neq 0$ is clear. Conversely, if $f \neq 0$, then there is $a \in A$ with $f(a) \neq 0$. Then $f(1) \cdot f(a) = f(1 \cdot a) = f(a)$ implies $f(1) = 1$. \square

THEOREM 2.12. *Let X be a smooth manifold and let \mathbb{K} be \mathbb{R} or \mathbb{C} . For each $x \in X$,*

$$\text{ev}_x: C^\infty(X, \mathbb{K}) \rightarrow \mathbb{K}, \quad f \mapsto f(x),$$

is a character. Conversely, any character is of this form. This yields a canonical bijection between the underlying set of X and the set of characters on $C^\infty(X, \mathbb{K})$.

PROOF. It is obvious that ev_x is a character. The converse direction is the interesting one. Our proof uses the Embedding Theorem 1.22 to reduce this to a problem about smooth functions on \mathbb{R}^N . The Embedding Theorem yields a proper embedding $f: X \rightarrow \mathbb{R}^N$ for some $N \in \mathbb{N}$. Let f_1, \dots, f_N be the coefficients of f . Let $\chi: C^\infty(X, \mathbb{K}) \rightarrow \mathbb{K}$ be a character and let $x_j = \chi(f_j)$. We claim that

$(x_1, \dots, x_N) = f(x)$ for some $x \in X$ and that $\chi = \text{ev}_x$. We may replace the functions f_j by $f_j - x_j$ without losing the proper embedding property. Thus we may assume without loss of generality that $x_1 = \dots = x_N = 0$, which we do from now on to simplify.

Let $h \in C^\infty(X, \mathbb{K})$; if there is a – necessarily unique – $x \in X$ with $f(x) = \vec{0}$, we also assume $h(x) = 0$. We want to show that $\chi(h) = 0$. The Extension Theorem 1.20 provides a smooth function $\bar{h}: \mathbb{R}^N \rightarrow \mathbb{K}$ with $h = \bar{h} \circ f$ and $\bar{h}(\vec{0}) = 0$.

LEMMA 2.13. *For any $\bar{h} \in C^\infty(\mathbb{R}^N)$ with $\bar{h}(\vec{0}) = 0$, there are smooth functions \bar{h}_j on \mathbb{R}^N with*

$$\bar{h}(\vec{y}) = y_1 \bar{h}_1(y) + \dots + y_N \bar{h}_N(y) \quad \text{for all } \vec{y} = (y_1, \dots, y_N) \in \mathbb{R}^N.$$

PROOF. We only sketch the proof of this technical lemma. The assertion is evident if \bar{h} is a polynomial. For general \bar{h} , we may construct \bar{h}_j recursively. We start with

$$\bar{h}_1(y_1, \dots, y_N) := \begin{cases} \frac{\bar{h}(y_1, \dots, y_N) - \bar{h}_1(0, y_2, \dots, y_N)}{y_1} & \text{for } y_1 \neq 0, \\ \partial_{y_1} \bar{h}(y_1, \dots, y_N) & \text{for } y_1 = 0. \end{cases}$$

An approximation of \bar{h} by Taylor polynomials shows that this function is smooth. The function $\bar{h} - y_1 \bar{h}_1$ depends only on y_2, \dots, y_N . Repeating the step above for this difference and the variable y_2 , we get a smooth function \bar{h}_2 so that $\bar{h} - y_1 \bar{h}_1 - y_2 \bar{h}_2$ depends only on y_3, \dots, y_N . We recursively build the functions $\bar{h}_3, \dots, \bar{h}_N$ so that $\bar{h} - y_1 \bar{h}_1 - \dots - y_N \bar{h}_N$ is constant. Since this function vanishes at 0, we get $\bar{h} = y_1 \bar{h}_1 + \dots + y_N \bar{h}_N$. \square

The lemma allows us to rewrite $h = \sum_{j=1}^N f_j \cdot (\bar{h}_j \circ f)$. Hence

$$\chi(h) = \sum_{j=1}^N \underbrace{\chi(f_j)}_0 \chi(\bar{h}_j \circ f) = 0.$$

If $\vec{0} \notin f(X)$, our computation yields $\chi(h) = 0$ for all $h \in C^\infty(X)$ – which is impossible because the zero map is not a character by convention. Thus $\vec{0} = f(x)$ for some $x \in X$ and $\chi(h) = 0$ if $h(x) = 0$. Now write an arbitrary function $g \in C^\infty(X)$ as $g = g(x) \cdot 1 + h$ with the constant function 1 and $h(y) := g(y) - g(x)$ for all $y \in X$. Then

$$\chi(g) = \chi(g(x) \cdot 1) + \chi(h) = g(x) \cdot \chi(1) + 0 = g(x) = \text{ev}_x(g)$$

because $h(x) = 0$ and χ is \mathbb{K} -linear and unital (Lemma 2.11). \square

EXERCISE 2.14. Let X be a smooth manifold of positive dimension. Prove that $C^\infty(X, \mathbb{R})$ has uncountable dimension.

Functional analysis is the study of such vector spaces of uncountable dimension. In our case, $C^\infty(X)$ carries a canonical topology that turns it into a topological algebra. If X has only countably many connected components, then this topology is metrisable, and we may as well describe it by its convergent sequences:

DEFINITION 2.15. A sequence $(f_n)_{n \in \mathbb{N}}$ in $C^\infty(X)$ converges to $f \in C^\infty(X)$ if and only if for each chart $\varphi: U \rightarrow \mathbb{R}^d$, each compact subset $L \subseteq \varphi(U)$, and each multi-index α , the sequence of functions $\partial^\alpha(f_n \circ \varphi^{-1})|_L: L \rightarrow \mathbb{K}$ converges uniformly towards $\partial^\alpha(f \circ \varphi^{-1})|_L$.

I will say more about this topology when it is needed seriously. Here I want to point out that using this topology usually *simplifies* matters. For instance, the proof that *continuous* characters on $C^\infty(X)$ are of the form ev_x for some $x \in X$ is elementary, using no more than local coordinates and partitions of unity:

PROPOSITION 2.16. *Any continuous character on $C^\infty(X)$ is equal to ev_x for some $x \in X$.*

PROOF. The proof is done by contradiction. Let $\chi: C^\infty(X, \mathbb{K}) \rightarrow \mathbb{K}$ be a character with $\chi \neq ev_x$ for all $x \in X$. We are going to prove that $\chi(1) = 0$, which is impossible.

By assumption, there is $f \in C^\infty(X)$ with $\chi(f) \neq f(x)$. Since $\chi(1) = 1$, we may subtract $\chi(f) \cdot 1$ from f to get $f_x \in C^\infty(X)$ with $\chi(f_x) = 0$ and $f_x(x) \neq 0$. Then $\chi(f_x \cdot \overline{f_x}) = 0$ as well. The function $|f_x|^2 = f_x \cdot \overline{f_x}$ is positive in a neighbourhood U_x of x . Any compact subset L of X is covered by finitely many such open subsets $(U_{x_i})_{1 \leq i \leq n}$. Then $\sum_{i=1}^n |f_{x_i}|^2$ is positive on L and in the kernel of χ . Multiplying with a suitable function, we get a smooth function g_L that is 1 on L and killed by χ . Let (L_n) be an increasing sequence of compact subsets with $\bigcup L_n = X$. Then the functions g_{L_n} converge to 1 in $C^\infty(X)$. Since χ is continuous, it follows that $\chi(1) = 0$. This is impossible. So there must be $x \in X$ with $\chi = ev_x$. \square

We needed some geometric input – the Embedding Theorem and the Extension Theorem – to get the same statement for potentially discontinuous characters.

Since the characters ev_x for $x \in X$ are clearly continuous, Theorem 2.12 shows that all characters on $C^\infty(X)$ are continuous. Such automatic continuity results are an important theme in functional analysis.

So far we have only recovered X as a set. Next we describe a canonical topology on the set of characters, which recovers the topology on X for $C^\infty(X)$.

DEFINITION 2.17. Let A be a commutative K -algebra and let \hat{A} be its set of characters. Let S be a subset of \hat{A} . Its *Zariski closure* \overline{S} is the set of all $\chi \in \hat{A}$ with $\chi(a) = 0$ for all $a \in A$ with $\omega(a) = 0$ for all $\omega \in S$ – that is, \overline{S} is the largest subset such that $\omega(a) = 0$ for all $\omega \in S$ implies $\chi(a) = 0$ for all $\chi \in \overline{S}$.

LEMMA 2.18. *This closure operation satisfies the Kuratowski closure axioms and therefore defines a topology on \hat{A} .*

This topology is called the *Zariski topology* or the *Jacobson topology*.

PROOF. For $S \subseteq \hat{A}$, let

$$I_S := \{a \in A : \chi(a) = 0 \text{ for all } \chi \in S\},$$

this is an ideal in A . By definition, $\chi \in \overline{S}$ if and only if χ vanishes on I_S . Since characters in S vanish on I_S , we get $S \subseteq \overline{S}$. Since $I_\emptyset = A$, we get $\overline{\emptyset} = \emptyset$. In addition, $I_S = I_{\overline{S}}$ by definition, so that the closure of \overline{S} is again \overline{S} .

It remains to verify $\overline{S_1 \cup S_2} = \overline{S_1} \cup \overline{S_2}$ for $S_1, S_2 \subseteq \hat{A}$. Clearly, $I_{S_1 \cup S_2} = I_{S_1} \cap I_{S_2}$. Hence a character in $\overline{S_1 \cup S_2}$ vanishes on $I_{S_1 \cup S_2}$, that is, $\overline{S_1 \cup S_2} \subseteq \overline{S_1} \cup \overline{S_2}$. Conversely, suppose that $\chi \notin \overline{S_1} \cup \overline{S_2}$. Hence there are $a_1 \in I_{S_1}$ with $\chi(a_1) \neq 0$ and $a_2 \in I_{S_2}$ with $\chi(a_2) \neq 0$. Then $\chi(a_1 \cdot a_2) \neq 0$ and

$$a_1 \cdot a_2 \in I_{S_1} \cdot I_{S_2} \subseteq I_{S_1} \cap I_{S_2} = I_{S_1 \cup S_2},$$

so that $\chi \notin \overline{S_1 \cup S_2}$. \square

LEMMA 2.19. *A subset of \hat{A} is Zariski closed if and only if it is of the form $\widehat{A/I} := \{\chi \in \hat{A} : \chi|_I = 0\}$ for some ideal $I \subseteq A$.*

PROOF. Let $S \subseteq \hat{A}$. Let $I_S := \{a \in A : \chi(a) = 0 \text{ for all } s \in S\}$ be its vanishing ideal. By definition, the Zariski closure of S is $\widehat{A/I_S}$, which has the asserted form. Let $I \subseteq A$ be an ideal. Then $\widehat{A/I}$ is equal to its own Zariski closure. \square

LEMMA 2.20. *The canonical bijection between X and the character space of $C^\infty(X)$ becomes a homeomorphism for the Zariski topology. That is, a point $x \in X$ lies in the closure of a subset S of X if and only if $\text{ev}_x(f) = 0$ for all $f \in C^\infty(X)$ with $f(y) = 0$ for all $y \in S$.*

PROOF. If x belongs to the closure of S , then $f(x) = 0$ for all continuous functions f with $f|_S = 0$. Conversely, if x does not belong to the closure of S , then there is a smooth function f that is supported in a small neighbourhood of x that does not meet S – this is an easy special case of the Extension Theorem 1.20. Then $f(x) \neq 0$ but $f|_S = 0$. \square

PROPOSITION 2.21. *A map $f: X \rightarrow Y$ is smooth if and only if $h \mapsto h \circ f$ defines a map $f^*: C^\infty(Y) \rightarrow C^\infty(X)$. This map f^* is automatically an algebra homomorphism. Any unital algebra homomorphism $C^\infty(Y) \rightarrow C^\infty(X)$ is of this form for a unique smooth map $f: X \rightarrow Y$. Thus smooth maps $f: X \rightarrow Y$ correspond bijectively to algebra homomorphisms $C^\infty(Y) \rightarrow C^\infty(X)$.*

In the language of category theory, $X \mapsto C^\infty(X, \mathbb{K})$ is a fully faithful functor from the category of smooth manifolds with smooth maps as morphisms to the category of unital \mathbb{K} -algebras with unital algebra homomorphisms as morphisms. Here \mathbb{K} may be \mathbb{R} or \mathbb{C} .

PROOF. It is clear that f^* maps $C^\infty(Y)$ to $C^\infty(X)$ and is a unital algebra homomorphism if f is smooth. Conversely, assume that f^* maps $C^\infty(Y)$ to $C^\infty(X)$. We have $\text{ev}_x \circ f^* = \text{ev}_{f(x)}$, and the map $\chi \mapsto \chi \circ f^*$ between the character spaces is automatically continuous for the Zariski topology. Since the latter agrees with the manifold topology, $f: X \rightarrow Y$ must be continuous.

Pick $x \in X$ and let $\varphi = (\varphi_1, \dots, \varphi_d): U \rightarrow \mathbb{R}^d$ be a chart in a neighbourhood of $f(x)$; let $\psi: Y \rightarrow \mathbb{R}$ be a smooth function supported in U with $\psi(f(x)) = 1$ in a neighbourhood of x . Then $\psi \cdot \varphi_j$ for $j = 1, \dots, d$ are elements of $C^\infty(Y)$, so that $f^*(\psi \cdot \varphi_j) \in C^\infty(X)$. Equivalently, $x \mapsto \psi(f(x)) \cdot \varphi(f(x))$ is a smooth function on X . Since $\psi = 1$ in a neighbourhood of $f(x)$, this function agrees with $\varphi \circ f$ in a neighbourhood of x – here we use that f is continuous. Thus $\varphi \circ f$ is smooth in a neighbourhood of x . Since the point x and the chart φ are arbitrary, f is smooth.

Now let $h: C^\infty(Y) \rightarrow C^\infty(X)$ be any unital algebra homomorphism. If $x \in X$, then $\text{ev}_x \circ h: C^\infty(Y) \rightarrow \mathbb{K}$ is a unital algebra homomorphism, that is, a character. By Theorem 2.12, $\text{ev}_x \circ h = \text{ev}_y$ for some $y \in Y$. Letting $f(x) := y$ defines a map $f: X \rightarrow Y$. By construction, $h(\psi)(x) = \psi(f(x))$ for all $x \in X$, that is, $h = f^*$. Hence f is smooth. \square

COROLLARY 2.22. *The topological space X and its smooth structure may be recovered from the algebra $C^\infty(X)$ of smooth functions. If $C^\infty(X)$ and $C^\infty(Y)$ are isomorphic algebras, then X and Y are diffeomorphic.*

PROOF. The last statement follows because C^∞ is a fully faithful functor; this is an instance of the Yoneda Lemma from category theory (see [12]). If $C^\infty(X) \cong C^\infty(Y)$ as algebras, then there are algebra isomorphisms in both directions, which must be of the form f^* and g^* for smooth maps $f: X \rightarrow Y$ and $g: Y \rightarrow X$. Since f^*g^* and g^*f^* are the identity homomorphisms, $g \circ f$ and $f \circ g$ are smooth maps that induce the identity morphism. Since the identity maps induce the same map on $C^\infty(X)$ and $C^\infty(Y)$, we get $g \circ f = \text{Id}_X$ and $f \circ g = \text{Id}_Y$ from the uniqueness of the map that induces an algebra homomorphism. Thus we may recover the diffeomorphism type of X from $C^\infty(X)$. Two smooth structures on X are considered the same if and only if the identity map is a diffeomorphism between them. So we also recover the smooth structure on X from $C^\infty(X)$. \square

REMARK 2.23. The claim that two smooth structures are the same if they have the same smooth functions is only literally true if we require the atlas that defines the smooth structure to be maximal. With our definition, the “same” smooth structure may be described by more than one atlas. We prefer to identify two smooth structures on a manifold X if the identity map on X is a diffeomorphism between them.

EXERCISE 2.24. Let $f: X \rightarrow Y$ be a smooth map between two smooth manifolds. Show that the induced algebra homomorphism $f^*: C^\infty(Y) \rightarrow C^\infty(X)$ is surjective if and only if f is a proper embedding, and injective if and only if f has dense range. Conclude that there is an injective algebra homomorphism $C^\infty(\mathbb{T}^2) \rightarrow C^\infty(\mathbb{R})$, where $\mathbb{T}^n = (\mathbb{S}^1)^{\times n}$ denotes the n -torus. Is there an injective algebra homomorphism $C^\infty(\mathbb{T}^2) \rightarrow C^\infty(\mathbb{T}^1)$ as well?

REMARK 2.25. In a similar way, we may recover a real-analytic manifold from its algebra of real-analytic functions, using an analogue of the embedding theorem in this category. The elementary argument about continuous characters works equally well for the *Banach* algebras of Lipschitz functions on a Lipschitz manifold or of C^k -functions on a C^k -manifold.

REMARK 2.26. It is hard to characterise those algebras that are of the form $C^\infty(X)$ for a manifold X . We know some necessary conditions – $C^\infty(X)$ must be commutative and unital – but these are far from sufficient.

The C^0 -case makes sense for any topological space. For a compact space X , the algebra $C(X, \mathbb{C})$ of continuous functions on X with the supremum norm and the pointwise complex conjugation is a C^* -algebra. The *Gelfand–Naimark Theorem* identifies the space of continuous characters on $C(X)$ with the underlying compact space X . We do not say more on C^* -algebras because they are treated in depth in another course. Instead, we focus on the differential geometric and algebraic aspects of noncommutative geometry. These do not require much analysis.

The goal of noncommutative geometry is to study algebras that are no longer commutative as if they were algebras of functions on some space.

3. Algebraic varieties

Algebraic geometry considers a different class of geometric objects. Roughly speaking, affine algebraic varieties are subsets of \mathbb{C}^n that can be defined by polynomial equations. We describe such varieties through their algebras of polynomial functions. And we also characterise which algebras arise in this fashion: they are exactly the commutative, finitely generated algebras with trivial radical (see Theorem 3.4). Algebraic varieties may also be described over the field \mathbb{R} . This makes an important difference because \mathbb{R} is no longer algebraically closed. Therefore, characters on an algebra over \mathbb{R} may take values in \mathbb{R} or in the larger field \mathbb{C} . We introduce maximal ideals in order to handle this more general concept of character. This paves the way for noncommutative generalisations because maximal ideals make sense also in noncommutative algebras – which typically admit no characters.

Algebraic geometry also considers more general objects than affine algebraic varieties. The problem is that there are no interesting examples of compact affine algebraic varieties over \mathbb{C} . To describe compact varieties, we must complete affine varieties in a suitable way. We briefly consider projective algebraic varieties. These are described by commutative graded algebras. Algebraic geometry also uses the idea of glueing affine algebraic varieties with suitable coordinate change maps – like for smooth manifolds. We do not explore this direction here, however. One reason is that our algebraic approach does not work directly for these kinds of

varieties. Namely, by Liouville's Theorem, the only polynomial functions on a compact complex algebraic variety are the constant functions. Algebraic geometry uses sheaves of rational functions with specified singularities to handle this. This does not yet have a fully developed noncommutative analogue. So we do not mention this in this course.

We also turn back to the algebra $C^\infty(X)$ and apply our new concept of maximal ideal to it. As it turns out, if X is compact, then all maximal ideals correspond to points in X as before. If X is non-compact, however, then there are more maximal ideals that live at ∞ . These maximal ideals are rather badly behaved. They contain the ideal $C_c^\infty(X)$ of smooth functions of compact support, so that they are dense in the standard topology on $C^\infty(X)$. The resulting quotient field of $C^\infty(X)$ has uncountable dimension. This shows that we really should treat $C^\infty(X)$ as a *topological algebra*.

DEFINITION 3.1. An *affine complex algebraic variety* is a subset V of \mathbb{C}^n that is defined by algebraic equations. That is, it is the solution set of a (usually finite) set of polynomial equations.

We may define the same variety by several different sets of polynomial equations; to avoid this ambiguity, we consider the set

$$I_V := \{p \in \mathbb{C}[x_1, \dots, x_n] : p|_V = 0\},$$

which is always an ideal in $\mathbb{C}[x_1, \dots, x_n]$. Since any polynomial equation that we may use to describe V is of the form $f = g$ with $f - g \in I_V$, our assumption on V is equivalent to

$$V = \{\vec{x} \in \mathbb{C}^n : p(\vec{x}) = 0 \text{ for all } p \in I_V\}.$$

Hilbert's Nullstellensatz asserts that an ideal in $\mathbb{C}[x_1, \dots, x_n]$ arises in this fashion if and only if it is *radical*, that is, if $p^n \in I_V$ for some $n \in \mathbb{N}$, then already $p \in I_V$. This yields a bijection between affine complex algebraic varieties and radical ideals in the polynomial ring $\mathbb{C}[x_1, \dots, x_n]$.

DEFINITION 3.2. The algebra of polynomial (usually called *regular*) functions on V is

$$\text{Pol}(V) := \mathbb{C}[x_1, \dots, x_n]/I_V.$$

The algebra $\text{Pol}(V)$ describes V independently of an embedding in \mathbb{C}^n .

EXAMPLE 3.3. Consider the variety in \mathbb{C}^2 defined by the equation $x^2 + y^2 = 1$. Its algebra of polynomial functions is the quotient of $\mathbb{C}[x, y]$ by the ideal generated by the polynomial $x^2 + y^2 - 1$. The classes of the functions $x^n y^k$ with $n \in \mathbb{N}$ and $k \in \{0, 1\}$ form a basis for this quotient (substitute $y^2 = 1 - x^2$ to reduce the exponent of y).

THEOREM 3.4. A \mathbb{C} -algebra is of the form $\text{Pol}(V)$ for an affine complex algebraic variety V if and only if it is commutative and finitely generated and its radical vanishes.

PROOF. Let A be a commutative, finitely generated \mathbb{C} -algebra, let a_1, \dots, a_n be a finite set of generators for A . Then there is a unique surjective homomorphism $\alpha: \mathbb{C}[x_1, \dots, x_n] \rightarrow A$ mapping x_j to a_j for all j . It descends to an isomorphism $\mathbb{C}[x_1, \dots, x_n]/\ker \alpha \cong A$. The ideal $\ker \alpha$ is radical because the radical of A vanishes by assumption. Hence it is of the form I_V for an affine complex algebraic variety $V \subseteq \mathbb{C}^n$. \square

EXAMPLE 3.5. The simplest example of a non-radical ideal is the ideal generated by the polynomial x^2 in $\mathbb{C}[x]$. The corresponding quotient ring $\mathbb{C}[x]/(x^2)$ has the basis $1 := 1 \bmod (x^2)$, $\varepsilon := x \bmod (x^2)$, where 1 is a unit element and $\varepsilon^2 = 0$. This

algebra is also called the algebra of *dual numbers*. The corresponding algebraic variety has only one point $0 \in \mathbb{C}$, which corresponds to the radical ideal (x) associated to (x^2) . The following exercise shows that certain computations in this algebra are related to differentiation of functions.

The dual numbers are a special case of the family of varieties $\mathbb{C}[x]/(x^2 - t)$ for $t \in \mathbb{C}$. For $t \neq 0$, this variety consists of two points $\pm\sqrt{t}$. These coincide for $t = 0$. The algebra $\mathbb{C}[x]/(x^2)$ remembers that the equation $x^2 = 0$ has a multiple solution at 0.

EXERCISE 3.6. Let A be the algebra of dual numbers of Example 3.5. Show that an element $a + b\varepsilon$ for $a, b \in \mathbb{C}$ is invertible if and only if $a \neq 0$, and compute its inverse. Let $f = \sum_{n=0}^{\infty} c_n x^n$ be a power series with infinite radius of convergence. Show that

$$f(a + b\varepsilon) := \sum_{n=0}^{\infty} c_n (a + b\varepsilon)^n$$

converges for all $a, b \in \mathbb{C}$ and describe the limit. In particular, check that

$$\exp(a + b\varepsilon) = \exp(a) + \exp(a)b\varepsilon.$$

A character on $\text{Pol}(V)$ is nothing but a character on $\mathbb{C}[x_1, \dots, x_n]$ that vanishes on I_V . It is easy to see that the character space of $\mathbb{C}[x_1, \dots, x_n]$ is just \mathbb{C}^n , given by evaluation homomorphisms. Hence the character set of $\text{Pol}(V)$ agrees with $V \subseteq \mathbb{C}^n$. In addition, any algebra homomorphism $\text{Pol}(V) \rightarrow \text{Pol}(V')$ is of the form f^* for some polynomial (regular) map $f^*: V' \rightarrow V$.

For algebraic varieties, the connections between algebra and geometry are even stronger than for smooth manifolds – but also less surprising. This is why we have discussed the example of smooth manifolds first: they are a more convincing example for the inevitability of algebraic methods in geometry because they appear, at first sight, to be much farther removed from algebra.

We have considered *complex* algebraic varieties above because the correspondence between algebra and geometry is particularly simple for them. More generally, the same things happen over any *algebraically closed* field instead of \mathbb{C} . Algebraic geometry over more general fields – such as the real numbers – is more complicated because the bijection between radical ideals and subsets of affine space breaks down (see the examples below). Thus the two possible definitions of an affine algebraic variety become different; it turns out that the more algebraic one works better:

DEFINITION 3.7. Let K be any field. An *affine algebraic variety* in K^n is a radical ideal I in $K[x_1, \dots, x_n]$. The algebra of *polynomials* (or regular functions) on this variety is the corresponding quotient algebra $\text{Pol}(V) := K[x_1, \dots, x_n]/I$.

EXAMPLE 3.8. Let $I \subseteq \mathbb{R}[x]$ be the ideal generated by $x^2 + 1$. The corresponding polynomial algebra is $\text{Pol}(V) := \mathbb{R}[x]/I \cong \mathbb{C}$. There is no character $\text{Pol}(V) \rightarrow \mathbb{R}$ because $x^2 + 1 = 0$ has no real solutions. There are two characters $\mathbb{C} \cong \text{Pol}(V) \rightarrow \mathbb{C}$: the identity map and complex conjugation. We interpret this statement as follows: the real algebraic variety defined by the equation $x^2 + 1 = 0$ has no real points and two complex points i and $-i$.

EXAMPLE 3.9. Now consider the real algebraic variety V described by the ideal in $\mathbb{R}[x, y]$ generated by $x^2 + y^2 - 1$. The \mathbb{R} -valued characters on $\text{Pol}(V) := \mathbb{R}[x, y]/(x^2 + y^2 - 1)$ correspond to the points on the unit circle in \mathbb{R}^2 . Characters $\text{Pol}(V) \rightarrow \mathbb{C}$ correspond to pairs $(x, y) \in \mathbb{C}^2$ with $x^2 + y^2 = 1$. Complex conjugation acts on the set of complex points. A complex point is real if and only if it is invariant under complex conjugation.

The correspondence between complex affine algebraic varieties and their polynomial algebras implies that a real affine algebraic variety V is determined by the set $V_{\mathbb{C}} \subseteq \mathbb{C}^n$ of its complex points, together with the complex conjugation map

$$\gamma: V_{\mathbb{C}} \rightarrow V_{\mathbb{C}}, \quad (x_1, \dots, x_n) \mapsto (\overline{x_1}, \dots, \overline{x_n}).$$

The real points of the variety are simply the γ -fixed points in $V_{\mathbb{C}}$. To recover the real affine variety according to Definition 3.7, we must recover the algebra $\text{Pol}(V)$. This is done as follows. Define a complex conjugation on $\text{Pol}(V_{\mathbb{C}})$ by $\overline{f}(x) := \overline{f(\overline{x})}$ or, equivalently, $\overline{\sum_{\alpha} c_{\alpha} x^{\alpha}} := \sum_{\alpha} \overline{c_{\alpha}} x^{\alpha}$, where the sum runs over multi-indices α and $c_{\alpha} \in \mathbb{C}$. Then $\text{Pol}(V)$ is the subalgebra of conjugation-invariant functions in $\text{Pol}(V_{\mathbb{C}})$.

To recover also the complex points of a real variety, we amend our definition of character:

DEFINITION 3.10. Let K be a field and let A be a K -algebra. A *character* on A is a surjective unital algebra homomorphism $A \rightarrow L$ for some field extension L of K . Two such characters $\varphi: A \rightarrow L$ and $\varphi': A \rightarrow L'$ are considered equivalent if there is an isomorphism $\lambda: L \rightarrow L'$ with $\lambda \circ \varphi = \varphi'$ and $\lambda|_K = \text{Id}_K$.

The surjectivity assumption ensures that an \mathbb{R} -valued character on a real affine algebraic variety is not counted as a \mathbb{C} -valued character as well.

DEFINITION 3.11. A *maximal* ideal in a K -algebra A is a proper ideal $I \subsetneq A$ such that there is no ideal strictly between I and A (that is, no ideal J with $I \subsetneq J \subsetneq A$).

A non-zero K -algebra A is called *simple* if it has no ideals besides $\{0\}$ and A .

Clearly, an ideal $I \subsetneq A$ is maximal if and only if the quotient algebra A/I is simple.

The commutative non-unital algebra K with zero multiplication is simple in the sense above, but there are good reasons for excluding it. There are interesting noncommutative examples of simple noncommutative algebras (such as the algebra of finite matrices introduced in (5.13)).

PROPOSITION 3.12. *Let A be a commutative, unital K -algebra. There is a canonical bijection between the set of maximal ideals in A and the set of equivalence classes of characters on A . It maps a character χ to its kernel $\ker \chi$.*

PROOF. First we claim that a commutative, unital algebra B is simple if and only if it is a field. It is clear that fields are simple. Conversely, let B be commutative, unital and simple and let $b \in B \setminus \{0\}$. Let (b) be the ideal generated by b . Since B is simple and $(b) \neq \{0\}$, $(b) = B$, so that $1 \in B$. Now $(b) = \{x \cdot b : x \in B\}$, so that there is $x \in B$ with $x \cdot b = 1$. Then $b \cdot x = 1$ as well because B is commutative. Thus any non-zero element of B is invertible, that is, B is a field.

If $I \subseteq A$ is a maximal ideal, then A/I is a commutative, unital, simple algebra, hence a field. The quotient map $A \rightarrow A/I$ is a character. Conversely, any character $\chi: A \rightarrow L$ induces – by the surjectivity requirement – an isomorphism $L \cong A/\ker \chi$. Since $A/\ker \chi$ is required to be a field, the ideal $\ker \chi$ is maximal. \square

Before we go on, we briefly mention some limitations of our approach of describing an algebraic variety or a smooth manifold through characters or maximal ideals on commutative algebras.

We have described an algebraic variety through a commutative unital ring. This only works for affine varieties, however. This is a serious restriction because affine complex algebraic varieties are only compact when they are finite, and many topological considerations only work for compact spaces. Therefore, we often want

to complete an affine algebraic variety $V \subseteq K^n$ to a *projective algebraic variety* $\bar{V} \subseteq K\mathbb{P}^n$ – a precise definition will not be needed for the following. The resulting object may be described by local coordinates in the same way as a smooth manifold; each of the coordinate neighbourhoods is a set-theoretic difference of two affine algebraic varieties, and the change of coordinate maps are rational functions with no poles in the subsets under consideration.

EXAMPLE 3.13. The projective space $\mathbb{C}\mathbb{P}^n$ is covered by the affine chart neighbourhoods $U_i := \{[x_0 : \cdots : x_n] \in \mathbb{C}\mathbb{P}^n : x_i \neq 0\}$, each of which is identified with the subspace $\mathbb{C}^n \setminus \{\vec{0}\}$ by the bijection

$$U_i \rightarrow \mathbb{C}^n, \quad [x_0 : \cdots : x_n] \mapsto x_i^{-1}(x_0, \dots, \hat{x}_i, \dots, x_n).$$

Liouville's Theorem shows that any holomorphic function on a compact complex manifold is constant. Thus it is unclear which commutative algebra of holomorphic functions to associate to a complex manifold. New concepts are needed to overcome this problem. One solution replaces algebras by differential graded algebras. Roughly speaking, the idea is to view holomorphic functions as solutions to a certain system of differential equations. The differential graded algebras in question encode the algebra of smooth functions together with this system of differential equations, so that holomorphy may be detected. At the moment, differential graded algebras and differential graded categories are a popular setting for noncommutative algebraic geometry. We will stick to smooth functions and limit our attention to algebras instead of differential graded algebras here.

For an affine algebraic variety over a field K that is not algebraically closed, we have suggested to use maximal ideals instead of characters. For instance, this gives us the complex points of a real algebraic variety. Before we turned to algebraic geometry, we studied characters on the algebra $C^\infty(X)$ of smooth functions on a smooth manifold X . We have shown that any character $C^\infty(X) \rightarrow \mathbb{C}$ is equal to the evaluation map $\text{ev}_x: f \mapsto f(x)$ for some $x \in X$. How about maximal ideals?

PROPOSITION 3.14. *Let X be a smooth manifold. If X is compact, then any ideal in $C^\infty(X)$ is contained in $I_x := \{f \in C^\infty(X) : f(x) = 0\}$ for some $x \in X$. If X is not compact, then $C^\infty(X)$ has maximal ideals that are not of this form.*

PROOF. Let X be compact and let $J \subseteq C^\infty(X)$ be an ideal that is not contained in I_x for any $x \in X$. We must show that $J = C^\infty(X)$. We proceed as in the proof of Proposition 2.16. For each $x \in X$, there is $f_x \in J$ with $f_x(x) \neq 0$. Then $f_x(y) \neq 0$ for y in an open neighbourhood U_x of x . The open subsets U_x for $x \in X$ cover X . Since X is compact, there is a finite sub-covering $X = U_{x_1} \cup \cdots \cup U_{x_n}$. Then

$$f := \sum_{j=1}^n |f_{x_j}|^2 = \sum_{j=1}^n f_{x_j} \cdot \overline{f_{x_j}} \in J.$$

The function f is everywhere positive and hence invertible. Thus $J = C^\infty(X)$.

Now assume that X is non-compact. Let $J := C_c^\infty(X)$ be the subset of compactly supported smooth functions; this is an ideal in $C^\infty(X)$ that is not contained in I_x for any $x \in X$. By Zorn's Lemma, J is contained in a maximal ideal. \square

A variant of the proof above shows that any proper ideal in $C_c^\infty(X)$ is contained in I_x for some $x \in X$. And any *closed* proper ideal in $C^\infty(X)$ is contained in I_x for some $x \in X$. Therefore, a maximal ideal I in $C^\infty(X)$ that differs from I_x for all $x \in X$ must be dense. The resulting quotient field $C^\infty(X)/I$ always has uncountable dimension. Since these quotient fields are so badly behaved, it is preferable to rule them out by restricting to closed ideals. By the way, a maximal ideal in a topological algebra is either dense or closed because its closure is again an ideal.

4. Representations and simple modules

So far, we have only studied commutative algebras. Now we turn to noncommutative algebras. What could replace the points of an ordinary space in this noncommutative setting? We have described points in a manifold or an algebraic variety through characters and maximal ideals. Characters are homomorphisms to commutative fields and, as such, not useful for noncommutative algebras. Maximal ideals work fine. In fact, it is better to replace maximal by primitive or prime ideals. Primitive ideals are the same as kernels of irreducible representations. So irreducible representations of algebras may play the role of “points” in noncommutative geometry. Whereas characters and maximal ideals correspond to each other bijectively, the irreducible representations do not correspond bijectively to their kernels any more. Instead, we show that they correspond to one-sided ideals up to a certain equivalence relation; this is complicated, and cannot be simplified much further. In general, the set of all irreducible representations of a noncommutative algebra is a wild object with little structure.

The points of a manifold carry a topology. This important structure also exists on the primitive ideal space and on the space of irreducible representations. In this course, however, we shall not study it. There is another, more fruitful way to study the interaction between the different irreducible representations of an algebra. Namely, we should study the category of all representations instead. This category of representations will occur frequently in the following.

Representations of algebras are the same thing as modules over it. And irreducible representations correspond to simple modules. The language of representations is more common for Hilbert space representations, whereas purely algebraic objects are more often called modules. We shall use both languages interchangeably.

The central examples for this section are the polynomial algebra $\mathbb{C}[x]$ and the algebra of upper triangular matrices. A module over $\mathbb{C}[x]$ is equivalent to a vector space V with a linear map $f: V \rightarrow V$. And a module over the algebra of upper triangular matrices is equivalent to two vector spaces V_0 and V_1 with a linear map $f: V_0 \rightarrow V_1$. Thus many important questions in linear algebra translate to questions about representations of these two algebras. In this way, the study of modules over more general algebras generalises linear algebra.

DEFINITION 4.1. Let K be a field and let V be a K -vector space. Let $\text{End}(V)$ be the space of K -linear maps $V \rightarrow V$; this is a K -algebra with respect to composition of maps as multiplication and pointwise addition and scalar multiplication. A *representation* of a K -algebra A on V is an algebra homomorphism $f: A \rightarrow \text{End}(V)$. The pair (V, f) is also called a (left) *A -module*.

A representation or module is called *faithful* if f is injective.

When we study noncommutative algebras, the first question is usually to understand their representations or, equivalently, their modules. Modules and representations are the same thing; the nuance is that the name “module” highlights V , while the name “representation” highlights f .

EXAMPLE 4.2. If $V = K$, then the representations of an algebra A on K are the K -valued characters and the zero map. The zero map is a representation on any vector space – but not an interesting one.

EXAMPLE 4.3. Let $A = \mathbb{C}[x]$ be the polynomial algebra in one generator. A representation $f: A \rightarrow \text{End}(V)$ of A on V is already determined by $f(x)$ because $f(x^n) = f(x)^n$ and f is \mathbb{C} -linear. Moreover, $f(x)$ may be any K -linear map $V \rightarrow V$. Therefore, representations of $\mathbb{C}[x]$ are the same as K -linear maps $V \rightarrow V$. Since A has infinite dimension, a linear map $V \rightarrow V$ can only give a faithful A -module if V is

infinite-dimensional. If a module is not faithful, then there is a polynomial $p \in \mathbb{C}[x]$ with $f(p) = 0$. Equivalently, the linear map satisfies some polynomial relation. Thus the module corresponding to a linear map $T: V \rightarrow V$ is faithful if and only if T satisfies no polynomial equation.

Various familiar constructions in linear algebra may be interpreted as natural operations with representations. For instance, the minimal polynomial of $T := f(x)$ is a normalised generator for the ideal $\ker f \triangleleft A$, which is principal because $\mathbb{C}[x]$ is a Euclidean ring (that is, there is a good notion of division with remainder in $\mathbb{C}[x]$). A vector $v \in V$ is an eigenvector if and only if the subspace $\mathbb{C} \cdot v$ it spans is an A -submodule of V .

EXAMPLE 4.4. Let A be a K -algebra. The *regular representation* of A on itself is the representation $\lambda: A \rightarrow \text{End}(A)$ defined by $\lambda_a(b) := a \cdot b$. This is a faithful representation for any unital K -algebra.

EXERCISE 4.5. Find a non-unital algebra for which the regular representation is not faithful. Show that any algebra, unital or not, has a faithful representation.

EXERCISE 4.6. An algebra A is simple if and only if any nonzero representation is faithful.

EXAMPLE 4.7. Let A be the unital K -algebra of upper triangular matrices in M_2K . We use the basis

$$1 := \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad p := \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad s := \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

A unital representation f of A on a K -vector space V is determined uniquely by its values $P := f(p)$ and $S := f(s)$ on these basis vectors – unitality forces $f(1) = \text{Id}_V$. The multiplicativity of f may also be checked on basis vectors: since both $f(x \cdot y)$ and $f(x) \cdot f(y)$ are K -bilinear maps $A \times A \rightarrow \text{End}(V)$, they agree for all $x, y \in A$ once they agree for $x, y \in \{1, p, s\}$. If $x = 1$ or $y = 1$, then $f(xy) = f(x)f(y)$ is trivial. Hence we remain with only four conditions

$$P^2 = P, \quad P \cdot S = 0, \quad S \cdot P = S, \quad S \cdot S = 0.$$

The last one is redundant as well because $S \cdot S = (S \cdot P) \cdot S = S \cdot (P \cdot S) = S \cdot 0 = 0$. Thus a unital representation of A on V is the same as a pair of operators P and S satisfying $P^2 = P$, $PS = 0$ and $SP = S$.

DEFINITION 4.8. An algebra element $p \in A$ is called *idempotent* if $p^2 = p$.

An idempotent operator $P \in \text{End}(V)$ gives rise to a direct sum decomposition $V = \ker(P) \oplus \ker(1 - P)$. Conversely, if $V = V_0 \oplus V_1$, then letting $P|_{V_0} := 0$ and $P|_{V_1} := \text{Id}_{V_1}$ defines an idempotent operator on V . Thus idempotent linear operators on V are in bijection with direct sum decompositions of V . The direct sum decomposition $V = V_0 \oplus V_1$ associated to P identifies $\text{End}(V)$ with an algebra of block matrices with entries

$$\begin{pmatrix} \text{Hom}(V_0, V_0) & \text{Hom}(V_1, V_0) \\ \text{Hom}(V_0, V_1) & \text{Hom}(V_1, V_1) \end{pmatrix}.$$

In this decomposition, P corresponds to the block matrix $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$, and $P \cdot S = S$ and $S \cdot P = 0$ means that

$$S = \begin{pmatrix} 0 & T \\ 0 & 0 \end{pmatrix}$$

for an operator $T: V_1 \rightarrow V_0$. Thus a unital representation of A on V is equivalent to a direct sum decomposition $V = V_0 \oplus V_1$ and an operator $V_1 \rightarrow V_0$. The vector spaces V_0 and V_1 determine V uniquely up to a canonical isomorphism. Hence we arrive at the following theorem:

THEOREM 4.9. *Let A be the unital K -algebra of upper triangular matrices in \mathbb{M}_2K . The category of unital A -modules with A -linear maps as morphisms is equivalent to the category whose objects are the K -linear maps between K -vector spaces and whose morphisms from T to T' are the commuting diagrams*

$$\begin{array}{ccc} V_0 & \xrightarrow{T} & V_1 \\ \downarrow \varphi_0 & & \downarrow \varphi_1 \\ V'_0 & \xrightarrow{T'} & V'_1, \end{array}$$

that is, pairs of maps (φ_0, φ_1) with $T' \circ \varphi_0 = \varphi_1 \circ T$.

Thus representations of $\mathbb{C}[x]$ and of the algebra of upper triangular 2×2 -matrices are both related to certain classes of linear maps – the first to endomorphisms, the second to all linear maps. It is remarkable that the first case is much more complicated, giving rise to an infinite-dimensional algebra $\mathbb{C}[x]$. This is because endomorphisms have a much richer structure (eigenvalues, diagonalisation). In contrast, linear maps between different vector spaces are so flexible that they are easy to classify: their only invariants are the dimensions of the vector spaces V_0 and V_1 and the rank of T :

EXERCISE 4.10. The A -modules associated to linear maps $T: V_0 \rightarrow V_1$ and $T': V'_0 \rightarrow V'_1$ are isomorphic as A -modules if and only if $\dim V_0 = \dim V'_0$, $\dim V_1 = \dim V'_1$, and $\text{rank } T = \text{rank } T'$. Here dimensions are to be taken as ordinal numbers, that is, the equalities mean that the vector spaces on both sides have bases of the same cardinality. If you find set theory confusing, restrict attention to the finite-dimensional case.

EXERCISE 4.11. Describe modules over the algebra $\mathbb{M}_2(\mathbb{C})$ and the algebra of upper triangular 3×3 -matrices in a similar fashion. More precisely, show that their categories of modules are equivalent to the category of invertible linear maps $V_0 \rightarrow V_1$ and to the category of pairs of composable linear maps $V_0 \rightarrow V_1 \rightarrow V_2$, respectively, with appropriate commuting diagrams as morphisms.

All relevant information about an algebra is encoded in its category of modules. That is, we should consider algebras as equivalent if they have equivalent categories of modules. This leads to the concept of Morita equivalence, which we will study in Section 11.

Although faithful representations always exist, they tend to be big and unwieldy. Therefore, we prefer to study irreducible representations:

DEFINITION 4.12. Let $(V, f: A \rightarrow \text{End}(V))$ be an A -module. An A -submodule is a vector subspace W of V for which $f(a)(w) \in W$ for all $a \in A, w \in W$; then f restricts to a map $A \rightarrow \text{End}(W)$ that turns W into an A -module. We call (V, f) with $f \neq 0$ a *simple A -module* or an *irreducible representation of A* if the only A -submodules are $\{0\}$ and V .

DEFINITION 4.13. An algebra is called (left) *primitive* if it has a faithful simple module. An ideal $I \triangleleft A$ is called *primitive* if A/I is primitive or, equivalently, there is a simple module (V, f) with $\ker f = I$.

PROPOSITION 4.14. *Let A be a unital K -algebra. Then it has a simple module. If A is simple, then it is primitive. Maximal ideals in A are primitive.*

PROOF. Subrepresentations of the regular representation of A (Example 4.4) are the same as left ideals in A . An increasing union of a directed set of proper left ideals is again a proper left ideal because it does not contain the unit of A ; this is where we use that A is unital. Then Zorn's Lemma gives a maximal proper left

ideal L in A . The corresponding quotient A/L carries an induced representation of A . Subrepresentations of this representation correspond to left ideals between L and A . Since L is a maximal proper left ideal, A/L carries an irreducible representation of A .

If A is a simple algebra, then any non-zero irreducible representation of A is faithful. Since such representations exist, A is primitive. An ideal is maximal if and only if the corresponding quotient is simple; since a simple algebra is primitive, maximal ideals are primitive. \square

DEFINITION 4.15. Let A be a K -algebra. The set of primitive ideals of A is called the *primitive ideal space* and denoted by $\text{Prim}(A)$. The set of isomorphism classes of simple A -modules is denoted by \hat{A} .

Both \hat{A} and $\text{Prim}(A)$ may claim to play the role of the set of points of a space. Since primitive ideals are exactly the kernels of irreducible representations, we get a surjective map $\hat{A} \rightarrow \text{Prim}(A)$, $(V, f) \mapsto \ker f$. There are many interesting examples where this map is a bijection, but there are even more where it is not. If the canonical map $\hat{A} \rightarrow \text{Prim}(A)$ is not bijective, then we usually need other tools to study the algebra A .

Now we have to justify a few implicit claims. The first is that we claimed \hat{A} to be a set – not just a class. The second is that we claim that for an algebra of functions on some space, \hat{A} and $\text{Prim}(A)$ agree with the underlying space. These claims all follow from a classification of the simple modules. This classification also shows that \hat{A} is an intrinsic property of the algebra A , which can be defined without introducing the concept of representation.

LEMMA 4.16. *Let A be a unital K -algebra over a field K . Any simple left A -module is isomorphic to A/L for a maximal proper left ideal $L \subseteq A$. The simple left A -modules A/L_1 and A/L_2 associated to two such left ideals are isomorphic if and only if there are $a, b \in A$ with $L_1a \subseteq L_2$, $L_2b \subseteq L_1$, $1 - ab \in L_1$, and $1 - ba \in L_2$.*

Roughly speaking, this asserts that the construction in the proof of Proposition 4.14 already yields all simple modules.

PROOF. Let (V, f) be a simple A -module. Pick $v \in V$ with $v \neq 0$ and consider the map $\varrho_v: A \rightarrow V$, $a \mapsto a \cdot v$. This is an A -module map, so that its range is an A -submodule of V . Since V is simple and the range of ϱ_v contains $v = 1 \cdot v \neq 0$, the map ϱ_v must be surjective. Thus $V \cong A/L$ with $L := \ker \varrho_v$. The subset L is a proper left ideal in A . It is a maximal proper left ideal if and only if V is simple. By linear algebra, the map ϱ_v descends to a vector space isomorphism $A/L \rightarrow V$; this map is still A -linear and hence an A -module isomorphism.

Now let $\varphi: A/L_1 \rightarrow A/L_2$ be an isomorphism of A -modules, where L_1 and L_2 are left ideals in A . Let $a \in A$ be any element with $a + L_2 = \varphi(1 + L_1)$. Then

$$\varphi(x + L_1) = \varphi(x \cdot (1 + L_1)) = x \cdot \varphi(1 + L_1) = x \cdot (a + L_2) = x \cdot a + L_2$$

for all $x \in A$. Since $\varphi(x + L_1) = a + L_2$ for all $x \in L_1$, we get $L_1a \subseteq L_2$. Similarly, there is $b \in A$ with $\varphi^{-1}(x + L_2) = x \cdot b + L_1$ and $L_2b \subseteq L_1$. Since $\varphi^{\pm 1}$ are inverse to each other, we get $xab \equiv x \pmod{L_1}$ and $xba \equiv x \pmod{L_2}$ for all $x \in A$. Since L_1 and L_2 are left ideals, this is equivalent to $1 - ab \in L_1$ and $1 - ba \in L_2$.

Conversely, if such a, b exist, then $x + L_1 \mapsto xa + L_2$ and $y + L_2 \mapsto yb + L_1$ are well defined A -module maps $A/L_1 \leftrightarrow A/L_2$ that are inverse to each other. \square

COROLLARY 4.17. *If A is a commutative unital K -algebra, then the map $\hat{A} \rightarrow \text{Prim}(A)$ is a bijection, and $\text{Prim}(A)$ is the set of maximal ideals in A , which is in bijection with the set of characters on A .*

PROOF. In the commutative case, any left ideal is two-sided, and the kernel of the representation $A \rightarrow \text{End}(A/I)$ for a two-sided ideal I is equal to I . Therefore, two different (left) ideals yield non-isomorphic representations. Hence $\text{Prim}(A)$ agrees with the set of maximal ideals in A . These are exactly the kernels of characters on A . \square

Summing up, we have seen that geometric objects such as manifolds or varieties may be recovered from suitable commutative algebras of functions on them. The different regularity conditions imposed on these functions (continuous, smooth, polynomial) amount to viewing the space as a topological space, a smooth manifold, or an algebraic variety. Representations are a good way to understand noncommutative algebras, relating them to algebras of linear maps. We have argued why isomorphism classes of simple modules and primitive ideals of a noncommutative algebra are analogous to the points of a space. Now it is time to examine some examples.

5. Endomorphism algebras and finite-dimensional algebras

What are the most basic examples of noncommutative algebras? The first examples of a noncommutative multiplication are symmetries in geometry. For instance, reflections and rotations in plane geometry do not commute, and neither of them commute with translations. Thus composition of maps is a prototype for noncommutative multiplications. This suggests to look at noncommutative groups as a source of examples. We will do this in Section 6.

Groups are not the simplest examples because in an algebra, we should be able to add and multiply elements. A group carries only one such structure. So we have to linearise it to turn a group into a group algebra. It seems simpler then to first look at linear maps on vector spaces, which carry an obvious addition and multiplication. The algebra of endomorphisms of a vector space already occurred in the concept of a representation. In this lecture, we first study this class of algebras. An important point is the ideal of finite-rank operators. This is the first important example of a non-unital algebra. In noncommutative geometry – unlike in ring theory – it is quite common to allow non-unital algebras because there are many important examples of non-unital algebras. We show that all irreducible representations of the algebra of finite-rank operators are equivalent to the obvious representation. This is not true for the algebra of all endomorphisms. It has representations that vanish on the ideal of finite-rank endomorphisms.

Another way to interpret basic examples is as “small” examples. This suggests to look at finite-dimensional algebras. They have a very rich structure, and their study goes back more than 100 years. I decided to only touch upon this story here. Its most interesting aspects occur only for algebras over fields that are not algebraically closed: such fields may admit non-trivial finite-dimensional skew-fields like the quaternions over the real numbers. We only mention the important structure theorems for finite-dimensional algebras without proof. An important idea here is semi-simplicity. This means that any representation is a direct sum of irreducible representations. An important insight is that a subalgebra of the algebra $M_n\mathbb{C}$ of $n \times n$ -matrices that is closed under taking adjoints is always semi-simple. This is one of the indications why C^* -algebras are so important. They are assumed to be closed under taking adjoints, and this makes them behave better than general algebras of operators. Their theory is so rich that it is better to treat them in a lecture course of their own. This course will therefore not mention them much.

5.1. Algebras of endomorphisms. We have made the point that the algebra $\text{End}(V)$ of linear maps on a vector space V is the prototypical example of an algebra and that other algebras should be represented as subalgebras of $\text{End}(V)$. Since

algebras of the form $\text{End}(V)$ can contain any other algebra, we cannot expect them to have particularly nice properties in general. But at least their ideal structure is rather simple.

First we study linear maps of finite rank. Since we will later generalise this to finite-rank maps between modules, it is useful at this point to work with vector spaces over a skew-field such as the algebra of quaternions. This forces us to distinguish between left and right vector spaces – this is good because it will become necessary later, anyway. But if you prefer, you may take $K = \mathbb{C}$ at this point and not worry about left and right.

DEFINITION 5.1. Let K be a skew-field and let V and W be right K -vector spaces. An operator $f: V \rightarrow W$ has *rank* n if its range $f(V) \subseteq W$ is of dimension n . Let $\text{Hom}_f(V, W) \subseteq \text{Hom}(V, W)$ be the subset of all endomorphisms of finite rank, and let $\text{End}_f(V) := \text{Hom}_f(V, V)$.

Clearly, $\text{End}_f(V)$ is an ideal in $\text{End}(V)$, so that $\text{End}(V)$ is not simple unless V is finite-dimensional.

Operators of rank 1 are particularly important because any operator of finite rank is a sum of operators of rank 1. We are going to introduce some notation to describe such operators. We let $V^* := \text{Hom}_K(V, K)$ be the *dual space* of V ; this is a left K -vector space if V is a right K -vector space by $(\kappa \cdot \hat{v})(v) := \kappa \cdot \hat{v}(v)$ for $\kappa \in K$, $\hat{v} \in V^*$, $v \in V$. Given $\hat{v} \in V^*$ and $w \in W$, we define K -linear maps

$$(5.2) \quad \begin{aligned} |w\rangle: K &\rightarrow W, & \kappa &\mapsto w \cdot \kappa, \\ \langle \hat{v} |: V &\rightarrow K, & v &\mapsto \hat{v}(v), \\ |w\rangle \langle \hat{v} |: V &\rightarrow W, & v &\mapsto |w\rangle (\langle \hat{v} | (v)) = w \cdot \hat{v}(v). \end{aligned}$$

LEMMA 5.3. *Any operator $V \rightarrow W$ of rank 1 is of the form $|w\rangle \langle \hat{v} |$ for some $w \in W \setminus \{0\}$ and some $\hat{v} \in V^* \setminus \{0\}$. And $|w_1\rangle \langle \hat{v}_1 | = |w_2\rangle \langle \hat{v}_2 |$ if and only if there is $\kappa \in K \setminus \{0\}$ with $w_1 = \kappa w_2$ and $\hat{v}_2 = \kappa \hat{v}_1$.*

Let $T: W \rightarrow W'$ and $S: V' \rightarrow V$ be K -linear maps. Then

$$T \circ |w\rangle \langle \hat{v} | \circ S = |T(w)\rangle \langle S^* \hat{v} |,$$

where $S^*: V^* \rightarrow (V')^*$ is defined by $S^* \hat{v}(v') := \hat{v}(S(v'))$. In particular,

$$|w\rangle \langle \hat{v} | \circ |v\rangle \langle \hat{x} | = |w \cdot \hat{v}(v)\rangle \langle \hat{x} | = |w\rangle \langle \hat{v}(v) \cdot \hat{x} |.$$

PROOF. Everything follows by short computations. □

PROPOSITION 5.4. *Let K be a skew-field and let V be a K -vector space. The algebra $\text{End}_f(V)$ is always simple. If V has countably infinite dimension, then the quotient $\text{End}(V)/\text{End}_f(V)$ is simple; in this case, $\{0\}$, $\text{End}_f(V)$, and $\text{End}(V)$ are the only ideals in $\text{End}(V)$.*

PROOF. Let $I \subseteq \text{End}_f(V)$ be some non-zero ideal. Pick $T \in I \setminus \{0\}$. Since $T \neq 0$, there is $v \in V$ with $Tv \neq 0$; let $\hat{v} \in V^* \setminus \{0\}$ be arbitrary. Since I is an ideal in $\text{End}_f(V)$, it contains $T \circ |v\rangle \langle \hat{v} | = |Tv\rangle \langle \hat{v} |$, which is an operator of rank 1. Now let $w \in V$ and $\hat{w} \in V^*$ be arbitrary non-zero elements. There are $\hat{x} \in V^*$ and $x \in V$ with $\hat{x}(Tv) = 1$ and $\hat{w}(x) = 1$. Since I is an ideal in $\text{End}_f(V)$, it contains

$$|w\rangle \langle \hat{x} | \circ |Tv\rangle \langle \hat{v} | \circ |x\rangle \langle \hat{w} | = |w\rangle \langle \hat{w} |.$$

That is, I contains all operators of rank 1. Since these operators span $\text{End}_f(V)$, we get $I = \text{End}_f(V)$. Thus $\text{End}_f(V)$ contains no ideals besides $\{0\}$ and $\text{End}_f(V)$.

Now let I be a non-zero ideal in $\text{End}(V)$. The same argument as above shows that $\text{End}_f(V) \subseteq I$. If $I \neq \text{End}_f(V)$, then it contains some operator $T: V \rightarrow V$ whose range is not finite-dimensional. Let $(v_j)_{j \in \mathbb{N}}$ be a basis for V . Since $T(V)$ is infinite-dimensional and contained in V , it also has a countable basis. Thus we may

find a sequence of vectors $(w_j)_{j \in \mathbb{N}}$ such that $T(w_j)$ is a basis for $T(V)$. Hence there is a unique linear map $T(V) \rightarrow V$ that maps $T(w_j)$ to v_j ; we may extend it to a linear map $a: V \rightarrow V$. In addition, there is a unique linear map $b: V \rightarrow V$ that maps $v_j \mapsto w_j$. The composite map aTb maps $v_j \mapsto w_j \mapsto T(w_j) \mapsto v_j$ for all j , so that $aTb = \text{Id}_V$. Since I is an ideal, $aTb \in I$. But an ideal that contains the unit element of an algebra must be the whole algebra. Thus $I = \text{End}(V)$. \square

Next we want to understand the representation theory of the algebra $\text{End}_f(V)$. Since it is non-unital, we must introduce a technical condition here.

DEFINITION 5.5. Let k be a field, let V be a k -vector space, and let A be a k -algebra. A representation $(V, f: A \rightarrow \text{End}(V))$ is called *non-degenerate* if elements of the form $f(a)(v)$ for $a \in A, v \in V$ span V or, briefly, $A \cdot V = V$.

LEMMA 5.6. *A module over a unital algebra is non-degenerate if and only if $1_A \cdot v = v$ for all $v \in V$, where $1_A \in A$ is the unit element.*

PROOF. If $1_A \cdot v = v$ for all $v \in V$, then $A \cdot V = V$. Since $1_A \cdot (a \cdot v) = (1_A \cdot a) \cdot v = a \cdot v$, it follows that $1_A \cdot v = v$ for all v in the closed linear span of $a \cdot v$ for $a \in A, v \in V$. Therefore, $1_A \cdot v = v$ for all $v \in V$ if V is non-degenerate. \square

PROPOSITION 5.7. *Let K be a skew-field and let $k \subseteq K$ be a field. The standard representation of $\text{End}_f(V)$ on V is irreducible. Up to isomorphism, this is the only k -linear irreducible representation of $\text{End}_f(V)$. Any non-degenerate k -linear representation of $\text{End}_f(V)$ is isomorphic to a direct sum of copies of the standard representation.*

The field k plays no significant role. The proof equips a non-degenerate $\text{End}_f(V)$ -module with a left K -vector space structure, but this is not quite canonical; only its restriction to the centre of K is independent of choices. We may therefore replace k by the centre of K .

PROOF. Let $W \subseteq V$ be a non-zero $\text{End}_f(V)$ -invariant subspace. Pick $w \in W \setminus \{0\}$. Then there is $\hat{v} \in V^*$ with $\hat{v}(w) = 1$. For any $v \in V$, $|v\rangle\langle\hat{v}|(w) = v$ belongs to W because W is $\text{End}_f(V)$ -invariant. Thus $W = V$, proving that V is irreducible. Next we prove that any non-degenerate representation is a direct sum of copies of the standard representation. This implies that any irreducible representation is isomorphic to the standard representation (our conventions ensure that degenerate representations cannot be irreducible).

Pick a non-zero element $v \in V$. Then there is $\hat{v} \in V^*$ with $\hat{v}(v) = 1$. Let $P := |v\rangle\langle\hat{v}|$. Then $P^2 = P$, and the subalgebra

$$P \text{End}_f(V) P := \{PTP : T \in \text{End}_f(V)\} = \{|v\kappa\rangle\langle\hat{v}| : \kappa \in K\}$$

is isomorphic to K as an algebra.

Let $f: \text{End}_f(V) \rightarrow \text{End}(W)$ be a non-degenerate representation. Let $W_0 := P(W)$. The action of $P \text{End}_f(V) P$ on W_0 turns W_0 into a left K -vector space. As such, it has a basis $(w_i)_{i \in I}$. Let

$$W_i := \text{End}_f(V) \cdot w_i = \{f(T)w_i : T \in \text{End}_f(V)\} = \{f(|x\rangle\langle\hat{v}|)(w_i) : x \in V\},$$

where we used $T|v\rangle\langle\hat{v}| = |Tv\rangle\langle\hat{v}|$. The map

$$S_i: V \rightarrow W_i, \quad x \mapsto f(|x\rangle\langle\hat{v}|)(w_i),$$

is linear and surjective by construction. In addition, $S_i(|x\rangle\langle\hat{x}|y) = f(|x\rangle\langle\hat{x}|)S_i(y)$ holds for all $x, y \in V, \hat{x} \in V^*$. That is, S_i is an $\text{End}_f(V)$ -module homomorphism. Thus its kernel is an $\text{End}_f(V)$ -submodule of V . But the only such submodules are $\{0\}$ and V itself. Since $S_i \neq 0$, the map S_i is injective. Thus the representation of $\text{End}_f(V)$ on W_i is isomorphic to the standard representation on V via S_i .

We claim that $W = \bigoplus_{i \in I} W_i$. This will finish the proof of the proposition. First we must check that the subspaces W_i span W . This follows because the representation is non-degenerate. Since the rank-1-operators span $\text{End}_f(V)$, the elements of the form $f(|x\rangle\langle\hat{x}|)(w)$ for $w \in W$, $x \in V$, $\hat{x} \in V^*$ span W . We may factor $|x\rangle\langle\hat{x}| = |x\rangle\langle\hat{v}| \circ |v\rangle\langle\hat{x}|$. Then $f(|v\rangle\langle\hat{x}|)(w) = f(|v\rangle\langle\hat{v}||v\rangle\langle\hat{x}|)(w) \in W_0$. Thus elements of the form $|x\rangle\langle\hat{v}|(w)$ with $x \in V$, $w \in W_0$ span W . Since the vectors w_i form a basis for W_0 , we get $\sum W_i = W$.

Now we consider a linear relation between the subspaces W_i . This is given by a finite subset of indices $F \subseteq I$ and $x_i \in V$ for $i \in F$ with $\sum_{i \in F} S_i(x_i) = 0$. Let $\hat{x} \in V^*$. Then

$$\begin{aligned} 0 &= \sum_{i \in F} |v\rangle\langle\hat{x}|(S_i(x_i)) = \sum_{i \in F} f(|v\rangle\langle\hat{x}||x_i\rangle\langle\hat{v}|)(w_i) \\ &= \sum_{i \in F} f(|v \cdot \hat{x}(x_i)\rangle\langle\hat{v}|)(w_i) = \sum_{i \in F} \hat{x}(x_i) \cdot w_i. \end{aligned}$$

This implies $\hat{x}(x_i) = 0$ for all $i \in F$ because the elements w_i are K -linearly independent. Since this holds for all $\hat{x} \in V^*$, this implies $x_i = 0$ for all $i \in I$. That is, the subspaces W_i are linearly independent, so that their sum is a direct sum. \square

COROLLARY 5.8. *The algebra $\text{End}(V)$ is primitive: its standard representation on V is irreducible and faithful. If V is countably infinite-dimensional, then its primitive ideal space has two points: the ideals $\{0\}$ and $\text{End}_f(V)$.*

We can refine Proposition 5.7 as follows:

THEOREM 5.9. *The category of non-degenerate representations of $\text{End}_f(V)$ is equivalent to the category of right K -vector spaces.*

PROOF. Proposition 5.7 asserts, essentially, that these two categories have the same isomorphism classes of objects: a K -vector space is a direct sum of copies of the standard vector space K , a non-degenerate representation of $\text{End}_f(V)$ is a direct sum of copies of the standard representation V . It remains to observe that the $\text{End}_f(V)$ -module homomorphisms $V \rightarrow V$ are exactly the maps of the form $v \mapsto v \cdot \kappa$ for $\kappa \in K$; this implies that both categories have the same morphisms. \square

We will interpret Theorem 5.9 in Section 11 as saying that $\text{End}_f(V)$ is ‘‘Morita equivalent’’ to K . Roughly speaking, these two algebras describe the same noncommutative space. This includes $\text{End}_f(K^n) \cong \mathbb{M}_n K$ as a special case.

EXERCISE 5.10. If V is finite-dimensional, show that any left ideal in $A := \text{End}(V)$ is of the form $L_W = \{S \in \text{End}(V) : S|_W = 0\}$ for a unique linear subspace W , that is, there is a bijection between left ideals in A and linear subspaces of V . When is the A -module A/L_W simple? When are two such quotient modules isomorphic?

EXERCISE 5.11. A non-degenerate representation $\text{End}_f(V) \rightarrow \text{End}(W)$ extends uniquely to a unital representation of $\text{End}(V)$ by $S(Tx) := (ST)x$ for all $S \in \text{End}(V)$, $T \in \text{End}_f(V)$, $x \in W$.

Let $f: \text{End}(V) \rightarrow \text{End}(W)$ be a representation. Let W_f be the linear span of $f(\text{End}_f(V))(W)$. The ideal $\text{End}_f(V)$ acts non-degenerately on this subspace. Proposition 5.7 and Exercise 5.11 imply that this representation is a direct sum of copies of the standard representation of $\text{End}(V)$ on V . Let

$$W_\infty := W/W_f.$$

The representation of $\text{End}(V)$ on W_∞ descends to one of $A := \text{End}(V)/\text{End}_f(V)$. Since this quotient is simple, any non-zero representation of it is faithful. But we

can say very little about the structure of such representations. In particular, it is hard to describe simple A -modules.

Left ideals in A correspond to left ideals in $\text{End}(V)$ that contain all finite-rank operators. Some of them are of the form

$$L_W := \{S \in \text{End}(V) : S|_W \text{ has finite rank}\}$$

for a subspace $W \subseteq V$. We have $L_{W_1} = L_{W_2}$ if and only if W_1 and W_2 are commensurable, that is, $(W_1 + W_2) / (W_1 \cap W_2)$ is finite-dimensional. In particular, $L_W = \text{End}(V)$ if and only if W has finite codimension in V , and $L_W = A$ if W has finite dimension.

EXERCISE 5.12. Let V be a vector space of countable dimension. Let $A := \text{End}(V) / \text{End}_f(V)$. Let W_1, W_2 be two subspaces of V . Show that $A/L_{W_1} \cong A/L_{W_2}$ if none of the spaces $W_1, W_2, V/W_1, V/W_2$ is finite-dimensional. Can you classify left ideals or isomorphism classes of left ideals in A ?

The algebra $\text{End}_f(V)$ is huge unless V is finite-dimensional: the dual space V^* of a vector space with a countable basis is isomorphic to $\prod_{n \in \mathbb{N}} K$, and this vector space does not have a countable basis. Therefore, we often prefer to work with the algebra of *finite matrices*

$$(5.13) \quad \mathbb{M}_\infty K := \bigcup_{n \in \mathbb{N}} \mathbb{M}_n K,$$

where we view $\mathbb{M}_n K$ as the space of all matrices $(\kappa_{ij})_{i,j \in \mathbb{N}}$ with $\kappa_{ij} = 0$ for $i > n$ or $j > n$. Thus $\mathbb{M}_\infty K$ consists of all matrices $(\kappa_{ij})_{i,j \in \mathbb{N}}$ with $\kappa_{ij} = 0$ for all but finitely many pairs $(i, j) \in \mathbb{N}^2$. If V is a K -vector space with a countable basis $(v_j)_{j \in \mathbb{N}}$, then we may identify $\mathbb{M}_\infty K$ with the subalgebra of $\text{End}_f(V)$ consisting of all finite rank maps that vanish on v_k for $k \gg 0$.

The algebra $\mathbb{M}_\infty K$ has the same representation theory as $\text{End}_f(V)$, that is, it is Morita equivalent to K as well. Roughly speaking, we can carry over all the arguments for $\text{End}_f(V)$ to $\mathbb{M}_\infty K$ as long as we take care to use the given basis whenever we need one.

5.2. Finite-Dimensional algebras. The structure of finite-dimensional algebras is quite well understood. We only mention the basic results here without proof and then discuss some examples. More examples will occur in Section 6.

We have seen that matrix algebras over division algebras are examples of simple algebras. Wedderburn's Theorem asserts that these are the only finite-dimensional examples:

THEOREM 5.14 (Wedderburn). *A finite-dimensional K -algebra for a field K is simple if and only if it is isomorphic to $\mathbb{M}_n D$ for some $n \in \mathbb{N}_{\geq 1}$ and some finite-dimensional division algebra D over K .*

The Frobenius Theorem classifies the finite-dimensional division algebras over the field of real numbers:

THEOREM 5.15 (Frobenius). *Any finite-dimensional division algebra over \mathbb{R} is isomorphic to \mathbb{R}, \mathbb{C} , or to the algebra \mathbb{H} of quaternions.*

In particular, \mathbb{C} is the only finite-dimensional division algebra over \mathbb{C} . This is a general feature of algebraically closed fields. There are many interesting examples of *infinite-dimensional* division algebras over \mathbb{C} . The most obvious examples are fields of rational functions, and finite-dimensional division algebras over those.

DEFINITION 5.16. Let A be an algebra. An ideal $I \triangleleft A$ is called *nilpotent* if there is $n \in \mathbb{N}$ such that all products $i_1 \cdots i_n$ with $i_1, \dots, i_n \in I$ vanish.

LEMMA 5.17. *Let $I_1, I_2 \triangleleft A$ be nilpotent ideals. Then*

$$I_1 + I_2 := \{i_1 + i_2 : i_1 \in I_1, i_2 \in I_2\}$$

is a nilpotent ideal as well. Therefore, any finite-dimensional algebra contains a maximal nilpotent ideal.

PROOF. Let $I_1^n = 0$ and $I_2^m = 0$. Let $k = n + m - 1$. We claim that $(I_1 + I_2)^k = 0$. Write a product of k factors in $I_1 + I_2$ as a sum of products of k factors in $I_1 \cup I_2$. If there are at least n factors in I_1 , then we may multiply the remaining factors in I_2 with neighbouring factors to see that the product belongs to $I_1^n = 0$. Otherwise, there are at most $n - 1$ factors in I_1 and hence at least $k - (n - 1) \geq m$ factors in I_2 . Therefore, the product belongs to $I_2^m = 0$. Thus $(I_1 + I_2)^k = 0$. If A is finite-dimensional, then any increasing chain of ideals is finite. Therefore, the result above implies that there is a largest nilpotent ideal in A . \square

DEFINITION 5.18. The maximal nilpotent ideal in an algebra is called its *nilradical*.

LEMMA 5.19. *Let A be a finite-dimensional algebra. The nilradical is contained in all primitive ideals of A .*

PROOF. Let $I \triangleleft A$ be a nilpotent ideal and let $f: A \rightarrow \text{End}(V)$ be an irreducible representation of A . We prove by contradiction that $f|_I = 0$. Otherwise, $I \cdot V = V$ by irreducibility. Then $I^k \cdot V = V$ for all $k \in \mathbb{N}$ by induction. This is impossible because $I^k = 0$ for some $k \in \mathbb{N}$. \square

DEFINITION 5.20. A finite-dimensional algebra is called *semi-simple* if its nilradical vanishes or, equivalently, if zero is its only nilpotent ideal.

THEOREM 5.21. *Any semi-simple finite-dimensional algebra is isomorphic to a finite direct product of simple algebras, that is, to a finite direct product of matrix algebras over division algebras.*

Let A be any finite-dimensional algebra, and let $\text{rad } A \triangleleft A$ denote its nilradical. Then $A/\text{rad } A$ is semi-simple and hence isomorphic to a finite direct product of matrix algebras over division algebras:

$$A/\text{rad } A \cong \bigoplus_{j=1}^n \mathbb{M}_{m_j} D_j$$

for some $n \in \mathbb{N}_{\geq 0}$, $m_j \in \mathbb{N}_{\geq 1}$, and division algebras D_j over K .

THEOREM 5.22. *The representations $A/\text{rad } A \rightarrow \mathbb{M}_{m_j} D_j$ for $j = 1, \dots, n$ form a set of representatives for the isomorphism classes of irreducible representations of A . Thus \hat{A} and $\text{Prim}(A)$ both consist of exactly n points. Any primitive ideal in A is a maximal two-sided ideal.*

PROOF. Let M be a simple A -module. Lemma 5.19 shows that $\text{rad } A \cdot M = 0$. So M is a module over $A/\text{rad } A$. This is isomorphic to a direct sum of matrix algebras. Let $p_j \in A/\text{rad } A$ be the image of the unit matrix in the summand $\mathbb{M}_{m_j} D_j \cdot M$ for $j = 1, \dots, n$. Each p_j is idempotent, and these idempotent elements are orthogonal and sum up to the unit element. Since they belong to the centre of $A/\text{rad } A$, the orthogonal subspaces $p_j M$ are A -submodules as well. Therefore, exactly one of them is non-zero. Then M becomes a simple module over the summand $\mathbb{M}_{m_j} D_j$ in $A/\text{rad } A$. We already know that all simple modules over such a matrix algebra are isomorphic to the standard module $D_j^{m_j}$. Thus there is exactly one isomorphism class of simple modules for $j = 1, \dots, n$. \square

LEMMA 5.23. *A finite-dimensional unital algebra A is semi-simple if and only if any unital A -module is a direct sum of simple A -modules.*

PROOF. If A is semi-simple, then it is isomorphic to a direct sum of matrix algebras. The unit element in one of these summands is a projection. These projections belong to the centre of A . They decompose a representation of A into a direct sum of pieces on which only one matrix algebra is acting. Then Proposition 5.7 shows that any representation of the matrix algebra $\mathbb{M}_n K$ is a direct sum of copies of the standard representation on K^n . And the latter are irreducible. Conversely, assume that any representation is a direct sum of irreducible representations. Then the radical acts trivially in any representation by Lemma 5.19. Since there is a faithful representation, namely, the regular representation, it follows that the radical vanishes. \square

Let A be a finite-dimensional algebra and let \hat{A} be its set of isomorphism classes of simple modules. Let d_π for $\pi \in \hat{A}$ be the dimension of the simple module. The family of simple modules provides an algebra homomorphism

$$F: A \rightarrow \bigoplus_{\pi \in \hat{A}} \mathbb{M}_{d_\pi} \mathbb{C}.$$

Theorem 5.22 and the structure theory of finite-dimensional algebras show that F is always surjective and that its kernel is the nilradical $\text{rad } A$. Hence F is an algebra isomorphism if and only if A is semi-simple. How can we detect whether this is the case?

THEOREM 5.24. *Let A be a finite-dimensional unital algebra and let $\langle _ | _ \rangle$ be an inner product on A with the property that for each $a \in A$, there is $a^* \in A$ with $\langle ax | y \rangle = \langle x | a^* y \rangle$. Then A is semi-simple.*

PROOF. The left regular representation $\lambda: A \rightarrow \text{End}(A)$ is faithful because A is unital. By construction, $\lambda(a)^* = \lambda(a^*)$ for all $a \in A$, where the adjoint T^* of an endomorphism T of A is defined by $\langle x | T^* y \rangle = \langle Tx | y \rangle$ for all $x, y \in A$.

Let $I \triangleleft A$ be a nilpotent ideal and let $x \in I$. Then $x^* x \in I$ as well because I is an ideal. Since I is nilpotent, $(x^* x)^n = 0$ and hence $\lambda(x^* x)^n = 0$ for some $n \in \mathbb{N}$. But the matrix $\lambda(x^* x) = \lambda(x)^* \lambda(x)$ is self-adjoint and therefore diagonalisable. Hence $(x^* x)^n = 0$ already implies $x^* x = 0$, so that $\langle x | x \rangle = \langle 1_A | x^* x \rangle = 0$ and hence $x = 0$. Therefore, $I = \{0\}$ is the only nilpotent ideal in A . \square

The same argument shows:

PROPOSITION 5.25. *Let $A \subseteq \mathbb{M}_n \mathbb{C}$ be a subalgebra with $x^* \in A$ for all $x \in A$; here the adjoint of a matrix is defined by $(x_{ij})^* := (\overline{x_{ji}})$. Then A is semi-simple.*

In the situation of Theorem 5.24, we know that $F: A \rightarrow \bigoplus_{\pi \in \hat{A}} \mathbb{M}_{d_\pi} \mathbb{C}$ is invertible. This can be proved directly, without using the structure theory of finite-dimensional algebras. We will only prove this in the more concrete case of group algebras; the ideas in the general case are the same.

6. Group algebras

The group algebra of a group is defined so that its unital representations are equivalent to representations of the group. It plays an important role in the representation theory of groups. We define the group algebra for all (discrete) groups in this section. Our study in this section is limited, however, to finite groups. We will consider a group algebra of a certain noncommutative group in Section 8. The structure of group algebras of general infinite groups still has many open problems.

The group algebra of a finite group is finite-dimensional. We show that it is semi-simple. Hence it is isomorphic to a direct sum of matrix algebras. We also prove this result directly by relating irreducible group representations to the regular representation. The Peter–Weyl Theorem describes the group algebra in terms of the set of irreducible representations and their dimensions. We first deduce this theorem from the general structure theorems for finite-dimensional algebras. Then we prove it by hand, using matrix coefficients of representations and the Schur orthogonality relations. The latter approach works for compact groups as well, which have a very similar representation theory. This requires some more analysis, however, and is therefore left to a course on C^* -algebras.

DEFINITION 6.1. Let G be a group and let K be a field. The *group algebra* $K[G]$ is the ring of all functions $f: G \rightarrow K$ with finite support – that is, $f(g) \neq 0$ for only finitely many $g \in G$ – and with the convolution product

$$(f_1 * f_2)(g) := \sum_{h,k \in G: hk=g} f_1(h)f_2(k) = \sum_{h \in G} f_1(h)f_2(h^{-1}g) = \sum_{h \in G} f_1(gh)f_2(h^{-1}).$$

The convolution is associative because the products in G and K are associative. Bilinearity of the convolution is evident as well. For $g \in G$, define $\delta_g \in K[G]$ by $\delta_g(g) = 1$ and $\delta_g(h) = 0$ for $h \neq g$. These functions form a basis for $K[G]$ as a K -vector space and satisfy $\delta_g * \delta_h = \delta_{gh}$. And δ_1 is the unit element of $K[G]$.

EXAMPLE 6.2. If $G = \mathbb{Z}$, then $\mathbb{C}[\mathbb{Z}] \cong \mathbb{C}[t, t^{-1}]$ is isomorphic to the algebra of Laurent polynomials. The isomorphism maps $t^n \mapsto \delta_n$ for all $n \in \mathbb{Z}$.

DEFINITION 6.3. A representation of a group G on a vector space V is a group homomorphism $\pi: G \rightarrow \text{Aut}(V)$, where $\text{Aut}(V)$ denotes the group of invertible linear maps on V . If V is a Banach space, then we allow only bounded linear maps on V .

PROPOSITION 6.4. *The group algebra $K[G]$ is a unital K -algebra.*

If $\pi: G \rightarrow \text{Aut}(V)$ is a representation of G on a K -vector space V , then

$$\bar{\pi}(f) := \sum_{g \in G} f(g)\pi(g)$$

defines a unital algebra homomorphism $\bar{\pi}: K[G] \rightarrow \text{End}(V)$. Conversely, any such unital algebra homomorphism comes from a unique group representation of G .

PROOF. Since the elements δ_g for $g \in G$ form a basis for $K[G]$, a linear map $K[G] \rightarrow \text{End}(V)$ is determined uniquely by its values on group elements, which may be specified arbitrarily. The product in $K[G]$ is defined so that $\bar{\pi}$ is a unital algebra homomorphism if and only if π is a group representation. \square

The group algebra $K[G]$ is finite-dimensional if and only if the group G is finite. Hence the structure theory of finite-dimensional algebras applies to $K[G]$ and provides some insights into the representation theory of G . The algebra $K[G]$ has the advantage over G that we may add its elements, making use of the linear structure that is present in group representations. For a field K of positive characteristic, the group algebra $K[G]$ need not be semi-simple (we need the order of G to be invertible in K). We only study the case $K = \mathbb{C}$ from now on.

THEOREM 6.5. *Let G be a finite group. Let \hat{G} be the set of isomorphism classes of irreducible representations of G and let d_π for $\pi \in \hat{G}$ be the dimension of the representation π . The group algebra $\mathbb{C}[G]$ is semi-simple, and there is an isomorphism*

$$\mathbb{C}[G] \cong \bigoplus_{\pi \in \hat{G}} \mathbb{M}_{d_\pi} \mathbb{C}.$$

PROOF. We equip $\mathbb{C}[G]$ with the inner product

$$(6.6) \quad \langle f_1 | f_2 \rangle := \sum_{g \in G} \overline{f_1(g)} f_2(g)$$

and define a conjugate-linear map $f \mapsto f^*$ on $\mathbb{C}[G]$ by

$$(6.7) \quad f^*(g) := \overline{f(g^{-1})} \quad \text{for } f \in \mathbb{C}[G], g \in G.$$

A straightforward computation yields $\langle f * x | y \rangle = \langle x | f^* * y \rangle$ for all $f, x, y \in \mathbb{C}[G]$. Now Theorem 5.24 yields that $\mathbb{C}[G]$ is semi-simple. The correspondence between group representations of G on \mathbb{C} -vector spaces and unital representations of $\mathbb{C}[G]$ preserves submodules – it is even an equivalence of categories. Hence it identifies irreducible representations of G with simple $\mathbb{C}[G]$ -modules. Thus $\mathbb{C}[G] \cong \bigoplus_{\pi \in \hat{G}} \mathbb{M}_{d_\pi} \mathbb{C}$. \square

COROLLARY 6.8. *Let G be a finite group. Then $\sum_{\pi \in \hat{G}} d_\pi^2 = |G|$.*

PROOF. The isomorphism in Theorem 6.5 implies that $\mathbb{C}[G]$ and $\bigoplus_{\pi \in \hat{G}} \mathbb{M}_{d_\pi} \mathbb{C}$ have the same dimension. \square

PROPOSITION 6.9. *Let G be a group, let K be a field, and let C be the set of all finite conjugacy classes in G . The characteristic functions of $\langle g \rangle \in C$ form a basis for the centre of $K[G]$. The number of isomorphism classes of irreducible \mathbb{C} -linear representations is the number of conjugacy classes in G .*

PROOF. An element $f \in K[G]$ is central if and only if $\delta_g * f = f * \delta_g$ for all $g \in G$. This is equivalent to $\delta_g * f * \delta_{g^{-1}} = f$ for all $g \in G$. Since $(\delta_g * f * \delta_{g^{-1}})(x) = f(gxg^{-1})$, the condition says that f is constant on conjugacy classes. So the dimension of the centre is the number of conjugacy classes. Let $K = \mathbb{C}$. Then

$$\mathbb{C}[G] \cong \bigoplus_{\pi \in \hat{G}} \mathbb{M}_{d_\pi} \mathbb{C}$$

implies that the dimension of the centre is also equal to the size of \hat{G} . \square

One of the main issues in representation theory is to describe the set \hat{G} of irreducible representations for specific groups. This is a very rich subject. We only consider two rather easy examples.

EXAMPLE 6.10. Let G be the symmetric group on three letters. This group has six elements and three conjugacy classes, namely, the classes of the trivial element and of the cycles of length two and three, respectively. By Proposition 6.9, G has three equivalence classes of irreducible representations. There are two group homomorphisms $G \rightarrow \{\pm 1\}$, namely, the trivial homomorphism and the sign homomorphism that maps each transposition to -1 . These two homomorphisms are 1-dimensional representations of G . These are two of the three irreducible representations of G . Corollary 6.8 implies that the third representation has dimension 2. The group G acts on \mathbb{C}^3 by permuting the basis vectors. This representation is reducible: the linear span of the vector $(1, 1, 1)$ is an invariant subspace. The complement is a 2-dimensional representation of G . It is easy to check that this complement contains neither the trivial representation nor the sign character as a subrepresentation. Therefore, it is irreducible. Thus we have found all irreducible representations of the group G .

THEOREM 6.11. *Let G be a finite commutative group. Then \hat{G} is equal to the set of characters of G . The Fourier transform is an isomorphism between the algebra $\mathbb{C}[G]$ with the convolution product and the algebra $\mathbb{C}[\hat{G}]$ with pointwise multiplication.*

For a finite cyclic group $G = \mathbb{Z}/n$ for $n \in \mathbb{N}_{\geq 1}$, any character of G is of the form $\chi_l: k \mapsto \exp(2\pi ikl/n)$ for some $l \in \mathbb{Z}/n$. So $\hat{G} \cong \mathbb{Z}/l$. The Fourier transform in this case is also called the discrete Fourier transform.

PROOF. The finite-dimensional algebra $\mathbb{C}[G]$ is commutative because G is Abelian. Hence it cannot surject onto a $d \times d$ -matrix algebra for $d \geq 2$. So all irreducible representations of $\mathbb{C}[G]$ are 1-dimensional. And the Fourier transform is an isomorphism from $\mathbb{C}[G]$ onto $\bigoplus_{g \in \hat{G}} \mathbb{C}$. This direct sum is isomorphic to $\mathbb{C}[\hat{G}]$ with the pointwise multiplication.

Let $G = \mathbb{Z}/n$ be cyclic. A character on G is determined by its value on the generator $1 \bmod n$. And this must be an n th root of unity. So any character has the form χ_l for some $l \in \mathbb{Z}/n$. \square

6.1. The orthogonality relations. The proof of Theorem 6.5 above appeals to the general structure theory, making it hard to see what goes on. We are going to construct the isomorphism between $\mathbb{C}[G]$ and $\bigoplus_{\pi \in \hat{G}} \mathbb{M}_{d_\pi} \mathbb{C}$ more directly. First we need some basic results from representation theory.

PROPOSITION 6.12. *Let $\pi: G \rightarrow \text{Aut}(V)$ be a group representation of G on a finite-dimensional vector space V . There is a G -invariant inner product on V , that is, an inner product with $\langle \pi_g \vec{v} | \pi_g \vec{w} \rangle = \langle \vec{v} | \vec{w} \rangle$ for all $g \in G$, $\vec{v}, \vec{w} \in V$.*

PROOF. Take any inner product (\vec{v}, \vec{w}) and make it G -invariant by averaging:

$$\langle \vec{v} | \vec{w} \rangle := \frac{1}{|G|} \sum_{g \in G} \langle \pi_g \vec{v} | \pi_g \vec{w} \rangle. \quad \square$$

An inner product is $\pi(G)$ -invariant if and only if π is unitary, that is, $\pi_g^{-1} = \pi_g^*$ for all $g \in G$.

Any representation (V, π) of G yields a unital algebra homomorphism $\bar{\pi}: \mathbb{C}[G] \rightarrow \text{End}(V) \cong \mathbb{M}_{d_\pi} \mathbb{C}$, where $d_\pi := \dim V$. Letting (V, π) run through the set \hat{G} of all irreducible representations, we get a unital algebra homomorphism

$$(6.13) \quad F: \mathbb{C}[G] \rightarrow \bigoplus_{\pi \in \hat{G}} \mathbb{M}_{d_\pi} \mathbb{C}.$$

This homomorphism is called the *Fourier transform* for G (this notation follows Theorem 6.11). Theorem 6.5 implies that F is an isomorphism. We will find a formula for the inverse of F . This gives an analogue of the Fourier inversion formula for representations of finite groups. First we explain why F is injective:

COROLLARY 6.14. *Any finite-dimensional representation V of G is a direct sum of irreducible representations. The canonical map $F: \mathbb{C}[G] \rightarrow \bigoplus_{\pi \in \hat{G}} \text{End}(V_\pi)$ is injective.*

PROOF. This is proved by induction on the dimension of V . If V is not itself irreducible, then it contains a non-trivial subrepresentation $W \subseteq V$, that is, a subspace W with $\pi_g(W) \subseteq W$ for all $g \in G$ and $1 \leq \dim W \leq \dim V - 1$. The orthogonal complement W^\perp of W for a G -invariant inner product is G -invariant as well. We get a direct sum decomposition $V = W \oplus W^\perp$. Since the dimensions of both W and W^\perp are smaller than that of V , induction shows that they are direct sums of irreducible representations. Hence so is V .

In particular, we may decompose the left regular representation on $\mathbb{C}[G]$ into a direct sum of irreducible representations. If $F(f) = 0$, then f acts by 0 in each irreducible representation and hence $\lambda(f) = 0$. Since λ is faithful, we get $f = 0$. \square

Corollary 6.14 together with Lemma 5.23 show once more that $\mathbb{C}[G]$ is semi-simple – but this result also uses the structure theory of finite-dimensional algebras.

Recall the standard inner product on $\mathbb{C}[G]$ in (6.6). If F is an isomorphism, we can transport this inner product to the target space $\bigoplus \mathbb{M}_{d_\pi} \mathbb{C}$. Then the inverse of F is equal to its adjoint with respect to these inner products. In fact, we will describe the correct inner product on $\mathbb{M}_{d_\pi} \mathbb{C}$ directly and check that the resulting adjoint F^* is an isometry. Since we already know that F is injective, F is unitary and we have established our Fourier Inversion Formula.

We define an inner product on $\text{End}(V_\pi)$ for $\pi \in \hat{G}$ by

$$\langle x | y \rangle := \frac{d_\pi}{|G|} \cdot \text{tr}(x^* y) \quad \text{for all } x, y \in \text{End}(V_\pi).$$

This is an inner product because $\text{tr}(x^* x) > 0$ for $x \neq 0$; we will see below that the normalisation factor $d_\pi/|G|$ makes F isometric. The inner product on $\bigoplus_{\pi \in \hat{G}} \text{End}(V_\pi)$ is defined so that the summands $\text{End}(V_\pi)$ for $\pi \in \hat{G}$ are orthogonal to each other and carry the inner products specified above.

Now we compute the adjoint of F . Let $x_\pi \in \text{End}(V_\pi)$ for $\pi \in \hat{G}$ and $f \in \mathbb{C}[G]$.

$$\langle F(f) | (x_\pi) \rangle = \sum_{\pi \in \hat{G}} \frac{d_\pi}{|G|} \text{tr}(\bar{\pi}(f)^* x_\pi) = \sum_{g \in G} \overline{f(g)} \sum_{\pi \in \hat{G}} \frac{d_\pi}{|G|} \text{tr}(\pi_g^* x_\pi).$$

Using $\text{tr}(\pi_g^* x_\pi) = \text{tr}(x_\pi \pi_g^*) = \text{tr}(x_\pi \pi_{g^{-1}})$, we get

$$F^*((x_\pi)_{\pi \in \hat{G}})(g) = \sum_{\pi \in \hat{G}} \frac{d_\pi}{|G|} \text{tr}(x_\pi \circ \pi_{g^{-1}}).$$

Before we establish that F^* is an isometry, we are going to reinterpret this definition in terms of matrix coefficients.

Let (V, π) be a finite-dimensional representation of G . By Proposition 6.12, there is a G -invariant inner product $\langle \cdot | \cdot \rangle$ on V . The map that maps $\vec{v} \in V$ to the linear functional $\langle \vec{v} | \cdot \rangle$ yields a vector space isomorphism between the dual space V^* of V and \bar{V} , the Abelian group V with the conjugate \mathbb{C} -vector space structure $\lambda \bullet x = \bar{\lambda} \cdot x$. The operators $|\vec{v}\rangle\langle\vec{w}|$ for $\vec{v}, \vec{w} \in V$ span the space $\text{End}(V)$; even more, there is a canonical vector space isomorphism

$$V \otimes \bar{V} \rightarrow \text{End}(V), \quad \vec{v} \otimes \vec{w} \mapsto |\vec{v}\rangle\langle\vec{w}|.$$

Our task is to define a map $\text{End}(V) \rightarrow \mathbb{C}[G]$ (for irreducible V). One of the obvious choices is the map

$$C: \text{End}(V) \rightarrow \mathbb{C}[G], \quad C(|\vec{v}\rangle\langle\vec{w}|) = c_{\vec{v}, \vec{w}} := \langle \pi_g \vec{w} | \vec{v} \rangle.$$

The function $c_{\vec{v}, \vec{w}}$ is called a *matrix coefficient* of the representation π because when we choose a basis for V and let \vec{v} and \vec{w} be basis vectors, then these functions are indeed the matrix coefficients of π represented in this basis.

Since $\langle \vec{v} | \vec{w} \rangle$ is the trace of the matrix $|\vec{v}\rangle\langle\vec{w}|$ and $\langle \vec{w} | \circ \pi_{g^{-1}} = \langle \vec{w} | \circ \pi_g^* = \langle \pi_g(\vec{w}) |$, we may rewrite C as

$$C(x)(g) = \text{tr}(x \circ \pi_{g^{-1}}) \quad \text{for all } x \in \text{End}(V).$$

Up to a normalisation factor, this is exactly our formula for F^* .

Since $\text{tr}(xy) = \text{tr}(yx)$, this satisfies $C(\pi_g \circ x \circ \pi_h) = \delta_g * C(x) * \delta_h$. More generally, if $f_1, f_2 \in \mathbb{C}[G]$ and $x \in \text{End}(V)$, then

$$C(\bar{\pi}(f_1) \circ x \circ \bar{\pi}(f_2)) = f_1 * C(x) * f_2.$$

That is, C is a bimodule homomorphism if we equip $\text{End}(V)$ and $\mathbb{C}[G]$ with the obvious $\mathbb{C}[G]$ -bimodule structures. This is equivalent to

$$c_{\pi_g \vec{v}, \pi_h \vec{w}} = \delta_g * c_{\vec{v}, \vec{w}} * \delta_h, \quad c_{\bar{\pi}(f_1) \vec{v}, \bar{\pi}(f_2) \vec{w}} = f_1 * c_{\vec{v}, \vec{w}} * f_2.$$

We want to prove that the maps F and F^* above are unitary. This requires some more theory. Let (V, π) be irreducible. Given $\vec{v} \in V$, we define operators

$$|\vec{v}\rangle: \mathbb{C}[G] \rightarrow V, \quad \langle\langle \vec{v} | : V \rightarrow \mathbb{C}[G]$$

by $|\vec{v}\rangle(f) := \sum_{g \in G} f(g) \pi_g(\vec{v})$ for $f \in \mathbb{C}[G]$ and $\langle\langle \vec{v} | (\vec{w})(g) := c_{\vec{w}, \vec{v}} = \langle \pi_g(\vec{v}) | \vec{w} \rangle$.

It is easy to check that $|\vec{v}\rangle^* = \langle\langle \vec{v} |$:

$$\langle\langle \vec{v} | (f) | \vec{w} \rangle = \sum_{g \in G} \langle f(g) \pi_g(\vec{v}) | \vec{w} \rangle = \sum_{g \in G} \overline{f(g)} \langle\langle \vec{v} | (\vec{w})(g) = \langle f | \langle\langle \vec{v} | (\vec{w}) \rangle$$

for all $\vec{v}, \vec{w} \in V$. The computations above show that $\langle\langle \vec{v} |$ and hence $|\vec{v}\rangle$ are G -equivariant, that is, $\langle\langle \vec{v} | \circ \pi_g = \lambda_g \circ \langle\langle \vec{v} |$ and $|\vec{v}\rangle \circ \lambda_g = \pi_g \circ |\vec{v}\rangle$ for all $g \in G$, $\vec{v} \in V$; here λ denotes the left regular representation

$$\lambda_g f(x) := f(g^{-1}x) = (\delta_g * f)(x)$$

for $g, x \in G$, $f \in \mathbb{C}[G]$.

Now let (V, π) and (W, ϱ) be two irreducible representations of G and let $\vec{v} \in V$, $\vec{w} \in W$. Then $|\vec{v}\rangle \langle\langle \vec{w} | : V \rightarrow W$ is a G -equivariant operator. Such operators must have a very simple form by Schur's Lemma:

LEMMA 6.15 (Schur). *Let (V, π) and (W, ϱ) be two finite-dimensional irreducible representations of G and let $T: V \rightarrow W$ be a G -equivariant operator. Then T is either invertible or zero. In addition, if $(V, \pi) = (W, \varrho)$, then T is a scalar multiple of the identity matrix, $T = \lambda \cdot \text{Id}_V$ for some $\lambda \in \mathbb{C}$.*

PROOF. The kernel of T is a subrepresentation of V . Since V is irreducible, T is either injective or zero. Similarly, since the range of T is a subrepresentation of the irreducible representation W , T is either surjective or zero. Thus T is either invertible or zero. Now assume that $(V, \pi) = (W, \varrho)$. Then $T - \lambda$ is G -equivariant and hence invertible or zero for all $\lambda \in \mathbb{C}$. Since T has at least one eigenvalue, there is λ for which $T - \lambda$ is not invertible. Then $T - \lambda = 0$ as desired. \square

THEOREM 6.16 (Orthogonality Relations). *Let (V, π) and (W, ϱ) be irreducible representations of a finite group G , equipped with G -invariant inner products; let $\vec{v}_1, \vec{v}_2 \in V$, $\vec{w}_1, \vec{w}_2 \in W$. If (V, π) and (W, ϱ) are not isomorphic, then $c_{\vec{v}_1, \vec{v}_2} \perp c_{\vec{w}_1, \vec{w}_2}$ in $\mathbb{C}[G]$, that is,*

$$0 = \sum_{g \in G} \overline{\langle \pi_g \vec{v}_2 | \vec{v}_1 \rangle} \langle \varrho_g \vec{w}_2 | \vec{w}_1 \rangle = \sum_{g \in G} \langle \vec{v}_1 | \pi_g \vec{v}_2 \rangle \langle \varrho_g \vec{w}_2 | \vec{w}_1 \rangle.$$

If $(V, \pi) = (W, \varrho)$ with the same inner product, then

$$\langle c_{\vec{v}_1, \vec{v}_2} | c_{\vec{w}_1, \vec{w}_2} \rangle = \frac{|G|}{d_\pi} \langle \vec{v}_1 | \vec{w}_1 \rangle \cdot \langle \vec{w}_2 | \vec{v}_2 \rangle.$$

PROOF. Assume first that (V, π) and (W, ϱ) are not isomorphic. By Schur's Lemma, the G -equivariant operator $|\vec{w}_1\rangle \langle\langle \vec{v}_1 |$ from V to W is either invertible or zero. But if it were invertible, then (V, π) and (W, ϱ) would be isomorphic. Hence $|\vec{w}_1\rangle \langle\langle \vec{v}_1 | = 0$. Now we compute

$$\langle c_{\vec{v}_1, \vec{v}_2} | c_{\vec{w}_1, \vec{w}_2} \rangle = \langle \langle\langle \vec{v}_2 | (\vec{v}_1) | \langle\langle \vec{w}_2 | (\vec{w}_1) \rangle = \langle |\vec{w}_2\rangle \langle\langle \vec{v}_2 | \vec{v}_1 | \vec{w}_1 \rangle = 0.$$

This proves the *First Orthogonality Relation* $\langle c_{\vec{v}_1, \vec{v}_2} | c_{\vec{w}_1, \vec{w}_2} \rangle = 0$ for not isomorphic (V, π) and (W, ϱ) .

Now assume that (V, π) and (W, ϱ) are equal and equipped with the same inner product. Then Schur's Lemma shows that $|\vec{w}_1\rangle \langle\langle \vec{v}_1 | = \alpha(\vec{w}_1, \vec{v}_1) \cdot \text{Id}_V$ for some scalar $\alpha(\vec{w}_1, \vec{v}_1)$. The same computation as above shows

$$\langle c_{\vec{v}_1, \vec{v}_2} | c_{\vec{w}_1, \vec{w}_2} \rangle = \overline{\alpha(\vec{w}_2, \vec{v}_2)} \langle \vec{v}_1 | \vec{w}_1 \rangle.$$

Recall the involution $f^*(g) := \overline{f(g^{-1})}$ on $\mathbb{C}[G]$. The formula for $c_{\vec{v}_1, \vec{v}_2}$ shows that $c_{\vec{v}_1, \vec{v}_2}^* = c_{\vec{v}_2, \vec{v}_1}$. Since $\langle f_1^* | f_2^* \rangle = \overline{\langle f_1 | f_2 \rangle}$ for $f_1, f_2 \in \mathbb{C}[G]$, we have

$$\langle c_{\vec{v}_1, \vec{v}_2} | c_{\vec{w}_1, \vec{w}_2} \rangle = \overline{\langle c_{\vec{v}_2, \vec{v}_1} | c_{\vec{w}_2, \vec{w}_1} \rangle}.$$

Together with the computation above, this implies $\alpha(\vec{w}_1, \vec{v}_1) = a \cdot \overline{\langle \vec{w}_1 | \vec{v}_1 \rangle}$ for some $a \in \mathbb{R}$. That is, $|\vec{v}_1\rangle\rangle\langle\langle \vec{v}_2| = a \cdot \langle \vec{v}_2 | \vec{v}_1 \rangle$ for all $\vec{v}_1, \vec{v}_2 \in V$. To compute a , we let $\vec{v}_1, \dots, \vec{v}_d$ be an orthonormal basis of V , $d = d_\pi$, and consider

$$a \cdot d = \sum_{j=1}^d |\vec{v}_j\rangle\rangle\langle\langle \vec{v}_j|.$$

We use $\sum_{j=1}^d |v_j\rangle\langle v_j| = \text{Id}_V$ to compute

$$\begin{aligned} \sum_{j=1}^d |\vec{v}_j\rangle\rangle\langle\langle \vec{v}_j|(\vec{w}) &= \sum_{j=1}^d \sum_{g \in G} \pi_g \vec{v}_j \langle \pi_g \vec{v}_j | \vec{w} \rangle = \sum_{g \in G} \sum_{j=1}^d \pi_g \circ |\vec{v}_j\rangle\langle \vec{v}_j| \circ \pi_g^*(\vec{w}) \\ &= \sum_{g \in G} \pi_g \circ \pi_g^*(\vec{w}) = \sum_{g \in G} \vec{w} = |G| \vec{w}. \end{aligned}$$

Thus $a = |G|/d$. □

THEOREM 6.17. *The Fourier transform*

$$F: \mathbb{C}[G] \rightarrow \bigoplus_{\pi \in \hat{G}} \text{End}(V_\pi), \quad f \mapsto (\widehat{\pi}(f))_{\pi \in \hat{G}},$$

is unitary. Its inverse is the unitary map

$$F^*: \bigoplus_{\pi \in \hat{G}} \text{End}(V_\pi) \rightarrow \mathbb{C}[G], \quad ((x_\pi)_{\pi \in \hat{G}})(g) \mapsto \sum_{\pi \in \hat{G}} \frac{d_\pi}{|G|} \text{tr}(x_\pi \circ \pi_{g^{-1}}).$$

PROOF. The Orthogonality Relations show that $F^*: \bigoplus_{\pi \in \hat{G}} \text{End}(V_\pi) \rightarrow \mathbb{C}[G]$ is an isometry. Its adjoint F is injective by Corollary 6.14. Hence F and F^* are unitaries and inverse to each other. □

We have found an explicit isomorphism between $\mathbb{C}[G]$ for a finite group G and a direct sum of matrix algebras. This isomorphism encodes basic results of representation theory.

7. Category algebras and quiver algebras

In this section, we associate convolution algebras to categories, generalising the construction of group algebras in Section 6. A special case are algebras associated to quivers. We discuss several examples that lead to well known finite-dimensional algebras such as full matrix algebras or algebras of lower triangular matrices. Category algebras for finite categories need not be semi-simple any more. The argument in the group case only extends to categories in which all morphisms are invertible – also called groupoids. Groupoids play an important role in noncommutative geometry. We shall not study them much in this course, however. The interesting groupoids are topological groupoids or even smooth groupoids. We will only look at groupoids coming from group actions on manifolds, and these examples can be handled without mentioning groupoids.

The examples treated in some detail in this section are algebra of matrices and upper triangular matrices. Much more could be done. In particular, a quiver gives rise to another algebra, called its Leavitt path algebra. These algebras have a rich structure, and they have received a lot of attention from operator algebraists as well because they are algebraic relatives of graph C^* -algebras. It would be nice to say

more about Leavitt path algebras. But time limitations make me stay away from this topic in this course.

DEFINITION 7.1. Let K be a field. Let \mathcal{C} be a small category. We write \cdot for its composition, \mathcal{C} or $\mathcal{C}^{(1)}$ for its morphism space, and $\mathcal{C}^{(0)}$ for its object space. Let $K[\mathcal{C}]$ be the free K -vector space over $\mathcal{C}^{(1)}$, that is, a K -vector space with a basis $(\delta_f)_{f \in \mathcal{C}^{(1)}}$. Equivalently, $K[\mathcal{C}]$ is the space of all functions $\mathcal{C} \rightarrow K$ with finite support. We define the multiplication on basis vectors of $K[\mathcal{C}]$ by

$$\delta_f * \delta_g := \begin{cases} \delta_{f \cdot g} & \text{if } f \cdot g \text{ is defined,} \\ 0 & \text{otherwise.} \end{cases}$$

This extends to a unique bilinear map

$$*: K[\mathcal{C}] \times K[\mathcal{C}] \rightarrow K[\mathcal{C}], \quad f * g(\alpha) := \sum_{\beta \cdot \gamma = \alpha} f(\beta) \cdot g(\gamma).$$

and turns $K[\mathcal{C}]$ into a K -algebra; associativity is easy to check.

The algebra $K[\mathcal{C}]$ is finite-dimensional if and only if \mathcal{C} is finite.

EXAMPLE 7.2. Let G be a group. View G as a category with one object. Then Definition 7.1 reproduces the group algebra of G .

EXAMPLE 7.3. Let X be a set. View X as a category with only identical morphisms, that is, $\mathcal{C}^{(0)} = \mathcal{C}^{(1)} = X$. Then $\mathbb{C}[X]$ carries the pointwise multiplication, that is, $\mathbb{C}[X] \cong \bigoplus_{x \in X} \mathbb{C}$.

EXAMPLE 7.4. Let $X_n = \{1, 2, \dots, n\}$ and let \mathcal{C}_n be the category with exactly one morphism $i \rightarrow j$ for each $i, j \in X_n$, that is, $\mathcal{C}_n = X_n \times X_n$; this dictates the composition in \mathcal{C}_n . We identify $K[\mathcal{C}_n]$ with the space of functions $X_n \times X_n \rightarrow K$ and let $\delta_{i,j}$ be the basis vector for the unique morphism $j \rightarrow i$ in \mathcal{C}_n . The multiplication in $K[\mathcal{C}_n]$ is given by

$$(f * g)(i, j) := \sum_{k, l, m, p} (f(k, l)g(m, p)\delta_{(k, l) \cdot (m, p)})(i, j) = \sum_{l=1}^n f(i, l)g(l, j).$$

This is exactly the matrix multiplication law. Hence $K[\mathcal{C}_n]$ is isomorphic to the algebra of $n \times n$ -matrices over K .

EXAMPLE 7.5. Let \mathcal{C} be the category with two objects 1 and 2 and three morphisms – the identity morphisms on 1 and 2 and a morphism f from 1 to 2. This is a subcategory of the category \mathcal{C}_2 from Example 7.4. Hence the resulting category algebra is a subalgebra of $K[\mathcal{C}_2] \cong \mathbb{M}_2 K$. More precisely, $K[\mathcal{C}]$ is isomorphic to the subalgebra of lower triangular matrices via

$$\begin{pmatrix} a_{11} & 0 \\ a_{21} & a_{22} \end{pmatrix} \mapsto a_{11}\delta_{\text{Id}_1} + a_{21}\delta_f + a_{22}\delta_{\text{Id}_2}.$$

The concept of a group representation generalises as follows:

DEFINITION 7.6. A (K -linear) *representation* of a category \mathcal{C} is a functor from \mathcal{C} to the category Vect_K of K -vector spaces. The morphisms in the category of representations are the natural transformations.

The following theorem generalises previous results for group representations (Proposition 6.4) and linear maps (Theorem 4.9).

THEOREM 7.7. *Let \mathcal{C} be a small category. The category of K -linear representations of \mathcal{C} is equivalent to the category of non-degenerate $K[\mathcal{C}]$ -modules.*

PROOF. Let $F: \mathcal{C} \rightarrow \text{Vect}_K$ be a representation. We want to construct a non-degenerate representation of $K[\mathcal{C}]$ on $V_F := \bigoplus_{x \in \mathcal{C}^{(0)}} F(x)$. A morphism $\alpha: x \rightarrow y$ in \mathcal{C} acts on V_F by $\iota_y \circ F(\alpha) \circ \pi_x$, where $\iota_y: F(y) \rightarrow V_F$ is the coordinate embedding and $\pi_x: V_F \rightarrow F(x)$ is the coordinate projection. This defines a K -linear map $\varrho_F: K[\mathcal{C}] \rightarrow \text{End}(V_F)$. It is easy to see that this map is an algebra homomorphism. In addition, for each $v \in V_F$ there is an idempotent $e \in K[\mathcal{C}]$ with $e \cdot v = v$ (see Remark 7.8). Thus (V_F, ϱ_F) is a non-degenerate $K[\mathcal{C}]$ -module.

Conversely, let (V, ϱ) be a non-degenerate $K[\mathcal{C}]$ -module. Let $P_x := \varrho(\text{Id}_x)$ for $x \in \mathcal{C}^{(0)}$. These are idempotent operators with $P_x P_y = 0$ for $x \neq y$. If $v \in V$, then $v = \sum_{j=1}^m \varrho(f_j) v_j$ for some $f_j \in K[\mathcal{C}]$, $v_j \in V$. There is a finite subset $S \subseteq \mathcal{C}^{(0)}$ such that all morphisms of \mathcal{C} that belong to the supports of the elements f_j have range in S . Then $\sum_{x \in S} \delta_{\text{Id}_x} * f_j = f_j$ for $j = 1, \dots, m$. This implies $\sum_{x \in S} P_x(v) = v$. Briefly, we may write $\sum_{x \in \mathcal{C}^{(0)}} P_x = \text{Id}_v$. Let $F(x)$ be the range of P_x for $x \in \mathcal{C}^{(0)}$. If $\alpha: x \rightarrow y$ is a morphism, then $\varrho(\delta_\alpha) = P_y \varrho(\delta_\alpha) P_x$. So we may view $\varrho(\delta_\alpha)$ as an operator $F(\alpha): F(x) \rightarrow F(y)$. This defines a representation of \mathcal{C} .

The two constructions above are functors between the categories of representations of \mathcal{C} and non-degenerate $K[\mathcal{C}]$ -modules. They are inverse to each other up to natural isomorphisms. Hence they form an equivalence of categories. If \mathcal{C} has more than one object, this equivalence of categories is not an isomorphism of categories, but only an equivalence because the direct sum of two vector spaces is not unique as a set – it is just unique up to a canonical isomorphism. \square

REMARK 7.8. The algebra $K[\mathcal{C}]$ is unital if and only if \mathcal{C} has only finitely many objects. Then $\sum_{x \in \mathcal{C}^{(0)}} \delta_{\text{Id}_x}$ is a unit element. In general, $K[\mathcal{C}]$ has *idempotent local units*: if $F \subseteq K[\mathcal{C}]$ is a finite subset, then there is $e \in K[\mathcal{C}]$ with $e^2 = e$ and $e \cdot x = x = x \cdot e$ for all $x \in F$. This is shown in the proof of Theorem 7.7.

Whereas group algebras are always semi-simple, this is no longer the case for category algebras, as Example 7.5 shows. Our proof that group algebras are semi-simple uses inversion in the group. Hence the following concept is relevant:

DEFINITION 7.9. A category is called a *groupoid* if each morphism is invertible.

THEOREM 7.10. *The category algebra $\mathbb{C}[\mathcal{C}]$ of a finite groupoid \mathcal{C} is semi-simple.*

PROOF. Define an inner product on $\mathbb{C}[\mathcal{C}]$ in the obvious way, declaring the basis $(\delta_\alpha)_{\alpha \in \mathcal{C}}$ to be orthonormal. Let $f^*(g) := f(g^{-1})$ for $f \in \mathbb{C}[\mathcal{C}]$. Then $\langle f^* * a | b \rangle = \langle a | f * b \rangle$ for all $a, b \in \mathbb{C}[\mathcal{C}]$. Hence $\mathbb{C}[\mathcal{C}]$ is semi-simple by Theorem 5.24. \square

We now discuss some category algebras with large nilpotent ideals.

EXAMPLE 7.11. Let (X, \leq) be a partially ordered set. Let $\mathcal{C}_{(X, \leq)}$ be the category with object set X and morphism set $\{(x, y) \in X : x \geq y\}$ with composition $(x, y) \cdot (y, z) := (x, z)$. That is, there is no morphism from x to y if $x > y$ and there is exactly one morphism (y, x) from x to y if $x \leq y$.

For instance, if $X = \{1, \dots, n\}$ with the usual ordering, then $K[\mathcal{C}_{(X, \leq)}]$ is isomorphic to the algebra of lower triangular matrices (compare Example 7.5). If we use the ordering where $x \leq y$ for all $x, y \in X$, then we get the algebra $\mathbb{M}_n K$.

We assume now that \leq is anti-symmetric, that is, $x \leq y$ and $y \leq x$ implies $x = y$. Let $I \triangleleft A := K[\mathcal{C}_{(X, \leq)}]$ be the subspace spanned by $\delta_{(x, y)}$ with $x \neq y$. This is an ideal because no morphism in $\mathcal{C}_{(X, \leq)}$ is invertible. We claim that I is nilpotent if X is finite. Since $A/I \cong \bigoplus_{x \in X} \mathbb{C}$ is semi-simple, I is the nilradical of A . Thus A is an extension of a semi-simple commutative algebra by a nilpotent algebra.

If $I^n \neq 0$, then there must be a chain of non-identity morphisms $\alpha_1, \dots, \alpha_n$ with $\delta_{\alpha_1} * \dots * \delta_{\alpha_n} \neq 0$. Equivalently, there is a chain of strict inequalities $x_0 < x_1 < \dots < x_n$. This is impossible if $n \geq |X|$. Thus I is nilpotent as asserted.

Example 7.11 can be generalised considerably.

DEFINITION 7.12. A *quiver* is a directed graph, that is, it consists of a set of objects Q^0 and a set of arrows Q^1 with range and source maps $Q^1 \rightrightarrows Q^0$ and no further structure.

Thus “quiver” is a synonym for “directed graph.” A *path* in a quiver is a finite sequence of composable arrows. By convention, there is an “empty path” v starting and ending at each object $v \in Q^0$. The paths in a quiver form a category with respect to concatenation of paths, called its *path category*.

EXERCISE 7.13. For any finite partially ordered set (X, \leq) , define a quiver whose path category is the category described in Example 7.11. Find an infinite partially ordered set whose category is not the path category of a quiver.

EXERCISE 7.14. The path category of a quiver is finite if and only if the quiver is finite and has no (directed) loops; a loop is a path with the same head and tail.

PROPOSITION 7.15. *Let \mathcal{C} be the path category of a finite quiver without loops. Then the nilradical $\text{rad } K[\mathcal{C}]$ is the linear span of δ_α for the non-empty paths α in \mathcal{C} , and the quotient $K[\mathcal{C}]/\text{rad } K[\mathcal{C}]$ is isomorphic to the direct sum $\bigoplus_{x \in \mathcal{C}^{(0)}} \mathbb{C}$.*

PROOF. Since concatenation of paths never produces an empty path, the linear span I of the non-empty paths is a two-sided ideal in $K[\mathcal{C}]$. The quotient $K[\mathcal{C}]/I$ is isomorphic to the commutative semi-simple algebra $\bigoplus_{x \in \mathcal{C}^{(0)}} K$. It only remains to check that I is nilpotent.

Let $I^n \neq \{0\}$. Then there are composable, non-empty paths $\alpha_1, \dots, \alpha_n$ in the quiver. The length of their composition is at least n . Since there are no loops, this composite path cannot visit an object of the quiver twice. So the quiver has at least $n + 1$ objects. Hence $I^{|\mathcal{C}^{(0)}|} = 0$. \square

COROLLARY 7.16. *All irreducible representations of a finite-dimensional quiver algebra are characters. These are in bijection with the vertex set $\mathcal{C}^{(0)}$ of the quiver. The character for $v \in \mathcal{C}^{(0)}$ maps $f \in K[\mathcal{C}]$ to $f(1_v)$.*

PROOF. This is a special case of Theorem 5.22. Since $K[\mathcal{C}]/\text{rad } K[\mathcal{C}]$ is commutative, all irreducible representations are characters. And these correspond to the projections to the summands K in $K[\mathcal{C}]/\text{rad } K[\mathcal{C}]$. \square

Although the simple modules are easy to understand, this does not tell us much about the category of all representations. In order to understand general representations, we should know the *indecomposable* modules instead of the simple modules. Recall that a module is called decomposable if it is a direct sum of two non-trivial modules. It can be shown that any module over a finite-dimensional algebra is a direct sum of indecomposable modules and that this decomposition is unique up to permutation. Indecomposable modules over a finite-dimensional algebra can, in general, not be classified, but there are some cases when this is possible.

EXERCISE 7.17. Let A be the algebra of upper (or lower) triangular 2×2 -matrices studied in Example 4.7. Describe three indecomposable A -modules that are not isomorphic and show that any other indecomposable A -module is isomorphic to one of them.

8. The group algebra of the dihedral group

We have already studied group algebras of finite groups in great detail. Group algebras of infinite groups are much more complicated. In this lecture, we examine a rather simple example, namely, the infinite dihedral group D_∞ . An important

feature of this example is that all its irreducible representations have dimension 1 or 2. The proof uses a version of Schur's Lemma. We classify the irreducible representations of D_∞ up to unitary equivalence. Let A be the group algebra of D_∞ . We show that $\hat{A} \cong \text{Prim}(A)$ and that all primitive ideals of A are maximal. We show that A is a subalgebra of finite codimension in a matrix algebra over a polynomial ring. Therefore, A is finitely generated as a module over its centre. Thus A is still very close to being commutative. We also show that all algebras that are finitely generated over their centre have similar properties.

The group algebra of D_∞ is also equal to the universal unital algebra generated by two projections. This algebra and related C^* -algebras have been studied by different authors, including Cuntz [5] in a C^* -algebraic setting.

DEFINITION 8.1. The *infinite dihedral group* D_∞ is the group of affine transformations of \mathbb{R} that is generated by translations τ_n for $n \in \mathbb{Z}$ and the reflection s at the origin.

Every element of D_∞ is either a translation τ_n or a reflection $\tau_n \circ s$ for some unique $n \in \mathbb{Z}$, and the multiplication table is determined by $\tau_n \tau_m = \tau_{n+m}$, $s^2 = 1$, and $s\tau_n = \tau_{-n}s$. The subgroup of translations $\mathbb{Z} \cong \{\tau_n\}$ in D_∞ is normal, and D_∞ is isomorphic to the semi-direct product group $\mathbb{Z} \rtimes \mathbb{Z}/2$, where $\mathbb{Z}/2$ acts on \mathbb{Z} by $n \mapsto -n$. Hence the group algebra $\mathbb{C}[D_\infty]$ has a basis $\delta_{(n,\epsilon)}$ with $n \in \mathbb{Z}$, $\epsilon \in \mathbb{Z}/2$.

The group D_∞ is generated by the two reflections s and $t := \tau_1 \circ s$ because $ts = \tau_1$, $st = \tau_{-1}$. These reflections satisfy no relations besides $s^2 = t^2 = 1$. Therefore, any two linear operators S and T with $S^2 = T^2 = 1$ determine a representation of D_∞ .

Assume from now on that we work over a field K whose characteristic is not 2. The representation theory of D_∞ is quite different in characteristic 2 because it contains elements of order 2.

Let $p := \frac{1}{2}(1 + s)$ and $q := \frac{1}{2}(1 + t) \in K[D_\infty]$. The relations $s^2 = t^2 = 1$ are equivalent to $p^2 = p$ and $q^2 = q$. Thus representations of D_∞ correspond to pairs of idempotent operators P and Q .

Similarly, unitary representations of D_∞ on a Hilbert space \mathcal{H} correspond to pairs of orthogonal projections $P, Q: \mathcal{H} \rightrightarrows \mathcal{H}$. Since orthogonal projections on Hilbert space correspond to subspaces, the representation theory of D_∞ is equivalent to the study of the relative position of two subspaces in a Hilbert space. Our treatment is purely algebraic and therefore has to disregard unitarity because this would lead to sine and cosine functions, which are not algebraic.

First we describe representations of dimensions 1 and 2. We will show later that this already contains all irreducible representations.

Representations of dimension 1 are characters. In our case, a character on D_∞ is determined by two signs $\chi(s)$ and $\chi(t)$, which may be arbitrary. Thus we get four 1-dimensional representations.

A representation of dimension 2 is specified by $S, T \in \mathbb{M}_2 K$ with $S^2 = T^2 = 1$. If either S or T is $\pm \text{Id}_{K^2}$, then the representation is a direct sum of characters, hence not interesting. Thus we assume that this is not the case. Equivalently, the idempotent operators P and Q both have rank 1.

By a change of basis, we may arrange that

$$P = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad S = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The idempotent matrix Q must have trace 1 because its trace is its rank. Hence it has the form

$$Q = \begin{pmatrix} x & a \\ b & 1-x \end{pmatrix}, \quad x, a, b \in K, \quad ab = x - x^2;$$

the condition $ab = x - x^2$ is equivalent to $Q^2 = Q$. The basis in which P and Q have this form is unique up to scaling the basis vectors. By rescaling the first basis vector, we may simplify Q further and arrange that either $a = 0$ or $a = 1$. If $a = 0$, then the second basis vector is a joint eigenvector of P and Q , so that our representation is reducible. If $a = 1$, then $b = x - x^2$. This case yields a 1-parameter family of 2-dimensional representations with

$$Q = \begin{pmatrix} x & 1 \\ x(1-x) & 1-x \end{pmatrix}, \quad T = \begin{pmatrix} 2x-1 & 2 \\ 2x(1-x) & 1-2x \end{pmatrix}.$$

EXERCISE 8.2. Show that the 2-dimensional representation described above is reducible if and only if $x = 0$ or $x = 1$.

We may also consider the above family of representations as an algebra homomorphism from $K[D_\infty]$ to the algebra of 2×2 -matrix valued polynomials in one variable, $\mathbb{M}_2K[x]$.

PROPOSITION 8.3. *The homomorphism $\varrho: K[D_\infty] \rightarrow \mathbb{M}_2K[x]$ that maps*

$$s \mapsto \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad t \mapsto \begin{pmatrix} 2x-1 & 2 \\ 2x(1-x) & 1-2x \end{pmatrix}$$

is injective, and its range is the subalgebra

$$\left\{ \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix} : f_{21}(0) = f_{21}(1) = 0 \right\}.$$

PROOF. Every element of $K[D_\infty]$ has a decomposition

$$pa_{11}p + pa_{12}(1-p) + (1-p)a_{21}p + (1-p)a_{22}(1-p)$$

with $a_{ij} \in K[D_\infty]$, which corresponds to the decomposition of a 2×2 -matrix into its entries. Any element of D_∞ is either 1 or of the form $s(ts)^n$, $(ts)^nt$, $(st)^{n+1}$, or $(ts)^{n+1}$ for a unique $n \geq 0$. This implies that the elements 1 and $p(qp)^n$, $(qp)^nq$, $(pq)^{n+1}$, $(qp)^{n+1}$ for $n \geq 0$ form a basis of $K[D_\infty]$. Then the elements $1-p$ and $p(qp)^np$, $p(qp)^nq(1-p)$, $(1-p)q(pq)^np$, $(1-p)q(pq)^n(1-p)$ for $n \geq 0$ form a basis of $K[D_\infty]$. Since $\varrho(pqp) = x \cdot \varrho(p)$ and $\varrho(qpq) = x \cdot \varrho(q)$, their images under ϱ are

$$\begin{aligned} \varrho(1-p) &= \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \\ \varrho(p(qp)^np) &= \varrho(p(pqp)^np) = \begin{pmatrix} x^n & 0 \\ 0 & 0 \end{pmatrix}, \\ \varrho(p(qp)^nq(1-p)) &= \varrho(p(pqp)^nq(1-p)) = \begin{pmatrix} 0 & 0 \\ x^n \cdot x(1-x) & 0 \end{pmatrix}, \\ \varrho((1-p)q(pq)^np) &= \varrho((1-p)q(pqp)^np) = \begin{pmatrix} 0 & x^n \\ 0 & 0 \end{pmatrix}, \\ \varrho((1-p)q(pq)^n(1-p)) &= \varrho((1-p)q(qpq)^n(1-p)) = \begin{pmatrix} 0 & 0 \\ 0 & x^n(1-x) \end{pmatrix} \end{aligned}$$

for $n \geq 0$. These matrix-valued polynomials are linearly independent and span the subspace of matrix-valued polynomials with $f_{12}(0) = f_{12}(1) = 0$. \square

COROLLARY 8.4. *The centre of $K[D_\infty]$ is isomorphic to $K[x]$; the isomorphism maps x to $pqp + (1-p)(1-q)(1-p) = 1 + (st + ts)/2$.*

Since $x \mapsto 1 + x/2$ induces an automorphism of $K[x]$, it makes no difference whether we use the generator $st + ts$ or $pqp + (1-p)(1-q)(1-p)$ for the centre.

PROOF. The centre of $\mathbb{M}_2K[x]$ is isomorphic to $K[x]$ and is generated by $\begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix}$. Since the latter is $\varrho(pqp + (1-p)(1-q)(1-p))$ and ϱ is faithful, the centre of $K[D_\infty]$ is isomorphic to the polynomial algebra in the generator $pqp + (1-p)(1-q)(1-p)$. \square

Corollary 8.4 also follows from Proposition 6.9, which identifies the centre of $K[D_\infty]$ with the K -linear span of the characteristic functions of finite conjugacy classes in D_∞ . The translations in D_∞ decompose into finite conjugacy classes $\{\tau_n, \tau_{-n}\}$ for $n \in \mathbb{N}$; the reflections $\tau_n s$ with even and odd n are all conjugate, so that they form two infinite conjugacy classes. Since $\tau_n = (ts)^n$ and $\tau_{-n} = (st)^n$, the centre of $K[D_\infty]$ is spanned by 1 and $(st)^n + (ts)^n$ for $n \geq 1$. The latter may be expressed as polynomials in $st + ts$.

THEOREM 8.5 (Schur's Lemma). *Let K be an algebraically closed uncountable field such as \mathbb{C} . Let A be a K -algebra of at most countable dimension. Let (V, f) be an irreducible representation of A . Let $T: V \rightarrow V$ be an A -module homomorphism, that is, a K -linear map that commutes with $f(A)$. Then $T = \kappa \cdot \text{Id}_V$ for some $\kappa \in K$.*

This is a generalisation of Lemma 6.15, which only deals with the special case of finite-dimensional representations of finite groups. There are several other variants of Schur's Lemma that yield the same conclusion under different assumptions.

PROOF. The operator T must be invertible or zero because its kernel and range are A -submodules and therefore either $\{0\}$ or V by irreducibility. Hence the algebra B of all A -module endomorphisms of (V, f) is a skew-field. Let $v \in V \setminus \{0\}$. The map $B \rightarrow V, T \mapsto T(v)$, is injective because operators in B are either invertible or zero, and the map $A \oplus K \rightarrow V, (a, \kappa) \mapsto a \cdot v + \kappa v$ is surjective. Since A has a countable basis, it follows that V and B have countable bases as well.

Thus we are reduced to proving that there are no skew-fields over K with a countable basis except K itself. Assume $B \neq K$ and pick $b \in B \setminus K$. The elements $(b - \kappa)^{-1} \in B$ for $\kappa \in K$ cannot be linearly independent because K is uncountable. Hence there are finitely many elements $\kappa_j \in K, c_j \in K^*$ with

$$\sum_{j=1}^n c_j (b - \kappa_j)^{-1} = 0.$$

Multiplying through by $\prod_{j=1}^n (b - \kappa_j)$ gives a non-zero polynomial $p \in K[x]$ with $p(b) = 0$. Since K is algebraically closed, p is a product of linear factors and we may write $p(b) = c \prod_{j=1}^k (b - \lambda_j)$. Since $b \notin K$, none of the factors is zero in B . So B has zero divisors and is no skew-field. We have reached the desired contradiction. \square

COROLLARY 8.6. *Let K be an uncountable, algebraically closed field. Then any irreducible representation of $A := K[D_\infty]$ is of dimension at most 2. The canonical map $\hat{A} \rightarrow \text{Prim}(A)$ is bijective and all primitive ideals are maximal.*

More precisely, the set of equivalence classes of irreducible representations is in bijection with $K \setminus \{0, 1\} \sqcup \mathbb{Z}/2 \times \mathbb{Z}/2$, where $y \in K \setminus \{0, 1\}$ corresponds to the representation $\text{ev}_y \circ \varrho$ and points $(n, m) \in \{0, 1\} \times \{0, 1\}$ correspond to the characters given by $s \mapsto (-1)^n, t \mapsto (-1)^m$.

PROOF. Let (V, f) be an irreducible representation of A . The centre $Z(A)$ of A acts on V by scalar multiples of the identity matrix because of Theorem 8.5. Hence there is a character $\chi: Z(A) \rightarrow K$ with $f(a) = \chi(a)\text{Id}_V$ for all $a \in Z(A)$. Now we identify A with a subalgebra of $\mathbb{M}_2 K[x]$ as in Proposition 8.3 and its centre with $K[x]$ as in Corollary 8.4. Characters on $Z(A)$ are of the form $\chi(a) = a(y)$ for some $y \in K$. Hence f vanishes on the two-sided ideal I_y generated by $(x - y) \cdot E$ in $A \subseteq \mathbb{M}_2 K[x]$. Then f descends to an irreducible representation of the quotient algebra A/I_y . These algebras are always finite-dimensional of dimension at most 4.

If $y \neq 0$ and $y \neq 1$, then the evaluation map $\mathbb{M}_2 K[x] \rightarrow \mathbb{M}_2 K$ induces an isomorphism $A/I_y \cong \mathbb{M}_2 K$. Hence there is, in this case, a single irreducible

representation with central character y – this is the representation $\text{ev}_y \circ \varrho$ with ϱ as in Proposition 8.3. If $y = 0$ or $y = 1$, the quotient algebra A/I_y is not semi-simple, and its semi-simple quotient is $\mathbb{C} \oplus \mathbb{C}$. Thus we get two 1-dimensional irreducible representations of A/I_y for both values of y . These are the four characters of A . This finishes the classification of the irreducible representations of A . Our proof also shows that the kernels of these irreducible representations are all different. Hence the map $\hat{A} \rightarrow \text{Prim}(A)$ is bijective. \square

The above proof technique can be extended quite a bit:

DEFINITION 8.7. Let K be a field. A unital K -algebra is called an algebra of *finite type* if it is finitely generated as a module over its centre. That is, there are finitely many elements $x_1, \dots, x_n \in A$ such that every element of A can be written as $x_1 z_1 + \dots + x_n z_n$ with central elements $z_1, \dots, z_n \in Z(A)$.

Proposition 8.3 says that $K[D_\infty]$ is an algebra of finite type.

THEOREM 8.8. *Let A be a unital algebra of finite type over K . Let $Z(A)$ be its centre. Assume that K is uncountable and algebraically closed and that A has a countable basis over K . Then the map $\text{Prim}(A) \rightarrow \hat{A}$ is bijective and all primitive ideals of A are maximal. Each irreducible representation of A restricts to a character on $Z(A)$, and the resulting map $\hat{A} \rightarrow \widehat{Z(A)}$ is finite-to-one.*

PROOF. As in the proof of Corollary 8.6, each irreducible representation restricts to $z \mapsto \chi(z) \cdot \text{Id}_V$ on the centre of A for some character χ of $Z(A)$. The quotient of A by the two-sided ideal generated by $\ker \chi$ is a finite-dimensional algebra over $Z(A) / \ker \chi \cong K$. Hence there are only finitely many irreducible representations for each $\chi \in \widehat{Z(A)}$, and they all have different kernels. \square

The following discussion uses some basic C^* -algebra theory. The results above allow us to classify pairs of closed subspaces in Hilbert spaces. Let \mathcal{H} be a Hilbert space and let $V, W \subseteq \mathcal{H}$. Let P and Q be the orthogonal projections onto V and W , respectively. These are self-adjoint idempotent operators. They generate a unitary representation of D_∞ or, equivalently, a representation of $\mathbb{C}[D_\infty]$ compatible with the involutions. Let $Z := PQP + (1-P)(1-Q)(1-P)$. This is a self-adjoint operator with $0 \leq Z \leq 1$, that is, its spectrum is contained in $[0, 1]$. Our computation of $\mathbb{C}[D_\infty]$ shows that Z commutes with P and Q . The operator Z encodes the relative position of P and Q .

By the spectral theorem, we may simultaneously diagonalise Z and P . Since their joint spectrum is contained in $X := [0, 1] \times \{0, 1\}$, there is a continuous field of Hilbert spaces $(\mathcal{H}_{z,\epsilon})_{(z,\epsilon) \in X}$ over X and a measure μ on X such that $\mathcal{H} \cong L^2(X, (\mathcal{H}_x)_{x \in X}, d\mu)$ and Z and P correspond to the multiplication operators $Zf(z, \epsilon) := z \cdot f(z, \epsilon)$ and $Pf(z, \epsilon) := \epsilon f(z, \epsilon)$. The operator Q is described by a function $[0, 1] \ni z \mapsto Q_z$ of projections on $\mathcal{H}_{z,0} \oplus \mathcal{H}_{z,1}$, which act on \mathcal{H} by pointwise multiplication. The operator $PQ_z(1-P) : \mathcal{H}_{z,0} \rightarrow \mathcal{H}_{z,1}$ is invertible for $z \neq 0, 1$, so that we may identify these two Hilbert spaces by a unitary operator (polar decomposition of PQ_z). After this identification, Q_z becomes the block matrix

$$Q_z = \begin{pmatrix} z & \sqrt{z(1-z)} \\ \sqrt{z(1-z)} & 1-z \end{pmatrix} \cdot \text{Id}_{\mathcal{H}_z} \in \mathbb{B}(\mathcal{H}_z \oplus \mathcal{H}_z).$$

For $z = 0, 1$, the Hilbert spaces $\mathcal{H}_{z,0}$ and $\mathcal{H}_{z,1}$ may be different; there Q_z is block diagonal, $Q_0 = P$ and $Q_1 = 1 - P$.

The 2-dimensional irreducible *unitary* representations of D_∞ are given by the pairs of projections (P, Q_z) with $z \in (0, 1)$. The rank-one projection Q_z belongs to the unit vector $(\sqrt{z}, \sqrt{1-z})$. Hence $\sqrt{z} = \cos(\alpha)$ where α is the angle between the range spaces of Q_z and P .

9. Involutive algebras

The class of C^* -algebras plays an important role in noncommutative geometry. They may be interpreted as algebras of functions on locally compact noncommutative spaces. A C^* -algebra is an involutive Banach algebra with an extra condition on the norm. We shall not consider them much because they deserve their own course. We briefly discuss involutive algebras in this lecture. We use the involution to define some classes of special elements such as projections and isometries. These are used to define interesting classes of noncommutative algebras, such as the Toeplitz algebra to be discussed in the next lecture.

We have seen that a subalgebra of $M_n\mathbb{C}$ is semi-simple if it is stable under the adjoint map $x \mapsto x^*$. And we have constructed such adjoint maps on group algebras for finite and infinite groups. The definition of an involutive algebra formalises some basic algebraic properties of this involution.

DEFINITION 9.1. An *involutive algebra*, briefly **-algebra*, is a \mathbb{C} -algebra A with a map $A \rightarrow A$, $a \mapsto a^*$, that is

- conjugate-linear – $(\lambda a + \mu b)^* = \bar{\lambda}a^* + \bar{\mu}b^*$ for all $\lambda, \mu \in \mathbb{C}$, $a, b \in A$;
- involutive – $(a^*)^* = a$ for all $a \in A$;
- anti-multiplicative $(a \cdot b)^* = b^* \cdot a^*$ for all $a \in A$.

EXAMPLE 9.2. The map $(x_{ij})^* := (\overline{x_{ji}})$ on $M_n\mathbb{C}$ turns $M_n\mathbb{C}$ into a *-algebra for $n \in \mathbb{N} \cup \{\infty\}$.

EXAMPLE 9.3. Let G be a group. Define $f^*(g) := \overline{f(g^{-1})}$ for $f \in \mathbb{C}[G]$, $g \in G$. This turns $\mathbb{C}[G]$ into a *-algebra.

EXAMPLE 9.4. Let X be a smooth manifold. Then $C^\infty(X)$ becomes a *-algebra by $f^*(x) := \overline{f(x)}$ for all $x \in X$, $f \in C^\infty(X)$.

EXAMPLE 9.5. Let \mathcal{H} be a Hilbert space and let $\mathbb{B}(\mathcal{H})$ be the algebra of bounded linear operators on \mathcal{H} . If $x \in \mathbb{B}(\mathcal{H})$, then there is $x^* \in \mathbb{B}(\mathcal{H})$ with $\langle xv | w \rangle = \langle v | x^*w \rangle$ for all $v, w \in \mathcal{H}$. The proof uses that any bounded linear functional $f: \mathcal{H} \rightarrow \mathbb{C}$ is of the form $f(v) = \langle w | v \rangle$ for a unique $w \in \mathcal{H}$ with $\|w\| = \|f\|$.

The involution $x \mapsto x^*$ turns $\mathbb{B}(\mathcal{H})$ into a *-algebra.

We are interested in certain special elements of *-algebras:

DEFINITION 9.6. Let A be a *-algebra. An element $a \in A$ is called *self-adjoint* if $a^* = a$, a *projection* if $a^* = a$ and $a^2 = a$, and a *partial isometry* if $aa^*a = a$ or, equivalently, $a^*aa^* = a^*$. If A is unital, then $a \in A$ is called *unitary* if $aa^* = a^*a = 1$, an *isometry* if $a^*a = 1$, and a *coisometry* if $aa^* = 1$.

If a is a partial isometry, then aa^* and a^*a are projections, called the *range and source projections* of a .

By definition, a is unitary if and only if a is both an isometry and a coisometry, and both isometries and coisometries are partial isometries. Projections are partial isometries as well. An element a is a coisometry if and only if a^* is an isometry – hence we rarely talk about coisometries; and if a is unitary or a partial isometry, so is a^* .

The classes of algebra elements in Definition 9.6 correspond to special classes of operators on Hilbert space:

EXERCISE 9.7. Let \mathcal{H} be a Hilbert space and let $a \in \mathbb{B}(\mathcal{H})$.

- a is a projection if and only if a is the orthogonal projection onto a closed subspace of \mathcal{H} ;
- a is an isometry if and only if $\|av\| = \|v\|$ for all $v \in \mathcal{H}$;

- a is a partial isometry if and only if there are closed subspaces \mathcal{H}_0 and \mathcal{H}_1 of \mathcal{H} such that $a|_{\mathcal{H}_0}$ is a unitary operator from \mathcal{H}_0 onto \mathcal{H}_1 and a vanishes on the orthogonal complement of \mathcal{H}_0 . The range and source projections of a are the orthogonal projections onto \mathcal{H}_1 and \mathcal{H}_0 , respectively.

DEFINITION 9.8. A **-homomorphism* between two *-algebras A and B is an algebra homomorphism $\varphi: A \rightarrow B$ with $\varphi(a^*) = \varphi(a)^*$ for all $a \in A$.

A **-subalgebra* of a *-algebra is, of course, a subalgebra closed under the involution $x \mapsto x^*$. The (unital) *-subalgebra of a *-algebra generated by a subset is the smallest (unital) *-subalgebra containing that subset. We are interested in the unital *-subalgebra of a unital *-algebra generated by a self-adjoint element, a projection, a unitary, or an isometry.

Let A be a unital *-algebra and let $a \in A$ be a self-adjoint element. We get a unital *-homomorphism

$$\mathbb{C}[x] \rightarrow A, \quad \sum c_n x^n \mapsto \sum c_n a^n;$$

here we equip $\mathbb{C}[x]$ with the involution $(\sum c_n x^n)^* := \sum \overline{c_n} x^n$. Its range is the unital *-subalgebra of A generated by a . Since any unital *-homomorphism $\mathbb{C}[x] \rightarrow A$ is of the above form for a unique self-adjoint element, we get a bijection between self-adjoint elements of A and unital *-homomorphisms $\mathbb{C}[x] \rightarrow A$. We say briefly that $x \in \mathbb{C}[x]$ is the *universal self-adjoint element of a unital *-algebra*.

Similarly, the projection $(1, 0)$ in the algebra $\mathbb{C} \oplus \mathbb{C}$ with the involution $(x, y)^* := (\overline{x}, \overline{y})$ is the universal projection in a unital *-algebra, that is, unital *-homomorphisms $\mathbb{C} \oplus \mathbb{C} \rightarrow A$ correspond to projections in A by evaluation at $(1, 0)$. Once again, the unital *-subalgebra of A generated by a projection is the range of the corresponding homomorphism $\mathbb{C} \oplus \mathbb{C} \rightarrow A$.

PROPOSITION 9.9. Equip $\mathbb{C}[t, t^{-1}] \cong \mathbb{C}[\mathbb{Z}]$ with the involution $(\sum_{n \in \mathbb{Z}} c_n t^n)^* := \sum_{n \in \mathbb{Z}} \overline{c_n} t^{-n} = \sum_{n \in \mathbb{Z}} \overline{c_{-n}} t^n$. The element $t \in \mathbb{C}[t, t^{-1}]$ is the universal unitary element of a unital *-algebra. The unital *-subalgebra of a unital *-algebra A generated by a unitary element is the range of the corresponding unital *-homomorphism $\mathbb{C}[t, t^{-1}] \rightarrow A$.

PROOF. Let A be a unital *-algebra. Mapping a homomorphism $f: \mathbb{C}[t, t^{-1}] \rightarrow A$ to $f(t)$ gives a bijection between homomorphisms f and invertible elements in A . The homomorphism f is a *-homomorphism if and only if $f(t)^* = f(t^{-1})$. This says that $f(t)$ is unitary. \square

EXERCISE 9.10. An element a of a unital *-algebra is called *normal* if $aa^* = a^*a$. Let $A = \mathbb{C}[x, x^*]$, polynomials in two variables with the involution defined by $(x)^* = x^*$ and $(x^*)^* = x$. Show that $x \in A$ is the universal normal element in a unital *-algebra.

The universal partial isometry is rather complicated to describe, so that we do not discuss it here. The universal isometry generates an interesting *-algebra – the Toeplitz algebra. Unlike the algebras constructed above, it is no longer commutative because isometries are not normal.

9.1. The Toeplitz algebra. We will first study a particularly important isometry – the unilateral shift – and the *-algebra it generates. We will see that this is the universal isometry.

DEFINITION 9.11. The *unilateral shift* is the operator S on the Hilbert space $\ell_2(\mathbb{N})$ that shifts every basis vector one to the right: $S\delta_n := \delta_{n+1}$ for all $n \in \mathbb{N}$. The *Toeplitz algebra* \mathcal{T} is the *-subalgebra of $\mathbb{B}(\ell_2\mathbb{N})$ generated by S .

The range of S is the linear span of the basis vectors δ_n with $n > 0$. Since the orthogonal complement of this range is spanned by δ_0 , the operator $1 - SS^*$ is the rank-1-projection onto δ_0 , which we denote by E_{00} . Moreover,

$$(9.12) \quad \begin{aligned} E_{nm} &:= |\delta_n\rangle\langle\delta_m| = |S^n(\delta_0)\rangle\langle S^m(\delta_0)| = S^n|\delta_0\rangle\langle\delta_0|(S^m)^* \\ &= S^n(1 - SS^*)(S^m)^* = S^n(S^*)^m - S^{n+1}(S^*)^{m+1}. \end{aligned}$$

Thus the $*$ -subalgebra \mathcal{T} of $\mathbb{B}(\ell_2\mathbb{N})$ generated by S contains the algebra $\mathbb{M}_\infty\mathbb{C}$ of finite matrices in the basis $(\delta_n)_{n \in \mathbb{N}}$. It is easy to see that $\mathbb{M}_\infty\mathbb{C}$ is an ideal in \mathcal{T} .

Next we want to understand the quotient $\mathcal{T}/\mathbb{M}_\infty\mathbb{C}$. Since $S^*S = 1$ and $SS^* \equiv 1 \pmod{\mathbb{M}_\infty\mathbb{C}}$, the images of S and S^* in $\mathcal{T}/\mathbb{M}_\infty\mathbb{C}$ become inverse to each other.

LEMMA 9.13. *Define a unital linear $*$ -map $\varrho: \mathbb{C}[t, t^{-1}] \rightarrow \mathcal{T}$ by $\varrho(t^n) = S^n$ for $n \geq 0$ and $\varrho(t^n) = (S^*)^{-n}$ for $n < 0$. This induces a $*$ -algebra isomorphism $\mathbb{C}[t, t^{-1}] \rightarrow \mathcal{T}/\mathbb{M}_\infty\mathbb{C}$.*

PROOF. The map $\mathbb{C}[t, t^{-1}] \rightarrow \mathcal{T}/\mathbb{M}_\infty\mathbb{C}$ is a $*$ -algebra homomorphism because $SS^* \equiv 1 \equiv S^*S$ modulo the $*$ -ideal $\mathbb{M}_\infty\mathbb{C}$. It is surjective because its range contains the generator S of \mathcal{T} . To check injectivity, take $f \in \mathbb{C}[t, t^{-1}] \setminus \{0\}$. Then $\varrho(f)(\delta_n) \neq 0$ for sufficiently large n . But all elements of $\mathbb{M}_\infty\mathbb{C}$ vanish on δ_n for sufficiently large n . Hence $\varrho(f) \notin \mathbb{M}_\infty\mathbb{C}$. \square

THEOREM 9.14. *The isometry $S \in \mathcal{T}$ is the universal isometry in a unital $*$ -algebra in the sense that for any unital $*$ -algebra A , unital $*$ -homomorphisms $\varrho: \mathcal{T} \rightarrow A$ correspond bijectively to isometries in A via $\varrho \mapsto \varrho(S)$.*

PROOF. We have seen above that the elements E_{nm} for $n, m \in \mathbb{N}$, S^n and $(S^*)^n$ for $n \in \mathbb{N}_{\geq 1}$, and Id form a basis for \mathcal{T} . Using (9.12), this implies that the elements $S^n(S^*)^m$ for $n, m \in \mathbb{N}$ form a basis for \mathcal{T} as well. Let $v \in A$ be an isometry in a unital $*$ -algebra. Then there is a unique linear map $\varrho: \mathcal{T} \rightarrow A$ with $\varrho(S^n(S^*)^m) = v^n(v^*)^m$ for all $n, m \in \mathbb{N}$. This satisfies $\varrho(x)^* = \varrho(x^*)$ and $\varrho(1) = 1$. We must check $\varrho(xy) = \varrho(x)\varrho(y)$ for all $x, y \in \mathcal{T}$. Actually, it suffices to check this if x and y are basis vectors. Using $v^*v = 1$, we may simplify

$$v^{n_1}(v^*)^{m_1} \cdot v^{n_2}(v^*)^{m_2} = \begin{cases} v^{n_1+n_2-m_1}(v^*)^{m_2} & \text{if } n_2 \geq m_1, \\ v^{n_1}(v^*)^{m_1-n_2+m_2} & \text{if } n_2 \leq m_1, \end{cases}$$

and similarly for S instead of v . Hence $\varrho(xy) = \varrho(x)\varrho(y)$ for basis vectors. \square

We have decomposed the Toeplitz algebra into a matrix algebra and a commutative algebra. Such decompositions frequently occur in applications, so that we introduce a name for them:

DEFINITION 9.15. Let $I \subseteq A$ be an ideal in an algebra A . Let A/I be the quotient algebra and let $i: I \rightarrow A$ and $p: A \rightarrow A/I$ be the canonical maps. Then we call the diagram $I \xrightarrow{i} A \xrightarrow{p} A/I$ an *algebra extension* and we call A an extension of A/I by I . More generally, any diagram isomorphic to such a diagram is also called algebra extension.

LEMMA 9.16. *Let $i: K \rightarrow E$ and $p: E \rightarrow Q$ be algebra homomorphisms. They form an algebra extension $K \rightarrow E \rightarrow Q$ if and only if i is injective, p is surjective, and $\ker(p) = i(K)$.*

PROOF. The map i is injective if and only if it defines an isomorphism between K and $i(K) \subseteq E$. The map p induces an isomorphism $E/i(K) \cong Q$ if and only if p is surjective and $\ker(p) = i(K)$. \square

Thus the Toeplitz algebra is an extension of the algebra $\mathbb{C}[t, t^{-1}]$ of Laurent polynomials by the algebra $\mathbb{M}_\infty\mathbb{C}$ of finite matrices.

Unlike the group algebra of the dihedral group studied in Section 8, the Toeplitz algebra has trivial centre: only multiples of the identity element are central because a bounded linear operator on $\ell_2(\mathbb{N})$ that commutes with $\mathbb{M}_\infty\mathbb{C}$ must be a multiple of the identity map. Nevertheless, we can use the extension $\mathbb{M}_\infty\mathbb{C} \hookrightarrow \mathcal{T} \twoheadrightarrow \mathbb{C}[t, t^{-1}]$ to classify primitive ideals and irreducible representations of \mathcal{T} .

First we describe all ideals in \mathcal{T} . If $I \subseteq \mathcal{T}$ is an ideal, so is $I \cap \mathbb{M}_\infty\mathbb{C}$. Since $\mathbb{M}_\infty\mathbb{C}$ is simple, we either have $\mathbb{M}_\infty\mathbb{C} \subseteq I$ or $I \cap \mathbb{M}_\infty\mathbb{C} = \{0\}$.

In the first case $\mathbb{M}_\infty\mathbb{C} \subseteq I$, the ideal I is determined by its image in $\mathcal{T}/\mathbb{M}_\infty\mathbb{C} \cong \mathbb{C}[t, t^{-1}]$, which is still an ideal. Moreover, I is maximal or primitive if and only if its image in $\mathbb{C}[t, t^{-1}]$ is maximal or primitive, respectively.

The ideal structure of $\mathbb{C}[t, t^{-1}]$ is easy to describe: for any non-zero ideal in $\mathbb{C}[t, t^{-1}]$, there are finitely many points $\alpha_1, \dots, \alpha_k \in \mathbb{C} \setminus \{0\}$ and orders $m_j \in \mathbb{Z}_{\geq 1}$ for $j = 1, \dots, k$, such that the ideal consists of all functions that have a zero of order at least m_j at α_j for $j = 1, \dots, k$. Such an ideal is maximal if and only if it is primitive if and only if $k = 1$ and $m_1 = 1$, that is, the ideal is of the form

$$\{f \in \mathbb{C}[t, t^{-1}] : f(\alpha) = 0\}$$

for some $\alpha \in \mathbb{C}^\times$. The corresponding character of \mathcal{T} is the map $\mathcal{T} \rightarrow \mathbb{C}$ that sends $v \mapsto \alpha$, $v^* \mapsto 1/\alpha$. Notice that this character is only compatible with the involutions if $|\alpha| = 1$: there are more characters than $*$ -characters.

Summing up, the case $\mathbb{M}_\infty\mathbb{C} \subseteq I$ gives a copy of $\mathbb{C} \setminus \{0\}$ in the primitive ideal space. All these ideals are maximal and even kernels of characters, that is, the corresponding simple quotients are always \mathbb{C} .

Now consider the second case $I \cap \mathbb{M}_\infty\mathbb{C} = \{0\}$. Since any non-zero $x \in \mathcal{T}$ must map some $\delta_n \in \ell_2\mathbb{N}$ to a non-zero vector, there is $n \in \mathbb{N}$ with $x E_{nn} \neq 0$. But $x E_{nn}$ has rank one, so that $I \cap \mathbb{M}_\infty\mathbb{C} \neq \{0\}$ once $I \neq \{0\}$. Hence the second case only occurs for $I = \{0\}$. This ideal is not maximal, but primitive:

PROPOSITION 9.17. *Let $W \subseteq \ell_2\mathbb{N}$ be the algebraic linear span of the basis vectors $(\delta_n)_{n \in \mathbb{N}}$ and represent \mathcal{T} on W by $S \mapsto S|_W$. This representation is irreducible and faithful, so that \mathcal{T} is a primitive algebra.*

PROOF. The image of the representation on W contains the algebra $\mathbb{M}_\infty\mathbb{C}$ of finite matrices. But a non-zero subspace of W that is invariant under $\mathbb{M}_\infty\mathbb{C}$ must be all of W . \square

PROPOSITION 9.18. *Any irreducible representation of \mathcal{T} that is non-zero on $\mathbb{M}_\infty\mathbb{C}$ is isomorphic to the standard representation in Proposition 9.17. The other irreducible representations of \mathcal{T} are characters on the quotient $\mathcal{T}/\mathbb{M}_\infty\mathbb{C} \cong \mathbb{C}[t, t^{-1}]$. The canonical map $\hat{\mathcal{T}} \rightarrow \text{Prim}(\mathcal{T})$ is a bijection.*

PROOF. Let $f: \mathcal{T} \rightarrow \text{End}(W)$ be an irreducible representation. The subspace $f(\mathbb{M}_\infty\mathbb{C})(W) \subseteq W$ is \mathcal{T} -invariant. Assume $f|_{\mathbb{M}_\infty\mathbb{C}} \neq 0$. Then $f(\mathbb{M}_\infty\mathbb{C})(W)$ must be equal to W . On this subspace, the representation of \mathcal{T} is uniquely determined by $f(x)f(y)w = f(x \cdot y)w$ for all $x \in \mathcal{T}$, $y \in \mathbb{M}_\infty\mathbb{C}$, $w \in W$. So f is irreducible if and only if $f|_{\mathbb{M}_\infty\mathbb{C}}$ is irreducible. And we already know that this happens if and only if $f|_{\mathbb{M}_\infty\mathbb{C}}$ is the standard representation on $\mathbb{C}[\mathbb{N}]$ (see Proposition 5.7).

Now assume $f|_{\mathbb{M}_\infty\mathbb{C}} = 0$. Then f descends to a representation of $\mathcal{T}/\mathbb{M}_\infty\mathbb{C} \cong \mathbb{C}[t, t^{-1}]$. Since this algebra is commutative, we already know that all its irreducible representations are characters. Both cases together show that the canonical map $\hat{\mathcal{T}} \rightarrow \text{Prim}(\mathcal{T})$ is a bijection. \square

The representation of \mathcal{T} on $\ell_2(\mathbb{N})$ is not irreducible in the algebraic sense that we use because it contains the non-trivial invariant subspace W . Nevertheless, our

analysis of \mathcal{T} has implications for the structure of isometries on Hilbert space. Any non-degenerate Hilbert space $*$ -representation of $M_\infty\mathbb{C}$ is a direct sum of copies of the standard representation on $\ell_2\mathbb{N}$. Any isometry generates a representation of \mathcal{T} and hence a representation of $M_\infty\mathbb{C}$, which may be degenerate. The non-degenerate part must be a sum of copies of the standard representation of $M_\infty\mathbb{C}$. But the only extension of the standard representation of $M_\infty\mathbb{C}$ to one of \mathcal{T} is the irreducible faithful representation generated by the unilateral shift (compare Proposition 9.18). On the part where $M_\infty\mathbb{C}$ acts by zero, the isometry V must satisfy $VV^* = 1$ and hence be unitary. Hence we get the following statement:

THEOREM 9.19. *Any isometry on a Hilbert space can be decomposed as a direct sum of a unitary and of copies of the unilateral shift.*

9.2. Leavitt path algebras. The Toeplitz algebra is a special case of a Leavitt path algebra (see [1]). These algebras have received a lot of attention recently. They are noncommutative algebras that are quite tractable, but show some curious properties. They are algebraic analogues of graph C^* -algebras. Here we define Leavitt path algebras and exhibit the Toeplitz algebra as a special case. The ideal structure of Leavitt path algebras is known, but we do not discuss it here. The representation theory is usually complicated.

Let Γ be a directed graph or quiver. We describe it through a set of vertices Γ_0 , a set of edges Γ_1 , and the range and source maps $s, r: \Gamma_1 \rightrightarrows \Gamma_0$. A vertex $v \in \Gamma_0$ is called *regular* if $r^{-1}(v)$ is non-empty and finite. We chose a set $\Gamma'_0 \subseteq \Gamma_0$ of regular vertices. We get the “absolute” Leavitt path algebra if we let Γ'_0 be the set of all regular vertices.

DEFINITION 9.20. The *Leavitt path algebra* $L(\Gamma, \Gamma'_0)$ of Γ relative to Γ'_0 is the universal $*$ -algebra with generators S_v for $v \in \Gamma_0$ and S_α for $\alpha \in \Gamma_1$, subject to the following relations:

- (V) $S_v \cdot S_w = 0$ for $v \neq w$ and $S_v^2 = S_v = S_v^*$;
- (E) $S_\alpha S_{s(\alpha)} = S_\alpha$ and $S_{r(\alpha)} S_\alpha = S_\alpha$ for all α ;
- (CK1) $S_\alpha^* S_\beta = 0$ if $\alpha \neq \beta$ and $S_\alpha^* S_\alpha = S_{s(\alpha)}$;
- (CK2) $\sum_{\alpha \in r^{-1}(v)} S_\alpha S_\alpha^* = S_v$ for all $v \in \Gamma'_0$.

The relation (CK2), called also second Cuntz–Krieger relation, is only imposed for vertices in the chosen subset Γ'_0 . This makes our construction relative to Γ'_0 .

The relations (V) and (E) imply

$$S_\alpha S_v = S_\alpha S_{s(\alpha)} S_v = \delta_{s(\alpha), v} S_\alpha, \quad S_v S_\alpha = S_{r(\alpha)} S_\alpha S_v = \delta_{r(\alpha), v} S_\alpha$$

for all $\alpha \in \Gamma_1, v \in \Gamma_0$. Taking adjoints, this implies

$$S_\alpha^* S_v = \delta_{r(\alpha), v} S_\alpha^*, \quad S_v S_\alpha^* = \delta_{s(\alpha), v} S_\alpha^*.$$

So (V) and (E) imply $S_\alpha^* S_\beta = S_\alpha^* S_{r(\alpha)} S_{r(\beta)} S_\beta = 0$ if $r(\alpha) \neq r(\beta)$. Thus it suffices to impose the relation (CK1) only for $\alpha, \beta \in \Gamma_1$ with $r(\alpha) = r(\beta)$.

REMARK 9.21. If Γ_0 is finite, then the sum of S_v over all $v \in \Gamma_0$ is a unit element in the Leavitt path algebra. If Γ_0 is infinite, then the Leavitt path algebra only has an idempotent local unit as in Remark 7.8, formed by the projections $\sum_{v \in F} S_v$ for finite subsets $F \subseteq \Gamma_0$.

PROPOSITION 9.22. *Let Γ have two vertices a, b and two edges $v: a \rightarrow b$ and $w: b \rightarrow b$. Let $\Gamma'_0 = \{b\}$. The Leavitt path algebra $L(\Gamma, \Gamma'_0)$ is isomorphic as a $*$ -algebra to the Toeplitz algebra \mathcal{T} .*

PROOF. Let $Y := S_v + S_w$. We compute

$$\begin{aligned} Y^*Y &= S_v^*S_v + S_w^*S_w = S_a + S_b = 1, \\ YY^* &= S_vS_v^* + S_wS_w^* = S_b \neq 1, \\ YS_b &= S_vS_b + S_wS_b = S_w, \\ Y(1 - S_b) &= S_vS_a + S_wS_a = S_v. \end{aligned}$$

The subalgebra generated by Y^* and Y also contains $YS_b = S_w$ and $Y(1 - S_b) = S_v$. Thus Y^* and Y generate the Leavitt path algebra. The relation $Y^*Y = 1$ implies that there is a unital $*$ -homomorphism $\varrho: \mathcal{T} \rightarrow L(\Gamma, \Gamma'_0)$ that maps S to Y . The computations above show that the elements

$$\hat{S}_a := 1 - SS^*, \quad \hat{S}_b := SS^*, \quad \hat{S}_v := S^2S^*, \quad \hat{S}_w := S - S^2S^*$$

are ϱ -preimages of S_a , S_b , S_v and S_w , respectively. Some more computations show that they satisfy the defining relations of $L(\Gamma, \Gamma'_0)$. Hence there is a unital $*$ -homomorphism

$$\varphi: L(\Gamma, \Gamma'_0) \rightarrow \mathcal{T}, \quad S_a \mapsto \hat{S}_a, \quad S_b \mapsto \hat{S}_b, \quad S_v \mapsto \hat{S}_v, \quad S_w \mapsto \hat{S}_w.$$

By construction, $\varrho \circ \varphi = \text{Id}_{L(\Gamma, \Gamma'_0)}$. And $\varphi \circ \varrho(S) = \hat{S}_v + \hat{S}_w = S$ implies $\varphi \circ \varrho = \text{Id}_{\mathcal{T}}$. \square

EXAMPLE 9.23. Let Γ be the graph with one vertex a and with n edges. Let $\Gamma'_0 = \Gamma$. Let \mathcal{O}_n denote the resulting Leavitt path algebra. It is the universal $*$ -algebra generated by n isometries S_1, \dots, S_n with the relation $\sum_{j=1}^n S_j S_j^* = 1$. This is because $S_a = 1$ and because $\sum_{j=1}^n S_j S_j^* = 1$ implies that the range projections $S_j S_j^*$ of the isometries S_j for $j = 1, \dots, n$ are orthogonal. We may also phrase all these defining relations as follows: the row matrix with entries S_1, \dots, S_n in $\mathbb{M}_{1 \times n} \mathcal{O}_n$ is unitary. In other words, \mathcal{O}_n is the universal $*$ -algebra with a unitary element in $\mathbb{M}_{1 \times n} A$. The general theory of Leavitt path algebras shows that \mathcal{O}_n is simple.

10. Crossed products

Let G be a discrete group and let A be an algebra. Given an action of G on A by automorphisms, we define the crossed product $A \rtimes G$. This is a very important method to construct noncommutative algebras. It goes back to Emmy Noether. The crossed product for the trivial action of G on the ground field is the group algebra. Like the group algebra, the crossed product may be characterised by a universal property that describes its representations. These are covariant representations, namely, a representation of the group and a representation of the algebra that satisfy a certain compatibility condition.

We have studied the representation theory of finite groups by studying the structure of its group algebra. In this section, we classify the irreducible representations of the crossed product $C^\infty(X) \rtimes G$ for an action of a finite group G on a smooth manifold X : they are given by a point $x \in X$ and an irreducible representation of the stabiliser subgroup $G_x \subseteq G$ of x . The same result remains true for proper actions of infinite discrete groups. We do not prove this more general result, but we examine a few simple special cases, namely, homogeneous spaces where G acts on G/H for a subgroup H in G . These examples will motivate the study of Morita equivalence. Another important example that we treat is the structure of the group algebra of a semidirect product group $G = N \rtimes H$: this is isomorphic to the crossed product algebra $\mathbb{C}[N] \rtimes H$ for the canonical action of H on the group algebra $\mathbb{C}[N]$ of N . This is particularly useful if N is Abelian and H is finite. In particular, we apply this to the dihedral group studied in Section 8.

DEFINITION 10.1. Let G be a group, let A be an algebra over a field K , and let $\alpha: G \rightarrow \text{Aut}(A)$ be a group homomorphism, where $\text{Aut}(A)$ denotes the group of algebra automorphisms of A . Then the *crossed product* $G \rtimes_{\alpha} A = A \rtimes_{\alpha} G$ is the vector space $K[G] \otimes_K A$ of functions $G \rightarrow A$ with finite support, equipped with the *convolution product*

$$(f_1 * f_2)(g) := \sum_{h \in G} f_1(h) \cdot \alpha_h(f_2(h^{-1}g)).$$

It is easy to check that the convolution product is associative.

For $g \in G$, $a \in A$, let $a\delta_g \in G \rtimes_{\alpha} A$ denote the function $G \rightarrow A$ with $a\delta_g(h) = 0$ for $h \neq g$ and $a\delta_g(g) = a$. Every element of $G \rtimes A$ decomposes uniquely as a sum $\sum_{g \in F} a(g)\delta_g$ for some finite subset $F \subseteq G$.

REMARK 10.2. The convolution product may also be characterised by the rule

$$a\delta_g * b\delta_h = a\alpha_g(b)\delta_{gh}.$$

Roughly speaking, when we move a group element past an element of A , we must act on the latter by α . As a consequence, if A is unital, then $1_A\delta_1$ is a unit element in $G \rtimes_{\alpha} A$.

Next we relate representations of the crossed product to representations of A and G .

DEFINITION 10.3. A *covariant representation* of (A, G, α) is a pair of representations $f: A \rightarrow \text{End}(V)$, $U: G \rightarrow \text{Aut}(V)$ on the same vector space V that satisfy the covariance condition $U_g f(a) U_g^{-1} = f(\alpha_g(a))$ for all $g \in G$, $a \in A$.

PROPOSITION 10.4. Let A be a unital algebra, let G be a group, and let $\alpha: G \rightarrow \text{Aut}(A)$ be a group homomorphism. Then the category of unital representations of $A \rtimes_{\alpha} G$ is isomorphic to the category of unital covariant representations of (A, G, α) ; the morphisms in the latter are G -equivariant A -linear maps.

PROOF. Define $i_A: A \rightarrow A \rtimes_{\alpha} G$ and $i_G: G \rightarrow A \rtimes_{\alpha} G$ by $i_A(a) := a\delta_1$ and $i_G(g) := 1_A \cdot \delta_g$. If ϱ is a representation of $A \rtimes_{\alpha} G$, then $\varrho \circ i_A$ is a representation of A and $\varrho \circ i_G$ is a representation of G . These are covariant because $\delta_g a \delta_g^{-1} = \alpha_g(a)\delta_1$.

Conversely, a covariant representation (f, U) determines a representation φ of the crossed product by $\varphi(h) := \sum_{g \in G} f(h(g))U_g$ for all $h \in A \rtimes_{\alpha} G$; this even works without a unit element in A .

The two constructions above are inverse to each other, that is, we get a bijection between unital covariant representations of (A, G, α) and unital representations of $A \rtimes_{\alpha} G$ on V . Moreover, a map is $A \rtimes_{\alpha} G$ -linear if and only if it is both A -linear and G -equivariant. \square

REMARK 10.5. Proposition 10.4 breaks down if A is not unital. For instance, the zero map on A is covariant with respect to any group representation of G . This gives many different covariant representations that all produce the zero representation on $A \rtimes G$. We must add a non-degeneracy condition to extend Proposition 10.4 to a non-unital algebra. We do not discuss this here.

EXAMPLE 10.6. Let $\lambda: A \rightarrow \text{End}(A)$ be the left regular representation, $\lambda_a(b) := a \cdot b$. Then the pair (λ, α) is a covariant representation:

$$\alpha_g \lambda_a \alpha_{g^{-1}}(b) = \alpha_g(a \cdot \alpha_{g^{-1}}(b)) = \alpha_g(a)b = \lambda_{\alpha_g(a)}(b)$$

for all $g \in G$, $a, b \in A$. It generates a left A -module structure

$$A \rtimes_{\alpha} G \rightarrow \text{End}(A), \quad f * a := \sum_{g \in G} f(g) \cdot \alpha_g(a).$$

Similarly, we may turn A into a right $A \rtimes_{\alpha} G$ -module by

$$a * f := \sum_{g \in G} \alpha_g^{-1}(a \cdot f(g)) \quad \text{for } a \in A, f \in A \rtimes_{\alpha} G.$$

Let A be a $*$ -algebra and $\alpha: G \rightarrow \text{Aut}(A, *)$ an action by $*$ -algebra automorphisms. That is, $\alpha_g(a^*) = \alpha_g(a)^*$ for all $g \in G, a \in A$. Then $G \rtimes_{\alpha} A$ is a $*$ -algebra as well. The involution is defined by

$$(a\delta_g)^* = \delta_g^* a^* = \delta_{g^{-1}} a^* = \alpha_{g^{-1}}(a^*) \delta_{g^{-1}}$$

or, equivalently, if $f: G \rightarrow A$ has finite support, then

$$f^*(g) = \alpha_g(f(g^{-1})^*).$$

The above crossed product construction is closely related to semi-direct products of groups. Let N and H be groups and let $\alpha: H \rightarrow \text{Aut}(N)$ be a group homomorphism, where $\text{Aut}(N)$ denotes the group of group automorphisms of N . The *semi-direct product* $H \rtimes_{\alpha} N = N \rtimes_{\alpha} H$ is the set $N \times H$ with the multiplication

$$(n_1, h_1) \cdot (n_2, h_2) := (n_1 \alpha_{h_1}(n_2), h_1 h_2) \quad \text{for all } n_1, n_2 \in N, h_1, h_2 \in H.$$

This defines a group structure on $N \times H$.

EXAMPLE 10.7. The isometry group of \mathbb{R}^n is the semi-direct product of the group $N = \mathbb{R}^n$ of translations and the orthogonal group $H = \text{O}(n)$ with respect to the canonical linear action $\alpha: \text{O}(n) \rightarrow \text{Aut}(\mathbb{R}^n)$.

LEMMA 10.8. *Let N and H be groups and let $\alpha: H \rightarrow \text{Aut}(N)$ be a group homomorphism. This induces a group homomorphism $\bar{\alpha}: H \rightarrow \text{Aut}(K[N])$ by $\bar{\alpha}(\delta_n) := \delta_{\alpha(n)}$ for all $n \in N$. The map $\delta_n \delta_h \mapsto \delta_{(n,h)}$ is an algebra isomorphism $K[N] \rtimes_{\bar{\alpha}} H \cong K[N \rtimes_{\alpha} H]$.*

PROOF. This is a trivial computation. \square

EXAMPLE 10.9. The dihedral group studied in Section 8 is a semi-direct product $D_{\infty} = \mathbb{Z} \rtimes_{\alpha} \mathbb{Z}/2$, where $\alpha_n(m) = (-1)^n \cdot m$ for $m \in \mathbb{Z}, n \in \mathbb{Z}/2$. Hence

$$K[D_{\infty}] \cong K[\mathbb{Z}] \rtimes \mathbb{Z}/2 \cong K[t, t^{-1}] \rtimes \mathbb{Z}/2.$$

Many results about $K[D_{\infty}]$ extend to crossed products of finite group actions on commutative algebras. The finiteness of the group is crucial here.

PROPOSITION 10.10. *Let A be a unital \mathbb{C} -algebra, let G be a finite group, and let $\alpha: G \rightarrow \text{Aut}(A)$ be a group homomorphism. Let G act on $B := A \otimes \text{End}(\mathbb{C}[G])$ by $g \cdot (a \otimes x) := \alpha_g(a) \otimes \lambda_g x \lambda_g^{-1}$, where $\lambda_g f(x) := f(g^{-1}x)$ denotes the left regular representation of G on $\mathbb{C}[G]$; this is an action by algebra automorphisms. Then $A \rtimes_{\alpha} G$ is naturally isomorphic to the fixed point subalgebra of the G -action on B .*

PROOF. Define $\varphi: A \rtimes_{\alpha} G \rightarrow B$ by

$$\varphi(a\delta_h) := \sum_{g \in G} \alpha_g(a) \otimes |\delta_g\rangle\langle\delta_{gh}|$$

for $a \in A, h \in G$. Elements of B are linear combinations of $a \otimes |\delta_g\rangle\langle\delta_h|$ for $a \in A, g, h \in G$. We may project B to its fixed-point subalgebra by averaging. This maps $a \otimes |\delta_g\rangle\langle\delta_h|$ to $|G|^{-1} \sum_{x \in G} \alpha_x(a) \otimes |\delta_{xg}\rangle\langle\delta_{xh}|$. Therefore, φ is a vector space isomorphism from $A \rtimes_{\alpha} G$ onto the fixed point subalgebra of B . We compute

$$\begin{aligned} \varphi(a\delta_h)\varphi(b\delta_k) &= \sum_{g \in G} \sum_{l \in G} \alpha_g(a)\alpha_l(b) \otimes |\delta_g\rangle\langle\delta_{gh}||\delta_l\rangle\langle\delta_{lk}| \\ &= \sum_{g \in G} \alpha_g(a)\alpha_{gh}(b) \otimes |\delta_g\rangle\langle\delta_{ghk}| = \varphi(a\alpha_h(b)\delta_{hk}). \end{aligned}$$

Thus φ is an algebra isomorphism. \square

LEMMA 10.11. *Let G be a countably infinite group. Let $A = \mathbb{C}[G]$ with pointwise multiplication and let $\alpha = \lambda$ be the left regular representation. Then $A \rtimes_{\alpha} G \cong \mathbb{M}_{\infty}\mathbb{C}$.*

PROOF. The elements δ_g for $g \in G$ form a basis of $A = \mathbb{C}[G]$. Then $A \rtimes_{\alpha} G$ is spanned by $\delta_g \otimes \delta_h$ for $g, h \in G$. The multiplication is

$$(\delta_g \otimes \delta_h) \cdot (\delta_k \otimes \delta_l) = \delta_{g, h \cdot k} \delta_g \otimes \delta_{h \cdot l}$$

Since G is countably infinite, we may label elements of G by natural numbers. This identifies $\mathbb{M}_{\infty}\mathbb{C}$ with the algebra that has a basis $\delta_g \otimes \delta_h$ for $g, h \in G$ with the multiplication

$$(\delta_g \otimes \delta_h) \cdot (\delta_k \otimes \delta_l) = \delta_{h, k} \delta_g \otimes \delta_l.$$

The bijection on basis vectors that maps $\delta_g \otimes \delta_h \in A \rtimes G$ to $\delta_g \otimes \delta_{h^{-1}g} \in \mathbb{M}_{\infty}\mathbb{C}$ is an algebra homomorphism. \square

Now we are going to study the representation theory and ideal structure of crossed products $C^{\infty}(X) \rtimes_{\alpha} G$ for a smooth compact manifold X and a finite group G .

First we consider the simplest possible case $G \rtimes C^{\infty}(G/H)$ for a subgroup H in a finite group G , where G acts on G/H by left translation. Since G/H is discrete, $C^{\infty}(G/H)$ is a direct sum of finitely many copies of \mathbb{C} .

LEMMA 10.12. *There is an algebra isomorphism*

$$G \rtimes C^{\infty}(G/H) \cong \mathbb{C}[H] \otimes \mathbb{M}_{|G/H|}\mathbb{C}.$$

Thus $G \rtimes C^{\infty}(G/H)$ is semi-simple and its irreducible representations correspond bijectively to irreducible representations of H .

PROOF. We identify $\mathbb{M}_{|G/H|}\mathbb{C} \cong \text{End}(\mathbb{C}[G/H])$. Choose a set of representatives $R \subseteq G$ for G/H . We define a linear map

$$\varphi: \mathbb{C}[H] \otimes \text{End}(\mathbb{C}[G/H]) \rightarrow C^{\infty}(G/H) \rtimes G, \quad \delta_h \otimes |\delta_x\rangle\langle\delta_y| \mapsto \delta_{xH} \otimes \delta_{xhy^{-1}},$$

for $h \in H, x, y \in R$. This is a vector space isomorphism. And it is also multiplicative:

$$\begin{aligned} \varphi(\delta_h \otimes |\delta_x\rangle\langle\delta_y|) \varphi(\delta_k \otimes |\delta_z\rangle\langle\delta_w|) &= \delta_{xH} \cdot \delta_{xhy^{-1}zH} \otimes \delta_{xhy^{-1}} \delta_{zkw^{-1}} \\ &= \delta_{y,z} \delta_{xH} \otimes \delta_{xhkw^{-1}} = \varphi(\delta_{y,z} \delta_{hk} \otimes |\delta_x\rangle\langle\delta_w|) = \varphi((\delta_h \otimes |\delta_x\rangle\langle\delta_y|) \cdot (\delta_k \otimes |\delta_z\rangle\langle\delta_w|)); \end{aligned}$$

here we have used that $xH = xhy^{-1}zH$ if and only if $yH = zH$. Hence φ is an algebra isomorphism.

The group algebra $\mathbb{C}[H]$ is isomorphic to a direct sum of matrix algebras indexed by the irreducible representations of H . Since $\mathbb{M}_n\mathbb{C} \otimes \mathbb{M}_m\mathbb{C} \cong \mathbb{M}_{n \cdot m}\mathbb{C}$, this carries over to $\mathbb{C}[H] \otimes \text{End}(\mathbb{C}[G/H])$. Hence $C^{\infty}(G/H) \rtimes_{\alpha} G$ is semi-simple and its irreducible representations correspond to irreducible representations of the group H . \square

Explicitly, the irreducible representation of $C^{\infty}(G/H) \rtimes G$ corresponding to $(V, \pi) \in \hat{H}$ is constructed as follows. Let

$$W := \{f: G \rightarrow V : f(gh^{-1}) = \pi_h f(g) \text{ for all } g \in G, h \in H\}.$$

We let $C^{\infty}(G/H)$ act on W by pointwise multiplication and G by left translations. This is a covariant representation. It generates a representation of $C^{\infty}(G/H) \rtimes_{\alpha} G$ by $f_1 \delta_g * f_2(x) := f_1(g) f_2(g^{-1}x)$. Notice that a function in W is determined uniquely by its values on a set of representatives for the orbit space G/H . Hence we may identify $W \cong \mathbb{C}[G/H] \otimes V$. But the action of G becomes more complicated on $\mathbb{C}[G/H] \otimes V$.

EXERCISE 10.13. Check that the irreducible representation of $C^\infty(G/H) \rtimes_\alpha G$ corresponding to $(v, \pi) \in \hat{H}$ is the one on W described above. (Hint: Identify $C^\infty(G/H) \rtimes_\alpha G \cong \mathbb{C}[H] \otimes \text{End}(\mathbb{C}[G/H])$ and check that the induced action of H on the range of the rank-1-projection $|\delta_H\rangle\langle\delta_H|$ in $\text{End}(\mathbb{C}[G/H])$ on W is equivalent to (V, π) .)

THEOREM 10.14. *Let X be a smooth compact manifold, let G be a finite group, and let α be an action of G on X by diffeomorphisms. This induces an action of G on $C^\infty(X)$ by algebra automorphisms. Let $A := C^\infty(X) \rtimes_\alpha G$. Then the map $\hat{A} \rightarrow \text{Prim}(A)$ is bijective and any primitive ideal in A is maximal and closed in the natural topology. There is a canonical bijection between \hat{A} and $\bigsqcup_{x \in G \backslash X} \hat{G}_x$, where $G \backslash X$ denotes the orbit space, $G_x \subseteq G$ denotes the stabiliser of $x \in X$, and \hat{G}_x denotes the set of irreducible representations of G_x .*

The same assertions hold for non-compact X if we replace $C^\infty(X)$ by $C_c^\infty(X)$ or if we restrict attention to *closed* primitive ideals and *continuous* irreducible representations of $C^\infty(X) \rtimes G$.

PROOF. For each $x \in G \backslash X$, restricting to the orbit $G \cdot x \cong G/G_x$ defines a quotient mapping $A = C^\infty(X) \rtimes_\alpha G \rightarrow C^\infty(G/G_x) \rtimes_\alpha G$. Hence any irreducible representation of $C^\infty(G/G_x) \rtimes_\alpha G$ yields an irreducible representation $(V_{x,\pi}, \varrho_{x,\pi})$ of A . Since $C^\infty(G/G_x) \rtimes_\alpha G$ is semi-simple and the quotient map $A \rightarrow C^\infty(G/G_x) \rtimes_\alpha G$ is continuous, the kernels of the resulting irreducible representations of A are both maximal ideals and closed ideals, and the correspondence between these irreducible representations and their kernels is bijective. It remains to show that any irreducible representation of A is of this form for some $x \in X$ and some irreducible representation of $C^\infty(G/G_x) \rtimes_\alpha G$; the latter are described in Lemma 10.12.

Any simple A -module is of the form A/L for some left ideal L . We claim that there is a quotient mapping of A -modules $A/L \rightarrow V_{x,\pi}$ for some $x \in X$ and some irreducible representation π of $C^\infty(G/G_x) \rtimes_\alpha G$. Simplicity yields $A/L \cong V_{x,\pi}$, and we are done. The proof of this claim is similar to the proof of Proposition 3.14. Let L be a left ideal and assume that there is no surjective module homomorphism $A/L \rightarrow V_{x,\pi}$ for any x, π . We must show that $L = A$.

For $x \in X$, let J be the image of L in $A_x := C^\infty(Gx) \rtimes G$; this is a left ideal in A_x . If it is not all of A_x , then there are quotient maps $A/L \rightarrow A_x/J \rightarrow V_{x,\pi}$ for some π . Hence $J = A_x$. So there is $f_x \in L$ whose image in A_x is the unit element.

Next we use Proposition 10.10 to identify A with the subalgebra of G -fixed elements in $C^\infty(X, \text{End}(\mathbb{C}[G]))$. Hence we view elements of $C^\infty(X) \rtimes G$ as matrix-valued smooth functions. We observe that the G -fixed point subalgebra A is closed under taking adjoints and under taking inverses of functions whose values are everywhere invertible.

By construction, $f_x(x)$ is an invertible matrix. Since the invertible matrices form an open subset of $\mathbb{M}_{|G|}\mathbb{C}$, $f_x(y)$ is invertible in some open neighbourhood U_x of x . Since X is compact, there are finitely many such open neighbourhoods U_{x_1}, \dots, U_{x_n} that cover X . Let

$$f := f_x^* f_x + \dots + f_y^* f_y \in L.$$

The values of this function f are strictly positive matrices, so that f is invertible in $C^\infty(X) \otimes \text{End}(\mathbb{C}[G]) \cong C^\infty(X, \mathbb{M}_{|G|}\mathbb{C})$ and hence in $C^\infty(X) \rtimes_\alpha G$. Since the left ideal L contains an invertible element, we get $L = C^\infty(X) \rtimes_\alpha G$ as asserted. \square

11. Morita equivalence

11.1. Motivation. Let X be a smooth compact manifold, let G be a group, and let $\alpha: G \rightarrow \text{Diffeo}(X)$ be a group action by diffeomorphisms on X . A basic paradigm

of noncommutative geometry is that the noncommutative algebra $C^\infty(X) \rtimes_\alpha G$ is a good substitute for the orbit space $G \backslash X$ and should be studied instead of the latter.

This is particularly interesting for infinite groups G , where the orbit space $G \backslash X$ is usually very badly behaved.

EXAMPLE 11.1. Let $X := \mathbb{T} := \{z \in \mathbb{C} : |z| = 1\}$ be a circle and let $G := \mathbb{Z}$ act on X by rotations: $n \bullet z := \exp(2\pi i \vartheta n) \cdot z$ for all $n \in \mathbb{Z}$, for some parameter $\vartheta \in \mathbb{R}$. If ϑ is rational, $\vartheta = p/q$, then $q \bullet z = z$ for all $z \in X$, so that the action of G factors through an action of the finite group $\mathbb{Z}/q\mathbb{Z}$. Hence $G \backslash X = (\mathbb{Z}/q\mathbb{Z}) \backslash X$ is again a circle. If ϑ is irrational, then the orbit $Gz := \{n \bullet z : n \in \mathbb{Z}\}$ is dense in X for each $z \in X$. Hence any G -invariant continuous function on X is constant. The orbit space $G \backslash X$ carries no useful topology and is certainly not a smooth manifold.

Spaces of the form $G \backslash X$ for *finite* G are typical examples of *orbifolds*. They need not be smooth manifolds any more because fixed points of the G -action on X provide singularities. But these singularities are easy enough to understand, so that many results about smooth manifolds can be extended to orbifolds. When studying orbifolds, it is crucial to take into account the stabilisers of points as well – this is the information that makes the singularities tractable. This point of view is built into the noncommutative geometry approach of studying the algebra $C^\infty(X) \rtimes_\alpha G$ instead of $G \backslash X$.

Theorem 10.14 shows that isomorphism classes of irreducible representations of $C^\infty(X) \rtimes_\alpha G$ correspond to pairs (x, ϱ) with $x \in G \backslash X$ and an irreducible representation ϱ of the stabiliser G_x of x . This is still reasonably close to $G \backslash X$. There is a canonical map $\text{Prim}(C^\infty(X) \rtimes_\alpha G) \rightarrow G \backslash X$, which is surjective and finite-to-one. The representations ϱ encode some information about the stabilisers of points in X .

If the group G is finite and the group action α is free, that is, $g \cdot x = x$ implies $g = 1$, then $G \backslash X$ is again a smooth manifold. Since we want $C^\infty(X) \rtimes_\alpha G$ to model the orbifold $G \backslash X$, our paradigm leads to the following *correspondence principle*: if a finite group G acts freely on a smooth manifold X , then noncommutative geometry should not distinguish between the algebras $C^\infty(X) \rtimes_\alpha G$ and $C^\infty(G \backslash X)$.

More precisely, this principle means that we should only study those invariants of noncommutative algebras that do not distinguish between $C^\infty(X) \rtimes_\alpha G$ and $C^\infty(G \backslash X)$ if G is finite and acts freely on X . Theorem 10.14 shows that the primitive ideal space and the set of isomorphism classes of irreducible representations pass this test: they yield the same answer for $C^\infty(X) \rtimes_\alpha G$ and $C^\infty(G \backslash X)$ provided G is finite and acts freely on X .

EXAMPLE 11.2. The set of characters does not pass the test: since all irreducible representations of $C^\infty(X) \rtimes_\alpha G$ have dimension $|G|$, the crossed product $C^\infty(X) \rtimes_\alpha G$ has no characters, whereas $C^\infty(G \backslash X)$ has lots of characters.

The following example forces us to strengthen our correspondence principle:

EXAMPLE 11.3. Let G be a finite group and let $\alpha : G \rightarrow \text{Diffeo}(X)$ be a group action on a smooth manifold X . Let $H \subseteq G$ be a normal subgroup that acts freely. Then the quotient group G/H acts on the orbit space $H \backslash X$. Since the actions of G on X and of G/H on $H \backslash X$ describe the same orbifold. Therefore, noncommutative geometry should not distinguish between the algebras $C^\infty(H \backslash X) \rtimes G/H$ and $C^\infty(X) \rtimes G$.

Even this formulation of our correspondence principle is still too weak, of course, because it only applies to crossed products by finite groups. For instance, in the situation of irrational rotations in Example 11.1, the rotations with angles $2\pi\vartheta$ and

$2\pi/\vartheta$ generate “the same” orbit space

$$\mathbb{T}/2\pi\vartheta\mathbb{Z} \cong \mathbb{R}/2\pi(\mathbb{Z} + \vartheta\mathbb{Z}) \xrightarrow[x \mapsto \vartheta^{-1}x]{\cong} \mathbb{R}/2\pi(\vartheta^{-1}\mathbb{Z} + \mathbb{Z}) \cong \mathbb{T}/2\pi\vartheta^{-1}\mathbb{Z}.$$

Therefore, noncommutative geometry should not distinguish between the crossed product algebras $C^\infty(\mathbb{T}) \rtimes_{\vartheta} \mathbb{Z}$ for ϑ and ϑ^{-1} .

The examples above are all examples of the concept called Morita equivalence. The interesting invariants in noncommutative geometry are those that are Morita invariant. Morita equivalence for unital algebras is defined very simply by an equivalence between the categories of modules over two unital algebras. We characterise when this happens using balanced tensor products with bimodules and using corners in algebras defined by projections.

11.2. The concept of Morita equivalence. We now define Morita equivalence. This is the right concept to understand the examples discussed above. Until further notice, we tacitly assume all rings to be unital.

DEFINITION 11.4 (Kiiti Morita [11]). Let R be a (unital) ring. Let \mathfrak{Mod}_R be the category with left R -modules as objects, module homomorphisms as arrows, and the usual composition. Two rings R and S are *Morita equivalent* if \mathfrak{Mod}_R and \mathfrak{Mod}_S are equivalent categories.

We will see later that it makes no difference whether we use left or right unital modules in this definition.

If a functor $\mathfrak{Mod}_R \rightarrow \mathfrak{Mod}_S$ is an equivalence, then it preserves both limits and colimits (see [12, Lemma 3.3.6]). We are going to show that functors between module categories that preserve colimits have a special form. This leads to the well known description of Morita equivalence using bimodules.

DEFINITION 11.5. Let R and S be two rings, let Q be an S, R -bimodule, and M an R -module. The *R -balanced tensor product* $Q \otimes_R M$ is the quotient of $Q \otimes M$ by the subgroup generated by $q \cdot r \otimes m - q \otimes r \cdot m$ for all $q \in Q$, $r \in R$, $m \in M$. We still write $q \otimes m$ for the image of $q \otimes m \in Q \otimes M$ in $Q \otimes_R M$. The group $Q \otimes_R M$ carries a unique S -module structure with $s \cdot (q \otimes m) := (s \cdot q) \otimes m$ all $s \in S$, $q \in Q$, $m \in M$. If M is an R, T -module for a third ring T , then $Q \otimes_R M$ carries a unique right T -module structure with $(q \otimes m) \cdot t := q \otimes (m \cdot t)$ for all $q \in Q$, $m \in M$, $t \in T$. This makes $Q \otimes_R M$ an S, T -bimodule.

LEMMA 11.6. *The multiplication map $r \otimes m \mapsto r \cdot m$ defines an isomorphism $R \otimes_R M \cong M$ for any left R -module M . Similarly, $n \otimes r \mapsto n \cdot r$ defines an isomorphism $N \otimes_R R \cong N$ for any right R -module N .*

PROOF. The multiplication map $r \otimes m \mapsto r \cdot m$ annihilates $q \cdot r \otimes m - q \otimes r \cdot m$ for all $q, r \in R$, $m \in M$. So it defines a map $R \otimes_R M \rightarrow M$. There is also a map $M \rightarrow R \otimes_R M$, $m \mapsto 1 \otimes m$. The composite map $M \rightarrow R \otimes_R M \rightarrow M$ maps $m \mapsto 1 \otimes m \mapsto 1 \cdot m = m$. The composite map $R \otimes_R M \rightarrow M \rightarrow R \otimes_R M$ maps $r \otimes m \mapsto r \cdot m \mapsto 1 \otimes r \cdot m \sim r \otimes m$. Thus the two maps are inverse to each other. The second isomorphism is proven similarly. \square

Let R and S be rings and Q an S, R -bimodule. Then $Q \otimes_{R, \square}$ defines a functor $Q \otimes_{R, \square}: \mathfrak{Mod}_R \rightarrow \mathfrak{Mod}_S$. We shall see that a functor preserves colimits if and only if it is of this form for an essentially unique bimodule Q . Even more, any such functor has a right adjoint. We describe this adjoint first. It will be useful to characterise which bimodules can occur in a Morita equivalence.

Let M be an S -module. Then there is a left action of R on $\text{Hom}_S(Q, M)$ defined by $(r \cdot f)(q) := f(q \cdot r)$ for all $r \in R$, $f \in \text{Hom}_S(Q, M)$, $q \in Q$.

LEMMA 11.7. *Let R and S be rings and Q an S, R -bimodule. Then there are natural isomorphisms*

$$\mathrm{Hom}_S(Q \otimes_R M, N) \cong \mathrm{Hom}_R(M, \mathrm{Hom}_S(Q, N))$$

for all R -modules M and S -modules N , which are natural in M and N .

PROOF. The isomorphism sends a map $f: Q \otimes_R M \rightarrow N$ to the map $M \rightarrow \mathrm{Hom}_S(Q, N)$ that maps $m \in M$ to the map $q \mapsto f(q, m)$. Some computations show that this is a well defined isomorphism and natural in M and N (see [9, Theorem 3.1]). \square

The following theorem will be used to describe a Morita equivalence through bimodules. Its formulation and proof use quite a bit of category theory. If you are unfamiliar with this, then you may skip this proof.

THEOREM 11.8. *Let R and S be rings and let $T: \mathfrak{Mod}_R \rightarrow \mathfrak{Mod}_S$ be a functor. The following are equivalent:*

- (1) *there is an S, R -bimodule Q such that T is naturally isomorphic to the functor $Q \otimes_R \square$;*
- (2) *T has a right adjoint functor;*
- (3) *T preserves colimits;*
- (4) *T preserves direct sums and is right-exact.*

Let T_1, T_2 be functors that satisfy this and let Q_1 and Q_2 be S, R -bimodules for them as in (1). There is a natural bijection between bimodule maps $Q_1 \rightarrow Q_2$ and natural transformations $T_1 \Rightarrow T_2$. It maps $f: Q_1 \rightarrow Q_2$ to the natural transformation consisting of the maps $f \otimes_R M: Q_1 \otimes_R M \rightarrow Q_2 \otimes_R M$.

PROOF. Lemma 11.7 shows that (1) implies (2). This implies (3) by [12, Theorem 4.5.3]. Coproducts and cokernels in module categories are special cases of colimits; the cokernel of $f: M \rightarrow N$ is the colimit of the coequaliser diagram formed by the pair of maps $f, 0: M \rightrightarrows N$. So a functor that preserves colimits preserves coproducts and cokernels. Coproducts in \mathfrak{Mod}_R are the same as direct sums. A functor $\mathfrak{Mod}_R \rightarrow \mathfrak{Mod}_S$ is additive if and only if it preserves finite direct sums (see [2, Proposition 1.3.4]). By definition, an additive functor is right-exact if and only if it preserves cokernels (the analogous statement for left-exact functors is [12, Proposition 4.5.10]). So (3) implies (4). The main point of the proof is that (4) implies (1).

Recall that an R -module is called free if it is a direct sum of copies of R with the obvious module structure. Any R -module M is a quotient of a free module. A natural choice is to take the direct sum $F := \bigoplus_{m \in M} R$ and map $F \rightarrow M$ by mapping $(x_m)_{m \in M}$ to $\sum_{m \in M} x_m \cdot m$. This is indeed an R -module homomorphism. Let M' be the kernel of this quotient map. This is another R -module, so it is also a quotient of a free module. This gives a map $d: F' \rightarrow F$ between two free modules whose cokernel is isomorphic to M . Any right-exact functor T satisfies

$$T(M) \cong T(\mathrm{coker} d) \cong \mathrm{coker} T(d).$$

Using this natural isomorphism, the whole functor T – including its action on arrows – is determined by its restriction to free modules. Even more, if $T_1, T_2: \mathfrak{Mod}_R \rightrightarrows \mathfrak{Mod}_S$ are two right-exact functors, then a natural transformation between their restrictions to the full subcategories of free R -modules extends uniquely to a natural transformation $T_1 \Rightarrow T_2$.

If the functor T also preserves direct sums, then there are natural isomorphisms $T(\bigoplus_{i \in I} R) \cong \bigoplus_{i \in I} T(R)$ for all sets I . Then the restriction of T to the single module R determines the restriction of T to free modules – including the action on arrows – and hence the functor T itself. Even more, if two functors T_1, T_2 are

right-exact and preserve direct sums, then any natural transformation between their restrictions to the full subcategory with only R as an object extends uniquely to a natural transformation $T_1 \Rightarrow T_2$.

Let $Q := T(R)$. This is some left S -module. Right multiplication with $r \in R$ is a left module homomorphism $R \rightarrow R$, $x \mapsto x \cdot r$. Since T is a functor, this induces a left S -module homomorphism on Q , which we denote multiplicatively. The distributive law $x \cdot (r_1 + r_2) = x \cdot r_1 + x \cdot r_2$ holds because T is additive (since it preserves finite direct sums). The functoriality of T implies $(x \cdot r_1) \cdot r_2 = x \cdot (r_1 \cdot r_2)$. Thus Q becomes an S, R -bimodule.

The multiplication map $Q \otimes_R R \rightarrow Q$ is an isomorphism between the restrictions of the functors $Q \otimes_{R \sqcup}$ and T to the full subcategory with only R as an object. It is natural because it is a right R -module homomorphism and $\text{Hom}_R(R, R) \cong R$. We have already shown that (1) implies (4). So both functors $Q \otimes_{R \sqcup}$ and T preserve direct sums and are right-exact. Hence the natural isomorphism on the single module R extends uniquely to a natural isomorphism between $Q \otimes_{R \sqcup}$ and T on all R -modules.

The proof also gives a natural bijection between natural transformations $T_1 \Rightarrow T_2$ and S, R -bimodule homomorphisms $f: Q_1 \rightarrow Q_2$ for the bimodules $Q_j := T_j(R)$ for $j = 1, 2$. The maps $f \otimes_R \text{Id}_M$ for all R -modules M form a natural transformation $Q_1 \otimes_{R \sqcup} \Rightarrow Q_2 \otimes_{R \sqcup}$ that restricts to f for $M = R$. Hence this is the unique natural transformation extending f . \square

Theorem 11.8 was discovered simultaneously by Eilenberg, Gabriel, and Watts around 1960 (see [15]).

EXAMPLE 11.9. A ring homomorphism $f: S \rightarrow R$ induces a functor

$$f^*: \mathfrak{Mod}_R \rightarrow \mathfrak{Mod}_S,$$

which maps an R -module M to the same Abelian group with the S -module structure $s \cdot m := f(s) \cdot m$. The functor f^* is exact and preserves direct sums and products because it does not change the underlying Abelian group. By Theorem 11.8, it must be of the form $Q \otimes_{R \sqcup}$ for some S, R -bimodule Q . The proof of the theorem shows that Q is R as a right R -module, with the left S -module structure $s \cdot r := f(s)r$ for all $s \in S, r \in R$.

THEOREM 11.10. *Two rings R and S are Morita equivalent if and only if there are an S, R -bimodule Q and an R, S -bimodule P with bimodule isomorphisms $Q \otimes_R P \cong S$ and $P \otimes_S Q \cong R$.*

PROOF. By Theorem 11.8, an equivalence of categories $\mathfrak{Mod}_R \xrightarrow{\cong} \mathfrak{Mod}_S$ and its inverse are of the form $Q \otimes_{R \sqcup}$ and $P \otimes_{S \sqcup}$ for an S, R -bimodule Q and an R, S -bimodule P . Since these functors are inverse to each other up to natural isomorphisms, the functors $Q \otimes_R (P \otimes_{S \sqcup})$ and $P \otimes_S (Q \otimes_{R \sqcup})$ are naturally isomorphic to the identity functors. There are obvious natural isomorphisms $Q \otimes_R (P \otimes_{S \sqcup}) \cong (Q \otimes_R P) \otimes_{S \sqcup}$ and $P \otimes_S (Q \otimes_{R \sqcup}) \cong (P \otimes_S Q) \otimes_{R \sqcup}$. So these composite functors come from the bimodules $Q \otimes_R P$ and $P \otimes_S Q$, respectively. The identity functors come from the bimodules R and S , respectively, by Example 11.9. By Theorem 11.8, the natural isomorphisms between our tensor product functors are equivalent to bimodule isomorphisms $Q \otimes_R P \cong S$ and $P \otimes_S Q \cong R$. \square

REMARK 11.11. Theorem 11.10 works in the same way for left instead of right modules. This shows that it makes no difference whether we use left or right modules in Definition 11.4.

EXAMPLE 11.12. Let R be a unital ring. Then R is Morita equivalent to $M_n(R)$. We build appropriate bimodules. Let $P = R^n$ and $Q = R^n$. Make P an $R, M_n(R)$ -bimodule and Q an $M_n(R), R$ -bimodule by matrix-vector multiplication, treating elements of P as row vectors and elements of Q as column vectors. Then matrix multiplication between $1 \times n$ - and $n \times 1$ -matrices gives the required bimodule isomorphisms $Q \otimes_R P \cong S$ and $P \otimes_S Q \cong R$.

EXERCISE 11.13. Consider the three categories of \mathbb{R} -vector spaces, \mathbb{C} -vector spaces, and \mathbb{H} -vector spaces, where \mathbb{H} denotes the quaternions. Show that no two among these three categories are equivalent. That is, the \mathbb{R} -algebras \mathbb{C} , \mathbb{R} , and \mathbb{H} are not Morita equivalent.

Lemma 10.12 and Exercise 11.12 imply that the crossed product $C^\infty(G/H) \rtimes_\alpha G$ is Morita equivalent to the group algebra $\mathbb{C}[H]$ for any subgroup H in a finite group G .

Now we turn to proving Morita equivalence for some examples of rings. We usually prove this by constructing appropriate bimodules. These are easier to handle than the corresponding functors between module categories. The bimodules that we need below have a rather simple form. They are constructed from a single idempotent element.

DEFINITION 11.14. An idempotent element $p \in A$ in a ring A is called *full* if it generates A as an ideal, that is, elements of the form apb with $a, b \in A$ span A .

THEOREM 11.15. Let $p \in A$ be a full idempotent in a unital ring A . Then A is Morita equivalent to pAp .

The subring pAp in Theorem 11.15 for a (full) idempotent element is also called a (full) *corner*. It is unital with unit p .

PROOF. The right ideal pA generated by p is a pAp, A -bimodule, the left ideal Ap generated by p is an A, pAp -bimodule. The multiplication in A gives natural maps

$$\begin{aligned} pA \otimes_A Ap &\rightarrow pAp, & [pa \otimes bp] &\mapsto pa \cdot bp, \\ Ap \otimes_{pAp} pA &\rightarrow A, & [ap \otimes pb] &\mapsto ap \cdot pb = apb. \end{aligned}$$

These maps are bimodule maps. We claim that they are both invertible. Then Theorem 11.10 shows that A and pAp are Morita equivalent.

The first map is clearly surjective. We prove that it is injective. First, Lemma 11.6 shows that $A \otimes_A Ap \cong Ap$. Since $A \cong pA \oplus (1-p)A$ and the balanced tensor product is additive for direct sums of modules, we get

$$(pA \otimes_A Ap) \oplus ((1-p)A \otimes_A Ap) \cong A \otimes_A Ap.$$

Hence the multiplication map $pA \otimes_A Ap \rightarrow A \otimes_A Ap \cong Ap$ is injective. Its image is pAp . So $pA \otimes_A Ap \cong pAp$ as pAp -bimodules.

Since p generates A as an ideal, we may write $1 = \sum_{i=1}^n a_i p b_i$ with $a_i, b_i \in A$. We define maps

$$\begin{aligned} v: A &\rightarrow (Ap)^n, & x &\mapsto (x a_i p)_{i=1, \dots, n}, \\ w: (Ap)^n &\rightarrow A, & (y_i p)_{i=1, \dots, n} &\mapsto \sum_{i=1}^n y_i p \cdot b_i. \end{aligned}$$

These are left A -module homomorphisms with $w \cdot v = \text{Id}_A$ by assumption. Hence A is a direct summand of $(Ap)^n$ as a left A -module. Let q be the projection from $(Ap)^n$ onto A . As above, the additivity of balanced tensor products and Lemma 11.6 yield

$$Ap \otimes_{pAp} pA \cong q(Ap \otimes_{pAp} (pAp)^n) \cong q(Ap)^n \cong A.$$

Thus A is Morita equivalent to pAp . \square

THEOREM 11.16. *Let M be a smooth compact manifold and let α be a free group action of a finite group G on M . Then $C^\infty(M) \rtimes_\alpha G$ and $C^\infty(G \backslash M)$ are Morita equivalent.*

PROOF. Let $A := C^\infty(M) \rtimes_\alpha G$. Define $p \in \mathbb{C}[G] \subseteq A$ by

$$p := \frac{1}{|G|} \sum_{g \in G} \delta_g.$$

A computation shows that $p^2 = p$. Moreover, if V is a representation of G , equipped with the canonical $\mathbb{C}[G]$ -module structure from the group action, then $p(V) \subseteq V$ is the subspace of G -invariant elements. We compute

$$\begin{aligned} Ap &= \left\{ \sum_{g \in G} f \delta_g : f \in C^\infty(M) \right\} \cong C^\infty(M), \\ pA &= \left\{ \sum_{g \in G} \alpha_g(f) \delta_g : f \in C^\infty(M) \right\} \cong C^\infty(M), \\ pAp &= \left\{ \sum_{g \in G} f \delta_g : f \in C^\infty(M) \text{ with } \alpha_g(f) = f \text{ for all } g \in G \right\}. \end{aligned}$$

A function on M is G -invariant if and only if it factors through the projection $\pi: M \rightarrow G \backslash M$. Since the latter is a covering map, a function f on $G \backslash M$ is smooth if and only if $f \circ \pi$ is a smooth function on M . Hence $pAp \cong C^\infty(G \backslash M)$.

In order to apply Theorem 11.15, it remains to check that the idempotent p is full: this yields the desired Morita equivalence between pAp and A . In the proof of Theorem 10.14, we have shown that any proper ideal in $A = C^\infty(M) \rtimes_\alpha G$ is contained in the kernel of an irreducible representation of A . In our case, this simply means that any proper ideal is contained in the kernel of the restriction map $A \rightarrow C^\infty(Gx) \rtimes_\alpha G \cong \mathbb{M}_{|G|} \mathbb{C}$ for some $x \in M$ because $\mathbb{M}_{|G|} \mathbb{C}$ is simple; here we have used Lemma 10.12.

As a consequence, the ideal generated by p cannot be proper because the image of p in $C^\infty(Gx) \rtimes_\alpha G$ is non-zero for each $x \in M$. Hence p must be a full idempotent, and we are done. \square

EXERCISE 11.17. The proofs above show that there must be $a_i, b_i \in C^\infty(M)$ with $\sum_{i=1}^n a_i p b_i = 1$ in $C^\infty(M) \rtimes_\alpha G$. Construct such elements $a_i, b_i \in C^\infty(M)$ using a partition of unity on M .

EXERCISE 11.18. Let M be a smooth compact manifold, G a finite group, and $\alpha: G \rightarrow \text{Diffeo}(M)$ a group action by diffeomorphisms on M . Let H be a normal subgroup of G that acts freely on M . Show that $A := C^\infty(M) \rtimes_\alpha G$ is Morita equivalent to $C^\infty(H \backslash M) \rtimes_\alpha G/H$ by constructing a full idempotent in A with $pAp \cong C^\infty(H \backslash M) \rtimes_\alpha G/H$. What happens if H is not normal in G ?

We are going to describe Morita equivalence using *linking rings*. This shows that the situation of Theorem 11.15 is not as special as it seems.

DEFINITION 11.19. Let A and B be two unital rings. Let P and Q be an A, B -bimodule and a B, A -bimodule. Let $\mu_{PQ}: P \otimes_B Q \rightarrow A$ and $\mu_{QP}: Q \otimes_A P \rightarrow B$ be bimodule homomorphisms. Assume also that the maps $\mu_{PQ} \otimes_A \text{Id}_P$ and $\text{Id}_P \otimes_B \mu_{QP}$ from $P \otimes_B Q \otimes_A P$ to P and $\mu_{QP} \otimes_B \text{Id}_Q$ and $\text{Id}_Q \otimes_A \mu_{PQ}$ from $Q \otimes_A P \otimes_B Q$ to Q are equal

The *linking ring* associated to $(A, B, P, Q, \mu_{PQ}, \mu_{QP})$ is the unital ring with underlying vector space $L := A \oplus P \oplus Q \oplus B$ and with the multiplication

$$\begin{pmatrix} a_1 & p_1 \\ q_1 & b_1 \end{pmatrix} \cdot \begin{pmatrix} a_2 & p_2 \\ q_2 & b_2 \end{pmatrix} := \begin{pmatrix} a_1 \cdot a_2 + \mu_{PQ}[p_1 \otimes q_2] & a_1 \cdot p_2 + p_1 \cdot b_2 \\ q_1 \cdot a_2 + b_1 \cdot q_2 & \mu_{QP}[q_1 \otimes b_2] + b_1 \cdot b_2 \end{pmatrix}$$

for all $a_1, a_2 \in A$, $p_1, p_2 \in P$, $q_1, q_2 \in Q$, $b_1, b_2 \in B$. Notice the similarity to matrix multiplication.

PROPOSITION 11.20. *The vector space L with this multiplication is an associative unital ring. The elements*

$$p_A := \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad p_B := \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

in L are idempotent elements with $p_A L p_A \cong A$ and $p_B L p_B \cong B$. The idempotent elements p_A and p_B are full if μ_{PQ} and μ_{QP} are bimodule isomorphisms.

Let A and B be Morita equivalent. Then there are bimodules P and Q and bimodule homomorphisms μ_{PQ} and μ_{QP} as in Definition 11.19; that is, the additional assumptions $\mu_{QP} \otimes_B \text{Id}_Q = \text{Id}_Q \otimes_A \mu_{PQ}$ and $\mu_{PQ} \otimes_A \text{Id}_P = \text{Id}_P \otimes_B \mu_{QP}$ can be satisfied. As a result, two rings are Morita equivalent if and only if they are both isomorphic to full corners in the same ring.

PROOF. The associativity of the multiplication in L is easy to check. Here the additional assumptions $\mu_{QP} \otimes_B \text{Id}_Q = \text{Id}_Q \otimes_A \mu_{PQ}$ and $\mu_{PQ} \otimes_A \text{Id}_P = \text{Id}_P \otimes_B \mu_{QP}$ are used. It is trivial to check that p_A and p_B are idempotent and that the corners $p_A L p_A$ and $p_B L p_B$ are canonically isomorphic to A and B , respectively. The ideal generated by p_A contains $p_A L p_A = A$, $(1 - p_A) L p_A = Q$, $p_A L (1 - p_A) = P$, and $(1 - p_A) L p_A L (1 - p_A) = Q \cdot P$. Since μ_{QP} is surjective $Q \cdot P$ – which denotes the linear span of products $q \cdot p$ – is equal to B . Thus p_A is full. The same argument works for p_B .

Now assume only that A and B are Morita equivalent. Fix the bimodules P and Q and the isomorphism μ_{QP} from a Morita equivalence. Then

$$\text{Id}_P \otimes_B \mu_{QP} \otimes_B \text{Id}_Q: P \otimes_B Q \otimes_A P \otimes_B Q \rightarrow P \otimes_B B \otimes_B Q$$

is an isomorphism as well. The isomorphisms $P \otimes_B B \cong P$ and $B \otimes_B Q \cong Q$ from Lemma 11.6 both yield the same isomorphism $P \otimes_B B \otimes_B Q \cong P \otimes_B Q$. Since $P \otimes_B Q \cong A$, the maps

$$(11.21) \quad \begin{aligned} \text{Hom}_{A,A}(P \otimes_B Q, A) &\rightarrow \text{Hom}_{A,A}(P \otimes_B Q \otimes_A P \otimes_B Q, P \otimes_B Q), \\ f &\mapsto f \otimes_A \text{Id}_P \otimes_B \text{Id}_Q, \text{Id}_P \otimes_B \text{Id}_Q \otimes_A f, \end{aligned}$$

are isomorphisms. Even more, identifying $\text{Hom}_{A,A}(P \otimes_B Q, A) \cong \text{Hom}_{A,A}(A, A)$ with the centre of A , we may see that these two maps are equal because left or right multiplication with a central element are the same map. So there is a unique A -bimodule map $\mu_{PQ}: P \otimes_B Q \rightarrow A$ with $\mu_{PQ} \otimes_A \text{Id}_P \otimes_B \text{Id}_Q = \text{Id}_P \otimes_B \mu_{QP} \otimes_B \text{Id}_Q$. Since tensoring with Id_Q is invertible as well, this is equivalent to $\mu_{PQ} \otimes_A \text{Id}_P = \text{Id}_P \otimes_B \mu_{QP}$. Since the two maps in (11.21) are equal, it follows also that $\text{Id}_Q \otimes_A \mu_{PQ} = \mu_{QP} \otimes_B \text{Id}_Q$. \square

PROPOSITION 11.22. *Let A and B be Morita equivalent unital rings. Then there are canonical bijections $\text{Prim}(A) \cong \text{Prim}(B)$ and $\hat{A} \cong \hat{B}$ and an isomorphism between the lattices $\mathbb{I}(A)$ and $\mathbb{I}(B)$ of ideals in A and B , respectively, which are part of a commuting diagram*

$$\begin{array}{ccccc} \hat{A} & \xrightarrow{\text{ann}} & \text{Prim}(A) & \hookrightarrow & \mathbb{I}(A) \\ \downarrow \cong & & \downarrow \cong & & \downarrow \cong \\ \hat{B} & \xrightarrow{\text{ann}} & \text{Prim}(B) & \hookrightarrow & \mathbb{I}(B). \end{array}$$

Here ann maps a simple module to its annihilator ideal.

PROOF. An equivalence between the categories of left unital modules over A and B must map simple modules to simple modules. And it preserves and respects isomorphism. Hence it induces a bijection $\hat{A} \cong \hat{B}$. The isomorphism of ideal lattices is proved most easily using linking rings. By Proposition 11.20, we may assume without loss of generality that $A = pBp$ for a projection $p \in B$ with $BpB = B$. Then the bijection $\hat{B} \rightarrow \hat{A}$ maps a simple B -module M to pM with the induced A -module structure. An ideal $I \subseteq B$ induces an ideal $I \cap A = pIp \subseteq A$. Conversely, an ideal $J \subseteq A$ generates an ideal $K := BJB = BpJpB$ in B . By construction, $K \cap A = pKp = AJA = J$ and $BpIpB \subseteq I$. The converse inclusion also holds:

$$BpIpB = Bp(BIB)pB = BIB = I$$

because $BpB = B$. Let M be a B -module and let $\text{ann}(M) = \{b \in B : bM = 0\}$ be its annihilator. Then

$$\text{ann}(M) \cap A = \{a \in A : aM = 0\} = \{a \in A : apM = 0\}$$

is the annihilator ideal of the A -module pM . Thus the maps from B -modules to A -modules and from ideals in B to ideals in A are compatible with taking annihilator ideals of modules. Therefore, an ideal I in B is the annihilator ideal of a simple B -module if and only if $I \cap A$ is the annihilator ideal of a simple A -module. In other words, the isomorphism of ideal lattices restricts to a bijection between the subsets of primitive ideals. \square

REMARK 11.23. It is more difficult to define Morita equivalence for non-unital rings. The problem is that we expect the ring $\mathbb{M}_\infty A$ of finite matrices over a ring A to be Morita equivalent to A . But the module categories of $\mathbb{M}_\infty \mathbb{C}$ and \mathbb{C} are not equivalent: only their categories of non-degenerate modules are equivalent. The category of non-degenerate modules may be quite badly behaved for general A ; for instance, the left and right regular representations on A may be degenerate. There is a well behaved Morita theory for rings with local units (compare Remark 7.8).

More generally, it suffices if the rings have the property that the multiplication map induces an isomorphism $A \otimes_A A \xrightarrow{\cong} A$; then the non-degeneracy condition for modules must be replaced by the condition that the multiplication map induces an isomorphism $A \otimes_A M \xrightarrow{\cong} M$. These rings are called *self-induced* in [6, 10], and the modules with the above property are called smooth. For a ring with idempotent local units, a left module M is smooth if and only if the multiplication map $A \otimes M \rightarrow M$ is surjective, if and only if for any element $m \in M$ there is an idempotent element $e \in A$ with $m = e \cdot m$.

12. Derivations as a noncommutative analogue of differentiation

We have already collected a number of examples of noncommutative algebras. Our tools to study them are, however, quite limited. We have only introduced the primitive and maximal ideal spaces and the space of irreducible representations. There is also the category of modules, but this is rather an equivalent way to look at the noncommutative algebra than a tool to study it: we have argued that two noncommutative algebras should be considered equivalent if they have equivalent categories of modules. In the commutative case, primitive ideals and irreducible representations both generalise the points of the underlying space. This is a very basic invariant. To determine a manifold, we should also describe the smooth structure. This requires differentiation.

In a manifold, differentiation is formulated through the tangent space, which is a vector bundle. Its points are directional derivatives. And its sections are vector fields on the manifold. The concept of a derivation generalises both directional derivatives and vector fields to noncommutative algebras. Here a derivation from

an algebra to a bimodule over it is a linear map that satisfies the Leibniz rule. We show that derivations for the algebra $C^\infty(M)$ on a smooth manifold and suitable target bimodules generalise directional derivatives and vector fields, respectively.

Derivations are an important concept and will be discussed in several sections. In this section, we introduce the concept of derivation and we describe derivations for algebras of smooth functions and matrix algebras. As we shall see, all derivations from a matrix algebra into a bimodule over it are inner. That is, they are built in a canonical way from elements in the target bimodule. This shows that the space of derivations fails to be Morita invariant because there are far too many derivations for matrix algebras. We will see later that the quotient of derivations modulo inner derivations is Morita invariant (see Theorem 15.10). Our proof that derivations on matrix algebras are inner foreshadows this proof.

Our description of derivations on matrix algebras links them to extensions of modules. We use this to classify all two-dimensional representations of the algebra of smooth functions on a compact manifold. Any such representation is an extension of two one-dimensional representations, and this provides a link to our computation of derivations.

12.1. Derivations on algebras of smooth functions. Let M be a smooth d -dimensional manifold, let $m \in M$, and let $f: M \rightarrow \mathbb{R}$ be a smooth function. Let $\varphi: U \rightarrow \mathbb{R}^d$ be a coordinate chart in a neighbourhood U of m . The smooth function $f \circ \varphi^{-1}: \mathbb{R}^d \rightarrow \mathbb{R}$ has *directional derivatives* $\partial_{\vec{v}}(f \circ \varphi^{-1})$ at $\varphi^{-1}(m)$ for vectors $\vec{v} \in \mathbb{R}^d$. These depend on the chosen coordinate chart. We may, however, remove the coordinate dependence: there is an abstract vector space $T_m M$ such that any coordinate chart induces a linear map $h_\varphi: T_m M \rightarrow \mathbb{R}^d$ with the property that the directional derivative $\partial_{h_\varphi(\vec{v})}$ for $\vec{v} \in T_m M$ does not depend on the chart φ .

There are several constructions of $T_m M$. The following one explains its name: *tangent space*. Fix a proper embedding $i: M \rightarrow \mathbb{R}^N$ with injective derivative – this exists by Theorem 1.22. Let $T_m M \subseteq \mathbb{R}^N$ be the space of all tangent vectors to the submanifold $i(M)$ at $i(m)$. This is the space of all $\vec{v} \in \mathbb{R}^N$ for which the line $t \mapsto i(m) + \vec{v}t$ is tangent to $i(M)$. This subspace $T_m M$ is equal to the range of the derivative of $i \circ \varphi^{-1}: \mathbb{R}^d \rightarrow \mathbb{R}^N$ at $\varphi^{-1}(m)$ for any coordinate chart φ near m . We let $h_\varphi: T_m M \rightarrow \mathbb{R}^d$ be the inverse of this derivative.

The disjoint union $\bigsqcup_{m \in M} T_m M$ of all tangent spaces may be turned into a smooth manifold as well, using the maps $\bigsqcup_{m \in U} T_m M \rightarrow \mathbb{R}^d \times \mathbb{R}^d$ induced by coordinate charts $\varphi: U \rightarrow \mathbb{R}^d$. For different charts, the coordinate change maps between these local tangent spaces are smooth and linear in the fibres. This makes $T_m M$ a *smooth vector bundle* over M .

A smooth *vector field* on M is a smooth section of the tangent bundle $TM \rightarrow M$, that is, a smooth map $X: M \rightarrow TM$ with $X(m) \in T_m M$ for all $m \in M$. The space of smooth vector fields on M is an \mathbb{R} -vector space, even a module over the algebra $C^\infty(M, \mathbb{R})$ by pointwise multiplication.

Our goal is to generalise tangent vectors and smooth vector fields to noncommutative algebras. The starting point is to identify a tangent vector $\vec{v} \in T_m M$ with the differentiation operator $f \mapsto Df_m(\vec{v})$ and a vector field X with the corresponding *first order differential operator* $C^\infty(M) \rightarrow C^\infty(M)$ mapping $f \in C^\infty(M)$ to the function $m \mapsto Df_m(X(m))$. The following theorem characterises the operators that arise in this way:

THEOREM 12.1. *Let $m \in M$. A linear map $l: C^\infty(M, \mathbb{R}) \rightarrow \mathbb{R}$ is of the form $f \mapsto Df_m(\vec{v})$ for some $\vec{v} \in T_m M$ if and only if it satisfies the product rule $l(f \cdot g) = f(m) \cdot l(g) + l(f) \cdot g(m)$ for all $f, g \in C^\infty(M)$.*

A linear map $l: C^\infty(M, \mathbb{R}) \rightarrow C^\infty(M, \mathbb{R})$ comes from a smooth vector field on M if and only if it satisfies $l(f \cdot g) = l(f) \cdot g + f \cdot l(g)$ for all $f, g \in C^\infty(M)$.

PROOF. Differentiation and first order differential operators clearly satisfy the product rule. Conversely, let $l: C^\infty(M) \rightarrow \mathbb{C}$ satisfy the product rule. Then $l(f \cdot g) = 0$ if $f(m) = g(m) = 0$. We claim that a smooth function h on M is a linear combination of such products $f \cdot g$ if (and only if) it has a zero at m of order at least two (the order may be defined in any local coordinate system). To see this, we may first reduce to the case $M = \mathbb{R}^d$ by replacing h by $h - h \cdot \chi$, where χ is a function that vanishes at m and is constant equal to 1 outside a chart neighbourhood around m . For functions on \mathbb{R}^d , we use the following variant of Lemma 2.13. We omit its proof.

LEMMA 12.2. Any smooth function $h: \mathbb{R}^d \rightarrow \mathbb{R}$ with a zero of order at least 2 at the origin may be written as $h = \sum_{ij} x_i x_j h_{ij}$ with smooth functions h_{ij} .

As a result, $l(h) = 0$ if h has a zero at m of order at least two. Therefore, l must be of the form

$$l(f) = \lambda_0 f(0) + \lambda_1 \frac{\partial}{\partial x_1} f(0) + \cdots + \lambda_d \frac{\partial}{\partial x_d} f(0)$$

with $\lambda_0, \dots, \lambda_d \in \mathbb{R}$ (in local coordinates). We have $\lambda_0 = l(1) = l(1 \cdot 1) = l(1) + l(1)$, hence $\lambda_0 = 0$. Therefore, $l(f)$ is a directional derivative as asserted.

Now consider a linear map $l: C^\infty(M) \rightarrow C^\infty(M)$ that satisfies the product rule. Then $\text{ev}_m \circ l$ satisfies the product rule at m and, by the first paragraph, is the derivative for some tangent vector $X(m) \in T_m M$. The map l is the first order differential operator associated to the vector field $m \mapsto X(m)$; this vector field is smooth because $l(f)$ is smooth for all smooth functions f . \square

DEFINITION 12.3. Let A be an algebra and let M be an A -bimodule. A *derivation* of A with values in M is a linear map $d: A \rightarrow M$ that satisfies the *Leibniz rule* $d(a \cdot b) = a \cdot d(b) + d(a) \cdot b$ for all $a, b \in A$. We write $\text{Der}(A, M)$ for the space of derivations $A \rightarrow M$.

We need M to be an A -bimodule in order for the Leibniz rule to make sense.

EXAMPLE 12.4. Let $A = C^\infty(M, \mathbb{R})$ for a smooth manifold M . Theorem 12.1 shows that $\text{Der}(A, A)$ is isomorphic to the space of smooth vector fields on M . And if $x \in M$, then $A/\ker \text{ev}_x$ is an A -bimodule and $\text{Der}(A, A/\ker \text{ev}_x)$ is isomorphic to the tangent space $T_x M$. Thus derivations generalise both smooth vector fields and tangent vectors on manifolds.

In this section, we work with \mathbb{R} -valued functions on smooth manifolds in order to get ordinary tangent vectors and vector fields. On $C^\infty(M, \mathbb{C})$, we would get the vectors in the complexified tangent bundle and the complex vector space generated by smooth vector fields. This is less geometric, and some things like integrating a vector field to a flow as in Section 13.1 no longer work.

EXERCISE 12.5. Let A be an algebra and M an A -bimodule. Let $d: A \rightarrow M$ be a derivation and $a, c \in A$. If a and c are central, then $a \cdot d \cdot c: b \mapsto a \cdot d(b) \cdot c$ is a derivation as well. Thus $\text{Der}(A, M)$ is a bimodule over the *centre* of A . In general, it is not a bimodule over A .

12.2. Derivations on matrix algebras. We now describe derivations into bimodules for matrix algebras $\mathbb{M}_n K$. As it turns out, they are all of the special form described in the next definition:

DEFINITION 12.6. Let A be an algebra, let M be an A -bimodule, and let $m \in M$. The associated *inner derivation* ad_m is defined by $\text{ad}_m(a) = [m, a] := m \cdot a - a \cdot m$ for all $a \in A$. Let $\text{Inn}(A, M) \subseteq \text{Der}(A, M)$ be the subspace of inner derivations.

The *centre* of M is the set of all $m \in M$ with $\text{ad}_m = 0$, that is, $a \cdot m = m \cdot a$ for all $a \in A$.

Inner derivations are indeed derivations because

$$\text{ad}_m(a)b + a \text{ad}_m(b) = mab - amb + amb - abm = \text{ad}_m(ab).$$

Hence there is a linear map

$$\text{ad}: M \rightarrow \text{Der}(A, M).$$

Its kernel is equal to the centre $Z(M)$ of M .

As a consequence, $\text{Der}(\mathbb{M}_n\mathbb{C}, \mathbb{M}_n\mathbb{C})$ is much bigger than $\text{Der}(\mathbb{C}, \mathbb{C}) = \{0\}$. Thus Morita equivalent unital algebras may have quite different spaces of derivations. We will see later that the quotient $\text{Der}(A, A) / \text{Inn}(A, A)$ is Morita invariant (Theorem 15.10). Since $\text{Der}(\mathbb{C}, \mathbb{C}) = \{0\}$, it follows that any derivation of $\mathbb{M}_n\mathbb{C}$ is inner. This may be checked directly, of course. Our proof relates derivations to certain module extensions. This allows to use the Morita equivalence between \mathbb{C} and $\mathbb{M}_n\mathbb{C}$.

Let $V \xrightarrow{i} X \xrightarrow{p} W$ be an extension of left A -modules, that is, the map $i: V \rightarrow X$ is injective and p induces an isomorphism $X/i(V) \cong W$. Choosing a basis for W , we can construct a linear map $s: W \rightarrow X$ with $p \circ s = \text{Id}_W$, which need not be a module homomorphism. We use s to split $X \cong V \oplus W$ as a vector space. For $a \in A$, the operator μ_a^X of left multiplication by a has the form

$$\mu_a^X = \begin{pmatrix} \mu_a^V & d_a \\ 0 & \mu_a^W \end{pmatrix}$$

where μ_a^V and μ_a^W denote the operators of left multiplication by a on V and W and $d_a: W \rightarrow V$ is some linear map. The condition $\mu_a^X \circ \mu_b^X = \mu_{ab}^X$ for a module structure is equivalent to $d_a \mu_b^W + \mu_a^V d_b = d_{ab}$. This means that $a \mapsto d_a$ is a derivation $A \rightarrow \text{Hom}(W, V)$ with respect to the canonical bimodule structure on $\text{Hom}(W, V)$ defined by $a \cdot x \cdot b := \mu_a^V \circ x \circ \mu_b^W$ for all $a, b \in A, x: W \rightarrow V$.

So far we have seen that derivations $A \rightarrow \text{Hom}(W, V)$ correspond bijectively to A -module structures μ_X on $V \oplus W$ for which the maps $V \rightarrow V \oplus W \rightarrow W$ are A -linear. This bijection depends on the choice of the section s . Another section is of the form $s + \delta$ for some linear map $\delta: W \rightarrow V$. The resulting derivation then takes the form $d + \text{ad}_\delta$. Since any linear map δ may appear, we get the following theorem:

THEOREM 12.7. *Let A be an algebra, let V and W be left A -modules, and equip $\text{Hom}(W, V)$ with the canonical A -bimodule structure. Let $\text{Ext}_A(W, V)$ be the set of equivalence classes of A -module extensions $V \rightarrow X \rightarrow W$, where two such extensions are considered equivalent if there is a commuting diagram*

$$\begin{array}{ccccc} V & \rightarrow & X_1 & \rightarrow & W \\ & & \parallel & \downarrow \cong & \parallel \\ V & \rightarrow & X_2 & \rightarrow & W. \end{array}$$

Then there is a natural bijection

$$\text{Ext}_A(W, V) \cong \frac{\text{Der}(A, \text{Hom}(W, V))}{\text{Inn}(A, \text{Hom}(W, V))}.$$

An extension splits by an A -module homomorphism if and only if the corresponding derivation is inner.

The map from extensions to derivations may also be understood as follows. A section $s: W \rightarrow X$ of an A -module extension defines an element $s \in \text{Hom}(W, X)$, which yields an inner derivation $\text{ad}_s: A \rightarrow \text{Hom}(W, X)$. Since $p \circ s = \text{Id}_W$, we have $p \circ \text{ad}_s = \text{ad}_{\text{Id}_W} = 0$. That is, ad_s maps A into $\text{Hom}(W, V)$. This is exactly the derivation described above.

In the example above, it turns out that any derivation $A \rightarrow \text{Hom}(W, V)$ becomes inner when we enlarge $\text{Hom}(W, V)$ to $\text{Hom}(W, X)$. This is a general feature:

EXERCISE 12.8. Let A be a unital algebra, let M be a unital A -bimodule, and let $d: A \rightarrow M$ be a derivation. Construct an A -bimodule extension $M \hookrightarrow X \twoheadrightarrow Y$ and $x \in X$ such that the inner derivation ad_x is a map to M and agrees with d .

THEOREM 12.9. *Let $A = \mathbb{M}_n\mathbb{C}$. Then any derivation in $\text{Der}(A, A)$ is inner. More generally, any derivation $A \rightarrow M$ for a unital A -bimodule M is inner.*

PROOF. Since $\mathbb{M}_n\mathbb{C} \cong \text{Hom}(\mathbb{C}^n, \mathbb{C}^n)$ with the canonical bimodule structure, the assertion is equivalent to the statement that any A -module extension $\mathbb{C}^n \hookrightarrow X \twoheadrightarrow \mathbb{C}^n$ splits by an A -module homomorphism. This follows from the Morita equivalence between $\mathbb{M}_n\mathbb{C}$ and \mathbb{C} because any extension of vector spaces splits by a linear map, which is the same as a \mathbb{C} -module homomorphism. The assertion extends to derivations into arbitrary bimodules because these are direct sums of copies of A . \square

EXERCISE 12.10. Let A be a finite-dimensional, semi-simple algebra and let M be an A -bimodule, not necessarily unital. Show that any derivation $A \rightarrow M$ is inner.

12.3. Derivations and representations of smooth functions. Let M be a smooth compact manifold. Let $\varrho: C^\infty(M) \rightarrow \mathbb{M}_2\mathbb{C}$ be a two-dimensional representation. The irreducible representations of $C^\infty(M)$ are precisely the characters. So ϱ is reducible. Choose a basis in \mathbb{C}^2 so that the first basis vector spans an invariant subspace. This defines a 1-dimensional $C^\infty(M)$ -submodule $V \subseteq (\mathbb{C}^2, \varrho)$. The quotient $W := (\mathbb{C}^2, \varrho)/V$ is another 1-dimensional $C^\infty(M)$ -module, and we have got a module extension

$$V \hookrightarrow (\mathbb{C}^2, \varrho) \twoheadrightarrow W.$$

Since V and W are 1-dimensional, they are given by characters. That is, V and W are \mathbb{C} with $C^\infty(M)$ acting by ev_{x_0} and ev_{x_1} for some $x_0, x_1 \in M$, respectively. Theorem 12.9 shows that the representation is determined up to equivalence by these two points x_0, x_1 and a derivation $C^\infty(M) \rightarrow \text{Hom}(W, V)$. Here $\text{Hom}(W, V)$ is isomorphic to \mathbb{C} with the bimodule structure by $f_0 \cdot \lambda \cdot f_1 := f_0(x_0)\lambda f_1(x_1)$. If $x_0 = x_1 = x$, then such a derivation is equivalent to a vector in $T_x M$ by Theorem 12.1. If $x_0 \neq x_1$, then the extension above always splits. To see this, choose $f \in C^\infty(M)$ with $f(x_0) \neq f(x_1)$. Then $\varrho(f)$ has the two different eigenvalues $f(x_0)$ and $f(x_1)$. Then $\varrho(f)$ is diagonalisable. Since the two eigenvalues are different, any matrix that commutes with $\varrho(f)$ also becomes diagonal in the eigenbasis of $\varrho(f)$. This shows how the representation ϱ is a direct sum of two 1-dimensional representations. By Theorem 12.1, this says that any derivation into $\text{Hom}(W, V)$ is inner in this case.

The equivalence relation in Theorem 12.1 only allows isomorphisms of representations that induce the identity on V and W . This is why two representations as above may be isomorphic. In the first case, the representations for pairs of points (x_0, x_1) and (x'_0, x'_1) are isomorphic if and only if $\{x_0, x_1\} = \{x'_0, x'_1\}$. This gives the orbit space $(M \times M)/\mathbb{Z}/2$ with $\mathbb{Z}/2$ acting on $M \times M$ by $(x_0, x_1) \mapsto (x_1, x_0)$. The representation for $x_0 = x_1 = x$ and the tangent vector $0 \in T_x M$ is already covered by this. A representation given by a non-trivial tangent vector $\xi \in T_x M$ is not isomorphic to any of these because its corresponding module extension does not

split. The representations for ξ and $c \cdot \xi$ for $c \in \mathbb{C}^*$ are equivalent: this corresponds to rescaling one of the basis vectors for \mathbb{C}^2 . And this is the only case when the representations for two tangent vectors are equivalent. Therefore, these representations are classified up to equivalence by the projective bundle $\mathbb{P}TM$ of the tangent bundle of M . Summing up, equivalence classes of two-dimensional representations of $C^\infty(M)$ are in bijection with the disjoint union $(M \times M)/\mathbb{Z}/2 \sqcup \mathbb{P}TM$.

13. More on derivations, automorphisms and Lie algebra structure

Roughly speaking, the space of derivations $\text{Der}(A, A)$ is the tangent space of the space $\text{Aut}(A)$ of algebra automorphisms $A \rightarrow A$. We are going to make this statement more precise in Section 13.1. We will also relate inner derivations and inner automorphisms in a similar fashion. We also briefly discuss some physical ideas related to the integration and differentiation of 1-parameter groups.

The ‘‘Lie algebra’’ of the automorphism group is not well defined because $\text{Aut}(A)$ is not a Lie group. We should have some topology on A to talk about smooth 1-parameter groups of automorphisms and their generators. In examples, however, it is clear enough what this means. And the generators of such smooth 1-parameter groups are indeed derivations. The integration of derivations is even more problematic: it often fails to exist. This problem already occurs for smooth manifolds because there are vector fields that do not integrate to a flow that exists for all times.

We introduce the concept of a Lie algebra in Section 13.2, and we show that derivations carry a Lie bracket. This should be expected since they are, roughly speaking, the Lie algebra of a group.

13.1. Derivations and automorphisms. We first consider the case where A is $C^\infty(M)$ for a smooth manifold M . By Theorem 12.1, $\text{Der}(A, A)$ is naturally isomorphic to the space of vector fields on M .

DEFINITION 13.1. A *flow* or a *1-parameter group* of diffeomorphisms on M is a group homomorphism $\Phi: \mathbb{R} \rightarrow \text{Diffeo}(M)$, $t \mapsto \Phi_t$, that is smooth in the sense that $(t, m) \mapsto \Phi_t(m)$ is a smooth map $\mathbb{R} \times M \rightarrow M$. The *generator* of the flow is the vector field $X: M \rightarrow TM$ defined by

$$X(m) := \left. \frac{\partial}{\partial t} \Phi_t(m) \right|_{t=0}.$$

The name ‘‘generator’’ suggests that we may reconstruct a flow from its generating vector field. This is indeed the case for smooth compact manifolds:

THEOREM 13.2. *Let M be a smooth compact manifold. For any smooth vector field X there is a unique flow Φ whose generator is X .*

PROOF. We only sketch the proof. By assumption, a flow satisfies $\Phi_t \Phi_s = \Phi_{t+s}$ for all $s, t \in \mathbb{R}$. The chain rule shows that

$$\dot{\Phi}_{t_0}(m) := \left. \frac{\partial}{\partial t} \Phi_t(m) \right|_{t=t_0} = X(\Phi_{t_0}(m)).$$

Hence the function $\Phi(t, m) := \Phi_t(m)$ solves the initial value problem $\frac{\partial}{\partial t} \Phi = X \circ \Phi$ and $\Phi(0, m) = m$. This differential equation has a unique local solution by the Picard–Lindelöf Theorem. That is, there is a neighbourhood U of $\{0\} \times M \subseteq \mathbb{R} \times M$ on which there is a unique smooth solution $\Phi: U \times M \rightarrow M$. Since M compact, U contains $M \times (-\varepsilon, \varepsilon)$ for some $\varepsilon > 0$. We may extend the local flow on $M \times (-\varepsilon, \varepsilon)$ to $M \times (-2\varepsilon, 2\varepsilon)$ by $\Phi(2t, m) = \Phi(t, \Phi(t, m))$ – this also solves the initial value problem and extends the previous local solution on $M \times (-\varepsilon, \varepsilon)$ by uniqueness of the local solution. Repeating this step, we extend our flow to all of $M \times \mathbb{R}$. \square

If M is not compact, then the uniqueness part of Theorem 13.2 still holds, but the trick to extend Φ to all of $\mathbb{R} \times M$ fails without additional assumptions on X .

EXERCISE 13.3. Show that Theorem 13.2 still holds if M is non-compact, but X has compact support. That is, there is a unique flow with generator X provided $X(m) = 0$ outside a compact subset of M .

Now we turn to a general algebra A . We assume given an algebra $C^\infty(\mathbb{R}, A)$ of smooth functions $\mathbb{R} \rightarrow A$. Of course, we let $C^\infty(\mathbb{R}, C^\infty(M)) := C^\infty(\mathbb{R} \times M)$ for a smooth manifold M . If A is a topological algebra, then there is a well defined notion of smooth A -valued function, which we may use to define $C^\infty(\mathbb{R}, A)$. If A carries no further structure, then we may declare that a function $f: \mathbb{R} \rightarrow A$ is smooth near $t \in \mathbb{R}$ if there is a neighbourhood U of t and a finite-dimensional subspace $V \subseteq A$ such that f restricts to a function from U to V that is smooth in the usual sense of calculus.

In the following, we do not care how $C^\infty(\mathbb{R}, A)$ is defined, we only assume it to have the following basic properties. We must be able to evaluate smooth functions at points, that is, we require algebra homomorphisms $\text{ev}_t: C^\infty(\mathbb{R}, A) \rightarrow A$, $f \mapsto f(t)$, for all $t \in \mathbb{R}$; we assume that $f = 0$ if $\text{ev}_t(f) = 0$ for all $t \in \mathbb{R}$. We need a derivation $\frac{\partial}{\partial t}: C^\infty(\mathbb{R}, A) \rightarrow C^\infty(\mathbb{R}, A)$, $f \mapsto f' = \frac{\partial f}{\partial t}$. We also assume that there are translation operators $\tau_s: C^\infty(\mathbb{R}, A) \rightarrow C^\infty(\mathbb{R}, A)$ that commute with the differentiation map $\frac{\partial}{\partial t}$ and satisfy $\tau_t \tau_s = \tau_{t+s}$ and $\text{ev}_t \tau_s = \text{ev}_{t-s}$ for all $s, t \in \mathbb{R}$. Finally, we assume that any $f \in C^\infty(\mathbb{R}, A)$ with $\frac{\partial f}{\partial t} = 0$ is constant.

DEFINITION 13.4. A *smooth 1-parameter group of automorphisms* of A is an algebra homomorphism $\alpha: A \rightarrow C^\infty(\mathbb{R}, A)$ such that the maps $\alpha_t := \text{ev}_t \circ \alpha: A \rightarrow A$ satisfy $\alpha_t \circ \alpha_s = \alpha_{t+s}$ for all $s, t \in \mathbb{R}$ and $\alpha_0 = \text{Id}_A$.

The *generator* of such a smooth 1-parameter group is the map

$$D\alpha: A \rightarrow A, \quad a \mapsto \text{ev}_0 \left(\frac{\partial}{\partial t} \alpha(a) \right).$$

The condition $\alpha_t \circ \alpha_s = \alpha_{t+s}$ in Definition 13.4 is equivalent to $\alpha \circ \alpha_s = \tau_s \circ \alpha$ for all $s \in \mathbb{R}$.

LEMMA 13.5. *The map $D\alpha: A \rightarrow A$ is a derivation.*

PROOF. This follows because α and ev_0 are algebra homomorphisms and $\partial/\partial t$ is assumed to be a derivation on $C^\infty(\mathbb{R}, A)$. \square

Let $d: A \rightarrow A$ be the generator of a smooth 1-parameter group of automorphisms $\alpha: A \rightarrow C^\infty(\mathbb{R}, A)$. Then

$$\text{ev}_t \circ \frac{\partial}{\partial t} \circ \alpha = \text{ev}_0 \circ \tau_t \circ \frac{\partial}{\partial t} \circ \alpha = \text{ev}_0 \circ \frac{\partial}{\partial t} \circ \alpha \circ \alpha_t = d \circ \alpha_t.$$

That is, α solves the differential equation $\dot{\alpha}_t = d \circ \alpha_t$. This implies by induction that

$$\frac{\partial^n}{\partial t^n} \alpha_t = d^n \circ \alpha_t.$$

Hence the formal *Taylor series* $\sum_{n=0}^{\infty} \frac{\alpha^{(n)}(0)}{n!} t^n$ of α at 0 is equal to the exponential series

$$\sum_{n=0}^{\infty} \frac{t^n d^n(a)}{n!} =: \exp(t \cdot d)(a).$$

Thus the problem of integrating a derivation d to a 1-parameter group of automorphisms is equivalent to the problem of defining linear operators $\exp(td): A \rightarrow A$ for $t \in \mathbb{R}$ with reasonable properties such as

$$\exp(td) \exp(sd) = \exp((t+s)d), \quad \frac{\partial}{\partial t} \exp(td) = d \exp(td).$$

In general, this may be impossible or there may be more than one solution. That is, the correspondence between derivations and smooth 1-parameter groups of automorphisms is no longer bijective as for smooth compact manifolds. Nevertheless, since a derivation determines at least the formal Taylor series of a 1-parameter group, we may consider $\text{Der}(A, A)$ as a formal tangent space of the automorphism group $\text{Aut}(A, A)$ at Id_A .

The following examples show some of the problems that appear in connection with derivations and 1-parameter automorphism groups of algebras.

EXAMPLE 13.6. For a smooth manifold M and $k \in \mathbb{N}$, let $C^k(M)$ be the algebra of k times continuously differentiable functions on M . Any 1-parameter group of diffeomorphisms of M generates a 1-parameter group of automorphisms of $C^k(M)$. But the latter is not smooth. The problem is that the generator of the diffeomorphism group – a vector field X on M – maps $C^k(M)$ only to $C^{k-1}(M)$. Thus we only get a derivation from $C^k(M)$ to the $C^k(M)$ -bimodule $C^{k-1}(M)$.

EXAMPLE 13.7. Let $A := C^\infty(\mathbb{R}, \mathbb{C})$ and let $d: A \rightarrow A$ be the derivation associated to the complex-valued vector field $X(t) := i\partial/\partial t$. On the subalgebra of holomorphic functions $\mathbb{C} \rightarrow \mathbb{C}$, the 1-parameter automorphism group $\tau_{it}f(s) := f(s + it)$ integrates this vector field by the Cauchy–Riemann equation. But this makes no sense for functions defined only on \mathbb{R} . In fact, the vector field above does not integrate to a smooth 1-parameter group of automorphism of A . This is because automorphisms of A all come from diffeomorphisms. Thus a smooth 1-parameter group of automorphisms of A must come from a flow, which has a real-valued vector field as its generator.

DEFINITION 13.8. Let A be a unital algebra and let $u \in A$ be invertible. Then we define an associated *inner automorphism* $\text{Ad}_u: A \rightarrow A$ by $\text{Ad}_u(a) := uau^{-1}$.

LEMMA 13.9. *The map Ad_u is an algebra automorphism. If A is $*$ -algebra, then Ad_u is a $*$ -automorphism of A if and only if u is unitary. We have $\text{Ad}_1 = \text{Id}_A$ and $\text{Ad}_{uv} = \text{Ad}_u \circ \text{Ad}_v$ for all $u, v \in A$, that is, $u \mapsto \text{Ad}_u$ is a group homomorphism from the group of invertible elements in A to the automorphism group of A . And $\text{Ad}_u = \text{Id}_A$ if and only if u belongs to the centre of A .*

PROOF. All claims are very simple computations. For instance,

$$\text{Ad}_u(x) \text{Ad}_u(y) = uxu^{-1}uyu^{-1} = uxyu^{-1} = \text{Ad}_u(xy). \quad \square$$

Let A^\times denote the group of invertible elements in A . A 1-parameter group in A is a group homomorphism $u: \mathbb{R} \rightarrow A^\times$. It is smooth if there is an element of $U \in C^\infty(\mathbb{R}, A)$ with $u(t) = \text{ev}_t(U)$. We define the *generator* of u to be the element $X := \text{ev}_0\left(\frac{\partial}{\partial t}U\right) \in A$. We have $u(t) = \exp(tX)$ in the sense that U has the formal power series expansion $\sum_{n=0}^{\infty} \frac{(tX)^n}{n!}$ at 0 and solves the differential equation $\dot{U}(t) = X \cdot U(t)$. We compute

$$\frac{\partial}{\partial t} \text{Ad}_U = \text{ad}_X.$$

That is, the generator of the inner automorphism group Ad_u generated by u is the inner automorphism associated to the generator of u .

Recall that quantum mechanics describes a physical system by its $*$ -algebra of observables A . (More precisely, the observables are the self-adjoint elements of this $*$ -algebra.) We assume that the system is closed; this means that there are $*$ -algebra automorphisms $U(s, t): A \rightarrow A$ for all $s, t \in \mathbb{R}$ with $U(s, t) = U(s, r) \circ U(r, t)$ for all $s, r, t \in \mathbb{R}$, which describe the time evolution from time s to time t .

If we assume, in addition, that the time evolution is not explicitly time dependent, then $U(s, t)$ depends only on $t - s$. Then $t \mapsto U(?, ? + t)$ is a 1-parameter group of automorphisms of A that describes the time evolution of the system.

The smoothness of the time evolution is not automatic. In fact, the time evolution is usually not smooth on all elements of A . We may, however, pass to a subalgebra of smooth elements on which the time evolution is a smooth 1-parameter group of automorphisms. Its generator is a derivation of A called *energy*. The assumption that the time evolution is not explicitly time-dependent is equivalent to energy conservation, that is, the energy is time-independent.

If our system has more symmetry like a translation or rotation symmetry, then we get further 1-parameter groups of automorphisms by translations along $t\vec{v}$ for some fixed $\vec{v} \in \mathbb{R}^3$ or rotations around some fixed axis by angle t . Again, these become smooth on a suitable dense subalgebra of the algebra of all observables. The generator of the translation $t\vec{v}$ describes the \vec{v} -component of the *momentum*; the generator of rotations around an axis describes the corresponding *angular momentum*.

One might expect energy, momenta, and angular momenta to be observables of the system and hence to belong to the algebra A . Mathematically, this would mean that these derivations are inner, given by some elements of A , which are determined uniquely up to adding an element of the centre of A . But these operators are usually unbounded: the energy of a system should be bounded below, but is usually not bounded from above. This means that the relevant observables are “unbounded.” The algebra A is often taken to consist only of bounded observables because products of unbounded observables are hard to define.

13.2. Derivations as a Lie algebra. A *Lie group* is a group and a smooth manifold at the same time, such that the multiplication and inversion maps are smooth. The tangent space \mathfrak{g} of a Lie group G is not just a vector space – it inherits a binary operation $[\cdot, \cdot]: \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$ called the *Lie bracket*.

It is constructed as follows. First, for each $g \in G$, the conjugation map $x \mapsto gxg^{-1}$ maps the unit element of G to itself. Hence its derivative yields a linear representation Ad of G on the vector space \mathfrak{g} . Any element X of \mathfrak{g} generates a 1-parameter group $\exp(tX)$, which is represented by a 1-parameter group $\text{Ad}(\exp tX)$ of linear maps $\mathfrak{g} \rightarrow \mathfrak{g}$. The generator of the latter is a linear map $\text{ad}_X: \mathfrak{g} \rightarrow \mathfrak{g}$. Finally, we let $[X, Y] := \text{ad}_X(Y)$.

EXAMPLE 13.10. Let $G = \text{GL}(n, \mathbb{R})$ be the Lie group of all invertible $n \times n$ -matrices. Then $\mathfrak{g} = \text{M}_n\mathbb{R}$, $\text{Ad}_g(X) = gXg^{-1}$, and $\text{ad}_X(Y) = XY - YX$. So $[X, Y]$ is the usual commutator bracket on $\text{M}_n\mathbb{R}$.

The properties of the Lie bracket are formalised in the concept of a Lie algebra:

DEFINITION 13.11. A *Lie algebra* over a field K is a K -vector space \mathfrak{g} with a map $\mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$, $(X, Y) \mapsto [X, Y]$, that is K -bilinear and anti-symmetric and satisfies the *Jacobi identity*

$$[X, [Y, Z]] + [Y, [Z, X]] + [Z, [X, Y]] = 0.$$

EXAMPLE 13.12. Any algebra A becomes a Lie algebra for the commutator bracket $[X, Y] := XY - YX$. The Jacobi identity for the commutator bracket is a routine computation.

Let A be an algebra. Since we view $\text{Der}(A, A)$ as a formal tangent space of the group $\text{Aut}(A)$ at the identity automorphism, we expect it to inherit a Lie bracket as well. This is indeed the case:

LEMMA 13.13. *Define*

$$[X, Y] := X \circ Y - Y \circ X : A \rightarrow A$$

for two derivations $X, Y \in \text{Der}(A, A)$. This is again a derivation. And this bracket turns $\text{Der}(A, A)$ into a Lie algebra. And $[\text{ad}_a, \text{ad}_b] = \text{ad}_{[a, b]}$ for all $a, b \in A$, that is, $a \mapsto \text{ad}_a$ is a Lie algebra homomorphism.

PROOF. We compute

$$\begin{aligned} [X, Y](a \cdot b) &= X(Y(a)b + aY(b)) - Y(X(a)b + aX(b)) \\ &= XY(a)b + Y(a)X(b) + X(a)Y(b) + aXY(b) \\ &\quad - (YX(a)b + X(a)Y(b) + Y(a)X(b) + aYX(b)) \\ &= [X, Y](a) \cdot b + a \cdot [X, Y](b). \end{aligned}$$

So $[X, Y]$ is a derivation. The equation $[\text{ad}_a, \text{ad}_b] = \text{ad}_{[a, b]}$ is another simple computation. It is equivalent to the Jacobi identity of the commutator bracket. \square

14. Representations and crossed products for Lie algebras

For a Lie algebra \mathfrak{g} , we are going to define representations of \mathfrak{g} . The universal enveloping algebra $U(\mathfrak{g})$ of \mathfrak{g} is the unique algebra so that a representation of \mathfrak{g} is equivalent to a representation of $U(\mathfrak{g})$. We define actions of Lie algebras on algebras. While Lie groups should act by automorphisms, Lie algebras should act by derivations. Given an action of \mathfrak{g} on an algebra A by derivations, we are going to define covariant representations and a crossed product $A \rtimes \mathfrak{g}$. This is analogous to the crossed product for group actions studied in Section 10.

The study of Lie algebra representations is a vast subject, which deserves a course of its own. We will not even scratch the surface of it. We mostly restrict our study to the example of the Lie algebra \mathbb{R} with 0 bracket. An action of \mathbb{R} on an algebra is a single derivation, and so we are going to define a crossed product of an algebra by a derivation. Then we examine one particular example in some detail, namely, the derivation $f \mapsto f'$ on the polynomial algebra $\mathbb{C}[x]$. The resulting crossed product is the Weyl algebra, the universal algebra generated by two elements subject to the relation $[p, q] = 1$. We define crossed products in greater generality for the benefit of readers who have already met Lie groups and Lie algebras.

DEFINITION 14.1. Let \mathfrak{g} be a Lie algebra and let V be a vector space. A *representation* of \mathfrak{g} on V is a linear map $\varrho : \mathfrak{g} \rightarrow \text{End}(V)$ such that

$$\varrho([X, Y]) = [\varrho(X), \varrho(Y)] := \varrho(X) \cdot \varrho(Y) - \varrho(Y) \cdot \varrho(X)$$

for all $X, Y \in \mathfrak{g}$.

The algebra $\text{End}(V)$ with the commutator bracket is a Lie algebra. A representation is a Lie algebra homomorphism to this Lie algebra.

EXAMPLE 14.2. If A is an algebra, then the bracket on derivations is defined so that $\text{Der}(A, A) \subseteq \text{End}(A)$ is a Lie subalgebra. Thus the canonical action of derivations on A is a representation of $\text{Der}(A, A)$ on A .

EXAMPLE 14.3. Let G be a Lie group and let $\alpha : G \rightarrow \text{End}(V)$ be a representation of G . We assume that a reasonable space $C^\infty(G, A)$ of “smooth” functions $G \rightarrow A$ is given. Then we call the representation *smooth* if there is an algebra homomorphism $\alpha : A \rightarrow C^\infty(G, A)$, $\alpha(a)(g) = \alpha_g(a)$. Let \mathfrak{g} be the Lie algebra of G . For each $X \in \mathfrak{g}$, we get a group homomorphism $\mathbb{R} \rightarrow G$, $t \mapsto \exp(tX)$, and a corresponding 1-parameter group of invertible maps on V . This is smooth if $C^\infty(G, A)$ is defined reasonably. So it has a generator $\alpha_X \in \text{End}(V)$. We claim

that the map $X \mapsto \alpha_X$ is a Lie algebra representation of \mathfrak{g} if $C^\infty(G, A)$ is defined reasonably. First, if $g \in G$, $X \in \mathfrak{g}$, then

$$\alpha_g \alpha_X \alpha_g^{-1} = \frac{\partial}{\partial t} \alpha_{g \exp(tX) g^{-1}} \Big|_{t=0} = \frac{\partial}{\partial t} \alpha_{\exp(t \operatorname{Ad}_g X)} \Big|_{t=0} = \alpha_{\operatorname{Ad}_g(X)}.$$

Secondly, when $g = \exp(sY)$ for $s \rightarrow 0$, then differentiation of this equation in the s -direction using the chain rule implies

$$\alpha_Y \alpha_X - \alpha_X \alpha_Y = \frac{\partial}{\partial s} \alpha_{\exp(sY)} \alpha_X \alpha_{\exp(-sY)} \Big|_{s=0} = \frac{\partial}{\partial s} \alpha_{\operatorname{Ad}_{\exp(sY)}(X)} \Big|_{s=0} = \alpha_{[Y, X]}.$$

The passage to Lie algebra representations above is very important to study the representation theory of Lie groups such as $\operatorname{Gl}(n, \mathbb{R})$. Typically, the interesting representations of G are continuous representations on a Banach space or even a Hilbert space. Then it follows that there is a dense subset of smooth vectors on which the representation becomes smooth. This dense subspace then carries a representation of the Lie algebra. This turns out to be easier to study. We shall not pursue this topic further in this course.

DEFINITION 14.4. Let \mathfrak{g} be a Lie algebra. Its *universal enveloping algebra* $U(\mathfrak{g})$ is the unital algebra generated by the set \mathfrak{g} with the relations that the map $i: \mathfrak{g} \rightarrow U(\mathfrak{g})$ is linear and satisfies $i([X, Y]) = i(X) \cdot i(Y) - i(Y) \cdot i(X)$.

Let V be a vector space. If $\rho: U(\mathfrak{g}) \rightarrow \operatorname{End}(V)$ is a representation of $U(\mathfrak{g})$, then $\rho \circ i: \mathfrak{g} \rightarrow \operatorname{End}(V)$ is a Lie algebra representation of \mathfrak{g} , and this map is a bijection between the two types of representations. This universal property is analogous to the universal property of the group ring.

The definitions above work for Lie algebras over any field. At the moment, our default field is the real numbers \mathbb{R} because there is a good geometric picture for derivations on $C^\infty(M, \mathbb{R})$. For the example of the Weyl, we will prefer to work over \mathbb{C} once again.

EXAMPLE 14.5. Let $\mathfrak{g} = \mathbb{R}^n$ with the zero bracket. A representation of \mathfrak{g} is equivalent to a family of n commuting operators on a vector space. Therefore, the universal enveloping algebra of \mathfrak{g} is isomorphic to the algebra of polynomials $\mathbb{R}[x_1, \dots, x_n]$. In particular, $\mathbb{R}[x]$ is the universal enveloping algebra of the Lie algebra \mathbb{R} with zero bracket.

A group should act on an algebra by algebra automorphisms. Similarly, a Lie algebra should act on an algebra by derivations:

DEFINITION 14.6. An *action* of a Lie algebra \mathfrak{g} on an algebra A is a Lie algebra homomorphism $\mathfrak{g} \rightarrow \operatorname{Der}(A, A)$, that is, a linear map $\alpha: \mathfrak{g} \rightarrow \operatorname{Der}(A, A)$ that satisfies $\alpha_{[X, Y]} = [\alpha_X, \alpha_Y]$.

DEFINITION 14.7. Let \mathfrak{g} be a Lie algebra, A a unital algebra, and $\alpha: \mathfrak{g} \rightarrow \operatorname{Der}(A, A)$ an action of \mathfrak{g} on A . A *covariant representation* of $(A, \mathfrak{g}, \alpha)$ on a vector space V is a pair (π, ρ) consisting of an algebra representation $\pi: A \rightarrow \operatorname{End}(V)$ and a Lie algebra representation $\rho: \mathfrak{g} \rightarrow \operatorname{End}(V)$, subject to the covariance condition

$$[\rho(X), \pi(a)] = \pi(\alpha_X(a))$$

for all $X \in \mathfrak{g}$, $a \in A$. The *crossed product* $A \rtimes_\alpha \mathfrak{g}$ is an algebra with the property that its representations are equivalent to covariant representations of $(A, \mathfrak{g}, \alpha)$.

EXAMPLE 14.8. Let G be a Lie group that acts smoothly on a unital algebra A by automorphisms. The construction in Example 14.3 gives an induced action of the Lie algebra \mathfrak{g} on A , which is an action by derivations. Given a smooth covariant representation of the Lie group action, the construction in Example 14.3 gives a covariant representation of the Lie algebra action on A .

EXAMPLE 14.9. Let \mathfrak{g} be \mathbb{R} with zero bracket. Then a Lie algebra action on an algebra A is equivalent to a single derivation d on A . A covariant representation becomes equivalent to a pair (π, X) consisting of a representation $\pi: A \rightarrow \text{End}(V)$ and a linear map $X \in \text{End}(V)$ that satisfies $[X, \pi(a)] = \pi(d(a))$ for all $a \in A$.

THEOREM 14.10. *Let \mathfrak{g} be a Lie algebra and A a unital algebra. Let $\alpha: \mathfrak{g} \rightarrow \text{Der}(A, A)$ be an action of \mathfrak{g} on A . There is a unique associative multiplication on $A \otimes U(\mathfrak{g})$ such that $i_A: A \rightarrow A \otimes U(\mathfrak{g})$, $a \mapsto a \otimes 1$, is an algebra homomorphism, $i_{\mathfrak{g}}: \mathfrak{g} \rightarrow A \otimes U(\mathfrak{g})$, $X \mapsto 1 \otimes X$, is a Lie algebra homomorphism, and $[i_{\mathfrak{g}}(X), i_A(a)] = i_A(\alpha_X(a))$ for all $X \in \mathfrak{g}$, $a \in A$. With this multiplication, $A \otimes U(\mathfrak{g})$ has the universal property of the crossed product $A \rtimes_{\alpha} \mathfrak{g}$.*

PROOF. We show that there is an associative multiplication as above by constructing a faithful representation of the algebra $A \otimes U(\mathfrak{g})$. Let (π, ϱ) be a covariant representation of $(A, \mathfrak{g}, \alpha)$ on a vector space V . It defines a representation $\bar{\varrho}: U(\mathfrak{g}) \rightarrow \text{End}(V)$ by the universal property of $U(\mathfrak{g})$. Then there is a map $\psi: A \otimes U(\mathfrak{g}) \rightarrow \text{End}(V)$, $a \otimes \omega \mapsto \pi(a)\bar{\varrho}(\omega)$. We claim that $\psi(A \otimes U(\mathfrak{g}))$ is a subalgebra of $\text{End}(V)$. The covariance condition implies $\bar{\varrho}(X) \cdot \pi(a) = \pi(\alpha_X(a)) + \pi(a) \cdot \bar{\varrho}(X)$ for $X \in \mathfrak{g}$, $a \in A$. Since \mathfrak{g} generates $U(\mathfrak{g})$, this implies $\bar{\varrho}(\omega) \cdot \pi(a) \in \psi(A \otimes U(\mathfrak{g}))$ for all $\omega \in U(\mathfrak{g})$, $a \in A$. Then it follows that $\psi(A \otimes U(\mathfrak{g}))$ is a subalgebra of $\text{End}(V)$. Assume for a moment that ψ is injective. Then the multiplication in $\text{End}(V)$ gives an associative multiplication in $A \otimes U(\mathfrak{g})$. It has the properties required in the theorem, and the proof shows that it is the only multiplication with these properties. In addition, the proof shows that $A \otimes U(\mathfrak{g})$ with this product has the universal property of the crossed product. So it remains to build a covariant representation for which ψ is injective.

We start with the regular representation λ of A in $\text{End}(A)$, defined by $\lambda_a(b) := a \cdot b$. The pair (λ, α) is covariant:

$$\lambda(\alpha_X(a))b = \alpha_X(a) \cdot b = \alpha_X(a \cdot b) - a \cdot \alpha_X(b) = [\alpha_X, \lambda(a)](b)$$

for all $X \in \mathfrak{g}$, $a, b \in A$. This representation is faithful on A because A is unital. But we need ψ to be faithful on $A \otimes U(\mathfrak{g})$. For this, we build a covariant representation on $A \otimes U(\mathfrak{g})$. It consists of the representation $\lambda \otimes 1$ of A and the representation $X \mapsto \alpha_X \otimes 1 + 1 \otimes \lambda_X$ of \mathfrak{g} , where λ_X means left multiplication by X . It is easy to check that this is a covariant representation. The proof that the resulting map ψ is injective is not hard if $\mathfrak{g} = \mathbb{R}$. For general Lie algebras, this follows from the Poincaré–Birkhoff–Witt Theorem, and we leave out all further details. \square

Now we specialise to the case where $\mathfrak{g} = \mathbb{R}$. Since $U(\mathbb{R}) = \mathbb{R}[t]$, the crossed product algebra $A \rtimes_d \mathbb{R}$ for a derivation d is $A \otimes \mathbb{R}[t]$ as a vector space. The multiplication turns out to be

$$(a \otimes t^m) \cdot (b \otimes t^n) = \sum_{k=0}^m \binom{m}{k} a \cdot d^k(b) \otimes t^{m-k+n}$$

for all $a, b \in A$, $n, m \in \mathbb{N}$. This formula follows by induction on m from the condition $t \cdot b = d(b) + b \cdot t$. The fact that the regular representation and d generate a representation of $A \rtimes_d \mathbb{R}$ on A implies the following formula for a derivation d :

$$(14.11) \quad d^m(a \cdot b) = \sum_{k=0}^m \binom{m}{k} d^k(a) \cdot d^{m-k}(b)$$

for all $a, b \in A$, $m \in \mathbb{N}$. This can also be proven directly by induction on m .

Inner derivations are trivial in some sense. One way to make this precise is the following proposition, which shows that the crossed product for an inner derivation is isomorphic to the crossed product for the zero derivation.

PROPOSITION 14.12. *Let $d = \text{ad}_x$ be an inner derivation on an algebra A . There is a unique algebra isomorphism $A \rtimes_d \mathbb{R} \cong A \otimes \mathbb{R}[t]$ that maps $a \otimes 1 \mapsto a \otimes 1$ and $1 \otimes t \mapsto 1 \otimes t - x \otimes 1$. Here the target $A \otimes \mathbb{R}[t]$ carries the obvious multiplication $(a \otimes t^m) \cdot (b \otimes t^n) := (ab) \otimes t^{m+n}$.*

PROOF. We built a faithful representation of $A \rtimes_d \mathbb{R}$ on $A \otimes \mathbb{R}[t]$. Its image is generated by $\lambda_a \otimes 1$ for $a \in A$ and $X := d \otimes 1 + 1 \otimes \lambda_t$, where λ_t means the operator of multiplication by t . Let $d = \text{ad}_x$ for some $x \in A$. Then $X - \lambda_x \otimes 1$ commutes with $\lambda(A) \otimes 1$. Hence there is an algebra homomorphism $A \otimes \mathbb{R}[t] \rightarrow \text{End}(A \otimes \mathbb{R}[t])$ that maps $a \otimes t^n \mapsto (\lambda_a \otimes 1) \cdot (X - \lambda_x \otimes 1)^n$. It is an isomorphism onto the image of $A \rtimes_d \mathbb{R}$ in $\text{End}(A \otimes \mathbb{R}[t])$. \square

Now we work once again over the complex numbers, and we examine a specific crossed product. Namely, let A be the crossed product of $\mathbb{C}[q]$ by the derivation $d(q^n) := i\hbar nq^{n-1}$ for some $\hbar \in \mathbb{C} \setminus \{0\}$. This algebra is called *Weyl algebra* by mathematicians and *Heisenberg algebra* by mathematical physicists. By definition, it is generated by a copy of $\mathbb{C}[q]$ and an extra generator p that satisfies $[p, f] = d(f)$ for all $f \in \mathbb{C}[q]$. Since q generates $\mathbb{C}[q]$, this relation follows from

$$[p, q] := pq - qp = \hbar i.$$

This equation is called the *canonical commutation relation*. It is the defining relation for the crossed product for the polynomial algebra with the derivation $f \mapsto i\hbar f'$. This relation plays a role in quantum mechanics. It is related to the Heisenberg uncertainty principle. By induction, the canonical commutation relation implies

$$[p^k, q] = \hbar i k p^{k-1}, \quad [p, q^k] = \hbar i k q^{k-1}.$$

Assume $\hbar \in \mathbb{R}$. Then there is a unique $*$ -algebra structure on A where p and q are self-adjoint. This needs the factor i : if p and q are self-adjoint, then $pq - qp$ is skew-adjoint.

The parameter \hbar above plays no significant role: if $[p, q] = 1$, then $[p, \hbar i q] = \hbar i$, so that these commutation relations generate the same algebra. It is, however, interesting to consider the *classical limit* $\hbar \rightarrow 0$. The limiting case is the algebra of polynomials in two variables with the usual commutative multiplication. As we shall see, the primitive ideal space and the space of irreducible representations differ drastically for $\hbar = 0$ and $\hbar \neq 0$. We will later meet other invariants of noncommutative algebras that yield the same result for $\hbar = 0$ and $\hbar \neq 0$. Roughly speaking, these are global topological invariants that are not affected by small-scale quantum effects.

As a vector space, the Weyl algebra is isomorphic to $\mathbb{C}[q] \otimes \mathbb{C}[p]$. Using this, we may write down a nice faithful representation:

EXAMPLE 14.13. Consider the operators

$$p, q: \mathbb{C}[x] \ni \mathbb{C}[x], \quad p(f) := x \cdot f, \quad q(f) := \frac{\hbar}{i} f' = -\hbar i f',$$

where f' denotes the derivative of f . The product rule for derivatives shows that p and q satisfy the canonical commutation relation. Hence they generate a representation of the Weyl algebra. We compute that $p^n q^m$ for $n, m \in \mathbb{N}$ is the operator $f \mapsto (-\hbar i)^m x^n f^{(m)}$. These operators are linearly independent. Therefore, the representation of the Weyl algebra defined above is faithful.

REMARK 14.14. The operators p and q above also act on the space $C_c^\infty(\mathbb{R})$ of smooth functions with compact support. This is a dense subspace of the Hilbert space $L^2(\mathbb{R})$, so that p and q become unbounded operators on $L^2(\mathbb{R})$. The closures of these unbounded operators are self-adjoint, and they are used in quantum mechanics

to describe the position and momentum of a 1-dimensional object. Many physicists use the name “Weyl algebra” for the algebra generated by the operators

$$\begin{aligned} \exp(2\pi itp): C_c^\infty(\mathbb{R}) &\rightarrow C_c^\infty(\mathbb{R}), & f &\mapsto (x \mapsto \exp(2\pi itx) \cdot f(x)), \\ \exp(2\pi itq): C_c^\infty(\mathbb{R}) &\rightarrow C_c^\infty(\mathbb{R}), & f &\mapsto (x \mapsto f(x - 2\pi \hbar t)), \end{aligned}$$

for all $t \in \mathbb{R}$. These extend to unitary operators on $L^2(\mathbb{R})$. This makes them easier to handle than the unbounded self-adjoint operators p and q .

PROPOSITION 14.15. *The Weyl algebra is simple.*

PROOF. Let $I \subseteq A$ be a non-zero ideal. We must show that $I = A$. Pick $g \in I$ with $g \neq 0$ and write $g = a_n q^n + a_{n-1} q^{n-1} + \cdots + a_0$ with $a_j \in \mathbb{C}[p]$ and $a_n \neq 0$. Since I is a two-sided ideal, it is invariant under the map

$$\text{ad}_x: A \rightarrow A, \quad y \mapsto [x, y] := xy - yx,$$

for all $x \in A$. Since $[p, a_j q^j] = \hbar i j a_j q^{j-1}$, we get $\text{ad}_p^n(g) = (\hbar i)^n n! a_n \in \mathbb{C}[p] \cap I$. Write $a_n = c_m p^m + \cdots + c_0$ with $c_m, \dots, c_0 \in \mathbb{C}$ and $c_m \neq 0$. The same reasoning as above shows that $\text{ad}_q^m(a_n) = (-\hbar i)^m m! c_m \in I$. Hence $1 \in I$, so that $I = A$. \square

PROPOSITION 14.16. *The algebra A is the unique unital algebra with two generators satisfying the canonical commutation relation. Let B be any unital algebra and let $P, Q \in B$ satisfy $[p, q] = \hbar i$. Then there is a unique unital algebra homomorphism $f: A \rightarrow B$ with $f(p) = P$ and $f(q) = Q$. And f is an algebra isomorphism onto the subalgebra of B generated by P and Q .*

PROOF. The existence of a unique algebra homomorphism $f: A \rightarrow B$ with $f(p) = P$ and $f(q) = Q$ follows easily from the universal properties for polynomial algebras and for the crossed product by a derivation. The range of f is the subalgebra of B generated by P and Q . Since A is simple and $f(1) = 1$, the ideal $\ker f$ must vanish. So f is injective. \square

The uniqueness statement in Proposition 14.16 is specific to a universal algebra with some generators and relations that is also simple. The Leavitt path algebras in Example 9.23 are also simple, so that they have a similar uniqueness property. The Toeplitz algebra, however, does not have this property. It is the universal unital $*$ -algebra generated by an isometry, but it is not simple. Some isometries – namely, those that are unitary – generate an algebra that is not isomorphic to the Toeplitz algebra.

PROPOSITION 14.17. *Let f be a polynomial and let $V_f \subseteq C^\infty(\mathbb{R})$ be the subspace of all functions of the form $g \exp(f)$ with a polynomial $g \in \mathbb{C}[x]$. Let p and q act on V_f by $p(g \exp f) := x \cdot g \exp f$ and*

$$q(g \exp f) := -i\hbar \frac{d}{dx}(g \exp f) = -i\hbar(g' - gf') \exp f.$$

This defines an irreducible representation of A on V_f . The irreducible representations V_{f_1} and V_{f_2} are only isomorphic if $f_2 = f_1 + c$ for a constant $c \in \mathbb{C}$.

PROOF. The algebra A clearly acts on $C^\infty(\mathbb{R})$ by differential operators, and the action of p and q on V_f is the restriction of that action. Hence the above formulas define an A -module structure on V_f . Let $W \subseteq V_f$ be a non-zero submodule. Since $p(W) \subseteq W$, the set of $g \in \mathbb{C}[x]$ with $g \exp(f) \in W$ must be an ideal in $\mathbb{C}[x]$. Hence

$$W = \{g g_0 \exp(f) : g \in \mathbb{C}[x]\}$$

for some $g_0 \in \mathbb{C}[x]$. Now $q(g_0 \exp f) \in W$ implies that g_0 divides g_0' . But this is impossible unless $g_0' = 0$, that is, g_0 is constant and $W = V_f$. Thus V_f is simple.

It is clear that $V_{f_2} = V_{f_1}$ if $f_2 - f_1$ is constant. Conversely, assume that $V_{f_2} \cong V_{f_1}$ for two polynomials f_1, f_2 . The isomorphism $\Phi: V_{f_1} \rightarrow V_{f_2}$ must map $g \exp(f_1)$ to $\varphi \cdot g \exp(f_2)$ for all $g \in \mathbb{C}[x]$ for some polynomial $\varphi \in \mathbb{C}[x]$ because it is p -linear. The condition $Q \circ \Phi(\exp f_1) = \Phi \circ Q(\exp f_2)$ means that φ satisfies the differential equation $\varphi' + \varphi f_2' = \varphi f_1'$, that is, $\varphi' = (f_2 - f_1)' \cdot \varphi$. Unless $(f_1 - f_2)' = 0$, this has no non-zero *polynomial* solution because the left hand side has smaller degree than the right hand side. Hence such an isomorphism cannot exist unless $f_1 - f_2$ is constant. \square

REMARK 14.18. In mathematical physics, we are mainly interested in self-adjoint solutions of the canonical commutation relation. This drastically changes the representation theory: if P, Q are (unbounded) self-adjoint operators on a Hilbert space \mathcal{H} that satisfy the canonical commutation relation $[P, Q] = i\hbar$, then there is a unitary operator on $\mathcal{H} \cong \bigoplus_{i \in I} L^2(\mathbb{R}, dx)$ that intertwines P, Q with a direct sum of copies of the multiplication operator $f \mapsto x \cdot f$ and the differentiation operator $f \mapsto -i\hbar f'$. In particular, the standard representation on $L^2(\mathbb{R}, dx)$ is the unique irreducible self-adjoint representation of the canonical commutation relation.

THEOREM 14.19. *Let A be the Weyl algebra. Any derivation $A \rightarrow A$ is inner.*

PROOF. We claim that any element of A may be written as a commutator $[x, p]$ for some $x \in A$. To see this, recall that the monomials $p^n q^m$ for $n, m \in \mathbb{N}$ form a basis of A and that $[p^n q^m, p] = -i\hbar m p^n q^{m-1}$. This shows that the map $x \mapsto -[x, p]$ is surjective and that its kernel is the subspace $\mathbb{C}[p]$ of polynomials in p . The same argument shows that any element of A may be written as $[y, q]$ for some $y \in A$. And we may choose $y \in \mathbb{C}[p]$ for elements of $\mathbb{C}[p]$.

Let $d: A \rightarrow A$ be a derivation. By the first paragraph, there is $x \in A$ with $[x, p] = d(p)$. The derivation $d' := d - \text{ad}_x$ satisfies $d'(p) = 0$ by construction. Since $d'(1) = 0$ for any derivation into a unital bimodule, the relation $[p, q] = i\hbar$ implies

$$0 = d'(pq - qp) = d'(p)q - qd'(p) + pd'(q) - d'(q)p = -[d'(q), p].$$

Since $d'(q)$ commutes with p , it is a polynomial in p . Hence $d'(q) = [y, q]$ for some $y \in \mathbb{C}[p]$. Then $[y, p] = 0 = d'(p)$. So the derivation $d - \text{ad}_{x+y}$ annihilates both generators p and q . This implies that it vanishes identically by the Leibniz rule. Thus $d = \text{ad}_{x+y}$. \square

REMARK 14.20. Theorem 14.19 is only about derivations $A \rightarrow A$. There are non-inner derivations from the Weyl algebra into bimodules over it.

LEMMA 14.21. *The only invertible elements of the Weyl algebra A are the constant multiples of the identity, which are central. So the only inner automorphism of A is the identity map.*

PROOF. We represent A faithfully as an algebra of linear differential operators on $C^\infty(\mathbb{R})$. If $D \in A$ is invertible, then the corresponding ordinary differential equation $D(f) = 0$ cannot have any non-zero solutions. This forces D to have order zero. That is, D must be a polynomial in p . So must the inverse of D . The only invertible polynomials are the non-zero constants. \square

Translations define a 1-parameter group of automorphisms

$$\tau_t: A \rightarrow A, \quad \tau_t(p^n q^m) := (p - t)^n q^m$$

for $t \in \mathbb{R}$, $n, m \in \mathbb{N}$ because $\tau_t(p) = p - t$ and $\tau_t(q) = q$ still satisfy the canonical commutation relation. The generator of this 1-parameter group is given by

$$\left. \frac{\partial}{\partial t} \tau_t(p^n q^m) \right|_{t=0} := \left. \frac{\partial}{\partial t} (p - t)^n q^m \right|_{t=0} = (-n)(p - t)^{n-1} q^m \Big|_{t=0} = -np^{n-1} q^m.$$

This is the inner derivation $\text{ad}_{q/i\hbar}$. Although the generator is inner, the automorphisms τ_t are not inner because $\exp(qt/i\hbar)$ cannot be defined as an element of A .

EXERCISE 14.22. Show that there is a unique automorphism of the Weyl algebra mapping $p \mapsto q$ and $q \mapsto -p$. Use this and the computations above to integrate the inner derivation ad_p to a 1-parameter group of automorphisms of A .

15. The universal derivation — Morita invariance

First, we describe derivations $A \rightarrow M$ as bimodule maps $\Omega^1(A) \rightarrow M$ for an A -bimodule $\Omega^1(A)$. The elements of this bimodule are called noncommutative differential forms. This bimodule will play a very important role for the definition of higher Hochschild cohomology and homology. It is also important to compute the space of derivations $A \rightarrow M$ for a given algebra A and bimodule M . Secondly, we learn how to compute the quotient of derivations modulo inner derivations using projective bimodule resolutions. This is a very powerful computational tool because there are often rather small projective bimodule resolutions. We will use this description of derivations modulo inner derivations to prove that this quotient is invariant under Morita equivalence.

15.1. The universal derivation. Let A be a unital algebra. We are going to describe a universal derivation $d: A \rightarrow \Omega^1(A)$, which is characterised by the following universal property: for any derivation $D: A \rightarrow M$ there is a unique A -bimodule homomorphism $f: \Omega^1(A) \rightarrow M$ with $f \circ d = D$. The bimodule $\Omega^1(A)$ here is a very important object. Its elements are called *noncommutative differential 1-forms* over A . Before we turn to arbitrary derivations, we first describe inner derivations through bimodule maps in a similar fashion.

Turn $A \otimes A$ into an A -bimodule in the obvious way: $a \cdot (b \otimes c) \cdot d := (a \cdot b) \otimes (c \cdot d)$ for all $a, b, c, d \in A$. Let $\text{Hom}_{A,A}(V, W)$ for two A -bimodules V, W denote the space of A -bimodule homomorphisms $V \rightarrow W$. Then

$$\text{Hom}_{A,A}(A \otimes A, M) \cong M$$

for any unital A -bimodule M by mapping $f \in \text{Hom}_{A,A}(A \otimes A, M)$ to $f(1 \otimes 1)$ and $m \in M$ to the map $a \otimes b \mapsto a \cdot m \cdot b$. Turn A into an A -bimodule by left and right multiplication, and let $\text{mult}: A \otimes A \rightarrow A$ be the multiplication map. This is a bimodule homomorphism.

Since mult is surjective, the induced map

$$\text{mult}^*: \text{Hom}_{A,A}(A, M) \rightarrow \text{Hom}_{A,A}(A \otimes A, M) \cong M$$

is injective. The range of mult^* in M is the centre of M :

$$\text{Hom}_{A,A}(A, M) \cong \{m \in M : a \cdot m = m \cdot a \text{ for all } a \in A\}.$$

Since $\text{ad}_m = 0$ if and only if m is central, we get

$$\text{Inn}(A, M) \cong \text{coker}(\text{mult}^*: \text{Hom}_{A,A}(A, M) \rightarrow \text{Hom}_{A,A}(A \otimes A, M)).$$

Here the *cokernel* of a map $f: X \rightarrow Y$ is defined to be the quotient $Y/f(X)$.

DEFINITION 15.1. Let $\Omega^1(A) := \ker(\text{mult}: A \otimes A \rightarrow A)$ and define

$$d: A \rightarrow \Omega^1(A), \quad d(a) := 1 \otimes a - a \otimes 1.$$

The map d is the inner derivation into $A \otimes A$ generated by $1 \otimes 1 -$ but it is usually not inner as a derivation into $\Omega^1(A)$ because $1 \otimes 1$ does not belong to $\Omega^1(A)$.

The bimodule $\Omega^1(A)$ is also called the *bimodule of noncommutative differential forms* on A . It is spanned by elements of the form $a \, db := a \cdot d(b)$ for $a \in A$, $b \in A/\mathbb{C} \cdot 1$ – recall that $d(1) = 0$ for any derivation d into a unital A -bimodule.

Thus $\Omega^1(A) \cong A \otimes (A/\mathbb{C} \cdot 1)$ as a left A -module. The right module structure is dictated by the Leibniz rule:

$$(adb) \cdot c := a d(bc) - ab dc.$$

PROPOSITION 15.2. *Let M be a unital A -bimodule and let $D: A \rightarrow M$ be a derivation. There is a unique bimodule homomorphism $f: \Omega^1(A) \rightarrow M$ with $f \circ d = D$. Roughly speaking, $d: A \rightarrow \Omega^1(A)$ is the universal derivation.*

PROOF. Given D , we define a linear map $f: A \otimes A \rightarrow M$ by $a \otimes b \mapsto a \cdot D(b)$. This is a left A -module homomorphism by definition, and if $a, b, c \in A$, then

$$f(a \otimes b \cdot c) = a \cdot D(bc) = a \cdot D(b) \cdot c + ab \cdot D(c) = f(a \otimes b) \cdot c + ab \cdot D(c)$$

Thus f is not a bimodule homomorphism on $A \otimes A$. But we only care about its restriction to $\Omega^1(A)$. On this submodule, the summand $ab \cdot D(c) = \text{mult}(a \otimes b) \cdot D(c)$ disappears. So $f(\omega \cdot c) = f(\omega) \cdot c$ for all $\omega \in \Omega^1(A)$, $c \in A$. It is also clear that any bimodule homomorphism $f: \Omega^1(A) \rightarrow M$ with $f \circ d = D$ must map $a \cdot d(b) = a \otimes b - ab \otimes 1$ to $a \cdot D(b) - ab \otimes D(1) = aD(b)$ for all $a, b \in A$. This forces the map to be $f|_{\Omega^1(A)}$. \square

REMARK 15.3. Let A be a non-unital K -algebra for a field K . Give $A^+ := A \oplus K$ the multiplication

$$(a_1, \lambda_1) \cdot (a_2, \lambda_2) := (a_1 \cdot a_2 + \lambda_1 \cdot a_2 + a_1 \cdot \lambda_2, \lambda_1 \cdot \lambda_2)$$

for $a_1, a_2 \in A$, $\lambda_1, \lambda_2 \in K$. This makes A^+ a unital K -algebra. The inclusion of A is a K -algebra homomorphism. Any homomorphism from A to a unital K -algebra B extends uniquely to a unital homomorphism $A^+ \rightarrow B$. Thus A^+ is the *universal unital K -algebra generated by A* . In particular, an A -module structure on a K -vector space is the same as a unital A^+ -module structure. If $D: A \rightarrow M$ is a derivation into an A -bimodule M , then D extends uniquely to a derivation $A^+ \rightarrow M$, $(a, \lambda) \mapsto D(a)$. This is easily seen to be a derivation. It is the only extension of D to a derivation on A^+ because any derivation $A^+ \rightarrow M$ must annihilate the unit element $(0, 1) \in A^+$. Roughly speaking, derivations from A to A -bimodules are the same as derivations from A^+ to A^+ -bimodules. Hence the universal derivation for A is the composite of the inclusion $A \hookrightarrow A^+$ with the universal derivation $d: A^+ \rightarrow \Omega^1(A^+)$.

Thus we define $\Omega^1(A) := \Omega^1(A^+)$ if A is non-unital. This is isomorphic to the tensor product $A^+ \otimes A$ by mapping $(a, \lambda) \otimes b \mapsto a d(b) + \lambda d(b)$ for $a, b \in A$, $\lambda \in K$. The above definition of $\Omega^1(A)$ for non-unital A causes some confusion for unital algebras because we may forget that A has a unit and redefine $\Omega^1(A)$ using A^+ . If this confusion is problematic, then one has to choose two different names for the two variants of Ω^1 . We shall, however, be sloppy about this point. Most algebras to which we apply Ω^1 below are unital.

By definition, we have a bimodule extension $\Omega^1(A) \hookrightarrow A \otimes A \twoheadrightarrow A$. When we apply the functor $\text{Hom}_{A,A}(\square, M)$ to it, we get the sequence

$$Z(M) \hookrightarrow M \xrightarrow{\text{ad}} \text{Der}(A, M).$$

Thus the quotient space

$$\text{HH}^1(A, M) := \frac{\text{Der}(A, M)}{\text{Inn}(A, M)} \cong \text{coker}(\text{ad}: M \rightarrow \text{Der}(A, M))$$

measures to what extent the map $\text{Hom}_{A,A}(A \otimes A, M) \rightarrow \text{Hom}_{A,A}(\Omega^1(A), M)$ is surjective; recall that the *cokernel* of a map $f: X \rightarrow Y$ is the quotient $Y/f(X)$. Similarly, the centre

$$\text{HH}^0(A, M) := Z(M) \cong \ker(\text{ad}: M \rightarrow \text{Der}(A, M))$$

measures the failure of the map $\text{Hom}_{A,A}(A \otimes A, M) \rightarrow \text{Hom}_{A,A}(\Omega^1(A), M)$ to be injective. Here $\text{HH}^0(A, A)$ is the centre of the algebra A . The names $\text{HH}^0(A, M)$ and $\text{HH}^1(A, M)$ are justified because these two groups are part of a general cohomology theory for A -bimodules called *Hochschild cohomology*.

COROLLARY 15.4. *The derivation $d: A \rightarrow \Omega^1(A)$ is inner if and only if all derivations into A -bimodules are inner.*

PROOF. If all derivations are inner, then so is d . Conversely, let d be inner. That is, there is $x \in \Omega^1(A)$ with $d = \text{ad}_x$. Let $D: A \rightarrow M$ be any derivation into an A -bimodule M . Then $D = f \circ d$ for a bimodule map $f: \Omega^1(A) \rightarrow M$ by Proposition 15.2. Then $D = f \circ \text{ad}_x = \text{ad}_{f(x)}$. So any derivation is inner. \square

EXERCISE 15.5. For $A = \mathbb{M}_n\mathbb{C}$, we already know that all derivations are inner. In particular, there must be $x \in \Omega^1(A)$ with $\text{ad}_x = d$. Find such an element x .

Why did we speak of “differential forms” in connection with $\Omega^1(A)$? There is a certain analogy between $\Omega^1(A)$ and differential forms on a smooth manifold. There are, however, also marked differences. We now discuss this. Let $A := C^\infty(M, \mathbb{R})$ for a smooth manifold M . The space $\Omega^1(M)$ of *smooth 1-forms* is the space of smooth sections of the cotangent bundle T^*M . A purely algebraic description is

$$(15.6) \quad \Omega^1(M) \cong \text{Hom}_{C^\infty(M)}(\mathfrak{X}(M), C^\infty(M)),$$

where $\mathfrak{X}(M)$ denotes the space of vector fields on M . In local coordinates – that is, on \mathbb{R}^n – elements of $\Omega^1(M)$ are formal linear combinations $\sum_{i=1}^n f_i dx_i$ with smooth functions $f_i: \mathbb{R}^n \rightarrow \mathbb{R}$. The canonical pairing with vector fields is

$$\left\langle \sum_{i=1}^n f_i dx_i, \sum_{j=1}^n g_j \frac{\partial}{\partial x_j} \right\rangle := \sum_{i=1}^n f_i g_i.$$

There is a canonical map $d_{\text{dR}}: C^\infty(M) \rightarrow \Omega^1(M)$ called *de Rham differential*, which maps $f \in C^\infty(M)$ to the module homomorphism $\mathfrak{X}(M) \rightarrow C^\infty(M)$, $X \mapsto X(f)$. It is a derivation in the sense that $d_{\text{dR}}(fg) = f d_{\text{dR}}g + g d_{\text{dR}}f$ for all $f, g \in C^\infty(M)$. In local coordinates, we have

$$d_{\text{dR}}(f) := \sum_{i=1}^n \frac{\partial f}{\partial x_i} dx_i.$$

Notice that (15.6) makes no sense for noncommutative algebras because the space $\text{Der}(A, A)$ of derivations, the analogue of $\mathfrak{X}(M)$, carries no canonical A -module structure. In fact, $\Omega^1(A)$ is not such a good noncommutative analogue of $\Omega^1(M)$ because $\Omega^1(C^\infty(M))$ is quite different from $\Omega^1(M)$.

Like $\Omega^1(A)$, we may characterise $\Omega^1(M)$ by a universal property. We may view any $C^\infty(M)$ -module V as a bimodule via $a \cdot v \cdot b := ab \cdot v$ for $a, b \in C^\infty(M)$, $v \in V$ – this works because $C^\infty(M)$ is commutative. Thus we can talk about derivations into $C^\infty(M)$ -modules instead of $C^\infty(M)$ -bimodules. The de Rham differential is an example of this.

PROPOSITION 15.7. *Let V be a $C^\infty(M)$ -module, viewed as a bimodule, and let $D: C^\infty(M) \rightarrow V$ be a derivation. Then there is a unique module homomorphism $\varphi: \Omega^1(M) \rightarrow V$ with $f \circ d_{\text{dR}} = D$.*

PROOF. Any element of $\Omega^1(M)$ may be written as a finite sum of elementary 1-forms $f dg$ with $f, g \in C^\infty(M)$. This follows, for instance, from the Embedding Theorem 1.22: if $h: M \rightarrow \mathbb{R}^N$ is an embedding, then any smooth differential 1-form may be written as $\sum_{j=1}^N f_j dh_j$. Hence there is at most one way to extend D to a module homomorphism $\Omega^1(M) \rightarrow V$ because $f dg$ must be mapped to $f Dg$. Conversely, there is a well defined map $\varphi: \Omega^1(M) \rightarrow V$ with $\varphi(f dg) = f Dg$. \square

COROLLARY 15.8. *Let M be a smooth manifold and $A := C^\infty(M)$. The A -module $\Omega^1(M)$ is isomorphic to the commutator quotient $\Omega^1(A) / [A, \Omega^1(A)]$ of $\Omega^1(A)$ by the linear span of $[a, \omega]$ for all $a \in A$, $\omega \in \Omega^1(A)$.*

PROOF. The quotient $\Omega^1(A)/[A, \Omega^1(A)]$ is an A -module with the same universal property as $\Omega^1(M)$. \square

REMARK 15.9. For topological algebras such as $A = C^\infty(M)$, it is usually better to complete $\Omega^1(A)$ to the kernel of the multiplication map $A \widehat{\otimes} A \rightarrow A$, where $A \widehat{\otimes} A$ denotes an appropriate completion of $A \otimes A$. For instance, it is useful to complete $C^\infty(M) \otimes C^\infty(M)$ to $C^\infty(M \times M)$. Then $\Omega^1(C^\infty M)$ becomes isomorphic to the quotient of $C^\infty(M \times M)$ by $C^\infty(M) \otimes 1$.

15.2. Morita invariance. Our next goal is to establish the following theorem:

THEOREM 15.10. *Let A and B be Morita equivalent unital algebras. Then $\mathrm{HH}^j(A, A) \cong \mathrm{HH}^j(B, B)$ for $j = 0, 1$.*

For $j = 0$, this asserts that Morita equivalent unital algebras have isomorphic centres. For $j = 1$, it asserts that both have isomorphic quotients $\mathrm{Der}/\mathrm{Inn}$. Derivations and inner derivations alone are not Morita invariant: this fails already for matrix algebras.

COROLLARY 15.11. *Let G be a finite group that acts freely on a smooth manifold M . Then*

$$\begin{aligned} \mathrm{HH}^0(C^\infty(M) \rtimes G, C^\infty(M) \rtimes G) &\cong C^\infty(G \backslash M), \\ \mathrm{HH}^1(C^\infty(M) \rtimes G, C^\infty(M) \rtimes G) &\cong \mathfrak{X}(G \backslash M). \end{aligned}$$

PROOF. Theorem 11.16 shows that $C^\infty(M) \rtimes G$ and $C^\infty(G \backslash M)$ are Morita equivalent. Being commutative, the latter algebra is its own centre and has no inner derivations. Its derivations correspond to smooth vector fields on $G \backslash M$ by Theorem 12.1. \square

Recall that a Morita equivalence between two unital algebras A and B is implemented by an A, B -bimodule P and a B, A -bimodule Q together with bimodule isomorphisms $P \otimes_B Q \cong A$ and $Q \otimes_A P \cong B$. These yield an equivalence between the categories of unital A -bimodules $\mathfrak{Bimod}(A)$ and of unital B -bimodules $\mathfrak{Bimod}(B)$ by

$$\varphi: \mathfrak{Bimod}(B) \rightarrow \mathfrak{Bimod}(A), \quad M \mapsto P \otimes_B M \otimes_B Q;$$

its inverse maps $M' \mapsto Q \otimes_A M' \otimes_A P$. In particular,

$$\varphi(B) := P \otimes_B B \otimes_B Q \cong P \otimes_B Q \cong A.$$

Since φ is an equivalence of categories, it maps extensions again to extensions. Since $\varphi(B) = A$, we get an A -bimodule extension

$$\varphi(\Omega^1 B) \mapsto \varphi(B \otimes B) \mapsto A.$$

We must compare it to the A -bimodule extension $\Omega^1 A \mapsto A \otimes A \mapsto A$. We have $\varphi(B \otimes B) \cong P \otimes Q$, which is usually not isomorphic to $A \otimes A$. But these two bimodules have a property in common that allows to compare the two extensions.

LEMMA 15.12. *Let P be an A -bimodule. The following assertions are equivalent:*

- (a) *the functor $\mathrm{Hom}_{A,A}(P, \square)$ is exact, that is, if $K \mapsto E \rightarrow Q$ is a bimodule extension, then the induced maps*

$$\mathrm{Hom}_{A,A}(P, K) \rightarrow \mathrm{Hom}_{A,A}(P, E) \rightarrow \mathrm{Hom}_{A,A}(P, Q)$$

form an extension of Abelian groups.

- (b) any surjective bimodule homomorphism $\pi: E \rightarrow P$ splits, that is, there is a bimodule homomorphism $\sigma: P \rightarrow E$ with $\pi \circ \sigma = \text{Id}_P$; we also call σ a section for π .
- (c) if $\pi: E \rightarrow Q$ is a surjective bimodule homomorphism, then any bimodule homomorphism $f: P \rightarrow Q$ lifts to a bimodule homomorphism $\hat{f}: P \rightarrow E$ (lifting means $\pi \circ \hat{f} = f$).

PROOF. Both (a) and (c) talk about a bimodule extension

$$(15.13) \quad K \hookrightarrow E \xrightarrow{\pi} Q$$

The statement in (c) says that the map $\text{Hom}_{A,A}(P, E) \rightarrow \text{Hom}_{A,A}(P, Q)$ is surjective. This is contained in (a). And a bimodule map $f: P \rightarrow K$ is equivalent to a bimodule map $\bar{f}: P \rightarrow E$ with $\pi \circ \bar{f} = 0$ because $K = \ker(\pi)$. Thus (a) and (c) are equivalent. The statement in (c) implies (b) by lifting the identity map on P to a map $\sigma: P \rightarrow E$. It remains to show that (b) implies (c). For this, take a bimodule extension as in (15.13) and a bimodule map $f: P \rightarrow Q$. Let

$$\hat{E} := \{(e, p) \in E \times P : \pi(e) = f(p)\}.$$

Since π is surjective, so is the second coordinate projection $\hat{\pi}: \hat{E} \rightarrow P$, $(e, p) \mapsto p$. By (b), there is a bimodule map $g: P \rightarrow \hat{E}$ with $\hat{\pi} \circ g = \text{Id}_P$. Hence $g(p) = (\hat{f}(p), p)$ for a bimodule map $\hat{f}: P \rightarrow E$. Since the image of g is contained in \hat{E} , $\pi \circ \hat{f} = f$. So \hat{f} lifts f as required in (c). \square

DEFINITION 15.14. An A -bimodule P is called *projective* if it has the equivalent properties in Lemma 15.12.

EXAMPLE 15.15. Let V be any vector space. Turn $A \otimes V \otimes A$ into an A -bimodule by $a \cdot (b \otimes v \otimes c) \cdot d := (ab) \otimes v \otimes (cd)$. This is called the *free* bimodule on V . It is characterised by the property that there is a natural isomorphism

$$\text{Hom}_{A,A}(A \otimes V \otimes A, M) \cong \text{Hom}(V, M)$$

for any A -bimodule M . Here $\text{Hom}_{A,A}(A \otimes V \otimes A, M)$ denotes A -bimodule maps $f: A \otimes V \otimes A \rightarrow M$ and $\text{Hom}(V, M)$ denotes linear maps $g: V \rightarrow M$. The bijection maps $f: A \otimes V \otimes A \rightarrow M$ to the map $V \rightarrow M$, $v \mapsto f(1 \otimes v \otimes 1)$, and it maps $g: V \rightarrow M$ to the map $A \otimes V \otimes A \rightarrow M$, $a_1 \otimes v \otimes a_2 \mapsto a_1 \cdot g(v) \cdot a_2$. It is not hard to see that this is a bijection as asserted. The bijection above implies that free bimodules are projective because any extension of vector spaces splits. In particular, $A \otimes A$ is a projective A -bimodule.

An equivalence of categories must preserve projective objects because they are defined in purely category theoretic terms. Hence both $A \otimes A$ and $\varphi(B \otimes B)$ are projective A -bimodules. For the same reason, P is projective as a left A -module and as a right B -module, and Q is projective as a right A -module and as a left B -module.

LEMMA 15.16 (Schanuel's Lemma). *Let*

$$K_1 \xrightarrow{i_1} E_1 \xrightarrow{p_1} Q \quad \text{and} \quad K_2 \xrightarrow{i_2} E_2 \xrightarrow{p_2} Q$$

be two bimodule extensions with the same quotient. Assume that E_1 and E_2 are projective. Then there are bimodule isomorphisms that make the following diagram commute:

$$\begin{array}{ccccc} K_1 \oplus K_2 & \xrightarrow{i_1 \oplus \text{Id}} & E_1 \oplus K_2 & \xrightarrow{(p_1, 0)} & Q \\ \downarrow \cong & & \downarrow \cong & & \parallel \\ K_1 \oplus K_2 & \xrightarrow{\text{Id} \oplus i_2} & K_1 \oplus E_2 & \xrightarrow{(0, p_2)} & Q. \end{array}$$

PROOF. The map $p_1: E_1 \rightarrow Q$ lifts to a map $\hat{p}_1: E_1 \rightarrow E_2$ with $p_2 \circ \hat{p}_1 = p_1$ because E_1 is projective. The map \hat{p}_1 restricts to a map $h: K_1 \rightarrow K_2$. The map $(\hat{p}_1, i_2): E_1 \oplus K_2 \rightarrow E_2$ is surjective because $E_2/i_2(K_2) \cong Q$ and p_2 is surjective; its kernel is isomorphic to K_1 via $(i_1, -h): K_1 \rightarrow E_1 \oplus K_2$. Thus we get a bimodule extension

$$K_1 \hookrightarrow E_1 \oplus K_2 \twoheadrightarrow E_2.$$

This extension splits because E_2 is projective. This produces a bimodule isomorphism $E_1 \oplus K_2 \cong K_1 \oplus E_2$. This isomorphism is compatible with the quotient maps to Q and hence restricts to a bimodule isomorphism $K_1 \oplus K_2 \xrightarrow{\cong} K_1 \oplus K_2$ between the kernels of this quotient map. \square

DEFINITION 15.17. Two bimodule homomorphisms $f_1: X_1 \rightarrow Y_1$ and $f_2: X_2 \rightarrow Y_2$ are called *stably isomorphic* if there are bimodules S_1 and S_2 and bimodule isomorphisms that make the following diagram commute:

$$\begin{array}{ccc} X_1 \oplus S_1 & \xrightarrow{f_1 \oplus \text{Id}} & Y_1 \oplus S_1 \\ \downarrow \cong & & \downarrow \cong \\ X_2 \oplus S_2 & \xrightarrow{f_2 \oplus \text{Id}} & Y_2 \oplus S_2 \end{array}$$

With this new notation, Schanuel's Lemma asserts that the maps i_1 and i_2 in Lemma 15.16 are stably isomorphic. Since stable isomorphism is a purely categorical concept, it is preserved by any functor. In particular, for any bimodule M , the induced maps

$$\text{Hom}_{A,A}(i_j, M): \text{Hom}_{A,A}(K_j, M) \leftarrow \text{Hom}_{A,A}(E_j, M), \quad j = 1, 2,$$

are stably isomorphic maps of Abelian groups. In addition, stably isomorphic maps have isomorphic kernels and cokernels. Now we have all the tools to establish the Morita invariance of $\text{HH}^0(A, A)$ and $\text{HH}^1(A, A)$:

PROOF OF THEOREM 15.10. Since $A \otimes A$ and $\varphi(B \otimes B)$ are both projective bimodules, Schanuel's Lemma applies to the bimodule extensions $\Omega^1(A) \hookrightarrow A \otimes A \twoheadrightarrow A$ and $\varphi(\Omega^1 B) \hookrightarrow \varphi(B \otimes B) \twoheadrightarrow A$ and shows that the inclusion maps $i_1: \Omega^1(A) \rightarrow A \otimes A$ and $i_2: \varphi(\Omega^1 B) \rightarrow \varphi(B \otimes B)$ are stably isomorphic. So are the induced maps $\text{Hom}_{A,A}(i_1, \varphi(M))$ and $\text{Hom}_{A,A}(i_2, \varphi(M))$ for any B -bimodule M . Hence the latter two maps have the same kernels and cokernels. For i_2 , we may simplify the result because φ is an equivalence of categories, that is, $\text{Hom}_{A,A}(\varphi X, \varphi Y) \cong \text{Hom}_{B,B}(X, Y)$. Thus the following two maps have isomorphic kernels and cokernels:

$$\begin{aligned} \text{Hom}_{A,A}(A \otimes A, \varphi(M)) &\rightarrow \text{Hom}_{A,A}(\Omega^1(A), \varphi(M)), \\ \text{Hom}_{B,B}(B \otimes B, M) &\rightarrow \text{Hom}_{B,B}(\Omega^1(B), M). \end{aligned}$$

Now recall that the kernels of these maps are $\text{HH}^0(A, \varphi(M))$ and $\text{HH}^0(B, M)$, respectively, and their cokernels are $\text{HH}^1(A, \varphi(M))$ and $\text{HH}^1(B, M)$. As a result,

$$(15.18) \quad \text{HH}^0(A, \varphi(M)) \cong \text{HH}^0(B, M), \quad \text{HH}^1(A, \varphi(M)) \cong \text{HH}^1(B, M).$$

This implies the theorem because $\varphi(B) = A$. \square

16. From deformation quantisations to Hochschild cohomology

A prototype of a deformation quantisation is the passage from the polynomial algebra $\mathbb{C}[p, q]$, which has the relation $[p, q] = 0$, to the Weyl algebra with the relation $[p, q] = -i\hbar$. Deformation theory studies this situation systematically. One problem is to classify the possible deformation quantisations of a given algebra. Another problem is how invariants of algebras behave under deformations.

Deformation theory has some physical significance. Let X be the phase space of a classical mechanical system. When we quantise this system, we may hope that there is a quantum theory for each value of the Planck constant \hbar and that the classical system is the limit of these quantum systems for $\hbar \rightarrow 0$. Then there would be a family of observable algebras for all \hbar . These should be related in a suitable way. The concept of a polynomial deformation quantisation is introduced to formalise this idea. Polynomial deformation quantisations are, however, hard to construct. Therefore, we replace them by formal deformation quantisations. Their theory is quite deep, and we will only scratch the surface. We will mostly limit our discussion to a low-order approximation of formal deformation quantisations. This is closely related to the second Hochschild cohomology group $\mathrm{HH}^2(A, A)$. If the latter group vanishes, then it implies that the algebra A has no non-trivial formal deformation quantisations. In general, however, the group $\mathrm{HH}^2(A, A)$ is not yet enough to classify formal deformation quantisations. The definition of $\mathrm{HH}^2(A, A)$ leads us to define Hochschild cohomology in all degrees. In the following sections, we will compute these invariants for some classes of algebras.

16.1. Deformation quantisations and their low-order approximations.

Our guiding example is the universal algebra A with three generators p, q, \hbar with relations $[p, q] = i\hbar$, $[p, \hbar] = 0$ and $[q, \hbar] = 0$. The monomials $\hbar^k p^m q^n$ for $k, m, n \in \mathbb{N}$ form a basis of A . The element \hbar belongs to the centre of A . Even more, the centre of A is equal to $\mathbb{C}[\hbar]$. As a $\mathbb{C}[\hbar]$ -module, A is isomorphic to $\mathbb{C}[\hbar] \otimes A_0$, where $A_0 = \mathbb{C}[p, q]$ is the space of polynomials in two variables. The multiplication in A is such that the isomorphism $A/(\hbar \cdot A) \cong A_0$ is an algebra homomorphism with respect to the usual product on $\mathbb{C}[p, q]$, that is, if f and g are two polynomials in p and q , then the difference between their products in A and A_0 is divisible by \hbar . These properties are formalised in the concept of a polynomial deformation quantisation:

DEFINITION 16.1. Let A_0 be an algebra. A *polynomial deformation quantisation* of A_0 is an associative multiplication m on $A_0[\hbar] := A_0 \otimes \mathbb{C}[\hbar]$ that is \hbar -bilinear and agrees with the multiplication in A_0 up to terms divisible by \hbar .

Since m is \hbar -bilinear, we have

$$m\left(\sum_{j=0}^{\infty} a_j \hbar^j, \sum_{k=0}^{\infty} b_k \hbar^k\right) = \sum_{j,k=0}^{\infty} \hbar^{j+k} m(a_j, b_k)$$

for all $a_j, b_k \in A_0$. Thus a polynomial deformation quantisation is specified by a bilinear map $m: A_0 \times A_0 \rightarrow A_0 \otimes \mathbb{C}[\hbar]$. We may further split $m(a, b) = \sum_{l=0}^{\infty} m_l(a, b) \hbar^l$. And $m_0(a, b) = a \cdot b$ because m is supposed to agree with the product in A_0 up to higher order terms.

What does associativity of the multiplication mean? We compute

$$\begin{aligned} m(a, m(b, c)) &= \sum_{l=0}^{\infty} \hbar^l \cdot \sum_{k=0}^l m_k(a, m_{l-k}(b, c)) \\ m(m(a, b), c) &= \sum_{l=0}^{\infty} \hbar^l \cdot \sum_{k=0}^l m_k(m_{l-k}(a, b), c). \end{aligned}$$

Hence associativity amounts to the conditions

$$(16.2) \quad \sum_{k=0}^l m_k(a, m_{l-k}(b, c)) = \sum_{k=0}^l m_k(m_{l-k}(a, b), c)$$

for all $a, b, c \in A_0$ and all $l \in \mathbb{N}$. This is a sequence of non-linear constraints on the maps m_l . For $l = 0$, (16.2) only involves m_0 and amounts to the associativity of the

multiplication in A_0 . For $l = 1$, we get the condition

$$(16.3) \quad a \cdot m_1(b, c) + m_1(a, b \cdot c) = m_1(a, b) \cdot c + m_1(a \cdot b, c).$$

More generally, the map m_l appears for the first time in the l th equation, and this takes the form

$$a \cdot m_l(b \cdot c) + \cdots + m_l(a, b \cdot c) = m_l(a, b) \cdot c + \cdots + m_l(a \cdot b, c),$$

where \cdots denotes expressions involving only m_1, \dots, m_{l-1} . This suggests that we may find solutions $(m_l)_{l \in \mathbb{N}}$ recursively, starting with (16.3).

Since the equation in each recursion step is quadratic, this recursive construction is still very difficult. It is usually quite hard to prove the existence of deformation quantisations with certain properties. Therefore, we will soon limit our discussion to the map m_1 only. Before we restrict to this limited problem, we briefly mention some important general issues.

For a polynomial deformation quantisation, $m(a, b) = \sum_{l=0}^{\infty} m_l(a, b) \hbar^l$ must be a polynomial. That is, for each $a, b \in A_0$ there must be $l_0 \in \mathbb{N}$ with $m_l(a, b) = 0$ for $l \geq l_0$. Since the recursive approach above does not provide such information, we replace polynomials by formal power series:

DEFINITION 16.4. Let A_0 be an algebra and let $A_0[[\hbar]]$ be the $\mathbb{C}[[\hbar]]$ -module of formal power series in one variable with coefficients in A_0 . A *formal deformation quantisation* of A_0 is an associative multiplication m on $A := A_0[[\hbar]]$ that is \hbar -linear in both variables and that agrees with the multiplication in A_0 up to terms divisible by \hbar .

Any sequence of bilinear maps $(m_l)_{l \in \mathbb{N}}$ that satisfies $m_0(a, b) = a \cdot b$ for all $a, b \in A_0$ and (16.2) for all $l \in \mathbb{N}$ yields a formal deformation quantisation of A_0 .

EXAMPLE 16.5. For $\hbar \in \mathbb{R}$, let $\alpha_{\hbar}: \mathbb{Z} \rightarrow \text{Aut}(\mathbb{C}[t, t^{-1}])$ be the automorphism that is induced by rotation with angle \hbar , that is, $\alpha_{\hbar}(t^n) := \exp(2\pi i \hbar n) t^n$ for all $n \in \mathbb{Z}$. The resulting crossed products $\mathbb{C}[t, t^{-1}] \rtimes_{\hbar} \mathbb{Z}$ give a formal deformation quantisation of the algebra $\mathbb{C}[t, t^{-1}] \rtimes_0 \mathbb{Z} \cong \mathbb{C}[t, t^{-1}, s, s^{-1}]$ of Laurent polynomials in two variables. More precisely, we let $B = \mathbb{C}[t, t^{-1}][[\hbar]]$ and define an \hbar -linear automorphism $\alpha \in \text{Aut}(B)$ by

$$\alpha(t^n \hbar^m) = t^n \exp(2\pi i \hbar n) \hbar^m := t^n \cdot \sum_{j=0}^{\infty} \frac{(2\pi i n)^j}{j!} \hbar^{m+j}.$$

The crossed product $B \rtimes_{\alpha} \mathbb{Z}$ is a formal deformation quantisation of $\mathbb{C}[t, t^{-1}] \rtimes_0 \mathbb{Z}$. This deformation quantisation is not polynomial because it involves exponential functions.

There is a trivial way to modify formal deformation quantisations:

DEFINITION 16.6. Two formal deformation quantisations m and m' are *equivalent* if there is an invertible \hbar -linear map $\Psi: A_0[[\hbar]] \rightarrow A_0[[\hbar]]$ with $m(\Psi(a), \Psi(b)) = \Psi(m'(a, b))$ for all $a, b \in A_0[[\hbar]]$ and $\Psi(a) \equiv a \pmod{\hbar}$ for all $a \in A_0$.

The map Ψ is of the form

$$(16.7) \quad \Psi \left(\sum_{j=0}^{\infty} a_j \hbar^j \right) = \sum_{l=0}^{\infty} \sum_{j=0}^{\infty} \Psi_{l-j}(a_j) \hbar^l$$

with linear maps $\Psi_l: A_0 \rightarrow A_0$ for $l \in \mathbb{N}$ and $\Psi_0 = \text{Id}_A$. Comparing coefficients, we see that (m_l) and (m'_l) are equivalent if and only if

$$(16.8) \quad \sum_{j=0}^l \sum_{k=0}^{l-j} m_{l-j-k}(\Psi_j(a), \Psi_k(b)) = \sum_{j=0}^l \Psi_j(m'_{l-j}(a, b))$$

for all $l \in \mathbb{N}$, $a, b \in A_0$. This condition is rather intransparent for $l \gg 0$. The condition is trivial for $l = 0$. For $l = 1$, it says that

$$(16.9) \quad m_1(a, b) + \Psi_1(a) \cdot b + a \cdot \Psi_1(b) = m'_1(a, b) + \Psi_1(a \cdot b).$$

Any family of linear maps $\Psi_l: A_0 \rightarrow A_0$ with $\Psi_0 = \text{Id}_A$ gives an invertible map Ψ on $A_0[[\hbar]]$. And then $m'(a, b) := \Psi^{-1}(m(\Psi(a), \Psi(b)))$ is a formal deformation quantisation. It is equivalent to m by construction. Roughly speaking, m and m' only differ by a change of coordinates.

16.2. Square-zero extensions and the second Hochschild cohomology.

Now we study only first order approximations of deformation quantisations, that is, we forget about the higher order terms m_l and Ψ_l for $l \geq 2$. This amounts to studying associative multiplications on $A_0[[\hbar]]/(\hbar^2) \cong A_0 \oplus \hbar \cdot A_0$. Equations (16.3) and (16.9) lead to the following concepts:

DEFINITION 16.10. Let A be an algebra and let M be an A -bimodule. An M -valued *Hochschild 2-cocycle* on A is a bilinear map $\omega: A \times A \rightarrow M$ satisfying

$$a \cdot \omega(b, c) + \omega(a, b \cdot c) = \omega(a, b) \cdot c + \omega(a \cdot b, c)$$

for all $a, b, c \in A$. A bilinear map $\omega: A \times A \rightarrow M$ is called a *Hochschild 2-coboundary* if it is of the form

$$\partial\psi(a, b) := a \cdot \psi(b) + \psi(a) \cdot b - \psi(a \cdot b)$$

for some linear map $\psi: A \rightarrow M$. Any Hochschild 2-coboundary is a Hochschild 2-cocycle because

$$\begin{aligned} a \cdot \partial\psi(b, c) + \partial\psi(a, b \cdot c) - \partial\psi(a, b) \cdot c - \partial\psi(a \cdot b, c) \\ = ab\psi(c) + a\psi(b)c - a\psi(bc) + a\psi(bc) + \psi(a)bc - \psi(abc) \\ - a\psi(b)c - \psi(a)bc + \psi(ab)c - ab\psi(c) - \psi(ab)c + \psi(abc) = 0. \end{aligned}$$

Hence we may form a quotient group

$$\text{HH}^2(A, M) := \frac{\{\text{Hochschild 2-cocycles}\}}{\{\text{Hochschild 2-coboundaries}\}}.$$

PROPOSITION 16.11. *There is a bijection between $\text{HH}^2(A, A)$ and equivalence classes of associative multiplications on $A[[\hbar]]/(\hbar^2) \cong A \oplus \hbar \cdot A$.*

PROOF. Equation (16.3) asserts that the map $m_1: A \times A \rightarrow A$ in a deformation quantisation is a Hochschild 2-cocycle. Equation (16.9) asserts that two maps m_1, m'_1 are equivalent if and only if $m_1 - m'_1$ is a Hochschild 2-coboundary. \square

We are going to derive a similar interpretation for Hochschild 2-cocycles with values in other bimodules.

DEFINITION 16.12. An algebra extension $I \hookrightarrow E \twoheadrightarrow Q$ is called a *square-zero extension* of Q by I if $i_1 \cdot i_2 = 0$ for all $i_1, i_2 \in I$.

Let $I \hookrightarrow E \twoheadrightarrow Q$ be a square-zero extension. Since I is an ideal, we get multiplication maps $E \times I \rightarrow I$ and $I \times E \rightarrow I$. These maps descend to maps $Q \times I \rightarrow I$ and $I \times Q \rightarrow I$ because $I \cdot I = 0$. This turns I into a Q -bimodule. Hence we speak of square-zero extensions of the algebra Q by a Q -bimodule in the following.

THEOREM 16.13. *Let A be an algebra and let M be an A -bimodule. There is a natural bijection between $\text{HH}^2(A, M)$ and equivalence classes of square-zero*

extensions of A by M . Here two square-zero extensions of A by M are considered equivalent if there is a commuting diagram

$$\begin{array}{ccccc} M & \longrightarrow & E_1 & \longrightarrow & A \\ \parallel & & \downarrow \cong & & \parallel \\ M & \longrightarrow & E_2 & \longrightarrow & A \end{array}$$

PROOF. Let $M \hookrightarrow E \xrightarrow{p} A$ be a square-zero extension. Choosing a basis for A , we define a linear map $s: A \rightarrow E$ with $p \circ s = \text{Id}_A$. Define

$$\omega_s: A \times A \rightarrow E, \quad \omega_s(a, b) := s(a \cdot b) - s(a) \cdot s(b).$$

Since s is section for p and p is an algebra homomorphism, $p \circ \omega_s = \omega_{p \circ s} = 0$, that is, the values of ω_s lie in $M \subseteq E$.

Together with the algebra structure on A and the A -bimodule structure on M , ω_s determines our square-zero extension up to equivalence: we must have $E \cong M \oplus A$ as a vector space, and the multiplication is

$$\begin{aligned} (m_1 + s(a_1)) \cdot (m_2 + s(a_2)) &= m_1 \cdot m_2 + s(a_1) \cdot m_2 + m_1 \cdot s(a_2) + s(a_1)s(a_2) \\ &= m_1 \cdot m_2 + a_1 m_2 + m_1 a_2 - \omega_s(a_1, a_2) + s(a_1 a_2) \end{aligned}$$

for all $a_1, a_2 \in A$, $m_1, m_2 \in M$.

The following computation shows that ω_2 is a Hochschild 2-cocycle:

$$\begin{aligned} a \cdot \omega_s(b, c) - \omega_s(ab, c) + \omega_s(a, bc) - \omega_s(a, b)c \\ &= s(a) \cdot \omega_s(b, c) - \omega_s(ab, c) + \omega_s(a, bc) - \omega_s(a, b)s(c) \\ &= s(a)s(bc) - s(a)s(b)s(c) - s(abc) + s(ab)s(c) \\ &\quad + s(abc) - s(a)s(bc) - s(ab)s(c) + s(a)s(b)s(c) = 0. \end{aligned}$$

These computations may be reversed: if ω is a Hochschild 2-cocycle, then

$$(m_1, a_1) \cdot (m_2, a_2) := (a_1 \cdot m_2 + m_1 \cdot a_2 - \omega(a_1, a_2), a_1 \cdot a_2)$$

for $m_1, m_2 \in M$, $a_1, a_2 \in A$ defines an associative multiplication μ on $M \oplus A$ and yields a square-zero extension $M \hookrightarrow (M \oplus A, \mu) \twoheadrightarrow A$.

The map ω_s depends on the section s . Another section differs from s by a linear map $\psi: A \rightarrow M$. We compute

$$\begin{aligned} \omega_{s+\psi}(a, b) &= (s + \psi)(ab) - (s + \psi)(a) \cdot (s + \psi)(b) \\ &= s(ab) + \psi(ab) - s(a)s(b) - s(a)\psi(b) - \psi(a)s(b) - \psi(a)\psi(b) \\ &= \omega_s(a, b) + \psi(ab) - a\psi(b) - \psi(a)b. \end{aligned}$$

Thus ω_s and $\omega_{s+\psi}$ differ by the Hochschild 2-coboundary of ψ . Conversely, if ω_s and $\omega_{s'}$ differ by the coboundary of $\psi: A \rightarrow M$, then $\omega_{s'} = \omega_{s+\psi}$. Hence the class of ω_s in $\text{HH}^2(A, M)$ depends only on the equivalence class of the square-zero extension and determines the latter uniquely. \square

The class of a square-zero extension in $\text{HH}^2(A, M)$ vanishes if and only if there is a section with $\omega_s = 0$. Then we say that the extension *splits* (by an algebra homomorphism). Our analysis shows that, up to equivalence, there is a unique split square-zero extension of A by M . The resulting algebra $A \ltimes M$ with multiplication

$$(a_1, m_1) \cdot (a_2, m_2) := (a_1 \cdot a_2, a_1 \cdot m_2 + m_1 \cdot a_2)$$

for all $a_1, a_2 \in A$, $m_1, m_2 \in M$ is also called the *crossed product algebra* of A by the bimodule M .

We now return to the study of formal deformation quantisations. In general, studying $\mathrm{HH}^2(A, A)$ is not enough to classify them. We only have the following strong result in case $\mathrm{HH}^2(A, A) = 0$:

THEOREM 16.14. *Let A be an algebra with $\mathrm{HH}^2(A, A) = 0$. Then all formal deformation quantisations are trivial.*

PROOF. Let $m_j : A \otimes A \rightarrow A$ for $j \in \mathbb{N}$ define a formal deformation quantisation. In particular, $m_0(a, b) = a \cdot b$ is the usual multiplication. We are going to construct maps $\Psi_j : A \rightarrow A$ for $j \in \mathbb{N}$ that give an equivalence between the formal deformation quantisation given by (m_j) and the trivial formal deformation quantisation given by $m'_j = 0$ for $j \geq 1$. We must put $\Psi_0 = \mathrm{Id}_A$.

In the first step, m_1 is a Hochschild 2-cocycle. Since $\mathrm{HH}^2(A, A) = 0$, there is $\Psi_1 : A \rightarrow A$ with $\partial\Psi_1 = m_1$. In other words, (16.9) holds with $m'_1 = 0$. The maps $\Psi_0, \Psi_1, 0, 0, \dots$ define an \hbar -linear vector space isomorphism $\Psi^{(1)}$ on $A[[\hbar]]$ as in (16.7). This gives an equivalence between the given formal deformation quantisation (m_j) and a formal deformation quantisation $(m_j^{(1)})$ with $m_1^{(1)} = 0$.

Before the k th recursion step, we have already found an equivalence from the given formal deformation quantisation (m_j) to a formal deformation quantisation $(m_j^{(k)})$ with $m_j^{(k)} = 0$ for $j = 1, \dots, k-1$. We wish to find an equivalence $\Psi^{(k)}$ from $(m_j^{(k)})$ to a formal deformation quantisation $(m_j^{(k+1)})$ with $m_j^{(k+1)} = 0$ for $j = 1, \dots, k$. In fact, we also assume $\Psi_j^{(k)} = 0$ for $j \in \mathbb{N} \setminus \{0, k\}$. We must always choose $\Psi_0^{(k)} = \mathrm{Id}_A$, so that we only need to find one map $\Psi_k^{(k)} : A \rightarrow A$. This must verify (16.8) for $l = k$. Under our assumptions, this simplifies to

$$m_k^{(k)}(a, b) + \Psi_k^{(k)}(a) \cdot b + a \cdot \Psi_k^{(k)}(b) = \Psi_k^{(k)}(a \cdot b).$$

In other words, we must choose $\Psi_k^{(k)}$ with $\partial\Psi_k^{(k)} = m_k^{(k)}$. Since $m_j^{(k)} = 0$ for $j \in \{1, \dots, k-1\}$, the associativity condition (16.2) for the multiplication $(m_j^{(k)})$ for $l = k$ says that $m_k^{(k)}$ is a Hochschild 2-cocycle. Once again, the assumption $\mathrm{HH}^2(A, A) = 0$ allows us to choose $\Psi_k^{(k)}$ as desired. The equivalences from (m_j) to $(m_j^{(k)})$ converge for $k \rightarrow \infty$ in the sense that each component Ψ_j becomes constant after j steps. The limit is the desired equivalence from (m_j) to $(m_0, 0, 0, \dots)$. \square

This theorem only applies to *formal* deformation quantisations. There may still be non-trivial polynomial deformation quantisations – these are hard to classify. There is a special case where we know much more:

THEOREM 16.15. *Let A be a finite-dimensional vector space and let $(m_t)_{t \in [0,1]}$ be a continuous family of associative multiplications on A , that is, $t \mapsto m_t$ is a continuous map from $[0, 1]$ to the vector space of bilinear maps $A \times A \rightarrow A$ and each m_t is associative. If (A, m_0) is semi-simple, then there is $\varepsilon > 0$ such that $(A, m_0) \cong (A, m_t)$ for all $t \in [0, \varepsilon]$.*

We do not prove this result here. If (A, m_{\hbar}) is a polynomial deformation quantisation of a finite-dimensional semi-simple algebra (A, m_0) , then $(A, m_{\hbar}) \cong (A, m_0)$ for all but finitely many $\hbar \in \mathbb{C}$.

EXAMPLE 16.16. Consider $\mathbb{C}[x, \hbar] / (x^2 - (1 + \hbar)x)$ as a polynomial deformation quantisation of $\mathbb{C}[x] / (x^2 - x)$ or, equivalently, as an \hbar -parametrised family of continuous multiplications on the 2-dimensional vector space with basis $1, x$:

$$m_{\hbar}(1, 1) = 1, \quad m_{\hbar}(1, x) = m_{\hbar}(x, 1) = x, \quad m_{\hbar}(x, x) = x^2 \equiv (1 + \hbar)x.$$

By the Chinese Remainder Theorem, $(\mathbb{C}^2, m_{\hbar}) \cong \mathbb{C} \oplus \mathbb{C}$ for $\hbar \neq -1$ because $x^2 - (1 + \hbar)x = x \cdot (x - 1 - \hbar)$ is a product of two coprime linear polynomials. For

$\hbar = -1$, we get the algebra of dual numbers (see Exercise 3.6). As a result, the above polynomial deformation of $\mathbb{C} \oplus \mathbb{C}$ is non-trivial, although the corresponding formal deformation quantisation is trivial.

16.3. Higher Hochschild cohomology. The very definitions show that Hochschild 2-coboundaries are related to derivations: the coboundary $\partial\psi: A \times A \rightarrow M$ of a linear map $\psi: A \rightarrow M$ vanishes if and only if ψ is a derivation. The connection between the centre of a module and inner derivations is similar: the inner derivation ad_x associated to $x \in M$ vanishes if and only if x is central. This leads us to expect a third Hochschild cohomology group based on Hochschild 3-cocycles and 3-coboundaries, such that the coboundary of a bilinear map vanishes if and only if this map is a Hochschild 2-cocycle, and so on.

DEFINITION 16.17. Let A be an algebra and let M be an A -bimodule. For an n -linear map $\varphi: A^n \rightarrow M$, define its *Hochschild coboundary* $\partial\varphi: A^{n+1} \rightarrow M$ by

$$\begin{aligned} \partial\varphi(a_0, \dots, a_n) &:= a_0 \cdot \varphi(a_1, \dots, a_n) - \varphi(a_0 \cdot a_1, a_2, \dots, a_n) \\ &\quad + \varphi(a_0, a_1 \cdot a_2, a_3, \dots, a_n) - \varphi(a_0, a_1, a_2 \cdot a_3, a_4, \dots, a_n) \pm \dots \\ &\quad + (-1)^n \varphi(a_0, \dots, a_{n-2}, a_{n-1} \cdot a_n) + (-1)^{n+1} \varphi(a_0, \dots, a_{n-2}, a_{n-1}) \cdot a_n. \end{aligned}$$

We call ψ a *Hochschild cocycle* if $\partial\psi = 0$. The following lemma allows to define the n th *Hochschild cohomology* for the algebra A with coefficients in the A -bimodule M as the quotient group

$$\text{HH}^n(A, M) := \frac{\{\text{Hochschild } n\text{-cocycles}\}}{\{\text{Hochschild } n\text{-coboundaries}\}}.$$

Two Hochschild cocycles that differ by adding a coboundary are called *cohomologous*.

LEMMA 16.18. *If $\psi: A^n \rightarrow M$ is an n -linear map, then $\partial(\partial\psi) = 0$. So $\text{HH}^n(A, M)$ is well defined.*

PROOF. The map $\partial^2\psi$ is a sum of terms of the form

$$\begin{aligned} &\pm\varphi(a_0, \dots, a_j a_{j+1}, \dots, a_k a_{k+1}, \dots, a_{n+2}), \\ &\quad \pm a_0 \varphi(a_1, \dots, a_j a_{j+1}, \dots, a_{n+2}), \\ &\quad \pm\varphi(a_0, \dots, a_j a_{j+1} a_{j+2}, \dots, a_{n+2}), \\ &\quad \pm\varphi(a_0, \dots, a_j a_{j+1}, \dots, a_{n+1}) \cdot a_{n+2}, \\ &\quad \pm a_0 a_1 \varphi(a_2, \dots, a_{n+2}), \\ &\quad \pm\varphi(a_0, \dots, a_n) a_{n+1} a_{n+2}, \\ &\quad \pm a_0 \varphi(a_1, \dots, a_{n+1}) a_{n+2}. \end{aligned}$$

Inspection shows that each term occurs twice with opposite signs. \square

By convention, a 0-linear map $A^0 \rightarrow M$ is an element m of M , and its coboundary is the associated inner derivation $\text{ad}_m: A^1 \rightarrow M$, $a \cdot m - m \cdot a$. Thus $\text{HH}^0(A, M)$ and $\text{HH}^1(A, M)$ are the centre of M and the quotient of derivations by inner derivations, respectively. We have interpreted $\text{HH}^2(A, M)$ as the set of equivalence classes of square-zero extensions of A by M . The higher Hochschild cohomology groups $\text{HH}^n(A, M)$ have no special interpretation any more.

DEFINITION 16.19. Let A be a unital algebra and let M be a unital A -bimodule. A Hochschild n -cochain $\omega: A^n \rightarrow M$ is called *normalised* if $\omega(a_1, \dots, a_n)$ vanishes whenever $a_i = 1$ for some $i \in \{1, \dots, n\}$.

Recall that any derivation into a unital A -bimodule satisfies $d(1) = 0$, that is, all Hochschild 1-cocycles are normalised. This is no longer true for Hochschild n -cocycles for higher n . The following remains true, however:

LEMMA 16.20. *Let A be a unital algebra, M a unital A -bimodule, and $n \in \mathbb{N}$. Any Hochschild cocycle $\omega: A^n \rightarrow M$ is cohomologous to a normalised one. And if two normalised Hochschild cocycles $\omega_1, \omega_2: A^n \rightarrow M$ are cohomologous, then there is a normalised Hochschild cochain $\psi: A^{n-1} \rightarrow M$ with $\partial\psi = \omega_1 - \omega_2$. The Hochschild coboundary of a normalised Hochschild cochain is again normalised.*

Roughly speaking, it makes no difference for the Hochschild cohomology whether we restrict attention to normalised Hochschild cochains or not. The proof of Lemma 16.20 will be sketched on page 89.

17. Computing Hochschild cohomology with projective resolutions

The definition suggests that Hochschild n -cocycles get more and more complicated for large n , making $\mathrm{HH}^n(A, M)$ more and more difficult to compute. Often, however, this is not the case and $\mathrm{HH}^n(A, M)$ vanishes for sufficiently large n . To understand this, we need another recipe to compute Hochschild cohomology, which uses projective resolutions. We have already seen a glimpse of this in Section 15.2 when we proved the Morita invariance of $\mathrm{HH}^1(A, A)$. Now we need the full machinery of chain complexes and projective resolutions. The methods developed in this section will be applied to examples in the following sections.

DEFINITION 17.1. A *chain complex* C_\bullet of A -modules is a sequence of A -modules C_n for $n \in \mathbb{Z}$ with *boundary maps* $d_n: C_n \rightarrow C_{n-1}$ such that $d_{n-1} \circ d_n = 0$ for all $n \in \mathbb{Z}$. An n -*cycle* in C_\bullet is $x \in C_n$ with $d_n(x) = 0$, and an n -*boundary* in C_\bullet is an element of the form $d_{n+1}(x)$ with $x \in C_{n+1}$. Two n -cycles are called *homologous* if their difference is an n -boundary. The n th *homology* of a chain complex is the quotient of n -cycles by n -boundaries, $H_n(C_\bullet) := \ker(d_n) / d_{n+1}(C_{n+1})$. A chain complex is called *exact* if all its cycles are boundaries or, equivalently, $H_n(C_\bullet) = 0$ for all $n \in \mathbb{Z}$.

Chain complexes of bimodules are defined similarly.

A *cochain complex* C^\bullet is like a chain complex, except that its *coboundary maps* $d^n: C^n \rightarrow C^{n+1}$ increase degrees. The analogues of cycles, boundaries, and homology for cochain complexes are called *cocycles*, *coboundaries* and *cohomology*, and the n th cohomology is denoted by $H^n(C_\bullet)$.

EXAMPLE 17.2. The Hochschild cochains $A^\bullet \rightarrow M$ form a cochain complex (of vector spaces) with respect to the Hochschild coboundary. Its cohomology is, by definition, the Hochschild cohomology $\mathrm{HH}^n(A, M)$.

DEFINITION 17.3. A *projective resolution* of an A -bimodule M is a chain complex P_\bullet of projective A -bimodules with $P_n = 0$ for $n < 0$ and with an *augmentation map* $d_0: P_0 \rightarrow M$ with $d_0 \circ d_1 = 0$, such that the augmented chain complex

$$\cdots \rightarrow P_3 \rightarrow P_2 \rightarrow P_1 \rightarrow P_0 \rightarrow M \rightarrow 0 \rightarrow 0 \rightarrow \cdots$$

is exact. We often describe a projective resolution by writing down this augmented chain complex, although M is not part of the complex P_\bullet .

View an algebra A as a bimodule over itself in the usual way. The main result of this section asserts that we can compute Hochschild cohomology using projective bimodule resolutions of A :

THEOREM 17.4. *Let A be a unital algebra and let M be an A -bimodule. Let $P_\bullet \rightarrow A$ be a projective resolution of A by unital A -bimodules, with boundary maps d_n^P . Let $\mathrm{Hom}_{A,A}(P_\bullet, M)$ be the cochain complex with A -bimodule maps $f: P_n \rightarrow M$ as n -cochains and the coboundary map $d^n(f) := (-1)^{n+1} f \circ d_{n+1}^P$. Then*

$$\mathrm{HH}^n(A, M) \cong H^n(\mathrm{Hom}_{A,A}(P_\bullet, M)).$$

The sign $(-1)^{n+1}$ in d^n has no effect on the cohomology. It is there to ensure consistency with a more general construction in Definition 17.8.

The proof requires some preparation. First, we exhibit a standard projective bimodule resolution of A called the *bar resolution* and observe that the cochain complex $\text{Hom}_{A,A}(P_\bullet, M)$ for the bar resolution is the Hochschild cochain complex for A with coefficients in M . Then we show that projective bimodule resolutions are unique in a certain sense and that equivalent projective resolutions produce the same cohomology $H^n(\text{Hom}_{A,A}(P_\bullet, M))$. Hence any projective bimodule resolution of A produces the same cohomology as the bar resolution, and we are done.

First for the *bar resolution*. Its definition is explicit but complicated, and its boundary map looks somewhat like the Hochschild coboundary map. Let $\bar{A} := A/\mathbb{C} \cdot 1$ and let $\text{Bar}_n(A) := A \otimes \bar{A}^{\otimes n} \otimes A$ for $n \geq 0$; here it is understood that $\text{Bar}_0(A) := A \otimes A$. These are free and hence projective A -bimodules (see Example 15.15). We define the boundary map $b': \text{Bar}_n(A) \rightarrow \text{Bar}_{n-1}(A)$ by

$$b'(a_0 \otimes \cdots \otimes a_{n+1}) := \sum_{j=0}^n (-1)^j a_0 \otimes \cdots \otimes a_{j-1} \otimes a_j \cdot a_{j+1} \otimes a_{j+2} \otimes \cdots \otimes a_{n+1}.$$

We must check that this is well defined, that is, $b'(a_0 \otimes \cdots \otimes a_{n+1}) = 0$ if $a_j = 1$ for some $j \in \{1, \dots, n\}$. This is because if $a_j = 1$, then the $(j-1)$ th and j th summands cancel, and all the other summands still contain 1 as an entry. We augment the bar resolution by the multiplication map $b': A \otimes A \rightarrow A$, $a_0 \otimes a_1 \mapsto a_0 \cdot a_1$.

Next we must check that b' is a chain complex, that is, $b' \circ b' = 0$. This composite boundary map sends the monomial $a_0 \otimes \cdots \otimes a_{n+1}$ to a sum of monomials of the form $a_0 \otimes \cdots \otimes a_j a_{j+1} \otimes \cdots \otimes a_k a_{k+1} \otimes \cdots \otimes a_n$ or $a_0 \otimes \cdots \otimes a_j a_{j+1} a_{j+2} \otimes \cdots \otimes a_n$. Each such term appears exactly twice with opposite signs. So the sum vanishes.

To prove exactness, we use what is called a *contracting homotopy*. They will reappear soon when we study the equivalence of different resolutions. We define maps $s_n: \text{Bar}_n(A) \rightarrow \text{Bar}_{n+1}(A)$ by

$$(17.5) \quad s_n(a_0 \otimes \cdots \otimes a_{n+1}) := 1 \otimes a_0 \otimes \cdots \otimes a_{n+1};$$

this is understood to vanish if $a_0 = 1$. So $s_n \circ s_{n-1} = 0$. We also define $s_{-1}: A \rightarrow A \otimes A$ by the same formula, $s_{-1}(a_0) := 1 \otimes a_0$. When we compare $s_{n-1} \circ b'$ and $b' \circ s_n$, we notice that all the terms in $s_{n-1} \circ b'(a_0 \otimes \cdots \otimes a_{n+1})$ appear in $b' \circ s_n$ as well, with opposite signs. This shows that

$$(s_{n-1} \circ b' + b' \circ s_n)(a_0 \otimes \cdots \otimes a_{n+1}) = a_0 \otimes \cdots \otimes a_{n+1}.$$

LEMMA 17.6. *Let C_\bullet be a chain complex with boundary map d and let $s_n: C_n \rightarrow C_{n+1}$ be maps with $s_{n-1}d_n + d_{n+1}s_n = \text{Id}_{C_n}$ for all $n \in \mathbb{Z}$. Then C_\bullet is exact.*

PROOF. Let $x \in \ker(d_n)$ be an n -cycle and let $y := s_n(x)$. Then

$$d_{n+1}(y) = d_{n+1}s_n(x) = (\text{Id} - s_{n-1}d_n)(x) = x$$

because $d_n(x) = 0$. Thus any cycle in C_\bullet is a boundary. \square

The lemma shows that the augmented chain complex $\text{Bar}_\bullet(A) \rightarrow A$ is exact. Thus the bar resolution is a projective bimodule resolution of A as claimed.

Now let M be another A -bimodule. Since $\text{Bar}_n(A)$ is the free A -bimodule on $\bar{A}^{\otimes n}$, Example 15.15 shows that

$$\text{Hom}_{A,A}(\text{Bar}_n(A), M) \cong \text{Hom}(\bar{A}^{\otimes n}, M)$$

is isomorphic to the space of n -linear maps $\bar{A}^n \rightarrow M$, that is, normalised Hochschild cochains. An n -linear map $\omega: \bar{A}^n \rightarrow M$ induces the A -bimodule homomorphism

$$\bar{\omega}: \text{Bar}_n(A) \rightarrow M, \quad a_0 \otimes \cdots \otimes a_{n+1} \mapsto a_0 \cdot \omega(a_1, \dots, a_n) \cdot a_{n+1}.$$

Its coboundary $\bar{\omega} \circ b'$ maps

$$a_0 \otimes \cdots \otimes a_{n+2} \mapsto a_0 a_1 \cdot \omega(a_2, \dots, a_{n+1}) \cdot a_{n+2} - a_0 \cdot \omega(a_1 a_2, \dots, a_{n+1}) \cdot a_{n+2} \\ \pm \cdots + (-1)^{n+1} a_0 \cdot \omega(a_1, \dots, a_n) \cdot a_{n+1} a_{n+2}.$$

This is the A -bimodule homomorphism associated to the Hochschild coboundary of ω . Thus $\text{Hom}_{A,A}(\text{Bar}_\bullet(A), M)$ is naturally isomorphic to the normalised Hochschild cochain complex for A with coefficients in M . The proof that the bar resolution is a resolution also reproves Lemma 16.18.

This completes the first step of the proof of Theorem 17.4: we have found one projective A -bimodule resolution P_\bullet of A with $\text{HH}^n(A, M) \cong \text{H}^n(\text{Hom}_{A,A}(P_\bullet, M))$. It remains to prove that all other projective bimodule resolutions yield the same cohomology. For this, we must first define an equivalence relation on chain complexes. This is an algebraic analogue of homotopy equivalence for topological spaces. Thus we first define a concept of homotopy for chain maps, which then leads to the concept of homotopy equivalence for chain complexes.

DEFINITION 17.7. A *chain map* between two chain complexes of A -modules C_\bullet and D_\bullet is a sequence of A -module maps $f_n: C_n \rightarrow D_n$ with $f_n \circ d^C = d^D \circ f_{n+1}$ for all $n \in \mathbb{Z}$. That is, the following diagram commutes:

$$\begin{array}{ccccccccccc} \cdots & \longrightarrow & C_{n+2} & \xrightarrow{d_{n+2}^C} & C_{n+1} & \xrightarrow{d_{n+1}^C} & C_n & \xrightarrow{d_n^C} & C_{n-1} & \longrightarrow & \cdots \\ & & \downarrow f_{n+2} & & \downarrow f_{n+1} & & \downarrow f_n & & \downarrow f_{n-1} & & \\ \cdots & \longrightarrow & D_{n+2} & \xrightarrow{d_{n+2}^D} & D_{n+1} & \xrightarrow{d_{n+1}^D} & D_n & \xrightarrow{d_n^D} & D_{n-1} & \longrightarrow & \cdots \end{array}$$

We may view chain maps as the 0-cycles of a chain complex:

DEFINITION 17.8. Let C_\bullet and D_\bullet be chain complexes. Let $\text{Hom}(C_\bullet, D_\bullet)$ be the chain complex whose k -chains for $k \in \mathbb{Z}$ are arbitrary sequences of maps $f_n: C_n \rightarrow D_{n+k}$ and whose boundary map maps (f_n) to the sequence of maps

$$d_{n+k}^D \circ f_n - (-1)^k f_{n-1} \circ d_n^C: C_n \rightarrow D_{n+k-1}.$$

The signs are dictated by the Koszul sign rule and ensure that $\text{Hom}(C_\bullet, D_\bullet)$ is a chain complex, that is, $d \circ d = 0$. By definition, the 0-cycles of $\text{Hom}(C_\bullet, D_\bullet)$ are the chain maps $C_\bullet \rightarrow D_\bullet$. The construction in Theorem 17.4 is a special case of this. In this case, the formula simplifies because the boundary map in D_\bullet vanishes.

DEFINITION 17.9. Two chain maps $f_\bullet, g_\bullet: C_\bullet \rightarrow D_\bullet$ are *chain homotopic* if they are homologous in $\text{Hom}(C_\bullet, D_\bullet)$, that is, there is a sequence of maps $s_n: C_n \rightarrow D_{n+1}$ with $d_{n+1}^D \circ s_n + s_{n-1} \circ d_n^C = f_n - g_n$. This sequence (s_n) is also called a *chain homotopy* between f_\bullet and g_\bullet .

A chain map is called *null homotopic* if it is homotopic to the zero map. A chain complex is called *contractible* if its identity map is homotopic to the zero map. A 1-chain s of $\text{Hom}(C_\bullet, C_\bullet)$ with $d(s) = \text{Id}_{C_\bullet}$ is called a *contracting homotopy* of C_\bullet .

For example, the maps $s_n: \text{Bar}_n(A) \rightarrow \text{Bar}_{n+1}(A)$ in (17.5) form a contracting homotopy of the bar complex. Lemma 17.6 shows that a contractible chain complex is exact. More generally:

LEMMA 17.10. *Let $f, g: C_\bullet \rightarrow D_\bullet$ be two homotopic chain maps. Then f and g induce the same map on homology.*

PROOF. Let $h: C_\bullet \rightarrow D_\bullet$ be a chain homotopy between f and g . Let $x \in C_n$ satisfy $d(x) = 0$. Then $dh(x) = dh(x) - hd(x) = g(x) - f(x)$. Thus $g(x)$ and $f(x)$ are homologous. \square

DEFINITION 17.11. A chain map $f_\bullet: C_\bullet \rightarrow D_\bullet$ is called a (chain) *homotopy equivalence* if there is a chain map $g: D_\bullet \rightarrow C_\bullet$ such that $f_\bullet \circ g_\bullet$ and $g_\bullet \circ f_\bullet$ are homotopic to the identity maps on D_\bullet and C_\bullet , respectively. Two chain complexes are called (chain) *homotopy equivalent* if such a chain homotopy exists.

EXAMPLE 17.12. A chain complex is chain homotopy equivalent to the zero complex if and only if it is contractible. In this case, the chain maps f_\bullet and g_\bullet and the chain homotopy on the zero complex are forced to be the zero map. So only the contracting homotopy remains.

For chain complexes of A -modules or A -bimodules, we restrict attention to chain maps that consist of A -module or A -bimodule homomorphisms. This leads to genuinely different concepts of homotopic chain maps and homotopy equivalence for chain complexes. We denote the chain complex of A -bimodule maps from C_\bullet to D_\bullet by $\text{Hom}_{A,A}(C_\bullet, D_\bullet)$, and similarly for A -module maps.

PROPOSITION 17.13. *Let $\varphi_\bullet: C_\bullet \rightarrow D_\bullet$ be a homotopy equivalence between two chain complexes of A -bimodules and let F be an additive functor on the category of A -bimodules. “Additive” means that $F(f_1 + f_2) = F(f_1) + F(f_2)$ if $f_1, f_2: M \rightarrow N$. Then $F(\varphi_\bullet)$ induces an isomorphism on homology, $H_n(F(C_\bullet)) \rightarrow H_n(F(D_\bullet))$.*

PROOF. Since $\varphi_\bullet: C_\bullet \rightarrow D_\bullet$ is a homotopy equivalence, there are chain maps $\psi_\bullet: D_\bullet \rightarrow C_\bullet$ and chain homotopies $\psi_\bullet \circ \varphi_\bullet \sim \text{Id}_{C_\bullet}$ and $\varphi_\bullet \circ \psi_\bullet \sim \text{Id}_{D_\bullet}$. Since F is additive, it preserves homotopy of chain maps: if s is a chain homotopy between f and g , then $F(s)$ is a chain homotopy between $F(f)$ and $F(g)$. Therefore, the chain maps $F(\varphi_\bullet)$ and $F(\psi_\bullet)$ are inverse to each other up to chain homotopy. Then Lemma 17.10 shows that the maps on homology induced by them are inverse to each other. \square

Since the maps s in (17.5) form a contracting homotopy for the bar resolution, it may seem that Proposition 17.13 implies that $\text{Hom}_{A,A}(\text{Bar}_\bullet(A), M)$ is again contractible, so that $\text{HH}^*(A, M)$ vanishes identically. This is clearly false. Proposition 17.13 does not apply because the maps s_n are not bimodule homomorphisms.

THEOREM 17.14. *Let M and M' be two A -bimodules and let $P_\bullet \rightarrow M$ and $P'_\bullet \rightarrow M'$ be projective A -bimodule resolutions. Then any bimodule homomorphism $f: M \rightarrow M'$ lifts to a bimodule homomorphism chain map $P_\bullet \rightarrow P'_\bullet$; this lifting is unique up to chain homotopy. In symbols:*

$$H_0(\text{Hom}_{A,A}(P_\bullet, P'_\bullet)) \cong \text{Hom}_{A,A}(M, M').$$

PROOF. Let $f: M \rightarrow M'$ be a bimodule homomorphism. Since P_0 is projective and $d'_0: P'_0 \rightarrow M'$ is surjective, the map $P_0 \xrightarrow{d_0} M \xrightarrow{f} M'$ lifts to a bimodule map $\bar{f}_0: P_0 \rightarrow P'_0$. It restricts to $f_1: \ker d_0 \rightarrow \ker d'_0$. We may view $(P_n)_{n \geq 1}$ as a projective bimodule resolution of $\ker d_0$, and similarly for $(P'_n)_{n \geq 1}$. Thus we may repeat our construction and lift f_1 to $\bar{f}_1: P_1 \rightarrow P'_1$, which again restricts to $f_1: \ker d_1 \rightarrow \ker d'_1$. Repeating the same construction over and over again, we construct a chain map $\bar{f}_\bullet: P_\bullet \rightarrow P'_\bullet$.

Let $f_\bullet, g_\bullet: P_\bullet \rightarrow P'_\bullet$ be two chain maps lifting the same map $M \rightarrow M'$. Then $f_\bullet - g_\bullet$ lifts the zero map. Thus we must show that a chain map $\bar{f}_\bullet: P_\bullet \rightarrow P'_\bullet$ that lifts the zero map is null homotopic. Again we check this by repeating a simple step. Since \bar{f}_\bullet lifts the zero map, \bar{f}_0 maps P_0 to $\ker d'_0$. Since P_0 is projective and $d_1: P_1 \rightarrow \ker d'_0$ is surjective, we may lift this to a map $s_0: P_0 \rightarrow P'_1$ with $d'_1 \circ s_0 = \bar{f}_0$. Subtracting the boundary of s_0 , we find that \bar{f}_0 is homotopic to a chain map \bar{f}'_0 with $\bar{f}'_0 = 0$. Then $(\bar{f}'_n)_{n \geq 1}$ is a chain map between the truncations of P_\bullet and P'_\bullet that lifts the zero map $\ker d_0 \rightarrow \ker d'_0$. We may repeat this step over and over again. This recursively builds a chain homotopy s_n between \bar{f}_\bullet and 0. \square

COROLLARY 17.15. *All projective A -bimodule resolutions of an A -bimodule M are homotopy equivalent as chain complexes of A -bimodules.*

PROOF. Let $P_\bullet \rightarrow M$ and $P'_\bullet \rightarrow M$ be two projective bimodule resolutions. Let $f_\bullet: P_\bullet \rightarrow P'_\bullet$ and $g_\bullet: P'_\bullet \rightarrow P_\bullet$ lift the identity map on M . Then $f_\bullet \circ g_\bullet$ and $g_\bullet \circ f_\bullet$ lift the identity map on M . So they are homotopic to the identity map by the uniqueness part of Theorem 17.14. \square

Corollary 17.15 and Proposition 17.13 show that $H^n(\text{Hom}_{A,A}(P_\bullet, M))$ does not depend on the choice of a projective bimodule resolution P_\bullet . This finishes the proof of Theorem 17.4. There is still one left-over business to finish, namely, the proof of Lemma 16.20. It says that the normalised and unnormalised Hochschild cochains define the the same cohomology.

PROOF OF LEMMA 16.20. The bar resolution defined above may also be called the *normalised* bar resolution because it uses \bar{A} in the interior tensor factors. Exactly the same formulas work when we replace \bar{A} by A everywhere above. This also defines a projective bimodule resolution of A . These two variants of the bar resolution are homotopy equivalent by Corollary 17.15. Therefore, they define the same cohomology. This is what Lemma 16.20 says. \square

EXAMPLE 17.16. As a first application of Theorem 17.4, we consider the case when the algebra A is a semi-simple, finite-dimensional algebra. Then A itself is projective as an A -bimodule. This gives a very short projective bimodule resolution with $P_n = 0$ for $n \geq 1$ and $P_0 = A$. We conclude that $\text{HH}^n(A, M)$ is the cohomology of the chain complex

$$\cdots 0 \rightarrow 0 \rightarrow 0 \rightarrow \text{Hom}_{A,A}(A, M) \rightarrow 0 \rightarrow 0 \rightarrow \cdots .$$

That is, $\text{HH}^n(A, M) = 0$ for $n \geq 1$ and $\text{HH}^0(A, M) \cong \text{Hom}_{A,A}(A, M)$; recall that this is isomorphic to the centre of M .

EXERCISE 17.17. Since $\mathbb{M}_n\mathbb{C}$ is a projective bimodule over itself, the multiplication map $\mathbb{M}_n\mathbb{C} \otimes \mathbb{M}_n\mathbb{C} \rightarrow \mathbb{M}_n\mathbb{C}$ must split by a bimodule homomorphism. Describe such a section explicitly. This proves once again that $\mathbb{M}_n\mathbb{C}$ is a projective $\mathbb{M}_n\mathbb{C}$ -bimodule because it realises $\mathbb{M}_n\mathbb{C}$ as a direct summand in the free bimodule $\mathbb{M}_n\mathbb{C} \otimes \mathbb{M}_n\mathbb{C}$, and direct summands and direct sums of projective bimodules remain projective.

18. Hochschild cohomology of algebras of polynomials

In this section, we compute the Hochschild cohomology for polynomial algebras. The main issue here is to find a small projective bimodule resolution. This bimodule resolution also has other uses. In particular, we use it to define the Taylor spectrum of an n -tuple of commuting operators on a vector space. Our computation for polynomial algebras is a special case of a more general result for smooth affine varieties over fields of characteristic 0, called the Hochschild–Kostant–Rosenberg Theorem. We do not discuss this here because we do not want to go that far into algebraic geometry. Instead, we will turn to algebras of smooth functions on manifolds in Section 19. In that case, we will get similar results and use similar resolutions, but the definitions of Hochschild cohomology and projective resolution must be changed a bit for this to come out.

18.1. Koszul resolutions. Koszul resolutions provide a general recipe for projective bimodule resolutions, which works for several algebras such as polynomial algebras or the Weyl algebra. The main ingredient are natural maps between exterior powers of a vector space.

Let V be a finite-dimensional \mathbb{C} -vector space and let $\Lambda^k V$ be its k th exterior power. This is understood to be \mathbb{C} for $k = 0$, and it vanishes for $k < 0$ or $k > \dim V$. Let b_1, \dots, b_n be a basis for V . Then the monomials $b_{i_1} \wedge \dots \wedge b_{i_k}$ with $1 \leq i_1 < i_2 < \dots < i_k \leq n$ form a basis of $\Lambda^k V$. So the latter vector space has dimension $\binom{n}{k}$. Let $V^* := \text{Hom}(V, \mathbb{C})$ be the dual space of V and let $\eta \in V^*$. This induces linear maps

$$i(\eta): \Lambda^k V \rightarrow \Lambda^{k-1} V,$$

$$v_1 \wedge \dots \wedge v_k \mapsto \sum_{j=1}^k (-1)^{j-1} \eta(v_j) v_1 \wedge \dots \wedge v_{j-1} \wedge v_{j+1} \wedge \dots \wedge v_k.$$

For $k = 1$, this means that $i(\eta): V \cong \Lambda^1 V \rightarrow \Lambda^0 V \cong \mathbb{C}$ maps $v \mapsto \eta(v)$. Notice that this map is well defined.

The antisymmetry of the \wedge -product implies $i(\eta) \circ i(\eta) = 0$. So

$$0 \rightarrow \Lambda^n V \xrightarrow{i(\eta)} \Lambda^{n-1} V \xrightarrow{i(\eta)} \Lambda^{n-2} V \rightarrow \dots \rightarrow \Lambda^1 V \xrightarrow{i(\eta)} \Lambda^0 V \rightarrow 0$$

is a chain complex. We call it the *Koszul complex* for V and η . If $\eta = 0$, then $i(\eta) = 0$ as well. For all other η , the Koszul complex is contractible:

LEMMA 18.1. *Let $v \in V$ satisfy $\eta(v) = 1$. Then the maps*

$$\lambda_v: \Lambda^k V \rightarrow \Lambda^{k+1} V, \quad \omega \mapsto v \wedge \omega,$$

form a contracting homotopy of the Koszul complex, that is, $i(\eta) \circ \lambda_v + \lambda_v \circ i(\eta) = \text{Id}$.

PROOF. Most terms in $i(\eta)(v \wedge v_1 \wedge \dots \wedge v_k)$ also appear in $v \wedge i(\eta)(v_1 \wedge \dots \wedge v_k)$ with the opposite sign. The only term in $i(\eta)(v \wedge v_1 \wedge \dots \wedge v_k) + v \wedge i(\eta)(v_1 \wedge \dots \wedge v_k)$ that survives this cancellation is $\eta(v)v_1 \wedge \dots \wedge v_n = v_1 \wedge \dots \wedge v_n$. \square

Thus the Koszul complex is contractible in a very explicit way whenever its boundary map is non-zero. This allows to compute the homology of chain complexes constructed out of the Koszul complex. Our first variant of the Koszul complex views the vector η as a free parameter and varies it.

Thus we consider $A = \mathbb{C}[\eta_1, \dots, \eta_n]$ and $V = \mathbb{C}^n$. We define the chain complex (C_\bullet, d_\bullet) with

$$C_k := A \otimes \Lambda^k V = \mathbb{C}[\eta_1, \dots, \eta_n] \otimes \Lambda^k \mathbb{C}^n,$$

which we view as the space of polynomial functions $\mathbb{C}^n \rightarrow \Lambda^k \mathbb{C}^n$, and

$$d_k: C_k \rightarrow C_{k-1}, \quad d_k(f \otimes \omega)(\eta) := f(\eta) i(\eta)(\omega)$$

for all $f \in A$, $\omega \in \Lambda^k V$, $\eta \in \mathbb{C}^n$. More explicitly, if e_1, \dots, e_n is the standard basis of \mathbb{C}^n , then

$$d_k(f \otimes \omega) = \sum_{j=1}^n f \cdot \eta_j \otimes i(e_j)(\omega).$$

We augment this chain complex by the map $d_0: A \cong C_0 \rightarrow \mathbb{C}$, $f \mapsto f(0)$, which clearly satisfies $d_0 \circ d_1 = 0$.

PROPOSITION 18.2. *The chain complex (C_\bullet, d_\bullet) above is a free A -module resolution of \mathbb{C} with the A -module structure $f \cdot x := f(0) \cdot x$ for all $f \in A$, $x \in \mathbb{C}$.*

We will use the following lemma to show that the would-be resolution is exact:

LEMMA 18.3. *Let C_\bullet be a chain complex and let $C_\bullet^{(k)}$ be an increasing filtration by subcomplexes, that is, $\partial(C_\bullet^{(k)}) \subseteq C_\bullet^{(k)}$ for all $k \in \mathbb{N}$. Assume that $\bigcup C_\bullet^{(k)} = C_\bullet$ and $C_\bullet^{(0)} = \{0\}$. If the quotient complexes $C_\bullet^{(k+1)} / C_\bullet^{(k)}$ are exact for all $k \in \mathbb{N}$, then C_\bullet is exact.*

PROOF. We prove by induction that $C_{\bullet}^{(k)}$ is exact for all $k \in \mathbb{N}$. This is trivial for $k = 0$. Assume the statement is true for k . Then we prove it for $k + 1$. There is a short exact sequence of chain complexes

$$C_{\bullet}^{(k)} \twoheadrightarrow C_{\bullet}^{(k+1)} \rightarrow C_{\bullet}^{(k+1)}/C_{\bullet}^{(k)}.$$

By assumption, $C_{\bullet}^{(k)}$ and $C_{\bullet}^{(k+1)}/C_{\bullet}^{(k)}$ are exact. We must show that $C_{\bullet}^{(k+1)}$ is exact. Take $x \in C_n^{(k+1)}$ with $d(x) = 0$. The image of x in $C_n^{(k+1)}/C_n^{(k)}$ is a boundary because this complex is exact. Thus there is $y \in C_{n+1}^{(k+1)}$ with $d(y) \equiv x \pmod{C_n^{(k)}}$. That is, $x - d(y) \in C_n^{(k)}$. Since $d(x - d(y)) = 0$ and $C_{\bullet}^{(k)}$ is exact, there is $z \in C_{n+1}^{(k)} \subseteq C_{n+1}^{(k+1)}$ with $x - d(y) = d(z)$. Thus $x = d(y + z)$.

Next we prove that C_{\bullet} is exact. Since $\bigcup C_{\bullet}^{(k)} = C_{\bullet}$, any $x \in C_{\bullet}$ with $d(x) = 0$ already belongs to $C_{\bullet}^{(k)}$ for some $k \in \mathbb{N}$. Since $C_{\bullet}^{(k)}$ is exact, there is $y \in C_{\bullet}^{(k)} \subseteq C_{\bullet}$ with $d(y) = x$. \square

PROOF OF PROPOSITION 18.2. It is clear that $(C_{\bullet}, d_{\bullet})$ is a chain complex of free A -modules and that d_0 is an A -module map. It remains to check that this chain complex is exact. For this, we use a filtration by subcomplexes. Namely, let $C_{\bullet}^{(k)}$ denote the complex above for $\mathbb{C}[\eta_1, \dots, \eta_k]$, that is, for polynomials in k variables. Embedding $\mathbb{C}[\eta_1, \dots, \eta_k] \subseteq \mathbb{C}[\eta_1, \dots, \eta_n]$ and $\Lambda^*(\mathbb{C}^k) \subseteq \Lambda^*(\mathbb{C}^n)$ in the obvious way for $k \leq n$, the complexes $C_{\bullet}^{(k)}$ for $k = 0, \dots, n$ become subcomplexes of $C_{\bullet}^{(n)}$. Now we want to use Lemma 18.3 to prove that the augmented complex $C_{\bullet}^{(n)}$ is exact. The complex $C_{\bullet}^{(0)}$ is just $\mathbb{C} \otimes \mathbb{C}$, the constant functions with values in $\Lambda^0(\mathbb{C}^n)$. When we augment it by the restriction of the augmentation map $C_0^{(n)} \rightarrow \mathbb{C}$ above, the augmented complex $C_{\bullet}^{(0)} \rightarrow \mathbb{C}$ becomes exact, as needed. The subquotients $C_{\bullet}^{(k+1)}/C_{\bullet}^{(k)}$ are not changed when we augment both complexes. We may describe this subquotient by decomposing polynomials and differential forms as follows:

$$\begin{aligned} \Lambda^* \mathbb{C}^{k+1} &\cong \Lambda^* \mathbb{C}^k \oplus (\Lambda^* \mathbb{C}^k \wedge e_{k+1}), \\ \mathbb{C}[\eta_1, \dots, \eta_{k+1}] &\cong \mathbb{C}[\eta_1, \dots, \eta_k] \oplus \eta_{k+1} \cdot \mathbb{C}[\eta_1, \dots, \eta_{k+1}]. \end{aligned}$$

As a result, the following subspace of $C_{\bullet}^{(k+1)}$ is a complement for $C_{\bullet}^{(k)}$:

$$\mathbb{C}[\eta_1, \dots, \eta_{k+1}] \otimes (\Lambda^* \mathbb{C}^k \wedge e_{k+1}) \oplus \mathbb{C}[\eta_1, \dots, \eta_{k+1}] \cdot \eta_{k+1} \otimes \Lambda^* \mathbb{C}^k.$$

We claim that this subquotient complex is contractible. Namely, we define the contracting homotopy to be 0 on the direct summand $\mathbb{C}[\eta_1, \dots, \eta_{k+1}] \otimes (\Lambda^* \mathbb{C}^k \wedge e_{k+1})$ and by $s^{(k+1)}(f \cdot \eta_{k+1} \otimes \omega) := f \otimes e_{k+1} \wedge \omega$ on the other direct summand. A computation shows that $d \circ s^{(k+1)} + s^{(k+1)} \circ d$ is the identity map on this subquotient complex. The computation is very easy for $\omega \in \mathbb{C}[\eta_1, \dots, \eta_{k+1}] \otimes (\Lambda^* \mathbb{C}^k \wedge e_{k+1})$ because $s^{(k+1)}$ vanishes there and all but one summand in $d(\omega)$ again belongs to the first summand. If $\omega \in \mathbb{C}[\eta_1, \dots, \eta_{k+1}] \cdot \eta_{k+1} \otimes \Lambda^* \mathbb{C}^k$, then all summands in $d(\omega)$ again belong to this summand. Then all terms in $s^{(k+1)} \circ d$ also occur in $d \circ s^{(k+1)}$ with an opposite sign, and the latter has one extra summand equal to ω . \square

One may use the the subcomplexes and the explicit contracting homotopies in the proof of Proposition 18.2 to write down an explicit contracting homotopy for the resolution. This homotopy is, however, rather complicated. It is only linear and not A -linear. This is to be expected because if a projective resolution of a module M has an A -linear contracting homotopy, then M must be projective by Lemma 15.12.

Next we modify the above construction to get a projective A -bimodule resolution of A . First, we may tensor C_k with another copy of A to get a chain complex of free A -bimodules $(C_k \otimes A, d_k \otimes \text{Id}_A)$. This is a resolution of A , but with the bimodule structure $f_1 \cdot f_2 \cdot f_3 = f_1(0) \cdot f_2 f_3$, which is not what we want. To remedy this, we

change coordinates from (x, y) to $(x - y, y)$. This still gives a projective bimodule resolution, and now we get the right bimodule structure on A .

THEOREM 18.4. *Let $A := \mathbb{C}[x_1, \dots, x_n]$. Let $C_k := A \otimes A \otimes \Lambda^k \mathbb{C}^n$ with the usual free bimodule structure and define*

$$d_k: A \otimes A \otimes \Lambda^k \mathbb{C}^n \rightarrow A \otimes A \otimes \Lambda^{k-1} \mathbb{C}^n, \quad d_k(f \otimes \omega)(x, y) := f(x, y) \otimes i(x - y)(\omega)$$

for all $f \in A \otimes A$, $\omega \in \Lambda^k \mathbb{C}^n$, $x, y \in \mathbb{C}^n$. With the augmentation map $d_0(f)(x) := f(x, x)$, which corresponds to the multiplication map $A \otimes A \rightarrow A$, this is a free A -bimodule resolution of A .

PROOF. This chain complex is isomorphic to the tensor product of C_\bullet with A and hence a resolution as desired. \square

We can use this resolution to compute the Hochschild cohomology for the polynomial algebra $A := \mathbb{C}[x_1, \dots, x_n]$. Let M be an A -bimodule. Then

$$\mathrm{Hom}_{A,A}(A \otimes A, M) \cong M \otimes \Lambda^k \mathbb{C}^n.$$

The induced boundary map $M \otimes \Lambda^k \mathbb{C}^n \rightarrow M \otimes \Lambda^{k+1} \mathbb{C}^n$ is

$$d_k^*(m \otimes \omega) = \sum_{j=1}^n \mathrm{ad}_{x_j}(m) \otimes i(e_j)(\omega).$$

In particular, for $M = A$ itself, this boundary map vanishes and we get

$$\mathrm{HH}^k(A, A) \cong A \otimes \Lambda^k \mathbb{C}^n.$$

We may view this as the space of polynomial k -vector fields on \mathbb{C}^n . Our next goal is to generalise this to algebras of smooth functions on smooth manifolds.

18.2. The Taylor spectrum of several commuting linear operators. We use Koszul resolutions to define a spectrum for several commuting operators.

DEFINITION 18.5. The *spectrum* of a single linear operator T on a \mathbb{C} -vector space V consists of all $\lambda \in \mathbb{C}$ for which $T - \lambda$ is *not* invertible.

Now consider an n -tuple $T := (T_1, \dots, T_n)$ of commuting linear operators on V . We would like to define a reasonable spectrum of T , which should be a subset of \mathbb{C}^n . For instance, we expect that the spectrum is the set of all joint eigenvalues of T if V is finite-dimensional. We call $\lambda = (\lambda_j) \in \mathbb{C}^n$ a *joint eigenvalue* of T if there is a vector $v \in V \setminus \{0\}$, called *joint eigenvector*, with $T_j v = \lambda_j v$ for $j = 1, \dots, n$. The following definition of a spectrum for linear operators on Banach spaces is due to Joseph L. Taylor [14].

In the single operator case, $T - \lambda$ is invertible if and only if the cochain complex

$$(18.6) \quad 0 \rightarrow V \xrightarrow{T-\lambda} V \rightarrow 0$$

is exact. We may view $T - \lambda$ as the generator of a $\mathbb{C}[t]$ -module structure on V . The cochain complex in (18.6) is $\mathrm{Hom}_{\mathbb{C}[t]}(P_\bullet, V)$ for the projective $\mathbb{C}[t]$ -module resolution

$$P_\bullet := (0 \rightarrow \mathbb{C}[t] \xrightarrow{t} \mathbb{C}[t] \xrightarrow{\mathrm{ev}_0} \mathbb{C}).$$

Thus the spectrum of T is related to a projective $\mathbb{C}[t]$ -module resolution of \mathbb{C} with the trivial $\mathbb{C}[t]$ -module structure. This suggests the following generalisation to several commuting operators.

Let $A := \mathbb{C}[t_1, \dots, t_n]$. An n -tuple of commuting linear operators on V is equivalent to an A -module structure on V , where T_j is the operator of multiplication by t_j . Given $\lambda \in \mathbb{C}^n$, we equip V with the A -module structure corresponding to the n -tuple $T - \lambda = (T_j - \lambda_j)$. Let $C_k := A \otimes \Lambda^k \mathbb{C}^n$ and d_k be the free A -module

resolution of \mathbb{C} as in Proposition 18.2. Then $\text{Hom}_A(C_k, V)$ is the cochain complex with entries $L_k := \Lambda^k \mathbb{C}^n \otimes V$ and the boundary map

$$d_k(e_{i_1} \wedge \cdots \wedge e_{i_k} \otimes v) := \sum_{j=1}^k (-1)^{j-1} e_{i_1} \wedge \cdots \wedge e_{i_{j-1}} \wedge e_{i_{j+1}} \wedge \cdots \wedge e_{i_k} \otimes (T_{i_j} - \lambda_{i_j})(v).$$

This chain complex is also called a Koszul complex.

DEFINITION 18.7. The *Taylor spectrum* of an n -tuple of commuting operators is the set of all $\lambda \in \mathbb{C}^n$ for which the Koszul complex above is *not* exact.

By design, this generalises the spectrum of a single operator. Our next goal is to compute the Taylor spectrum in the finite-dimensional case.

PROPOSITION 18.8. *Let $T = (T_1, \dots, T_n)$ be an n -tuple of commuting operators on a vector space V . Let $W \subseteq V$ be T_j -invariant for $j = 1, \dots, n$. Then T induces n -tuples $T|_W$ and $T|_{V/W}$ of commuting operators on W and V/W , respectively. The Taylor spectrum of T is contained in the union of the Taylor spectra of $T|_W$ and $T|_{V/W}$.*

PROOF. The Koszul complex C_W for $T|_W$ is a subcomplex in the Koszul complex C_V for T , and the quotient is isomorphic to the Koszul complex $C_{V/W}$ for $T|_{V/W}$. The proposition says that the complex for T is exact provided those for $T|_W$ and $T|_{V/W}$ are exact. This is a special case of Lemma 18.3, where the filtration is quite short, namely, $0 \subseteq C_W \subseteq C_V$. \square

PROPOSITION 18.9. *If V is finite-dimensional, then the Taylor spectrum of T is the set of joint eigenvalues of T .*

PROOF. First let λ be a joint eigenvalue of T . Then there is a joint eigenvector $v \in V$ with $(T_j - \lambda_j)v = 0$ for $j = 1, \dots, n$. Thus v is in the kernel of the map $d_0: V \rightarrow V^n$ in the Koszul resolution. So λ belongs to the Taylor spectrum. We prove the converse by induction on the dimension of V . The case where $\dim V = 0$ is trivial. If $\dim V = 1$, then the relevant complex is isomorphic to the Koszul complex in Lemma 18.1, and the claim follows from that lemma. Now assume the assertion to be known for all n -tuples of operators on W with $\dim W < \dim V$.

We claim that V has a 1-dimensional T -invariant subspace. To see this, we turn the n -tuple of commuting operators T into a module structure over the polynomial algebra $\mathbb{C}[x_1, \dots, x_n]$. Since V is a finite-dimensional module, it contains a simple submodule. This is proven by induction on the dimension of V . If V is not itself irreducible, it contains a proper invariant submodule, which has lower dimension. Therefore, it contains a simple submodule by induction assumption, and this is a simple submodule of V . Now we already know that all simple modules over commutative algebras are 1-dimensional because all primitive ideals are given by characters. So V has a T -invariant subspace $W \subseteq V$ of dimension 1. This is simply the span of a joint eigenvector $v \in T$ with $T_j v = \tau_j v$ with $\tau_j \in \mathbb{C}$ for $j = 1, \dots, n$. Now we consider $T|_W$ and $T|_{V/W}$.

By assumption, λ is not a joint eigenvalue of T . So $\tau \neq \lambda$. We claim that λ is not a joint eigenvalue of $T|_{V/W}$. Otherwise, there would be $\xi \in V$ with $\xi \notin W$ and $(T_j - \lambda_j)\xi \in W$ for $j = 1, \dots, n$. Then $(T_k - \tau_k)(T_j - \lambda_j)\xi = 0$ for all $j, k \in \{1, \dots, n\}$. So $(T_k - \tau_k)\xi$ is either 0 or a joint eigenvector of T with joint eigenvalue λ . Since λ is not a joint eigenvalue of T , it follows that $(T_k - \tau_k)\xi = 0$ for $k = 1, \dots, n$. This and $(T_j - \lambda_j)\xi \in W$ imply $\tau = \lambda$ or $\xi \in W$, which contradicts our assumptions. So λ is not a joint eigenvalue of $T|_{V/W}$.

Proposition 18.8 shows that the Taylor spectrum of T is contained in the union of the Taylor spectra of $T|_W$ and $T|_{V/W}$. Since these are lower-dimensional,

the induction hypothesis and the last claim imply that neither Taylor spectrum contains λ . Thus λ is not in the Taylor spectrum of T . \square

The following lemma describes another feature of the Taylor spectrum that one would expect of any reasonable definition of a joint spectrum:

LEMMA 18.10. *If $T_j - \lambda_j$ is invertible for some j , then λ is not in the Taylor spectrum of T .*

PROOF. The operator $(T_j - \lambda_j)^{-1}$ commutes with T_2, \dots, T_n . Then we compute as in Lemma 18.1 that $(e_j \wedge \square) \otimes (T_j - \lambda_j)^{-1}$ is a contracting homotopy for the Koszul complex for T and λ . Being contractible, this complex is exact. \square

19. Hochschild cohomology for algebras of smooth functions

We are going to compute the *continuous* Hochschild cohomology of the algebra of smooth functions $A := C^\infty(X)$ on a smooth manifold X . The continuity assumption means that we restrict attention to Hochschild cochains $A^n \rightarrow M$ that are continuous for the canonical topology on A described in Definition 2.15 and some topology on M , which is part of the data. Thus we do homological algebra with topological modules. Along the way, we explain why we cannot compute the purely algebraic Hochschild cohomology of $C^\infty(X)$. Only the continuity assumption makes the theory computable. In addition, we also change our concept of projective resolution. The effect of these two changes in the basic definitions is that the “projective resolutions” that we use to compute the continuous Hochschild cohomology are quite similar to the Koszul resolution that we used for the purely algebraic polynomial algebra.

DEFINITION 19.1. A *k*-vector field is a section of the *k*th exterior power of the tangent bundle of X . That is, $\pi(x) \in \Lambda^k T_x X$ for all $x \in X$.

THEOREM 19.2. *The *k*th continuous Hochschild cohomology of $C^\infty(X)$ with coefficients in $C^\infty(X)$ is naturally isomorphic to the space $\mathfrak{X}^k(X)$ of smooth *k*-vector fields on X . The isomorphism maps a *k*-vector field π to the Hochschild *k*-cocycle*

$$\begin{aligned} \Sigma_\pi : C^\infty(X)^k &\rightarrow C^\infty(X), \\ \Sigma_\pi(f_1, \dots, f_k)(x) &:= \langle \pi(x) \mid Df_1(x) \wedge \dots \wedge Df_k(x) \rangle. \end{aligned}$$

Here $Df_j(x) \in (T_x X)^*$ is the derivative of f_j at x and the pairing is the standard one between $\Lambda^k(T_x X)$ and $\Lambda^k(T_x X)^*$.

We will only prove that there is an isomorphism and omit the further details that describe the isomorphism as in the theorem.

We want to compute the continuous Hochschild cohomology using some kind of projective bimodule resolutions. This is possible, but requires some technical discussion. The main point is to complete tensor products. To understand this, we examine what happens when we replace polynomials by smooth functions on \mathbb{R}^n in the construction of a projective bimodule resolution of $\mathbb{C}[x_1, \dots, x_n]$ in Theorem 18.4.

Proposition 18.2 remains true if we replace $\mathbb{C}[x_1, \dots, x_n] \otimes \Lambda^k \mathbb{C}^n$ by $C_k := C^\infty(\mathbb{R}^n, \Lambda^k \mathbb{C}^n)$. The same proof continues to work. Namely, the subcomplexes for $\mathbb{R}^k \subseteq \mathbb{R}^n$ give a filtration of the complex C_\bullet for \mathbb{R}^n , such that the subquotients of the filtration are contractible with an explicit contracting homotopy, which also turns out to be bounded. So we get a free $C^\infty(\mathbb{R}^n)$ -module resolution of the trivial representation. Then the chain complex $C_\bullet \otimes C^\infty(\mathbb{R}^n)$ is a free bimodule resolution of $C^\infty(\mathbb{R}^n)$ – but with the useless bimodule structure where the left multiplication factors through evaluation at 0. But the change of basis map $(x, y) \mapsto (x - y, y)$ does not act on the algebraic tensor product $C^\infty(\mathbb{R}^n) \otimes C^\infty(\mathbb{R}^n)$. So we do not get a free bimodule resolution of $C^\infty(\mathbb{R}^n)$.

These problems appear because the algebraic tensor product $C^\infty(\mathbb{R}^n) \otimes C^\infty(\mathbb{R}^n)$ is not a nice space of functions. If we *complete* this space to $C^\infty(\mathbb{R}^n \times \mathbb{R}^n)$, everything works exactly as above. But $C^\infty(\mathbb{R}^n \times \mathbb{R}^n)$ is not projective because the diagonal restriction map $C^\infty(\mathbb{R}^n \times \mathbb{R}^n) \rightarrow C^\infty(\mathbb{R}^n)$ does not lift to a bimodule map $C^\infty(\mathbb{R}^n \times \mathbb{R}^n) \rightarrow C^\infty(\mathbb{R}^n) \otimes C^\infty(\mathbb{R}^n)$. We must modify our definitions to make $C^\infty(\mathbb{R}^n \times \mathbb{R}^n)$ projective.

This requires the theory of complete topological tensor products, a rather specialised branch of functional analysis (see [7]). The complete projective topological tensor product of two complete, locally convex topological vector spaces is defined by a universal property. Namely, a linear map from it to another complete, locally convex topological vector space is equivalent to a (jointly) continuous bilinear map. We will mainly use the following special case:

THEOREM 19.3. *Let M and N be smooth manifolds and V a complete locally convex topological vector space. A continuous bilinear map $b: C^\infty(M) \times C^\infty(N) \rightarrow V$ extends uniquely to a continuous linear map $l: C^\infty(M \times N) \rightarrow V$. The complete projective topological tensor product of $C^\infty(M)$ and $C^\infty(N)$ is naturally isomorphic to $C^\infty(M \times N)$.*

The two statements in Theorem 19.3 are equivalent because of the universal property of the complete projective topological tensor product.

Even slightly more is true: if V and W are complete locally convex topological vector spaces, then any continuous trilinear map $C^\infty(X) \times C^\infty(X) \times V \rightarrow W$ extends to a continuous bilinear map $C^\infty(X \times X) \times V \rightarrow W$. Thus a topological $C^\infty(X)$ -bimodule structure on V , that is, a continuous trilinear map $C^\infty(X) \times V \times C^\infty(X) \rightarrow W$, is equivalent to a topological $C^\infty(X \times X)$ -module structure, that is, a continuous bilinear map $C^\infty(X \times X) \times V \rightarrow W$. Thus the category of complete locally convex topological $C^\infty(X)$ -bimodules is equivalent to the category of complete locally convex topological $C^\infty(X \times X)$ -modules. This is exactly what we need: we want to replace $C^\infty(X) \otimes C^\infty(X)$ by $C^\infty(X \times X)$.

In particular, this implies that $C^\infty(\mathbb{R}^n \times \mathbb{R}^n, V)$ for a finite-dimensional vector space V is projective because

$$\mathrm{Hom}_{C^\infty(\mathbb{R}^n \times \mathbb{R}^n)}(C^\infty(\mathbb{R}^n \times \mathbb{R}^n, V), M) \cong \mathrm{Hom}(V, M).$$

It is crucial here that V is finite-dimensional. If V is an infinite-dimensional complete locally convex topological vector space as well, then the functor $\mathrm{Hom}(V, M)$ need not be exact.

To avoid this new problem, we must refine our concept of exact chain complex and resolution: we only allow exact chain complexes and resolutions with a *continuous* contracting homotopy and call these *admissible*. An exact chain complex of vector spaces is automatically contractible. But there are exact chain complexes of topological vector spaces without a *continuous* contracting homotopy. If a chain complex C_\bullet has a continuous contracting homotopy, then the chain complex $\mathrm{Hom}(V, C_\bullet)$ of continuous linear maps $V \rightarrow C_\bullet$ is again contractible, hence exact.

For any smooth manifold X and any complete locally convex topological vector space V , there is a complete locally convex topological vector space $C^\infty(X, V)$ of smooth functions $X \rightarrow V$ with the expected properties. In particular, it is a topological $C^\infty(X)$ -module, and $\mathrm{Hom}_{C^\infty(X)}(C^\infty(X, V), M) \cong \mathrm{Hom}(V, M)$ for any complete topological $C^\infty(X)$ -module M . Hence $C^\infty(X, V)$ is *relatively projective*, that is, the functor $\mathrm{Hom}(C^\infty(X, V), \square)$ maps continuously contractible chain complexes to exact chain complexes. In addition, $C^\infty(X, C^\infty(Y)) \cong C^\infty(X \times Y)$.

After these technical changes, Theorem 17.4 has a nice analogue for locally convex topological bimodules:

THEOREM 19.4. *Let A be a complete, locally convex topological unital algebra and let M be a complete, locally convex topological unital A -module. Let P_\bullet be a chain complex with an augmentation $P_0 \rightarrow A$. Suppose that the bimodules P_n are relatively projective and that the augmented chain complex is admissibly exact. Then $\mathrm{HH}_{\mathrm{cont}}^n(A, M) \cong \mathrm{H}^n(\mathrm{Hom}_{A,A}(P_\bullet, M))$.*

In the case of interest, we find that the chain complex that we get from Theorem 18.4 by replacing $\mathbb{C}[x_1, \dots, x_n, y_1, \dots, y_n]$ by $\mathrm{C}^\infty(\mathbb{R}^n \times \mathbb{R}^n)$ satisfies all requirements and hence may be used to compute the continuous Hochschild cohomology of $\mathrm{C}^\infty(\mathbb{R}^n)$. In particular, for the bimodule $M = \mathrm{C}^\infty(\mathbb{R}^n)$, we get

$$\mathrm{HH}_{\mathrm{cont}}^k(\mathrm{C}^\infty(\mathbb{R}^n), \mathrm{C}^\infty(\mathbb{R}^n)) \cong \mathrm{C}^\infty(\mathbb{R}^n, \Lambda^k(\mathbb{C}^n)).$$

This is a special case of Theorem 19.2.

Now we want to extend this result to the algebra $A := \mathrm{C}^\infty(X)$ of smooth functions on a general smooth manifold X . Since we expect to see k -vector fields in the Hochschild cohomology and since the passage to cohomology involves taking duals, we try to build a free bimodule resolution with $C_k = \mathrm{C}^\infty(X \times X, \Lambda^k \mathrm{T}^* X)$; more precisely, if $f \in C_k$, then $f(x, y) \in \Lambda^k \mathrm{T}_x^* X$, that is, we use the cotangent bundle of X along the first coordinate direction (it would make no difference to use the second coordinate). To define an analogue of the boundary map d_k , we need a vector field $\delta: X \times X \rightarrow \mathrm{TX}$ that associates to $(x, y) \in X \times X$ a tangent vector $\delta(x, y) \in \mathrm{T}_x X$. For \mathbb{R}^n , we take $\delta(x, y) := x - y$. The vector field δ should have certain properties which, unfortunately, cannot always be achieved. This makes the construction somewhat more involved (see also [4]).

We follow an indirect route that avoids writing down a projective bimodule resolution. As a preparation, we consider an easier intermediate case where we replace $\mathrm{C}^\infty(X \times X)$ by $\mathrm{C}^\infty(\mathrm{TX})$ and view $\mathrm{C}^\infty(X)$ as a $\mathrm{C}^\infty(\mathrm{TX})$ -bimodule via restriction to the zero section: $(f_1 \cdot g \cdot f_2)(x) := f_1(x, 0)g(x)f_2(x, 0)$ for all $f_1, f_2 \in \mathrm{C}^\infty(\mathrm{TX})$, $g \in \mathrm{C}^\infty(X)$, $x \in X$. In this case, there is a canonical vector field, the identical vector field that maps $\xi \in \mathrm{T}_x X \subseteq \mathrm{TX}$ to the tangent vector ξ . This leads to a chain complex with $C_k := \mathrm{C}^\infty(\mathrm{TX}, \Lambda^k \mathrm{T}^* X)$ and

$$(19.5) \quad d_k(f)(x, \xi) := i(\xi)(f(x, \xi))$$

for all $f \in \mathrm{C}^\infty(\mathrm{TX}, \Lambda^k \mathrm{T}^* X)$, $x \in X$, $\xi \in \mathrm{T}_x X$.

LEMMA 19.6. *The chain complex (C_k, d_k) is contractible with a continuous contracting homotopy.*

PROOF. If the tangent bundle is trivial, $\mathrm{TX} \cong X \times \mathbb{R}^n$, then the lemma follows exactly as for \mathbb{R}^n . In general, the tangent bundle is only local trivial. So there is an open covering $X = \bigcup_{i \in I} U_i$ such that the formula in the proof of Theorem 18.4 provides a continuous contracting homotopy h_i for the restriction of our complex to U_i . There is a smooth partition of unity $(w_i)_{i \in I}$ subordinate to this covering. Let $w_i h_i$ denote the operator on C_\bullet that first restricts a function on TX to $\mathrm{T}U_i$, then applies h_i , then multiplies with w_i , and then extends by 0 outside U_i . This is well defined because after multiplication with w_i , a function on $\mathrm{T}U_i$ extended by 0 outside $\mathrm{T}U_i$ is a smooth function on TX . Since the multiplication operators w_i commute with the boundary map, we compute $[d, w_i h_i] = w_i [d, h_i] = w_i$. So $[d, \sum w_i h_i] = \sum w_i = 1$. \square

We have now constructed a free $\mathrm{C}^\infty(\mathrm{TX})$ -module resolution of $\mathrm{C}^\infty(X)$. This is exactly what we need for $X = \mathbb{R}^n$ because $\mathrm{T}\mathbb{R}^n \cong \mathbb{R}^n \times \mathbb{R}^n$; more precisely, the passage from $\mathrm{T}\mathbb{R}^n$ to $\mathbb{R}^n \times \mathbb{R}^n$ involves a change of coordinates $(x, \xi) \mapsto (x, x + \xi)$.

In general, TX is not so far from $X \times X$ because of the following special case of the Tubular Neighbourhood Theorem:

THEOREM 19.7. *Let X be a smooth manifold. Then there is a diffeomorphism from TX onto an open subset of $X \times X$ that maps $0 \in \mathrm{T}_x X$ to (x, x) for all $x \in X$.*

Hence $X \times X$ looks like TX near the diagonal $X \subseteq X \times X$. Since the bimodule $\mathrm{C}^\infty(X)$ is supported on this diagonal, it is irrelevant what happens away from the diagonal. This is made precise by the following theorem:

THEOREM 19.8. *View $\mathrm{C}^\infty(\mathrm{TX})$ as a $\mathrm{C}^\infty(X \times X)$ -bimodule using a tubular neighbourhood embedding $\mathrm{TX} \rightarrow X \times X$ (Theorem 19.7). Then the chain complex $\mathrm{C}^\infty(\mathrm{TX}, \Lambda^k \mathrm{T}^* X)$ with the boundary map d_k in (19.5) is homotopy equivalent, as a chain complex of topological $\mathrm{C}^\infty(X \times X)$ -modules, to a relatively projective, admissible $\mathrm{C}^\infty(X \times X)$ -module resolution of $\mathrm{C}^\infty(X)$.*

This means that we may use the resolution constructed above to compute the Hochschild cohomology of $\mathrm{C}^\infty(X)$ although it is not relatively projective.

PROOF. We may view $\mathrm{C}^\infty(\mathrm{TX})$ -modules as $\mathrm{C}^\infty(X \times X)$ -modules using the open embedding $\mathrm{TX} \rightarrow X \times X$, which induces an algebra homomorphism $\mathrm{C}^\infty(X \times X) \rightarrow \mathrm{C}^\infty(\mathrm{TX})$. We must show that some projective $\mathrm{C}^\infty(X \times X)$ -module resolution of $\mathrm{C}^\infty(X)$ is homotopy equivalent to some projective $\mathrm{C}^\infty(\mathrm{TX})$ -module resolution of $\mathrm{C}^\infty(\mathrm{TX})$. Since projective resolutions are unique up to chain homotopy, the latter is homotopy equivalent to the resolution (C_k, d_k) . Throughout, homotopy equivalences are tacitly required $\mathrm{C}^\infty(X \times X)$ -linear or $\mathrm{C}^\infty(\mathrm{TX})$ -linear, whatever is appropriate.

First we write down some standard projective resolutions. They are related to the bar resolution because the multiplication map on $\mathrm{C}^\infty(X)$ extends to the map $m: \mathrm{C}^\infty(X \times X) \rightarrow \mathrm{C}^\infty(X)$, $m(f)(x) = f(x, x)$. Let Y be a smooth manifold that contains X as a smooth submanifold – we need the cases of TX and $X \times X$. Let $C_k(Y, \mathrm{C}^\infty(X)) := \mathrm{C}^\infty(Y^{k+1} \times X)$ for $k \geq 0$. Define boundary maps $d_k: C_k(Y, \mathrm{C}^\infty(X)) \rightarrow C_{k-1}(Y, \mathrm{C}^\infty(X))$ by

$$(d_k f)(y_0, \dots, y_{k-1}, x) := \sum_{j=0}^{k-1} (-1)^j f(y_0, \dots, y_j, y_j, \dots, y_{k-1}, x)$$

for $f: Y^k \times X \rightarrow \mathbb{C}$. That is, the entry y_j is doubled. We augment this chain complex by the map

$$d_0: C_0(Y, \mathrm{C}^\infty(X)) := \mathrm{C}^\infty(Y \times X) \rightarrow \mathrm{C}^\infty(X), \quad d_0(f)(x) := f(x, x).$$

We view $C_k(Y, \mathrm{C}^\infty(X))$ as a $\mathrm{C}^\infty(Y)$ -module by multiplication in the first variable:

$$(f_1 \cdot f_2)(y_0, \dots, y_k, x) := f_1(y_0) \cdot f_2(y_0, \dots, y_k, x)$$

for $f_1 \in \mathrm{C}^\infty(Y)$, $f_2 \in \mathrm{C}^\infty(Y^{k+1} \times X)$. Then $C_k(Y, \mathrm{C}^\infty(X))$ is relatively projective. A computation as for the bar resolution shows that the maps $s_k: C_{k-1} \rightarrow C_k$, $s_k f(y_0, \dots, y_k, x) := f(y_1, \dots, y_k, x)$ form a contracting homotopy of the augmentation of the complex (C_k, d_k) . Hence we have defined an admissible resolution of $\mathrm{C}^\infty(X)$.

Restriction from $X \times X$ to TX defines a chain map $C_\bullet(X \times X, \mathrm{C}^\infty(X)) \rightarrow C_\bullet(\mathrm{TX}, \mathrm{C}^\infty(X))$. We claim that it is a chain homotopy equivalence. The chain homotopy inverse $\Phi_\bullet: C_\bullet(\mathrm{TX}, \mathrm{C}^\infty(X)) \rightarrow C_\bullet(X \times X, \mathrm{C}^\infty(X))$ is of the form $\Phi_n f := f \cdot \omega_n$ for a sequence of smooth functions $\omega_k: Y^{k+1} \times X \rightarrow [0, 1]$ with certain properties. We need $\omega_{k+1}(y_0, \dots, y_j, y_j, \dots, y_k, x) = \omega_k(y_0, \dots, y_j, \dots, y_k, x)$ for all $y_0, \dots, y_k \in Y$, $x \in X$, and $\omega_0(x, x) = 1$ for all $x \in X$; this ensures that f_\bullet is a chain map that lifts the identity map on $\mathrm{C}^\infty(X)$. And we need that the support of ω_k is contained in $(\mathrm{TX})^{k+1} \times X$; this ensures that multiplication by ω_k and extension by 0 outside $(\mathrm{TX})^{k+1} \times X$ maps $\mathrm{C}^\infty((\mathrm{TX})^{k+1} \times X)$ to $\mathrm{C}^\infty((X \times X)^{k+1} \times X)$. There are indeed smooth functions with these properties.

The result is a chain map $f_\bullet: C_\bullet(TX, C^\infty(X)) \rightarrow C_\bullet(X \times X, C^\infty(X))$. Its composition with the restriction map in the opposite direction yields chain maps on $C_\bullet(TX, C^\infty(X))$ and $C_\bullet(X \times X, C^\infty(X))$ that lift the identity map on $C^\infty(X)$. Since the lifting of a module homomorphism to projective resolutions is unique up to chain homotopy, f_\bullet and the restriction map are inverse to each other up to chain homotopy. This finishes the proof. \square

As a result, the continuous Hochschild cohomology $\mathrm{HH}_{\mathrm{cont}}^*(C^\infty(X), C^\infty(X))$ is the cohomology of the cochain complex of continuous $C^\infty(X \times X)$ -module homomorphisms from $C^\infty(TX, \Lambda^k T^*X)$ to $C^\infty(X)$. Any such map is of the form $\omega \mapsto \langle \omega|_X | \pi \rangle$ for a section π of the dual bundle $\Lambda^k TX$ on X ; that is, π is a k -vector field on X . The boundary map vanishes because the space of k -vector fields is a symmetric $C^\infty(X)$ -bimodule. Hence

$$\mathrm{HH}_{\mathrm{cont}}^k(C^\infty(X), C^\infty(X)) \cong C^\infty(X, \Lambda^k TX).$$

20. Quasi-free algebras and their Hochschild cohomology

One way to measure the complexity of an algebra is by the length of a projective bimodule resolution of A . By Theorem 17.4, $\mathrm{HH}^n(A, M) = 0$ for all A -bimodules M and all $n > k$ if A has a projective bimodule resolution of length k . In fact, the converse of this is also true. The easiest case is when $\mathrm{HH}^n(A, M) = 0$ for all $n > 0$. Equivalently, all derivations into A -bimodules are inner. This is pretty rare. We have seen in Example 17.16 that this happens for semi-simple finite-dimensional algebras. Now we examine the case when $\mathrm{HH}^n(A, M) = 0$ for all $n > 1$. As it turns out, this is equivalent to the vanishing of $\mathrm{HH}^2(A, M)$ for all bimodules M . We choose this as our definition of quasi-freeness:

DEFINITION 20.1. An algebra is called *quasi-free* if any square-zero algebra extension $I \hookrightarrow E \twoheadrightarrow A$ splits by an algebra homomorphism. By Theorem 16.13, this is equivalent to $\mathrm{HH}^2(A, M) = 0$ for all A -bimodules M .

THEOREM 20.2. *If A is quasi-free, then any formal deformation quantisation of A is equivalent to the trivial one with $m(a, b) = a \cdot b$ for all $a, b \in A$.*

PROOF. This follows from Theorem 16.14 and $\mathrm{HH}^2(A, A) = 0$. \square

We have already met a number of quasi-free algebras. We are going to prove that \mathbb{C} , $\mathbb{C}[p]$, the group algebra of the dihedral group, $\mathbb{C}[t, t^{-1}]$, the Toeplitz algebra, and quiver algebras are quasi-free. We will also find several equivalent characterisations of quasi-free algebras. Namely, A is quasi-free if and only if all nilpotent extensions of A split – this is an analogue of Theorem 16.14. And A is quasi-free if and only if $\Omega^1(A)$ is a projective A -bimodule.

The examples of quasi-free algebras mentioned above are unital. To prove that they are quasi-free, the following lemma is very helpful:

LEMMA 20.3. *A unital algebra A is quasi-free if and only if $\mathrm{HH}^2(A, M) = 0$ for all unital A -bimodules, if and only if any square-zero extension $I \hookrightarrow E \twoheadrightarrow A$ with unital E splits by a unital algebra homomorphism.*

PROOF. Let M be an A -bimodule. Let

$$\begin{aligned} M_{11} &:= 1_A \cdot M \cdot 1_A, \\ M_{01} &:= (\mathrm{Id} - 1_A) \cdot M \cdot 1_A, \\ M_{10} &:= 1_A \cdot M \cdot (\mathrm{Id} - 1_A), \\ M_{00} &:= (\mathrm{Id} - 1_A) \cdot M \cdot (\mathrm{Id} - 1_A). \end{aligned}$$

These are A -subbimodules of M with

$$M \cong M_{11} \oplus M_{01} \oplus M_{10} \oplus M_{00}.$$

By construction, M_{11} is a unital A -bimodule, M_{10} is unital as a left A -module and M_{01} is unital as a right A -module, and all other multiplications are zero; that is, $A \cdot M_{01} = 0 = M_{10} \cdot A = A \cdot M_{00} \cdot A$. Routine computations show that all Hochschild 2-cocycles into M_{00} , M_{01} , M_{10} are Hochschild 2-coboundaries. For instance, if $\omega: A \times A \rightarrow M_{01}$ is a Hochschild 2-cocycle, then $\omega = \partial\psi$ with $\psi(a) := -\omega(1, a)$. Hence $\mathrm{HH}^2(A, M) \cong \mathrm{HH}^2(A, M_{11})$. Therefore, $\mathrm{HH}^2(A, M)$ vanishes for all bimodules once it vanishes for unital bimodules.

Let $I \twoheadrightarrow E \twoheadrightarrow A$ be a square-zero extension. Give I the induced A -bimodule structure. The proof of Theorem 16.13 shows that $\mathrm{HH}^2(A, M) = 0$ for all unital A -bimodules if and only if all those square-zero extensions split where I is a unital A -bimodule. For the second claim, we must show that E is unital if and only if I is a unital A -bimodule. If E is unital, then the quotient map $E \twoheadrightarrow A$ is unital. And then I is a unital A -bimodule. Conversely, assume I to be a unital A -bimodule. We have described E through a Hochschild 2-cocycle $\omega: A \otimes A \rightarrow I$ in the proof of Theorem 16.13. Only the class of ω up to Hochschild coboundaries matters. By Lemma 16.20, ω is cohomologous to a normalised Hochschild 2-cocycle $\omega': A \otimes A \rightarrow I$. Then $(1_A, 0) \in A \oplus I$ is a unit element for the multiplication on $A \oplus I$ defined by ω' . Thus E is unital. \square

PROPOSITION 20.4. *The field \mathbb{C} is quasi-free. If $I \twoheadrightarrow E \twoheadrightarrow A$ is a square-zero extension, then any idempotent element $p \in A$ lifts to an idempotent element $\hat{p} \in E$.*

PROOF. Assume first that $I \twoheadrightarrow E \twoheadrightarrow \mathbb{C}$ is a square-zero extension with unital E . Then the homomorphism $\mathbb{C} \rightarrow E$, $\lambda \mapsto \lambda \cdot 1_E$, is a section. By Lemma 20.3, it follows that \mathbb{C} is quasi-free. That is, a square-zero extension $I \twoheadrightarrow E \twoheadrightarrow \mathbb{C}$ also splits by a homomorphism if E is not unital. Now let $I \twoheadrightarrow E \xrightarrow{q} A$ be any square-zero extension and let $p \in A$ be idempotent. Let $\hat{E} := \{x \in E : q(x) \in \mathbb{C} \cdot p\}$. Then $I \twoheadrightarrow \hat{E} \xrightarrow{q} \mathbb{C} \cdot q$ is a square-zero extension of \mathbb{C} . A section for it is of the form $\lambda \mapsto \lambda \cdot \hat{p}$ for an idempotent element $\hat{p} \in E$ that lifts p . \square

PROPOSITION 20.5. *The group algebra of the infinite dihedral group D_∞ is quasi-free.*

PROOF. We have seen in Section 8 that the group algebra $\mathbb{C}[D_\infty]$ is the universal unital algebra generated by two idempotent elements p, q . That is, if B is a unital algebra, then there is a bijection between unital algebra homomorphisms $\mathbb{C}[D_\infty] \rightarrow B$ and pairs of idempotent elements (P, Q) in B . Now let $I \twoheadrightarrow E \twoheadrightarrow \mathbb{C}[D_\infty]$ be a square-zero extension with unital E . By Proposition 20.4, the idempotent elements p, q in $\mathbb{C}[D_\infty]$ lift to idempotent elements \hat{p}, \hat{q} in E . These generate a unital homomorphism $\mathbb{C}[D_\infty] \rightarrow E$. It is a section of the extension because \hat{p}, \hat{q} lift p, q . By Lemma 20.3, it follows that $\mathbb{C}[D_\infty]$ is quasi-free. \square

EXAMPLE 20.6. The polynomial algebra $\mathbb{C}[p]$ is quasi-free, even free: any algebra extension $I \twoheadrightarrow E \twoheadrightarrow \mathbb{C}[p]$ with unital E splits by a unital algebra homomorphism $\mathbb{C}[p] \rightarrow E$: lift the generator p to some $e \in E$ and map $p^n \mapsto e^n$ for $n \in \mathbb{N}$.

EXAMPLE 20.7. The polynomial algebra $\mathbb{C}[p, q]$ is not quasi-free because the Weyl algebra deformation provides a non-split square-zero extension.

In the world of noncommutative algebras, the free unital algebra on two generators is not $\mathbb{C}[p, q]$: it is the algebra of all non-commuting polynomials in two generators p and q . This free algebra is quasi-free for the same reason as $\mathbb{C}[p]$.

The polynomial algebra $\mathbb{C}[p]$ in Example 20.6 is the quiver algebra of the quiver with one object and one arrow (see Section 7). The above example generalises as follows:

PROPOSITION 20.8. *Let Q be a quiver with countably many vertices. Its quiver algebra is quasi-free.*

PROOF. Recall that the quiver is described by a pair of maps $s, r: Q^1 \rightrightarrows Q^0$. The quiver algebra has the following universal property. Let B be an algebra. An algebra homomorphism $\mathbb{C}[Q] \rightarrow B$ is equivalent to maps $p: Q^0 \rightarrow B$ and $t: Q^1 \rightarrow B$ such that $p_v p_w = \delta_{v,w} p_v$ for all $v, w \in Q^0$ and $t_e = p_{r(e)} t_e = t_e p_{s(e)}$ for all $e \in Q^1$. In other words, the elements p_v are orthogonal idempotents and t_e are arbitrary elements of $p_{r(e)} B p_{s(e)}$. Let $I \rightrightarrows E \rightrightarrows \mathbb{C}[Q]$ be a square-zero extension. Since Q^0 is countable, we may enumerate it: $Q_0 = \{v_i : i \in \mathbb{N}\}$ or $Q_0 = \{v_i : i = 0, 1, \dots, n\}$.

First, we recursively lift the orthogonal idempotents in Q_0 . By Proposition 20.4, we may lift p_{v_0} to an idempotent element \hat{p}_{v_0} of E . Assume that lifts \hat{p}_{v_j} of p_{v_j} for $j = 0, \dots, n-1$ have been constructed and that they are orthogonal idempotents. Let $q_n := \sum_{j=0}^{n-1} \hat{p}_{v_j}$. This is idempotent. Let $E_n := \{e \in E : q_n e = 0 = e q_n\}$. If $e \in E$ is arbitrary, then $e - q_n e - e q_n + q_n e q_n$ belongs to E_n . Using this, we show that p_{v_n} belongs to the image of E_n in $\mathbb{C}[Q]$. Therefore, Proposition 20.4 allows to lift p_{v_n} to an idempotent element $\hat{p}_{v_n} \in E_n$. In particular, \hat{p}_{v_n} is orthogonal to \hat{p}_{v_j} for $j < n$.

For each $e \in Q^1$, $\hat{p}_{r(e)} E \hat{p}_{s(e)}$ surjects onto $p_{r(e)} \mathbb{C}[Q] p_{s(e)}$. Hence we may lift $t_e \in p_{r(e)} \mathbb{C}[Q] p_{s(e)}$ to an element $\hat{t}_e \in \hat{p}_{r(e)} E \hat{p}_{s(e)}$. Now the elements \hat{p}_v and \hat{t}_e for $v \in Q^0$, $e \in Q^1$ generate a homomorphism $\mathbb{C}[Q] \rightarrow E$ that is a section for our extension. \square

PROPOSITION 20.9. *The algebra $\mathbb{C}[t, t^{-1}]$ of Laurent polynomials is quasi-free.*

PROOF. Let $I \rightrightarrows E \rightrightarrows \mathbb{C}[t, t^{-1}]$ be a square-zero extension with unital E and a unital algebra homomorphism $E \rightarrow \mathbb{C}[t, t^{-1}]$. Let $e \in E$ lift t . We must show that e is invertible – then $t^n \mapsto e^n$ for $n \in \mathbb{Z}$ is the required algebra homomorphism section. Let $f \in E$ be any lifting of t^{-1} . Then $1_E, ef$, and fe all lift tt^{-1} . So $1 - ef \in I$ and $1 - fe \in I$ have square zero. The equations $1 - 2ef + efef = 0$ and $1 - 2fe + fefe = 0$ imply that ef and fe are invertible with inverses $2 - ef$ and $2 - fe$, respectively. Hence e is both left and right invertible. Then e is invertible. \square

PROPOSITION 20.10. *The Toeplitz algebra is quasi-free.*

PROOF. Let $I \rightrightarrows E \rightrightarrows \mathcal{T}$ be a unital square-zero extension of the Toeplitz algebra. Let $\hat{s} \in E$ lift the generating isometry of \mathcal{T} . An argument as in the proof of Proposition 20.9 shows that \hat{s} is left-invertible. Let v be a left inverse of \hat{s} . Then there is a unital homomorphism $\mathcal{T} \rightarrow E$ that maps $s \mapsto \hat{s}$ and $s^* \mapsto v$. This is a section for the extension. Lemma 20.3 shows that \mathcal{T} is quasi-free. \square

DEFINITION 20.11. An algebra extension $I \rightrightarrows E \rightrightarrows Q$ is called *nilpotent* if there is $k \in \mathbb{N}$ with $I^k = 0$.

THEOREM 20.12. *Let A be a unital algebra. The following assertions are equivalent:*

- (1) *any square-zero extension $I \rightrightarrows E \rightrightarrows A$ splits;*
- (2) *for any square-zero extension $I \rightrightarrows E \rightrightarrows Q$, any algebra homomorphism $A \rightarrow Q$ lifts to an algebra homomorphism $A \rightarrow E$;*
- (3) *any nilpotent extension $I \rightrightarrows E \rightrightarrows A$ splits;*
- (4) *for any nilpotent extension $I \rightrightarrows E \rightrightarrows Q$, any algebra homomorphism $A \rightarrow Q$ lifts to an algebra homomorphism $A \rightarrow E$.*

PROOF. It is clear that (4) implies (2) and (3) and that (2) or (3) implies (1). We will establish the implications (1) \implies (2) \implies (4).

Let $I \twoheadrightarrow E \xrightarrow{p} Q$ be a square-zero extension and let $f: A \rightarrow Q$ be an algebra homomorphism. We shall build a commuting diagram of algebra homomorphisms

$$\begin{array}{ccccc} I & \twoheadrightarrow & E' & \xrightarrow{p'} & A \\ \parallel & & \downarrow f' & \swarrow s & \downarrow f \\ I & \twoheadrightarrow & E & \xrightarrow{p} & Q \end{array}$$

Let

$$E' := \{(e, a) \in E \times A : p(e) = f(a)\}$$

and let $f': E' \rightarrow E$ and $p': E' \rightarrow A$ be the restrictions of the coordinate projections. The map p' is a surjection with kernel $\{(i, 0) : i \in I\} \cong I$, so that $I \twoheadrightarrow E' \twoheadrightarrow A$ is a square-zero extension of A . By (1), it splits by an algebra homomorphism $s: A \rightarrow E'$. Then $f' \circ s: A \rightarrow E$ lifts $f: A \rightarrow Q$. Thus (1) implies (2).

Now assume (2). We prove (4) by induction on the number $k \in \mathbb{N}$ with $I^k = 0$. The case $k = 1$ is trivial, and $k = 2$ is the assumption (2). For the induction step, we consider an extension $I \twoheadrightarrow E \twoheadrightarrow Q$ with $I^k = 0$ and an algebra homomorphism $f: A \rightarrow Q$. Then $I/I^2 \twoheadrightarrow E/I^2 \twoheadrightarrow Q$ is a square-zero extension. So f lifts to an algebra homomorphism $f': A \rightarrow E/I^2$ by (2). The extension $I^2 \twoheadrightarrow E \twoheadrightarrow E/I^2$ satisfies $(I^2)^{k-1} = 0$. So f' lifts to an algebra homomorphism $f'': A \rightarrow E$ by the induction assumption. This finishes the induction step. \square

THEOREM 20.13. *Let A be a unital algebra. The following are equivalent:*

- (i) $\Omega^1(A)$ is a projective A -bimodule;
- (ii) A has a projective A -bimodule resolution of length 1;
- (iii) A is quasi-free, that is, $\mathrm{HH}^2(A, M) = 0$ for all A -bimodules M .

PROOF. (i) \implies (ii): By definition of $\Omega^1(A)$, there is a bimodule extension $\Omega^1(A) \twoheadrightarrow A \otimes A \twoheadrightarrow A$. Thus $0 \rightarrow \Omega^1(A) \rightarrow A \otimes A \rightarrow A$ is a projective bimodule resolution if $\Omega^1(A)$ is projective as an A -bimodule.

(ii) \implies (iii): follows immediately from Theorem 17.4.

(iii) \implies (i): Assume A to be quasi-free. The following diagram shows part of the bar resolution of A and the kernels of the boundary maps:

$$\begin{array}{ccccccc} A & \xleftarrow{b'_0} & A \otimes A & \xleftarrow{b'_1} & A \otimes \bar{A} \otimes A & \xleftarrow{b'_2} & A \otimes \bar{A} \otimes \bar{A} \otimes A \xleftarrow{b'_3} \dots \\ & & \swarrow & \dashrightarrow \gamma_1 & \swarrow & \dashrightarrow \gamma_2 & \\ & & \ker b'_0 & & \ker b'_1 & & \end{array}$$

There are unique dotted arrows that make the diagram commute because $b'_j \circ b'_{j+1} = 0$. These satisfy $\gamma_j \circ b'_{j+1} = 0$, and they are surjective because the bar resolution is exact. By definition, $\Omega^1(A) := \ker b'_0$. By the universal properties of free modules and tensor products, γ_2 comes from a bilinear map $\tilde{\gamma}_2: \bar{A} \times \bar{A} \rightarrow \ker b'_1$. The latter is a Hochschild 2-cocycle because $\gamma_2 \circ b'_3 = 0$; this is how we showed above that $\mathrm{Hom}_{A,A}(\mathrm{Bar}_\bullet(A), M)$ is naturally isomorphic to the normalised Hochschild cochain complex for A with coefficients in M . By assumption, any Hochschild cocycle is a coboundary. Thus there is an A -bimodule map $\xi: A \otimes \bar{A} \otimes A \rightarrow \ker b'_1$ with $\xi \circ b'_2 = \gamma_2$. This is equivalent to $\xi(x) = x$ for all $x \in \ker b'_1$. So we may view ξ as a projection from $A \otimes \bar{A} \otimes A$ onto $\ker b'_1$. Then $1 - \xi$ is a projection onto an A -subbimodule of $A \otimes \bar{A} \otimes A$ that is complementary to $\ker b'_1$. This makes the image of $1 - \xi$ isomorphic to the image of b'_1 . The latter is equal to the kernel of b'_0 , which is $\Omega^1(A)$. Thus $\Omega^1(A)$ is a direct summand in the free A -bimodule $A \otimes \bar{A} \otimes A$. Since

the functor $\text{Hom}_{A,A}(_, M)$ is additive and free bimodules are projective, this implies that $\Omega^1(A)$ is projective. \square

THEOREM 20.14. *Let Γ be a directed graph with countably many vertices and let $\Gamma'_0 \subseteq \Gamma_0$ be a set of regular vertices. Then the relative Leavitt path algebra $L(\Gamma, \Gamma'_0)$ is quasi-free.*

PROOF. Let $L := L(\Gamma, \Gamma'_0)$. Let $I \rightrightarrows E \rightarrow L$ be a square-zero extension. Let $\mathbb{C}[\Gamma]$ be the quiver algebra of Γ . This embeds in a canonical way into L as the subalgebra generated by S_v for $v \in \Gamma_0$ and S_α for $\alpha \in \Gamma_1$, without the generators S_α^* . The quiver algebra is quasi-free by Proposition 20.8. By Theorem 20.12, the inclusion $\mathbb{C}[\Gamma] \rightarrow L$ lifts to an algebra homomorphism $\varphi: \mathbb{C}[\Gamma] \rightarrow E$. Now we claim that φ extends to an algebra homomorphism on L . To build this extension, we view elements of L as left multiplication operators on L . We restrict the operator of left multiplication by S_α to a map $S_{s(\alpha)}L \rightarrow S_{r(\alpha)}L$ and the operator of left multiplication by S_α^* to a map $S_{r(\alpha)}L \rightarrow S_{s(\alpha)}L$. We combine all S_α with $r(\alpha) = v$ into an operator $I_v: \bigoplus_{\alpha \in r^{-1}(v)} S_{s(\alpha)}L \rightarrow S_vL$, $I_v(x_\alpha) := \sum_{\alpha \in r^{-1}(v)} S_\alpha(x_\alpha)$. Similarly, we combine all S_α^* with $r(\alpha) = v$ into an operator $I_v^*: S_vL \rightarrow \bigoplus_{\alpha \in r^{-1}(v)} S_{s(\alpha)}L$, $I_v^*(x) := (S_\alpha^*(x))_{\alpha \in r^{-1}(v)}$. Then $I_v^* \circ I_v: \bigoplus_{\alpha \in r^{-1}(v)} S_{s(\alpha)}L \rightarrow \bigoplus_{\alpha \in r^{-1}(v)} S_{s(\alpha)}L$ is the operator described by the block matrix $S_\alpha^* S_\beta = 0$ for $\alpha, \beta \in r^{-1}(v)$. So the relation (CK1) in Definition 9.20 says that $I_v^* I_v$ is the identity map on $\bigoplus_{\alpha \in r^{-1}(v)} S_{s(\alpha)}L$. The other composite $I_v I_v^*$ is equal to the sum $\sum_{\alpha \in r^{-1}(v)} S_\alpha S_\alpha^*$. So the other relation (CK2) in Definition 9.20 for v says that $I_v I_v^*$ is the identity map on S_vL provided $v \in \Gamma'_0$.

Now we apply the homomorphism $\varphi: \mathbb{C}[Q] \rightarrow E$. We may combine $\varphi(S_\alpha)$ for $\alpha \in r^{-1}(v)$ to an operator $\bigoplus_{\alpha \in r^{-1}(v)} \varphi(S_{s(\alpha)})E \rightarrow \varphi(S_v)E$. Let f_α be arbitrary lifts of $S_\alpha^* \in S_{s(\alpha)}L S_{r(\alpha)}$ to elements of $S_{s(\alpha)}E S_{r(\alpha)}$ and form a homomorphism $F_v: \varphi(S_v)E \rightarrow \bigoplus_{\alpha \in r^{-1}(v)} \varphi(S_{s(\alpha)})E$, $x \mapsto (f_\alpha x)_{\alpha \in r^{-1}(v)}$. Then $F_v I_v$ lifts the identity map on $\bigoplus_{\alpha \in r^{-1}(v)} S_{s(\alpha)}L$. And if $v \in \Gamma'_0$, then $I_v F_v$ lifts the identity map on S_vL . Therefore, $F_v I_v$ minus the identity is a matrix with entries in the square-zero algebra I . This makes it invertible as in the proof of Proposition 20.9. Let $I_v^* := (F_v I_v)^{-1} F_v$. Then $I_v^* I_v$ is the identity map. In addition, if $v \in \Gamma'_0$, then $I_v I_v^*$ still lifts the identity map on S_vL . Then it is also invertible. Then I_v is both left and right invertible, hence invertible. And the inverse must be I_v^* . The entries of I_v^* are given by left multiplication with elements $\tilde{f}_\alpha \in S_{s(\alpha)}E S_{r(\alpha)}$. The properties of I_v^* proven above say exactly that the elements $\varphi(S_v)$ for $v \in \Gamma_0$ and $\varphi(S_\alpha)$ and \tilde{f}_α for $\alpha \in \Gamma_1$ satisfy the relations of the Leavitt path algebra. Hence there is a unique homomorphism $L \rightarrow E$ that maps $S_v \mapsto \varphi(S_v)$, $S_\alpha \mapsto \varphi(S_\alpha)$ and $S_\alpha^* \mapsto \tilde{f}_\alpha$. \square

21. Hochschild cohomology of the Weyl algebra — De Rham cohomology

In the first subsection, we compute the Hochschild cohomology for the Weyl algebra. This is another application of a Koszul resolution. We merely observe that the “same” resolution works for polynomials and the Weyl algebra. There are more general results for the Hochschild cohomology of crossed products by derivations, but we limit our discussion to one interesting example.

Our computations so far show that Hochschild cohomology has several drawbacks. First, the computations for commutative algebras make it clear that it is not homotopy invariant. Secondly, it is also sensitive to deformations because it yields quite different results for the algebra of polynomials and the Weyl algebra.

Periodic cyclic cohomology is another cohomology theory for algebras that is homotopy invariant and, in many cases, is invariant under deformation quantisations. So it captures some purely topological information about a noncommutative algebra. In order to prepare for the study of periodic cyclic homology, we briefly discuss de Rham cohomology for smooth manifolds. This is what periodic cyclic homology will give for algebras of smooth functions. We also motivate de Rham cohomology and discuss it in the case of \mathbb{R}^3 .

21.1. Hochschild cohomology of the Weyl algebra. Fix $\hbar \in \mathbb{C}$ and let $A = \mathbb{C}\langle p, q \mid [p, q] = i\hbar \rangle$. We are going to write down a projective A -bimodule resolution of A . In fact, the resolution is just another instance of the Koszul resolution. Let

$$P_k := A \otimes \Lambda^k(\mathbb{C}^2) \otimes A,$$

that is, $P_0 \cong P_2 \cong A \otimes A$, $P_1 \cong (A \otimes A) \oplus (A \otimes A)$. Define the boundary maps $d_k: P_k \rightarrow P_{k-1}$ and the augmentation map $d_0: A \otimes A \rightarrow A$ by

$$\begin{aligned} d_2(a \otimes e_1 \wedge e_2 \otimes b) &:= a \cdot p \otimes e_2 \otimes b - a \otimes e_2 \otimes p \cdot b - a \cdot q \otimes e_1 \otimes b + a \otimes e_1 \otimes q \cdot b, \\ d_1(a \otimes e_1 \otimes b) &:= a \cdot p \otimes b - a \otimes p \cdot b, \\ d_1(a \otimes e_2 \otimes b) &:= a \cdot q \otimes b - a \otimes q \cdot b, \\ d_0(a \otimes b) &:= a \cdot b. \end{aligned}$$

For $\hbar = 0$, this is exactly the Koszul resolution for the polynomial algebra $\mathbb{C}[p, q]$ introduced in Section 18. The idea of the formulas above is that the boundary map replaces e_1 by an inner commutator with p and e_2 by an inner commutator with q .

THEOREM 21.1. *The chain complex above is a projective A -bimodule resolution of A . Let M be an A -bimodule. Then $\mathrm{HH}^n(A, M) = 0$ for $n \geq 3$ and*

$$\mathrm{HH}^2(A, M) \cong M [p, M] + [q, M].$$

In particular, $\mathrm{HH}^2(A, A) = 0$ if $\hbar \neq 0$.

PROOF. The A -bimodules P_k for $k \geq 0$ are free, hence projective by Example 15.15. And the maps d_k for $0 \leq k \leq 2$ are A -bimodule maps. It is easy to compute that $d_k \circ d_{k+1} = 0$ for $k = 0, 1$ and for all $\hbar \in \mathbb{C}$. This depends on the commutation relation $[p, q] = i\hbar$. We must show that the following complex C_\bullet is exact:

$$0 \rightarrow P_2 \xrightarrow{d_2} P_1 \xrightarrow{d_1} P_0 \xrightarrow{d_0} A \rightarrow 0.$$

We use a filtration of C_\bullet . We declare that p, q and e_1, e_2 each have degree 1. Let $C_\bullet^{(k)}$ be the subspace spanned by “products” of these generators with at most k factors. That is, $C_2^{(k)}$ is spanned by $p^n q^m \otimes e_1 \wedge e_2 \otimes p^a q^b$ with $n + m + 2 + a + b \leq k$, $C_1^{(k)}$ is spanned by $p^n q^m \otimes e_j \otimes p^a q^b$ with $n + m + 1 + a + b \leq k$, and $C_0^{(k)}$ is spanned by $p^n q^m \otimes p^a q^b$ with $n + m + a + b \leq k$. The boundary map replaces e_1 by an inner commutator with p and e_2 by an inner commutator with q . Therefore, it does not increase the degree. That is, each $C_\bullet^{(k)}$ is a subcomplex. Since degrees are non-negative, $C_\bullet^{(-1)} = 0$. We use a variant of Lemma 18.3, where we start counting at -1 . So to prove that C_\bullet is exact, it suffices to prove that each subquotient $C_\bullet^{(k)}/C_\bullet^{(k-1)}$ for $k \geq 0$ is exact. This can be done by hand, but it gets rather complicated. Therefore, we prefer an indirect argument.

If $\hbar = 0$, then the complex above is the Koszul resolution for the polynomial algebra $\mathbb{C}[p, q]$. We already know that this resolution is exact. In addition, if $\hbar = 0$, then the complex above is not just filtered, but graded by degree: the boundary map sends a monomial of degree k to a sum of monomials of degree k as well. So the Koszul resolution is isomorphic as a chain complex to the direct

sum $\bigoplus_{k=0}^{\infty} C_{\bullet}^{(k)} / C_{\bullet}^{(k-1)}$. If a direct sum is exact, then each summand is exact. So $C_{\bullet}^{(k)} / C_{\bullet}^{(k-1)}$ for $k \geq 0$ is exact if $\hbar = 0$. Now the subquotients $C_{\bullet}^{(k)} / C_{\bullet}^{(k-1)}$ for $k \geq 0$ are independent of \hbar : the effect of a non-zero \hbar are commutators, which have lower degree. So the boundaries of a monomial of degree k for different \hbar differ only in terms of lower degree, which become 0 in the quotient $C_{\bullet}^{(k)} / C_{\bullet}^{(k-1)}$. Since this quotient is exact for $\hbar = 0$, it is also exact for $\hbar \neq 0$.

This finishes the proof that the complex above is a projective A -bimodule resolution of A . Now we compute the Hochschild cohomology of A with this resolution as in Theorem 17.4. The cochain complex $\text{Hom}_{A,A}(P_{\bullet}, M)$ is

$$0 \rightarrow M \xrightarrow{\begin{pmatrix} \text{ad}_p \\ -\text{ad}_q \end{pmatrix}} M^2 \xrightarrow{\begin{pmatrix} \text{ad}_q & \text{ad}_p \end{pmatrix}} M,$$

where $\text{ad}_x(m) := [x, m]$ for $x = p, q$. Thus $\text{HH}^n(A, M) = 0$ for all $n \geq 3$ and all M , and

$$\text{HH}^2(A, M) \cong M / ([p, M] + [q, M]).$$

If $\hbar \neq 0$, then the monomial $p^n q^m$ may be written as a commutator with p or q . Hence the computation of $\text{HH}^2(A, M)$ above implies $\text{HH}^2(A, A) = 0$. \square

COROLLARY 21.2. *Any formal deformation quantisation of the Weyl algebra A with $\hbar \neq 0$ is equivalent to the trivial deformation quantisation with $m(a, b) = a \cdot b$ for $a, b \in A$.*

PROOF. This follows from Theorem 16.14 and the formula $\text{HH}^2(A, A) = 0$ proven in Theorem 21.1. \square

Varying the parameter \hbar defines a polynomial deformation quantisation of each individual Weyl algebra. We already know that all these Weyl algebras for $\hbar = 0$ are isomorphic. On the level of formal power series, the triviality of this formal deformation quantisation follows from the previous theorem. For $\hbar = 0$, however, the Weyl algebra is the polynomial algebra, which *has* non-trivial formal deformation quantisations.

We already know that $\text{HH}^1(A, A) = 0$ and $\text{HH}^0(A, A) = \mathbb{C} \cdot 1_A$. Thus the Weyl algebra has the same Hochschild cohomology with coefficients A as the algebra \mathbb{C} of complex numbers. In contrast, the Hochschild cohomology of the polynomial algebra $\mathbb{C}[p, q]$ is quite different. So deformation quantisation may change the Hochschild cohomology rather drastically.

21.2. De Rham cohomology of smooth manifolds. The de Rham boundary operator contains information about differentiation of smooth functions on a manifold. We work with real-valued functions instead of complex-valued functions here.

EXAMPLE 21.3. We first consider the trivial case of the real line. Here the de Rham boundary operator is equivalent to the differentiation map

$$C^{\infty}(\mathbb{R}) \rightarrow C^{\infty}(\mathbb{R}), \quad f \mapsto f'.$$

This map is surjective, and its kernel is the space of constant functions. Hence the de Rham cohomology of \mathbb{R} is $H_{\text{dR}}^1(\mathbb{R}) = 0$ and $H_{\text{dR}}^0(\mathbb{R}) = \mathbb{R}$. Integration provides explicit sections for the differentiation map of the form $(I_a f)(x) := \int_a^x f(t) dt$ for any $a \in \mathbb{R}$.

On a general smooth manifold M , there is no longer a unique derivative. Instead, each vector field X on M defines a derivation $f \mapsto \partial_X(f)$ on $C^{\infty}(M)$. These maps for different X are not independent: they combine to a derivation

$$(21.4) \quad d: C^{\infty}(M) \rightarrow \Omega^1(M), \quad (df)(X) := \partial_X(f),$$

where $\Omega^1(M)$ is the $C^\infty(M)$ -module of differential 1-forms on M discussed already in Section 15.1. The map d is the de Rham boundary map on $C^\infty(M)$. And (21.4) is the correct analogue of differentiation for smooth functions on \mathbb{R} .

Using a Riemannian metric, we may identify $\Omega^1(M)$ with the space of vector fields on M and d with the gradient map $V \mapsto \text{grad}(V)$. Then the operator $-d$ has the physical interpretation of mapping a potential function V to the resulting force field $-\text{grad}(V)$. Its kernel and image correspond to important physical questions. The kernel of d is related to the uniqueness of the potential associated to a force field. If $df = 0$, then the function f is locally constant, that is, constant on each connected component of M . This means that the potential is unique up to adding a constant on each connected component of M . The image of d describes which force fields are conservative, that is, have a potential. While d is surjective for \mathbb{R} , this fails for other smooth manifolds.

To understand the image of d , we first assume $M = \mathbb{R}^n$. Then $C^\infty(\mathbb{R}^n)^n \cong \Omega^1(\mathbb{R}^n)$ via the isomorphism

$$C^\infty(\mathbb{R}^n)^n \cong \Omega^1(\mathbb{R}^n), \quad (f_1, \dots, f_n) \mapsto f_1 dx_1 + f_2 dx_2 + \dots + f_n dx_n,$$

and d corresponds to the differential operator

$$C^\infty(\mathbb{R}^n) \rightarrow C^\infty(\mathbb{R}^n)^n, \quad f \mapsto \left(\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_n} \right).$$

The range of this operator consists of all n -tuples (f_1, \dots, f_n) of smooth functions that satisfy the system of linear partial differential equations

$$(21.5) \quad \frac{\partial f_i}{\partial x_j} = \frac{\partial f_j}{\partial x_i}$$

for all $1 \leq i < j \leq n$. Equivalently, (f_1, \dots, f_n) belongs to the kernel of the map

$$(21.6) \quad d: \Omega^1(\mathbb{R}^n) \rightarrow \Omega^2(\mathbb{R}^n),$$

$$\sum_{j=1}^n f_j dx_j \mapsto \sum_{i,j=1}^n \frac{\partial f_j}{\partial x_i} dx_i \wedge dx_j = \sum_{1 \leq i < j \leq n} \left(\frac{\partial f_j}{\partial x_i} - \frac{\partial f_i}{\partial x_j} \right) dx_i \wedge dx_j.$$

This is the de Rham boundary map on $\Omega^1(\mathbb{R}^n)$. We have used special properties of the basis of vector fields (∂x_j) on \mathbb{R}^n . So the extension of (21.6) to other smooth manifolds is still unclear.

Let M be a smooth manifold and let X, Y be two smooth vector fields on M . Then $\partial_{[X,Y]}(f) = \partial_X(\partial_Y(f)) - \partial_Y(\partial_X(f))$ defines another smooth vector field $[X, Y]$ on M . Therefore, if a differential 1-form ω is of the form df , then

$$\partial_X \langle \omega | Y \rangle - \partial_Y \langle \omega | X \rangle - \langle \omega | [X, Y] \rangle = 0.$$

Here $\langle \omega | X \rangle$ denotes the canonical pairing between vector fields and differential 1-forms. This leads to a version of (21.5) that makes sense for general manifolds. Namely, we define the de Rham boundary $d: \Omega^1(M) \rightarrow \Omega^2(M)$ by

$$d\omega(X, Y) := \partial_X \langle \omega | Y \rangle - \partial_Y \langle \omega | X \rangle - \langle \omega | [X, Y] \rangle.$$

Then $d(d(f)) = 0$ for all $f \in C^\infty(M)$. That is, $d(\omega) = 0$ is necessary for ω to be of the form $d(f)$ for some $f \in C^\infty(M)$. This condition need not be sufficient, however:

EXAMPLE 21.7. Let $M = \mathbb{R}/\mathbb{Z}$ be the circle. Then $C^\infty(M)$ is the space of \mathbb{Z} -periodic smooth functions on \mathbb{R} . Here $\Omega^2(M) = 0$, so that $d: \Omega^1(M) \rightarrow \Omega^2(M)$ vanishes. The smooth covector field dx is constant, hence \mathbb{Z} -periodic and belongs to $\Omega^1(M)$. Up to addition of constant functions, there is a unique smooth function on \mathbb{R} with $dV = dx$, namely, the identical function $V(x) = x$. But there is no \mathbb{Z} -periodic function V with $dx = dV$.

The quotient

$$H_{\text{dR}}^1(M) := \ker(d: \Omega^1(M) \rightarrow \Omega^2(M)) / d(C^\infty(M))$$

is called the *first de Rham cohomology of M* . It is an important invariant of a smooth manifold. It agrees with the singular cohomology of M with coefficients \mathbb{R} . It is closely related to the fundamental group $\pi_1(M)$: if M is connected, then $H_{\text{dR}}^1(M)$ is isomorphic to the space $\text{Hom}(\pi_1(M), \mathbb{R})$ of group homomorphisms $\pi_1(M) \rightarrow \mathbb{R}$.

Differential forms in the image of $d: \Omega^1(M) \rightarrow \Omega^2(M)$ also satisfy certain linear differential equations, which may be derived exactly as above. This leads to a differential $d: \Omega^2(M) \rightarrow \Omega^3(M)$. This process continues until we reach $\Omega^{\dim(M)+1}(M) = 0$, where we get no further obstructions. This yields the *de Rham cochain complex*

$$(21.8) \quad 0 \rightarrow \Omega^0(M) \xrightarrow{d} \Omega^1(M) \xrightarrow{d} \Omega^2(M) \rightarrow \dots \rightarrow \Omega^{n-1}(M) \xrightarrow{d} \Omega^n(M) \rightarrow 0$$

for a smooth manifold of dimension n , with $\Omega^0(M) := C^\infty(M)$. In local coordinates, that is, for $M = \mathbb{R}^n$, the de Rham boundary map becomes

$$d(f dx_{i_1} dx_{i_2} \dots dx_{i_k}) = \sum_{j=1}^n \frac{\partial f}{\partial x_j} dx_j dx_{i_1} dx_{i_2} \dots dx_{i_k}.$$

This works because the basic vector fields ∂x_j on \mathbb{R}^n commute ($[\partial x_j, \partial x_k] = 0$). The cohomology of the de Rham complex in (21.8) is called *de Rham cohomology of M* . It agrees with the singular cohomology of M and with the Čech cohomology of M – all reasonable cohomology theories with \mathbb{R} -coefficients agree for smooth manifolds.

Since de Rham cohomology is a cohomology theory, it is homotopy invariant. One easy consequence of homotopy invariance is the following theorem:

THEOREM 21.9 (Poincaré Lemma). $H_{\text{dR}}^k(\mathbb{R}^n) = 0$ for all $k \geq 1$, that is, the condition $d\omega = 0$ is both necessary and sufficient for ω to be of the form $d\eta$.

PROOF. Since \mathbb{R}^n is contractible and H_{dR} is homotopy invariant, \mathbb{R}^n and the one-point space have isomorphic de Rham cohomology. \square

In fact, the above proof is circular. We have not explained how to prove that de Rham cohomology agrees with singular cohomology. A direct verification of the Poincaré Lemma is one of the steps in standard proofs of this result.

EXAMPLE 21.10. Let M be an *oriented* smooth manifold of dimension 3. Then M admits a volume form $\omega \in \Omega^3(M)$, and $f \mapsto f \cdot \omega$ provides an isomorphism $C^\infty(M) \cong \Omega^3(M)$. We have $C^\infty(M) = \Omega^0(M)$, anyway, and we may identify $\Omega^2(M)$ and $\Omega^1(M)$ with the space $\mathfrak{X}(M)$ of smooth vector fields on M using a Riemannian metric on M . Thus the de Rham complex becomes isomorphic to a cochain complex of the form

$$0 \rightarrow C^\infty(M) \rightarrow \mathfrak{X}(M) \rightarrow \mathfrak{X}(M) \rightarrow C^\infty(M) \rightarrow 0.$$

On \mathbb{R}^3 , the first map $C^\infty(M) \rightarrow \mathfrak{X}(M)$ is the *gradient* map

$$\text{grad}(f) := \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z} \right);$$

the second map takes the *rotation* of a vector field:

$$\text{rot}(f_x, f_y, f_z) := \left(\frac{\partial f_y}{\partial z} - \frac{\partial f_z}{\partial y}, \frac{\partial f_z}{\partial x} - \frac{\partial f_x}{\partial z}, \frac{\partial f_x}{\partial y} - \frac{\partial f_y}{\partial x} \right);$$

the third map takes the *divergence*

$$\text{div}(f_x, f_y, f_z) := \frac{\partial f_x}{\partial x} + \frac{\partial f_y}{\partial y} + \frac{\partial f_z}{\partial z}.$$

Thus the natural differential operators grad , rot , and div between functions and vector fields on \mathbb{R}^3 are special cases of the de Rham boundary map. That the de Rham complex is a cochain complex means that $\text{rot grad}(V) = 0$ and $\text{div rot}(F) = 0$.

22. Noncommutative differential forms

In this section, we provide the ingredients needed to define periodic cyclic homology. These are the algebra of noncommutative differential forms over an algebra A and two canonical maps b and d on them. The bimodule of noncommutative differential forms is an analogue of $\Omega^1(A)$ that is universal for normalised Hochschild cocycles instead of derivations. We will also write its elements as formal differential forms $a_0 da_1 \dots da_n$. The difference to differential forms on a manifold is that there is no commutation relation for the factors da . The differential forms come with an obvious boundary map d . Another important map is the Hochschild boundary map b , which looks a bit like the boundary map b' in the bar resolution, but has an extra terms which makes its homology highly non-trivial. The homology of $(\Omega^\bullet A, b)$ is called the Hochschild homology of A . We explain how to compute it using projective bimodule resolutions, like Hochschild cohomology.

Let A be a unital algebra and let M be a unital A -bimodule. Then

$$\begin{aligned} \text{HH}^0(A, M) &:= Z(M) \cong \text{Hom}_{A,A}(A, M), \\ \text{Der}(A, M) &:= \text{Hom}_{A,A}(\Omega^1(A), M). \end{aligned}$$

We are going to define A -bimodules $\Omega^n(A)$ such that there is a natural bijection between $\text{Hom}_{A,A}(\Omega^n(A), M)$ and normalised Hochschild n -cocycles $A^n \rightarrow M$. In particular, $\Omega^0(A) := A$. And then we combine all these bimodules into a differential graded algebra over A .

DEFINITION 22.1. For $n \geq 1$, we let $\Omega^n(A)$ be the balanced tensor product

$$\Omega^1(A)^{\otimes_A n} := \underbrace{\Omega^1(A) \otimes_A \Omega^1(A) \otimes_A \cdots \otimes_A \Omega^1(A)}_{n \text{ factors}}.$$

This is an A -bimodule in a natural way. Elements of $\Omega^n(A)$ are called *noncommutative differential n -forms over A* .

We now describe this bimodule more concretely. Let $\bar{A} := A/\mathbb{C} \cdot 1_A$. Recall that $\Omega^1(A) \cong A \otimes \bar{A}$ as a left A -module via the map

$$A \otimes \bar{A} \rightarrow \Omega^1(A) := \ker(A \otimes A \xrightarrow{\text{mult}} A), \quad a_0 \otimes \bar{a}_1 \mapsto a_0 da_1 := a_0 \otimes a_1 - a_0 a_1 \otimes 1.$$

Thus $\Omega^1(A)$ is free as a left A -module. This allows us to simplify the above tensor products:

LEMMA 22.2. *There are natural left A -module isomorphisms $\Omega^n(A) \cong A \otimes \bar{A}^{\otimes n}$ for all $n \in \mathbb{N}_{\geq 1}$, such that the right A -module structure on $\Omega^n(A)$ becomes*

$$\begin{aligned} (a_0 \otimes a_1 \otimes \cdots \otimes a_n) \cdot b &:= a_0 \otimes a_1 \otimes \cdots \otimes (a_n \cdot b) \\ &- a_0 \otimes a_1 \otimes \cdots \otimes (a_{n-1} \cdot a_n) \otimes b + a_0 \otimes a_1 \otimes \cdots \otimes (a_{n-2} \cdot a_{n-1}) \otimes a_n \otimes b \\ &\mp \cdots + (-1)^n a_0 \cdot a_1 \otimes a_2 \otimes \cdots \otimes a_n \otimes b. \end{aligned}$$

We rewrite the above elementary tensors as differential forms

$$a_0 \otimes a_1 \otimes \cdots \otimes a_n \leftrightarrow a_0 da_1 \dots da_n.$$

PROOF. We prove the lemma by induction on n . We already know everything for $n = 1$. Assume the assertion holds for $\Omega^{n-1}(A)$. Then

$$\Omega^n(A) \cong \Omega^{n-1}(A) \otimes_A \Omega^1(A) \cong \Omega^{n-1}(A) \otimes_A A \otimes \bar{A} \cong \Omega^{n-1}(A) \otimes \bar{A} \cong A \otimes \bar{A}^n$$

as left A -modules by Lemma 11.6. In addition, the right multiplication with A becomes

$$\begin{aligned} a_0 da_1 \dots da_n \cdot b &= a_0 da_1 \dots da_{n-1} \otimes (da_n)b \\ &= a_0 da_1 \dots da_{n-1} \otimes (d(a_nb) - a_n db) \\ &= a_0 da_1 \dots da_{n-1} d(a_nb) - (a_0 da_1 \dots da_{n-1}) \cdot a_n db. \end{aligned}$$

Now we use the induction hypothesis about the right module structure on $\Omega^{n-1}(A)$ to get the formula for the right A -module structure on $\Omega^n(A)$. \square

PROPOSITION 22.3. *The map $d^{\cup n}: (a_1, \dots, a_n) \mapsto da_1 \dots da_n$ is the universal normalised Hochschild n -cocycle with values in a unital A -bimodule: it is a normalised Hochschild n -cocycle, and any other normalised Hochschild n -cocycle factors as $f \circ d^{\cup n}$ for a unique bimodule homomorphism f .*

PROOF. Since $\Omega^n(A)$ is isomorphic to the free A -module on $\bar{A}^{\otimes n}$, left A -module homomorphisms $f: \Omega^n(A) \rightarrow M$ correspond to n -linear maps $\omega: \bar{A}^n \rightarrow M$. A routine computation shows that f is a right module homomorphism as well if and only if ω is a normalised Hochschild n -cocycle. \square

By definition, $\Omega^1(A)$ is the kernel of the multiplication map $A \otimes A \rightarrow A$. This is the augmentation map of the bar resolution. The following exercise relates the bimodules $\Omega^n(A)$ to the bar resolution as well:

EXERCISE 22.4. Prove that the kernel of $b': \text{Bar}_{n+1}(A) \rightarrow \text{Bar}_n(A)$ is naturally isomorphic to $\Omega^{n+2}(A)$ if $n \geq 0$.

The bimodule structures on $\Omega^n(A)$ for $n \in \mathbb{N}$ extend to a multiplication

$$\Omega^n(A) \otimes \Omega^m(A) \rightarrow \Omega^{n+m}(A), \quad \omega \cdot a_0 da_1 \dots da_m := (\omega \cdot a_0) da_1 \dots da_m.$$

This multiplication of differential forms is associative.

DEFINITION 22.5. A *graded algebra* is an algebra A together with a decomposition $A \cong \bigoplus_{n \in \mathbb{N}} A_n$ such that $A_n \cdot A_m \subseteq A_{n+m}$.

A *graded derivation* on a graded algebra A is a linear map $d: A \rightarrow A$ with $d(A_n) \subseteq A_{n+1}$ that satisfies the *graded Leibniz rule*

$$d(a \cdot b) = d(a) \cdot b + (-1)^n a \cdot d(b)$$

for $a \in A_n, b \in A_m$.

A *differential graded algebra* is a graded algebra A with a graded derivation $d: A \rightarrow A$ that satisfies $d^2 = 0$.

It is clear that $\Omega(A) := \bigoplus_{n \in \mathbb{N}} \Omega^n(A)$ becomes a graded algebra with the multiplication defined above. We define a differential by the obvious formula

$$d: \Omega^n(A) \rightarrow \Omega^{n+1}(A), \quad a_0 da_1 \dots da_n \mapsto da_0 da_1 \dots da_n,$$

where it is understood that $d(1) = 0$.

LEMMA 22.6. *With this map d , $\Omega(A)$ becomes a differential graded algebra.*

PROOF. The property $d^2 = 0$ is built into our definition. It remains to check that d satisfies the graded Leibniz rule. Since $d^2 = 0$, the graded Leibniz rule implies $d(\omega \cdot d(\eta)) = d(\omega) \cdot d(\eta)$, which is certainly true. Let X be the set of $\eta \in \Omega(A)$ with $d(\omega \cdot \eta) = d(\omega) \cdot \eta + (-1)^n \omega d\eta$ for all n and all $\omega \in \Omega^n(A)$. This is a subalgebra of $\Omega(A)$ because if $\eta_1, \eta_2 \in X$, then

$$\begin{aligned} d(\omega \eta_1 \eta_2) &= d(\omega \eta_1) \eta_2 + (-1)^{|\omega \eta_1|} \omega \eta_1 d\eta_2 \\ &= d(\omega) \eta_1 \eta_2 + (-1)^{|\omega|} \omega ((d\eta_1) \eta_2 + (-1)^{|\eta_1|} \eta_1 d\eta_2) \\ &= d(\omega) \eta_1 \eta_2 + (-1)^{|\omega|} \omega d(\eta_1 \eta_2); \end{aligned}$$

here $|\omega|$ and $|\omega\eta_1| = |\omega| + |\eta_1|$ denote the degrees of ω and $\omega\eta_1$, respectively. Since X contains the range of d and $a_0 da_1 \dots da_n = a_0 \cdot d(a_1 da_2 \dots da_n)$, it remains to prove

$$d(\omega \cdot a) = d(\omega) \cdot a + (-1)^n \omega da$$

for $\omega \in \Omega^n(A)$, $a \in A$. We check this by induction on n . The case $n = 1$ is clear. In the induction step, we write $\omega = \eta db$ for some $b \in A$, $\eta \in \Omega^{n-1}(A)$. Then

$$\begin{aligned} d(\omega \cdot a) - d(\omega) \cdot a &= d(\eta \cdot d(ba) - \eta \cdot b da) - d(\eta db) \cdot a \\ &= d(\eta) \cdot d(ba) - d(\eta) \cdot b da - (-1)^{n-1} \eta db da - d(\eta) \cdot db \cdot a = (-1)^n \omega da. \end{aligned}$$

Thus d is a graded derivation as desired. \square

PROPOSITION 22.7. *The differential graded algebra $(\Omega(A), d)$ is the universal differential graded unital algebra over A in the sense that any unital algebra homomorphism from A to the degree-0-part of a differential graded unital algebra (B, D) extends uniquely to a differential graded algebra homomorphism from $(\Omega(A), d)$ to (B, D) .*

PROOF. If $f: A \rightarrow B_0$ is an algebra homomorphism, then we define maps $f_n: \Omega^n(A) \rightarrow B_n$ by mapping $a_0 da_1 \dots da_n$ to $f(a_0) \cdot Df(a_1) \cdots Df(a_n)$. This satisfies $f_{n+1} \circ d = D \circ f_n$ for all $n \in \mathbb{N}$. It is multiplicative because D satisfies the graded Leibniz rule, which dictates the multiplication of differential forms. And the maps f_n are the only maps that are compatible with the multiplication and the differentials. \square

22.1. The two boundaries on noncommutative differential forms. How can we define an analogue of the de Rham complex for a noncommutative algebra A ? We have already met the bimodules $\Omega^n(A)$ of noncommutative differential forms in connection with Hochschild cohomology. In some contexts, these bimodules play the role of the spaces $\Omega^n(M)$ of differential forms on a smooth manifold. They carry an obvious differential

$$d: \Omega^n(A) \rightarrow \Omega^{n+1}(A), \quad a_0 da_1 da_2 \dots da_n \mapsto da_0 da_1 da_2 \dots da_n,$$

which annihilates the *closed* forms $da_1 da_2 \dots da_n$. Thus $d \circ d = 0$ by definition, and we get a cochain complex $(\Omega^n(A), d)$. This complex is not interesting, however, because $H^0(\Omega^n(A), d) = \mathbb{C} \cdot 1$ and $H^n(\Omega^n(A), d) = 0$ for $n \neq 0$ for any unital algebra A . To see this, we use a splitting $A \cong \mathbb{C} \cdot 1 \oplus \bar{A}$ to define a contracting homotopy:

$$s(da_0 da_1 da_2 \dots da_n) := a_0 da_1 da_2 \dots da_n, \quad s(a_0 da_1 da_2 \dots da_n) := 0$$

for $a_0, \dots, a_n \in \bar{A}$. We have $ds + sd = \text{Id}$ on $\Omega^n(A)$ for $n \geq 1$, while $sd = ds + sd$ on $\Omega^0(A) = A$ is the projection onto \bar{A} .

We did not get an interesting cohomology theory because the noncommutative differential forms $\Omega^n(A)$ are an *analogue* but not a *generalisation* of the space of differential forms $\Omega^n(M)$ on a smooth manifold M . We get $\Omega^1(M)$ from $\Omega^1(\mathbb{C}^\infty M)$ by enforcing the relation $f \cdot \omega = \omega \cdot f$ for all $f \in \mathbb{C}^\infty(M)$, $\omega \in \Omega^1(\mathbb{C}^\infty M)$ (see Corollary 15.8). This suggests replacing $\Omega^n(A)$ by the commutator quotient $\Omega^n(A) / [\Omega^n(A), A]$ for all $n \in \mathbb{N}$. Although the solution is more complicated, understanding commutators is a good move. We define the *Hochschild boundary* $b: \Omega^{n+1}(A) \rightarrow \Omega^n(A)$ on differential forms by

$$\begin{aligned} (22.8) \quad b(a_0 da_1 da_2 \dots da_{n+1}) &:= (-1)^n [a_0 da_1 da_2 \dots da_n, a_{n+1}] \\ &= (-1)^n (a_0 da_1 da_2 \dots da_n) \cdot a_{n+1} + (-1)^{n+1} a_{n+1} \cdot a_0 da_1 da_2 \dots da_n. \end{aligned}$$

More explicitly,

$$\begin{aligned} b(a_0 da_1 da_2 \dots da_{n+1}) &= a_0 a_1 da_2 \dots da_{n+1} \\ &+ \sum_{j=1}^n (-1)^j a_0 da_1 da_2 \dots d(a_j a_{j+1}) \dots da_{n+1} \\ &+ (-1)^{n+1} a_{n+1} \cdot a_0 da_1 \dots da_n. \end{aligned}$$

This looks similar to the boundary map b' in the bar resolution. However, the extra term $a_{n+1} \cdot a_0 da_1 da_2 \dots da_n$ in b makes a big difference. We check that b is a boundary map, that is, $b^2 = 0$:

$$\begin{aligned} b^2(\omega da_n da_{n+1}) &= (-1)^n b(\omega d(a_n a_{n+1}) - \omega a_n da_{n+1} - a_{n+1} \omega d(a_n)) \\ &= -[\omega, a_n a_{n+1}] + [a_{n+1} \omega, a_n] + [\omega a_n, a_{n+1}] \\ &= -\omega a_n a_{n+1} + a_n a_{n+1} \omega + a_{n+1} \omega a_n - a_n a_{n+1} \omega \\ &\quad + \omega a_n a_{n+1} - a_{n+1} \omega a_n = 0. \end{aligned}$$

DEFINITION 22.9. The homology of the chain complex $(\Omega^n(A), b)$ is called the *Hochschild homology* of the unital algebra A and denoted by $\mathrm{HH}_n(A)$.

We are going to compute this homology theory using projective resolutions and to relate it to certain Hochschild cohomology groups.

DEFINITION 22.10. The *commutator quotient* of an A -bimodule M is the quotient of M by the linear span of $[a, m] := am - ma$ for all $a \in A$, $m \in M$.

THEOREM 22.11. Let $P_\bullet \rightarrow A$ be a projective A -bimodule resolution of A . Then the Hochschild homology of A is naturally isomorphic to the homology of the chain complex of commutator quotients $P_\bullet / [P_\bullet, A]$.

PROOF. Recall that any two projective bimodule resolutions are chain homotopy equivalent. Since taking commutator quotients is functorial, the homology of the chain complex $P_\bullet / [P_\bullet, A]$ does not depend on the choice of resolution. (Compare the proof of Theorem 17.4.) Hence it suffices to find a projective bimodule resolution of A whose associated commutator quotient complex is $(\Omega^n(A), b)$. In fact, the bar resolution used in the proof of Theorem 17.4 works fine here as well. The commutator quotient of $A \otimes \bar{A}^{\otimes n} \otimes A$ is $A \otimes \bar{A}^{\otimes n}$, mapping the class of $a_0 \otimes \dots \otimes a_{n+1}$ in $A \otimes \bar{A}^{\otimes n} \otimes A$ to $a_{n+1} a_0 \otimes a_1 \otimes \dots \otimes a_n$ in $A \otimes \bar{A}^{\otimes n}$ and back by

$$a_0 \otimes a_1 \otimes \dots \otimes a_n \mapsto [a_0 \otimes a_1 \otimes \dots \otimes a_n \otimes 1].$$

The boundary map b' of the bar resolution maps $[a_0 \otimes a_1 \otimes \dots \otimes a_n \otimes 1]$ to the alternating sum of $[a_0 \otimes \dots \otimes a_j a_{j+1} \otimes a_n \otimes 1]$ and $[a_0 \otimes \dots \otimes a_{n-1} \otimes a_n] = [a_n a_0 \otimes \dots \otimes a_{n-1} \otimes 1]$. Hence b' induces the boundary map $b: \Omega^n(A) \rightarrow \Omega^{n-1}(A)$ on the commutator quotient. \square

Thus Hochschild homology can be computed by the same tools that are used to compute Hochschild cohomology.

EXERCISE 22.12. Let A be a unital algebra and let M be an A -bimodule. Its commutator quotient is naturally isomorphic to $M \otimes_{A,A} A$, that is, to the quotient of $M \otimes A$ by the relations $a \cdot m \cdot b \otimes c \sim m \otimes b \cdot c \cdot a$ for all $m \in M$, $a, b, c \in A$.

EXERCISE 22.13. Let A be a \mathbb{C} -algebra. Show that the dual vector space of $\mathrm{HH}_n(A)$ is the Hochschild cohomology $\mathrm{HH}^n(A, A^*)$ with coefficients in the A -bimodule $A^* := \mathrm{Hom}(A, \mathbb{C})$ with the bimodule structure $a \cdot f \cdot b(c) := f(b \cdot c \cdot a)$.

When we are dealing with algebras of smooth functions, then we must complete the spaces of differential forms to spaces of smooth functions on suitable product spaces. More precisely, we let $\Omega^n(C^\infty M)$ be the quotient space of the space of all smooth functions $M^{n+1} \rightarrow \mathbb{C}$ by the closed linear span of the subspaces of functions that are constant in the j th coordinate direction for some $j \neq 0$. If we pick a base point $x \in M$, then we may identify this quotient space with the subspace of all smooth functions $f: M^{n+1} \rightarrow \mathbb{C}$ with $f(x_0, \dots, x_n) = 0$ if $x_j = x$ for some $j \neq 0$.

THEOREM 22.14. *Complete $\Omega^n(C^\infty M)$ as explained above. Then*

$$\mathrm{HH}_n(C^\infty M) = \mathrm{H}^n(\Omega^n(C^\infty M, b)) \cong \Omega^n(M)$$

is the space of differential n -forms on M .

PROOF. This is proven in the same way as Theorem 19.2 about the Hochschild cohomology of $C^\infty(M)$. \square

23. Towards periodic cyclic homology

This section explains an obvious but wrong Ansatz to generalise the de Rham complex to noncommutative algebras. This Ansatz is not a waste of time, however, because the homology groups defined in this way occur in a certain computational scheme for the periodic cyclic homology. We have seen above that the n th Hochschild homology of the algebra $A = C^\infty(M)$ is naturally isomorphic to the space $\Omega^n(M)$ of differential n -forms on M . Therefore, we seek an operator $\mathrm{HH}_n(A) \rightarrow \mathrm{HH}_{n+1}(A)$ for a noncommutative algebra A that could play the role of the de Rham boundary map for a smooth manifold. We cannot use the map d because it does not induce a map on the b -homology of $\Omega^n(A)$. That is, $b\omega = 0$ does not imply $bd\omega = 0$. We want, however, to use the two operators b and d to build it. So we carefully examine the algebra of operators on $\Omega(A)$ generated by these two operators. More precisely, we also implicitly use the grading operators that split $\Omega(A)$ into the direct sum $\bigoplus_{n \in \mathbb{N}} \Omega^n(A)$. We define a second boundary map $B: \Omega^n(A) \rightarrow \Omega^{n+1}(A)$ that anticommutes with b . As a result, the subspaces $\ker b$ and $\mathrm{im} b$ of $\Omega^n(A)$ are B -invariant. So B induces maps $B_*: \mathrm{HH}_n(A) \rightarrow \mathrm{HH}_{n+1}(A)$ for $n \in \mathbb{N}$. These satisfy $B_*^2 = 0$ because $B^2 = 0$. In the second subsection, we compute the cohomology of the cochain complex $(\mathrm{HH}_*(A), B_*)$ for the polynomial algebra in two generators and for the Weyl algebra. This is not so hard, based on our previous computations of Hochschild (co)homology for these two algebras.

23.1. The Karoubi operator. In this section, we examine the relations among the operators b and d . Both operators b and d are boundary maps, that is,

$$b^2 = 0, \quad d^2 = 0.$$

To measure to what extent they anti-commute, we introduce the *Karoubi operator*

$$(23.1) \quad \kappa := 1 - [d, b] = 1 - (db + bd).$$

Here and in the following, the brackets denote *graded* commutators. More explicitly,

$$(23.2) \quad \begin{aligned} \kappa(\omega dx) &= \omega dx - (-1)^n d([x, \omega]) - (-1)^{n+1} [x, d\omega] \\ &= \omega dx - (-1)^n [dx, \omega] = (-1)^{n-1} dx \cdot \omega \end{aligned}$$

for $\omega \in \Omega^{n-1}(A)$, $x \in A$. We have used here that d is a graded derivation.

We also define an operator $B: \Omega(A) \rightarrow \Omega(A)$ of degree $+1$ by

$$(23.3) \quad B := \sum_{j=0}^n \kappa^j \circ d \quad \text{on } \Omega^n(A)$$

or, more explicitly

$$B(x_0 dx_1 \dots dx_n) = \sum_{j=0}^n (-1)^{j^n} dx_j \dots dx_n dx_0 \dots dx_{j-1}.$$

Our next goal is to establish that the operators b and B satisfy

$$b^2 = 0, \quad B^2 = 0, \quad [b, B] := bB + Bb = 0.$$

Therefore, B maps the subspaces $\ker b$ and $\operatorname{im} b$ of $\Omega(A)$ into themselves and induces a map $B_*: \operatorname{HH}_n(A) \rightarrow \operatorname{HH}_{n+1}(A)$ on Hochschild homology. This map generalises the de Rham boundary map for algebras of smooth functions. First, we derive some relations satisfied by the operators b , d and κ .

To begin with, κ is chain homotopic to 1 with respect to either boundary map b or d by its very definition. Therefore, κ is a chain map:

$$\kappa d = d\kappa, \quad \kappa b = b\kappa.$$

Thus κ commutes with all operators that we can construct out of b and d . Then $d^2 = 0$ implies $B^2 = 0$. Notice also that κ is an operator of degree 0, that is, κ maps $\Omega^n(A)$ to $\Omega^n(A)$ for all $n \in \mathbb{N}$.

Equation (23.2) implies

$$\kappa(dx_1 dx_2 \dots dx_n) = (-1)^{n-1} dx_n dx_1 dx_2 \dots dx_{n-1}.$$

Hence κ^n acts identically on $dx_1 dx_2 \dots dx_n$. Then

$$d\kappa^{n+1} = \kappa^{n+1}d = d \quad \text{on } \Omega^n(A).$$

The first equality follows because κ and d commute. This equation explains why the sum defining the operator B runs from 0 to n . And it implies

$$\kappa B = B\kappa = B.$$

We also compute

$$\begin{aligned} \kappa^n(x_0 dx_1 dx_2 \dots dx_n) &= (-1)^{n-1} \kappa^{n-1}(dx_n \cdot x_0 dx_1 \dots dx_{n-1}) \\ &= \dots = (-1)^{(n-1)^n} dx_1 dx_2 \dots dx_n \cdot x_0 = dx_1 dx_2 \dots dx_n \cdot x_0. \end{aligned}$$

The definition of the Hochschild boundary as a commutator yields

$$b\kappa^n d(x_0 dx_1 dx_2 \dots dx_n) = (-1)^{n-n} b(dx_1 dx_2 \dots dx_n dx_0) = -[x_0, dx_1 dx_2 \dots dx_n].$$

A comparison of the last two computations shows

$$(23.4) \quad 1 + b\kappa^n d = \kappa^n \quad \text{on } \Omega^n(A).$$

If we compose (23.4) on the right with d and use $d^2 = 0$, we get $d = \kappa^n d$ on $\Omega^{n-1}(A)$, which we already know. If we compose (23.4) on the left with b and use $b^2 = 0$, we get the new relation

$$\kappa^n b = b\kappa^n = b \quad \text{on } \Omega^n(A).$$

If we compose (23.4) on either side with κ and use that κ commutes with b and d and $\kappa^{n+1}d = d$ on $\Omega^n(A)$, we get

$$\kappa^{n+1} = \kappa + b\kappa^{n+1}d = \kappa + bd = 1 - (bd + db) + bd = 1 - db.$$

So

$$(23.5) \quad db = 1 - \kappa^{n+1}, \quad bd = \kappa^{n+1} - \kappa \quad \text{on } \Omega^n(A).$$

Moreover, (23.4) and $d\kappa^{n+1} = d$ imply

$$(\kappa^n - 1)(\kappa^{n+1} - 1) = b\kappa^n d(\kappa^{n+1} - 1) = b\kappa^n (d\kappa^{n+1} - d) = 0 \quad \text{on } \Omega^n(A).$$

Any operator on $\Omega(A)$ of degree zero that we can construct out of b and d must be a polynomial in db and bd because $d^2 = 0$ and $b^2 = 0$. Equation (23.5) shows

that any such operator is already a polynomial in κ . In particular, Bb and bB are polynomials in κ :

$$\begin{aligned} Bb &= \sum_{j=0}^{n-1} \kappa^j db = \sum_{j=0}^{n-1} \kappa^j (1 - \kappa^{n+1}) = \sum_{j=0}^{n-1} \kappa^j - \sum_{j=n+1}^{2n} \kappa^j, \\ bB &= \sum_{j=0}^n bd\kappa^j = \sum_{j=0}^n (\kappa^{n+1} - \kappa)\kappa^j = -\sum_{j=1}^{n+1} \kappa^j + \sum_{j=n+1}^{2n+1} \kappa^j \end{aligned}$$

on $\Omega^n(A)$. So

$$Bb + bB = 1 - \kappa^n - \kappa^{n+1} + \kappa^{2n+1} = (1 - \kappa^n)(1 - \kappa^{n+1}) = 0 \quad \text{on } \Omega^n(A).$$

Thus we arrive finally at the relation $Bb + bB = 0$ that we wanted to establish.

Since B and b anti-commute, the subspaces $\ker b$ and $\operatorname{im} b$ of $\Omega^n(A)$ are B -invariant. So we get induced maps $B_*: \operatorname{HH}_n(A) \rightarrow \operatorname{HH}_{n+1}(A)$ for all $n \in \mathbb{N}$. These satisfy $B_*^2 = 0$ because $B^2 = 0$. Thus we get a cochain complex $(\operatorname{HH}_n(A), B_*)$. It generalises the de Rham complex for a smooth manifold.

23.2. Computations for polynomials and for the Weyl algebra. Now we are going to compute the cohomology of the cochain complex $(\operatorname{HH}_n(A), B_*)$ for the polynomial algebra in two generators and for the Weyl algebra.

We have seen that the Hochschild homology $\operatorname{HH}_n(A)$ generalises the space of differential forms on a smooth manifold, that is, $\operatorname{HH}_n(\mathbb{C}^\infty M) \cong \Omega^n(M)$ (provided tensor products are completed). Similarly, $\operatorname{HH}_n(\mathbb{C}[x, y])$ is the space of polynomial differential forms on \mathbb{R}^2 , that is,

$$\operatorname{HH}_n(\mathbb{C}[x, y]) \cong \begin{cases} \mathbb{C}[x, y], & \text{if } n = 0, \\ \mathbb{C}[x, y] dx \oplus \mathbb{C}[x, y] dy, & \text{if } n = 1, \\ \mathbb{C}[x, y] dx \wedge dy, & \text{if } n = 2, \\ 0, & \text{if } n > 2. \end{cases}$$

In order to determine the action of B_* on Hochschild homology, we must write down an explicit isomorphism between differential forms and $\ker b / \operatorname{im} b$ on $\Omega^n(\mathbb{C}[x, y])$.

The following recipe works for any smooth manifold: we map a classical differential form $f_0 df_1 \wedge df_2 \wedge \dots \wedge df_n$ to the noncommutative differential form

$$\natural(f_0 df_1 \wedge df_2 \wedge \dots \wedge df_n) := \sum_{\sigma \in S_n} (-1)^{|\sigma|} f_0 df_{\sigma(1)} df_{\sigma(2)} \dots df_{\sigma(n)},$$

where S_n denotes the symmetric group on n letters.

EXERCISE 23.6. Check that $b \circ \natural = 0$, so that \natural induces a map $\Omega^n(M) \rightarrow \operatorname{HH}_n(\mathbb{C}^\infty(M))$ for any smooth manifold M , and check that $B \circ \natural = \natural \circ d$, where d now denotes the de Rham boundary on differential forms.

In fact, \natural induces an isomorphism $\Omega^n(M) \cong \operatorname{HH}_n(\mathbb{C}^\infty(M))$. This is plausible because of Theorem 22.14.

Analogous computations work for algebras of polynomials instead of algebras of smooth functions. Hence the action of B_* on $\operatorname{HH}_n(\mathbb{C}[x, y])$ corresponds to the de Rham boundary on differential forms. We can use the above isomorphism to identify the action of B_* on differential forms because $B \circ \natural = \natural \circ d$, that is, our isomorphism \natural intertwines the de Rham differential and the map B . Now we recall the Poincaré Lemma, which asserts that the de Rham cohomology of \mathbb{R}^2 is trivial except in dimension 0, where we get the subspace of constant functions. This has an analogue for the polynomial algebra:

LEMMA 23.7. *The cohomology of the cochain complex*

$$0 \rightarrow \mathbb{C}[x, y] \xrightarrow{d} \mathbb{C}[x, y] dx \oplus \mathbb{C}[x, y] dy \xrightarrow{d} \mathbb{C}[x, y] dx \wedge dy \rightarrow 0$$

vanishes except in degree 0, where it is spanned by the constant functions.

PROOF. We use the basis of monomials

$$x^n y^m, \quad x^{n-1} y^m dx, \quad x^n y^{m-1} dy, \quad x^{n-1} y^{m-1} dx \wedge dy.$$

These four monomials for fixed $n, m \in \mathbb{N}$ span a subspace that is invariant under the de Rham differential, and we compute

$$\begin{aligned} d(x^n y^m) &= nx^{n-1} y^m dx + mx^n y^{m-1} dy, \\ d(x^{n-1} y^m dx) &= -mx^{n-1} y^{m-1} dx \wedge dy, \\ d(x^n y^{m-1} dy) &= nx^{n-1} y^{m-1} dx \wedge dy. \end{aligned}$$

Inspection shows that this 4-dimensional chain complex is exact for $n, m \neq 0$. If $n = 0$ or $m = 0$, then two of the above four monomials disappear. We remain with an exact 2-dimensional chain complex unless both $n, m = 0$, when we get just one monomial 1. \square

As a result, the homology of the chain complex $(\mathrm{HH}_*(\mathbb{C}[x, y]), B_*)$ is

$$H_n(\mathrm{HH}_*(\mathbb{C}[x, y]), B_*) \cong \begin{cases} \mathbb{C} & \text{if } n = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Now we compute this homology for the Weyl algebra A . We have determined a free bimodule resolution of A when we computed its Hochschild cohomology. Now we use the same resolution to compute the Hochschild homology $\mathrm{HH}_*(A)$. Theorem 22.11 shows that it is the homology of the chain complex

$$0 \rightarrow A \xrightarrow{(\mathrm{ad}_p, -\mathrm{ad}_q)} A \oplus A \xrightarrow{(\mathrm{ad}_q, \mathrm{ad}_p)} A \rightarrow 0,$$

which lives in degrees 0–2. Here we use that the map $A \otimes A \rightarrow A$, $a \otimes b \mapsto ba$, identifies the commutator quotient of $A \otimes A$ with A . Exactly the same chain complex already appeared in our Hochschild cohomology computations (such a phenomenon is called *duality*). Our previous computations yield

$$\mathrm{HH}_n(A) \cong \begin{cases} \mathbb{C} & \text{if } n = 2, \\ 0 & \text{otherwise.} \end{cases}$$

Since already the Hochschild homology is supported in only one degree, the boundary map B_* has to vanish. Hence

$$H_n(\mathrm{HH}_*(A), B_*) \cong \begin{cases} \mathbb{C} & \text{if } n = 2, \\ 0 & \text{otherwise.} \end{cases}$$

24. Periodic cyclic homology

Our computations for the polynomial algebra and the Weyl algebra show that the cochain complexes $(\mathrm{HH}_*(A), B_*)$ for them have non-isomorphic cohomology. Is there another cohomology theory for noncommutative algebras that is invariant under such deformation quantisations? In particular, it should give the same result for the Weyl algebra and the polynomial algebra. The homology of the cochain complex $(\mathrm{HH}_*(A), B_*)$ is not yet the right invariant.

When we replace a chain complex by its homology, then we forget much additional information. This is intended, of course: we want to get reasonably simple answers that we can understand. But when we do several-step constructions with chain complexes, then it is a bad idea to take homology in intermediate

steps because the “approximation errors” can add up and produce systematically wrong results. This is analogous to the problem with rounding errors in numerical computations. In that context, it is advisable to store intermediate results with higher precision. In our context, we get a better behaved homology theory for algebras by directly constructing a chain complex out of the operators b and B on $\Omega(A)$, without taking homology in intermediate steps. Namely, we form a complex with boundary map $b + B$. Since this map is inhomogeneous, we make $\Omega(A)$ 2-periodic and consider

$$\Omega^{\text{even}}(A) := \prod_{n=0}^{\infty} \Omega^{2n}(A), \quad \Omega^{\text{odd}}(A) := \prod_{n=0}^{\infty} \Omega^{2n+1}(A),$$

and the 2-periodic chain complex

$$(24.1) \quad \cdots \rightarrow \Omega^{\text{even}}(A) \xrightarrow{B+b} \Omega^{\text{odd}}(A) \xrightarrow{B+b} \Omega^{\text{even}}(A) \xrightarrow{B+b} \Omega^{\text{odd}}(A) \rightarrow \cdots .$$

DEFINITION 24.2. The homology of the chain complex (24.1) is called the *periodic cyclic homology* of the algebra A and denoted by $\text{HP}_*(A)$.

This theory is called “periodic” because it is 2-periodic by its very definition. Many computations and applications show that this is a good way to generalise de Rham cohomology from smooth manifolds to noncommutative algebras. We have motivated its definition through the preliminary study of the complex $(\text{HH}_*(A), B_*)$. We will compute periodic cyclic homology in only a few cases. Many more computations are known. Before we turn to examples, we must discuss one issue with the above definition. Why did we take direct products instead of direct sums to define $\Omega^{\text{even}}(A)$ and $\Omega^{\text{odd}}(A)$? The main reason is that this definition is the one that works. It can be shown that the corresponding chain complex with direct sums instead of direct products is contractible for any algebra A . Thus it is of no interest. In contrast, the direct product turns out to have an interesting homology.

A good way to explain the need to take direct products is that we would like to approximate the chain complex (24.1) by smaller ones. The following theorem shows how this can be done, and it works for the direct product, but not for the direct sum. Let

$$\mathcal{F}_n := b(\Omega^n(A)) \times \prod_{k=n}^{\infty} \Omega^k(A).$$

This subspace is b -invariant by definition and B -invariant because B increases degrees by 1.

THEOREM 24.3. *Assume that $\text{HH}_N(A) = 0$ for all $N \geq n$. Then the chain complex \mathcal{F}_n is contractible. So $\text{HP}_*(A)$ is isomorphic to the homology of the truncated chain complex*

$$\left(\prod_{k \in \mathbb{N}} \Omega^k(A) / \mathcal{F}_{n+1}, b + B \right) \cong \left(\prod_{k=0}^{n-1} \Omega^k(A) \times \frac{\Omega^n(A)}{b(\Omega^{n+1}(A))}, b + B \right).$$

PROOF. The subquotient $\mathcal{F}_N / \mathcal{F}_{N+1}$ is $b(\Omega^N A) \oplus \Omega^N A / b(\Omega^{N+1} A)$ with the boundary map b ; B vanishes on this subquotient because it maps $b(\Omega^N A)$ into $b(\Omega^{N+1} A)$. Hence the homology of $\mathcal{F}_N / \mathcal{F}_{N+1}$ is $\text{HH}_N(A) = 0$ for $N \geq n$. Then Lemma 18.3 implies that the homology of the quotient $\mathcal{F}_n / \mathcal{F}_N$ vanishes for all $N \geq n$. Next, we claim that the projective limit $\varprojlim \mathcal{F}_n / \mathcal{F}_N$ is exact as well. It is rather easy to prove that the product $\prod \mathcal{F}_n / \mathcal{F}_N$ is exact. Then there is an extension of chain complexes

$$\varprojlim \mathcal{F}_n / \mathcal{F}_N \twoheadrightarrow \prod_{N \geq n} \mathcal{F}_n / \mathcal{F}_N \twoheadrightarrow \prod_{N \geq n} \mathcal{F}_n / \mathcal{F}_N,$$

where the second map maps $(x_N)_{N \geq n} \in \prod \mathcal{F}_n/\mathcal{F}_N$ to $(q_N(x_{N+1}) - x_N)$, where $q_N: \mathcal{F}_n/\mathcal{F}_{N+1} \rightarrow \mathcal{F}_n/\mathcal{F}_N$ is the quotient map. This map can be shown to be surjective. Its kernel is the projective limit by definition. Now a variant of Lemma 18.3 shows that $\varprojlim \mathcal{F}_n/\mathcal{F}_N$ inherits completeness from $\prod_{N \geq n} \mathcal{F}_n/\mathcal{F}_N$. The projective limit $\varprojlim \mathcal{F}_n/\mathcal{F}_N$ is exactly \mathcal{F}_n with the boundary map $B + b$. Since \mathcal{F}_n is exact, dividing it out does not change homology by a variant of Lemma 18.3. \square

The filtration defined above establishes a close link between Hochschild homology and periodic cyclic homology. Readers familiar with “spectral sequences” will recall that a filtration on a chain complex induces such a spectral sequence that converges towards its homology. The following remark will only make sense to readers who are already familiar with spectral sequences and homological algebra.

REMARK 24.4. The filtration \mathcal{F}_n on the chain complex that defines periodic cyclic homology generates a spectral sequence. Its k th page consists of the homology groups of the subquotients $\mathcal{F}_n/\mathcal{F}_{n+k}$ for $n \in \mathbb{N}$. In particular, the first page of the spectral sequence contains the homologies of the subquotients $\mathcal{F}_n/\mathcal{F}_{n+1}$. The proof of Theorem 24.3 identifies the homology of $\mathcal{F}_n/\mathcal{F}_{n+1}$ with the n th Hochschild homology of our algebra A . The $k + 1$ -st page of a spectral sequence is gotten from the k th page by taking homology for the k th “differential”, which is a family of maps between the homology groups on the k th page. The first differential of the spectral sequence induced by the filtration \mathcal{F}_n turns out to be the family of maps $B_*: \text{HH}_n(A) \rightarrow \text{HH}_{n+1}(A)$ introduced in the previous section. The reason is the summand B in the boundary map $b + B$ on our chain complex. Thus the second page of our spectral sequence consists of the homology groups $H_n(\text{HH}_*(A), B_*)$. At least if $\text{HH}_N(A) = 0$ for $N \geq n$ for some $n \in \mathbb{N}$, then the spectral sequence above converges to $\text{HP}_*(A)$ by Theorem 24.3. In general, convergence may fail, but there is still a close link between $\text{HP}_*(A)$ and the spectral sequence. So the spectral sequence above clarifies how Hochschild homology and the homology of the chain complex $(\text{HH}_n(A), B_*)$ are the first and the second approximation to $\text{HP}_*(A)$. In many cases, the spectral sequence already becomes constant at the second page, that is,

$$\text{HP}_n(A) \cong \prod_{k \in \mathbb{Z}} H_{n+2k}(\text{HH}_*(A), B_*).$$

This happens, in particular, for the algebras of polynomials and for the Weyl algebra because the entire cohomology of the chain complex $(\text{HH}_*(A), B_*)$ is concentrated in one degree – this forces all higher differentials to vanish.

We conclude with a more down-to-earth way to relate Hochschild and periodic cyclic homology, which replaces the spectral sequence above by a long exact sequence and uses an intermediate theory, the (non-periodic) cyclic homology.

DEFINITION 24.5. Since $b^2 = 0$, $B^2 = 0$ and $bB + Bb = 0$, the squares in the diagram in Figure 2 anti-commute. So this diagram is a bicomplex up to this sign. We call it the *cyclic bicomplex*. The *cyclic homology* $\text{HC}_*(A)$ is the homology of the total complex of this anti-commuting bicomplex. The degree n space in this total complex is

$$\Omega^n A \times \Omega^{n-2} A \times \Omega^{n-4} \times \dots \times \Omega^{n \bmod 2} A,$$

where $n \bmod 2$ is 0 or 1 depending on whether n is even or odd, and the boundary map is $b + B$ on most summands, and just b on the first summand where B would not have the right target.

The first column in the cyclic bicomplex above is $(\Omega^n A, b)$. Its homology is, by definition, the Hochschild homology of A . When we leave out this column, we get a copy of the same bicomplex, but shifted up and left by 1. This gives a shift of

$$\begin{array}{ccccccc}
& \cdots & \longleftarrow & \cdots & \longleftarrow & \cdots & \longleftarrow & \cdots \\
& \downarrow & & \downarrow & & \downarrow & & \downarrow \\
\Omega^3(A) & \xleftarrow{B} & \Omega^2(A) & \xleftarrow{B} & \Omega^1(A) & \xleftarrow{B} & \Omega^0(A) \\
& \downarrow b & & \downarrow b & & \downarrow b & \\
\Omega^2(A) & \xleftarrow{B} & \Omega^1(A) & \xleftarrow{B} & \Omega^0(A) \\
& \downarrow b & & \downarrow b & \\
\Omega^1(A) & \xleftarrow{B} & \Omega^0(A) \\
& \downarrow b & \\
\Omega^0(A)
\end{array}$$

FIGURE 2. The cyclic bicomplex with the boundary maps b and B

degree 2 on the total complex. The resulting extension of chain complexes induces a long exact sequence of homology groups

$$(24.6) \quad \cdots \rightarrow \mathrm{HH}_n(A) \rightarrow \mathrm{HC}_n(A) \xrightarrow{S} \mathrm{HC}_{n-2}(A) \rightarrow \mathrm{HH}_{n-1}(A) \rightarrow \mathrm{HC}_{n-1}(A) \rightarrow \cdots$$

We have not explained this construction, except in the trivial case when two of the three complexes are exact. Then the third one is exact as well by Lemma 18.3. A more general result relates the homology groups of three complexes that form an extension.

The periodicity operator S is used to relate HC to HP . In the cases we are considering, there is an index n such that $S: \mathrm{HC}_{N+2}(A) \rightarrow \mathrm{HC}_N(A)$ is invertible for all $N \geq n$. This follows from the exact sequence above if $\mathrm{HH}_N(A) = 0$ for $N \geq n-1$. Then $\mathrm{HP}_{n \bmod 2}(A) \cong \mathrm{HC}_n(A)$ and $\mathrm{HP}_{1-n \bmod 2}(A) \cong \mathrm{HC}_{n+1}(A)$; this follows from Theorem 24.3.

EXAMPLE 24.7. Let A be the Weyl algebra. Then $\mathrm{HH}_n(A) = 0$ for $n \neq 2$ and $\mathrm{HH}_2(A) \cong \mathbb{C}$. And $\mathrm{HC}_n(A) = 0$ for $n < 0$ is trivial. Now the exact sequence in (24.6) allows to compute $\mathrm{HC}_n(A)$ recursively for all $n \in \mathbb{N}$. Namely, the exactness of (24.6) forces $\mathrm{HC}_n(A) = 0$ for $n \leq 1$ and $\mathrm{HC}_n(A) \cong \mathbb{C}$ for $n \geq 2$, with an invertible map $S: \mathrm{HC}_{n+2}(A) \rightarrow \mathrm{HC}_n(A)$ for $n \geq 2$. A similar computation works whenever Hochschild homology is concentrated in one degree.

THEOREM 24.8. *Let M be a smooth manifold. Let $n \in \mathbb{N}$. Then*

$$\begin{aligned}
\mathrm{HC}_n(\mathrm{C}^\infty M) &\cong \Omega^n M / \mathrm{d}(\Omega^{n-1} M) \oplus \mathrm{H}_{\mathrm{dR}}^{n-2}(M) \oplus \mathrm{H}_{\mathrm{dR}}^{n-4}(M) \oplus \mathrm{H}_{\mathrm{dR}}^{n-6}(M) \oplus \cdots, \\
\mathrm{HP}_n(\mathrm{C}^\infty M) &\cong \bigoplus_{k \in \mathbb{Z}} \mathrm{H}_{\mathrm{dR}}^{n-2k}(M).
\end{aligned}$$

Here we use completed tensor products in the definition of $\Omega^n(\mathrm{C}^\infty M)$.

PROOF. There are obvious maps

$$\natural: \Omega^n(\mathrm{C}^\infty M) \rightarrow \Omega^n(M), \quad \natural(f_0 \mathrm{d}f_1 \cdots \mathrm{d}f_n) := f_0 \mathrm{d}f_1 \wedge \cdots \wedge \mathrm{d}f_n.$$

They intertwine the vertical boundary map b on $\Omega^n(\mathrm{C}^\infty M)$ and the zero map on $\Omega^n(M)$, and they induce an isomorphism on b -homology. In addition, they intertwine the map d on $\Omega^n(\mathrm{C}^\infty M)$ and the de Rham boundary map on $\Omega^n(M)$. Therefore, they also intertwine the boundary map B with a multiple of the de Rham boundary map, where the constant factor depends on the degree n . Since \natural induces an isomorphism on the vertical homology, it follows that it induces an isomorphism on the homology

of the total complexes. This is a generalisation of Lemma 18.3. As a consequence, $\mathrm{HC}_*(C^\infty M)$ is isomorphic to the homology of the total complex of the much simpler bicomplex in Figure 3. This explains the formulas for $\mathrm{HC}_*(C^\infty M)$ in the theorem. This yields the assertion about $\mathrm{HC}_*(C^\infty M)$. If $n > \dim M$, then $\mathrm{HH}_n(C^\infty M) = 0$.

$$\begin{array}{ccccccc}
 & \dots & \longleftarrow & \dots & \longleftarrow & \dots & \longleftarrow & \dots \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 \Omega^3(M) & \xleftarrow{d} & \Omega^2(M) & \xleftarrow{d} & \Omega^1(M) & \xleftarrow{d} & \Omega^0(M) & \\
 \downarrow 0 & & \downarrow 0 & & \downarrow 0 & & & \\
 \Omega^2(M) & \xleftarrow{d} & \Omega^1(M) & \xleftarrow{d} & \Omega^0(M) & & & \\
 \downarrow 0 & & \downarrow 0 & & & & & \\
 \Omega^1(M) & \xleftarrow{d} & \Omega^0(M) & & & & & \\
 \downarrow 0 & & & & & & & \\
 \Omega^0(M) & & & & & & &
 \end{array}$$

FIGURE 3. A bicomplex computing the cyclic homology of smooth functions on a manifold

Therefore, the periodicity map $S: \mathrm{HC}_{n+2}(C^\infty M) \rightarrow \mathrm{HC}_n(C^\infty M)$ is invertible if $n > \dim M$. Therefore, $\mathrm{HC}_{\dim M+1}(C^\infty M)$ and $\mathrm{HC}_{\dim M}(C^\infty M)$ agree with the periodic cyclic homology of $C^\infty M$. \square

The two examples treated above show a general route how to compute the periodic cyclic homology of an algebra. First find a small projective bimodule resolution of the algebra. Use it to compute the Hochschild homology of the algebra. Then use the long exact sequence above or a map of bicomplexes to compute the cyclic homology from the Hochschild homology. Finally, the periodic cyclic homology is a kind of limit of cyclic homology groups.

24.1. Periodic cyclic cocycles, traces and K-theory. We end this course by sketching how to map the K-theory of an algebra to its periodic cyclic homology and how to construct interesting maps $\mathrm{HP}_*(A) \rightarrow \mathbb{C}$. This is an important application of periodic cyclic homology. In the case of smooth manifolds, the constructions above produce the Chern character from the K-theory of a smooth manifold to the de Rham cohomology of the manifold.

So far, we have defined periodic cyclic homology for unital algebras. And our definition is clearly functorial for unital algebra homomorphisms. We need units because we used $A := A/\mathbb{C} \cdot 1$ to define noncommutative differential forms over A . Now we extend the theory to non-unital algebras A . Namely, let $A^+ := A \oplus \mathbb{C}$ with the multiplication where the second summand behaves like constant multiples of the identity map. This is a unital algebra that contains A as an ideal, and the projection $A^+ \rightarrow \mathbb{C}$ is a unital algebra homomorphism. We define

$$\mathrm{HH}_*(A) := \ker(\mathrm{HH}_*(A^+) \rightarrow \mathrm{HH}_*(\mathbb{C})), \quad \mathrm{HC}_*(A) := \ker(\mathrm{HC}_*(A^+) \rightarrow \mathrm{HC}_*(\mathbb{C})),$$

and similarly for $\mathrm{HP}_*(A)$. This extends the three homology theories studied above to non-unital algebras. The same idea is used to extend the definition of K-theory from unital to non-unital algebras. The definition is functorial for algebra homomorphisms that need not preserve the unit.

In the unital case, we have seen that HH_* is Morita invariant and hence stable under tensoring with matrices. This is inherited by HC_* and HP_* . This follows from

the exact sequence (24.6), which is natural for non-unital algebra homomorphisms, and from the “Five Lemma” in homological algebra. For non-unital algebras, HH_* and HC_* may fail to be Morita invariant, even stability for matrices may fail:

EXAMPLE 24.9. Let A be a finite-dimensional vector space with the zero multiplication. Then the multiplication on $\mathbb{M}_n A$ is still zero. Since no element of A is a commutator, we see that $\mathrm{HH}_0(A) = A$ and $\mathrm{HH}_0(\mathbb{M}_n A) \cong \mathbb{M}_n A$. And $\mathrm{HH}_0 \cong \mathrm{HC}_0$. Hence neither HH_0 nor HC_0 is stable for matrix algebras.

This problem goes away for HP_* . There is, however, no time to explain this properly in this course. We merely state the result without proof.

PROPOSITION 24.10. *Morita equivalent unital algebras have isomorphic HP_* . For a non-unital algebra A , HP_* is still stable for matrix algebras, that is, if $n \geq 1$, then $\mathrm{HP}_*(\mathbb{M}_n A) \cong \mathrm{HP}_*(A)$.*

THEOREM 24.11. *Let A be a Banach algebra. There are canonical maps $K_*(A) \rightarrow \mathrm{HP}_*(A)$ for $* = 0, 1$.*

PROOF. By definition, elements in $K_0(A)$ are represented by formal differences of two idempotent elements $p, q \in \mathbb{M}_n A^+$ with the property that $p-1, q-1 \in \mathbb{M}_n A$. We may turn the idempotent elements p, q into non-unital algebra homomorphisms $\mathbb{C} \rightarrow A^+$. These induce maps $p_*, q_*: \mathrm{HP}_0(\mathbb{C}) \rightarrow \mathrm{HP}_0(A^+)$. Now a simple computation using (24.6) shows that $\mathrm{HP}_0(\mathbb{C}) \cong \mathbb{C}$. The difference $p_* - q_*$ takes values in $\mathrm{HP}_0(A) \subseteq \mathrm{HP}_0(A^+)$, and we map the K-theory class of the pair (p, q) to the image of $1 \in \mathbb{C}$ under the map $p_* - q_*: \mathbb{C} \cong \mathrm{HP}_0(\mathbb{C}) \rightarrow \mathrm{HP}_0(A)$. To prove that this gives a well defined map $K_0(A) \rightarrow \mathrm{HP}_0(A)$, we may proceed in two ways, depending on the equivalence relation on idempotents that we use. First, a Murray–von Neumann equivalence between two idempotents p, q in $\mathbb{M}_n A^+$ defines a homomorphism $\mathbb{M}_2 \mathbb{C} \rightarrow \mathbb{M}_n A^+$ that agrees with p_* and q_* on the two corners. Since $\mathrm{HP}_0(\mathbb{M}_2 \mathbb{C}) \cong \mathbb{C}$, the two corner embeddings $\mathbb{C} \rightarrow \mathbb{M}_2 \mathbb{C}$ induce the same map $\mathrm{HP}_0(\mathbb{C}) \rightarrow \mathrm{HP}_0(\mathbb{M}_2 \mathbb{C})$. Therefore, p_* and q_* induce the same map $\mathrm{HP}_0(\mathbb{C}) \rightarrow \mathrm{HP}_0(A^+)$ if p, q are Murray–von Neumann equivalent. Secondly, a smooth homotopy of idempotents $p_* \sim q_*$ induces a smooth homotopy between the resulting homomorphisms p_*, q_* . And it is known that smoothly homotopic homomorphisms induce the same map in HP_* .

By definition, an element of $K_1(A)$ is represented by an invertible element T in $\mathbb{M}_n A^+$ with $T - 1 \in \mathbb{M}_n A$ for some $n \in \mathbb{N}$. The element T creates a unital algebra homomorphism $\mathbb{C}[t, t^{-1}] \rightarrow \mathbb{M}_n A^+$. We have shown that the algebra $\mathbb{C}[t, t^{-1}]$ is quasi-free (see Proposition 20.9). This implies $\mathrm{HH}_n(\mathbb{C}[t, t^{-1}]) = 0$ for $n \geq 2$. And it is not hard to compute $\mathrm{HH}_n(\mathbb{C}[t, t^{-1}])$ for $n = 0, 1$. Then it follows that $\mathrm{HP}_1(\mathbb{C}[t, t^{-1}]) \cong \mathbb{C}$. Therefore, a homomorphism $\mathbb{C}[t, t^{-1}] \rightarrow \mathbb{M}_n A^+$ induces an element in $\mathrm{HP}_1(A)$ exactly as for $K_0(A)$. The smooth homotopy invariance of HP_* implies that smoothly homotopic invertible elements induce the same class in $\mathrm{HP}_1(A)$. This implies that the map $K_1(A) \rightarrow \mathrm{HP}_1(A)$ is well defined (on topological K-theory). \square

Next, we are going to describe a method to define maps $f: \Omega^n(A) \rightarrow \mathbb{C}$ that satisfy $f \circ b = 0$ and $f \circ B = 0$. As a result, f induces a map $\mathrm{HP}_n(A) \rightarrow \mathbb{C}$. Together with the construction in Theorem 24.11, we get linear maps $K_*(A) \rightarrow \mathbb{C}$ for $* \equiv n \pmod{2}$. Constructing such maps was the main motivation for Alain Connes to introduce (periodic) cyclic homology in the 1980s.

Let (C, ∂) be a differential graded unital algebra and let $\varphi: A \rightarrow C_0$ be a unital algebra homomorphism from A to the degree-zero part of C . This induces a differential graded algebra homomorphism $\varphi_*: \Omega(A) \rightarrow C$ by Proposition 22.7. Let $\tau: C_n \rightarrow \mathbb{C}$ be a linear map that is a closed graded trace. That is, $\tau(a \cdot b) = (-1)^{k \cdot (n-k)} \tau(b \cdot a)$ if $a \in C_k, b \in C_{n-k}$ and $\tau \circ \partial = 0$. Then $\tau \circ \varphi_*: \Omega^n(A) \rightarrow \mathbb{C}$ is a

closed graded trace as well because φ_* is a homomorphism of differential graded algebras. In particular, $\tau \circ d = 0$,

$$(\tau \circ \varphi_*) \circ b(\omega dx) = (-1)^n (\tau \circ \varphi_*)([\omega, x]) = 0$$

and $\tau \circ \varphi_* \circ \kappa = \tau \circ \varphi_*$ by (23.2). Thus $\tau \circ \varphi_* \circ B = 0$ as well.

EXAMPLE 24.12. Let M be a smooth, compact manifold (without boundary). The differential forms on M form a graded-commutative differential graded algebra $\Omega^*(M)$. Thus any linear map $\Omega^*(M) \rightarrow \mathbb{C}$ is a graded trace. If M is oriented, then integration over M is a closed graded trace on $\Omega^*(M)$ by Stokes' Theorem, which says here that $\int_M d\omega = 0$ for all $\omega \in \Omega^{\dim(M)-1}(M)$. So is $\omega \mapsto \int_M \omega \wedge \eta$ for any closed form η because $d\eta = 0$ implies $\int_M d(\omega) \wedge \eta = \int_M d(\omega \wedge \eta) = 0$. Poincaré duality for smooth compact manifolds implies that any linear functional $H_{\text{dR}}^k(M) \rightarrow \mathbb{C}$ on the k th de Rham cohomology is of the form $\omega \mapsto \int_M (\omega \wedge \eta)$ for some $\eta \in \Omega^{n-k}(M)$ with $d\eta = 0$. Since $\text{HP}_*(C^\infty(M))$ is $H_{\text{dR}}^*(M)$ made 2-periodic, the construction above gives all linear functionals on $\text{HP}_*(C^\infty(M))$.

References

- [1] Gene Abrams, Pere Ara, and Mercedes Siles Molina, *Leavitt path algebras*, Lecture Notes in Mathematics, vol. 2191, Springer, London, 2017. MR3729290
- [2] Francis Borceux, *Handbook of categorical algebra. 2*, Encyclopedia of Mathematics and its Applications, vol. 51, Cambridge University Press, Cambridge, 1994. Categories and structures. MR1313497
- [3] Glen E. Bredon, *Topology and Geometry*, Graduate Texts in Mathematics, vol. 139, Springer, New York, 1997. Corrected third printing of the 1993 original. MR1700700
- [4] Alain Connes, *Noncommutative differential geometry*, Inst. Hautes Études Sci. Publ. Math. **62** (1985), 257–360, available at http://www.numdam.org/item?id=PMIHES_1985__62__41_0. MR823176
- [5] Joachim Cuntz, *A survey of some aspects of non-commutative geometry*, Jahresbericht der DMV **95** (1993), 60–84.
- [6] Niels Grønbaek, *Morita equivalence for self-induced Banach algebras*, Houston J. Math. **22** (1996), no. 1, 109–140, available at <http://www.math.uh.edu/~hjm/restricted/archive/v022n1/0109GRONBAEK.pdf>. MR1434388
- [7] Alexandre Grothendieck, *Produits tensoriels topologiques et espaces nucléaires*, Mem. Amer. Math. Soc., vol. 16, 1955. MR0075539
- [8] Morris W. Hirsch, *Differential topology*, Graduate Texts in Mathematics, vol. 33, Springer, New York, 1976. MR0448362
- [9] Saunders MacLane, *Homology*, Grundlehren Math. Wiss., vol. 114, Springer, Berlin, 1963. MR1344215
- [10] Ralf Meyer, *Smooth and rough modules over self-induced algebras*, Münster J. Math. **4** (2011), 29–51, available at <http://nbn-resolving.de/urn:nbn:de:hbz:6-32449573355>. MR2869252
- [11] Kiiti Morita, *Duality for modules and its applications to the theory of rings with minimum condition*, Sci. Rep. Tokyo Kyoiku Daigaku Sect. A **6** (1958), 83–142. MR0096700
- [12] Emily Riehl, *Category theory in context*, Dover, Mineola, NY, 2016.
- [13] Ahmad Reza Haj Saeedi Sadegh and Nigel Higson, *Euler-like vector fields, deformation spaces and manifolds with filtered structure*, Doc. Math. **23** (2018), 293–325, DOI 10.25537/dm.2018v23.293-325.
- [14] Joseph L. Taylor, *A general framework for a multi-operator functional calculus*, Adv. Math. **9** (1972), 183–252, DOI 10.1016/0001-8708(72)90017-5. MR0328625
- [15] Charles E. Watts, *Intrinsic characterizations of some additive functors*, Proc. Amer. Math. Soc. **11** (1960), 5–8, DOI 10.1090/S0002-9939-1960-0118757-0. MR0118757