Lectures on quantum field theory

February 3, 2009

Chapter 1

Quantum theory

1.1 The quantum description

Quantum field theory aims to give an account of the ultimate constituents of matter and their interactions, but excluding gravitation from the picture except as a given fixed background. It should tell us, for example, how it is that light, though usually best described as electromagmetic waves, sometimes behaves like a stream of particles. The possibility of two such different descriptions of the same reality arises from the nature of quantum theory, so we must begin with the relation between quantum theory and classical physics.

Classical physics — that is to say, physics as it was understood at the beginning of the twentieth century — gives a simple and familiar description of the world. It presupposes that we know about "space" and "time", given to us as two smooth Riemannian manifolds of dimensions three and one. It tells us that the world consists at any time of a collection of point particles located in the space-manifold M, and that they interact by means of "fields" pervading space which are described (roughly) by a number of smooth functions on M. Once M is given, the possible configurations of the world at a given time are the points of an infinite-dimensional manifold $\mathcal{X}$, and classical physics tells us how the configuration evolves with time. The general mathematical form of the evolution is strikingly simple. The trajectory is determined by giving any one of its points and the instantaneous velocity with which it is moving, i.e. by giving a tangent vector to $\mathcal{X}$. Let us write $\mathcal{Y}$ for the bundle of all tangent vectors to $\mathcal{X}$. Classical physics can be summed up in the statement that there is a smooth closed 2-form ω on $\mathcal{Y}$, and a smooth function $H : \mathcal{Y} \to \mathbb{R}$ called the *Hamiltonian*, such that the evolution of $y \in \mathcal{Y}$ is determined by the differential equation — called *Hamilton's equation* —

$$\omega_y(\frac{\mathrm{d}y}{\mathrm{d}t}, \eta) = \mathrm{d}H(y; \eta),$$

for every tangent vector η to $\mathcal{Y}$ at y. (The right-hand side of this equation denotes the gradient of the function H at the point y in the direction η, while the left-hand side denotes the value of the 2-form ω at the point y on the pair of tangent vectors $\frac{dy}{dt}$ and η. It is important that ω is non-degenerate at each point of $\mathcal{Y}$, so that the tangent vector $\frac{dy}{dt}$ is completely determined by giving the left-hand side of the equation for each η.)

A manifold with a such a non-degenerate closed 2-form ω is called a *symplectic* manifold. The symplectic stucture is related to the quantum description of the world by non-commutative algebra, but, within the classical perspective, it is linked to the fact that the trajectories of classical physics are characterized by a *variational* property: the "principle of least action". I shall return to that in §x.

I have given this classical description in a non-relativistic language, treating time and space differently. As far as physics is concerned, there is no fundamental conceptual problem in going over to a relativistic picture with a unified four-dimensional space-time manifold on which the fields are defined, and in which the particles sit as "world-lines": the fact that space and time seem so very different to us can in practice be ignored and left to philosophy. The serious problem with the classical picture is one which at first sight looks merely "technical", but seems to be insuperable: it is the difficulty of mixing point particles with "fields". For the particles are supposed to move under the influence of a field which acts on them, and each particle helps generate the field, making a contribution which inevitably becomes infinite at its own position. But on the other hand the contributions of the individual particles to the total field cannot be separated out: there is just one agglomerated field which has to act on each particle, and is infinite at all the points where we need to know its value.

It is helpful to keep separate the ontological content of classical physics — that the world consists of particles and fields — from the mathematical structure $(\mathcal{Y}, \omega, H)$ consisting of a *symplectic manifold*, i.e. a smooth manifold $\mathcal{Y}$ with a non-degenerate closed 2-form ω, with a Hamiltonian function H. I shall refer to such triples as *Hamiltonian systems*. They can be used to describe not just the whole world but also any number of subsystems and idealizations of subsystems of it.

Quantum theory — as distinct from quantum field theory — is not an alternative to classical physics, as it has no ontological ingredient corresponding to the picture of particles in space-time interacting by means of fields. It retains only the notion of *time* as something understood in advance. Quantum theory is simply the idea that classical Hamiltonian systems are "really" approximations to *quantum* systems. At each time t a quantum system has a "state" which is supposed to be a ray L_t in a complex Hilbert space $\mathcal{H}$, and it has " observables" which are self-adjoint operators in $\mathcal{H}$. But all the information is contained[1] in the

[1] I shall return to this point below.

topological algebra $\mathcal{A}$ of operators generated by the observables, together with the linear map $\theta_t : \mathcal{A} \to \mathbb{C}$ given by

$$\theta_t(a) = \langle \psi \, , \; a \, \psi \, \rangle,$$

where $\psi \in L_t$ is a unit vector in the state-ray.

We can therefore adopt the following point of view, which enables us to put the classical and quantum pictures side-by-side.

A *classical system* is described by giving a symplectic manifold $\mathcal{Y}$ of possible states, and at each time t the *state* $y_t \in \mathcal{Y}$ of the system. The evolution of the system is described by a 1-parameter group $\{u_t\}_{t \in \mathbb{R}}$ of diffeomorphisms of $\mathcal{Y}$, so that $u_t(y_{t'}) = y_{t+t'}$.

A *quantum system* is described by giving a topological $*$-algebra $\mathcal{A}$ of observables [2], and at each time t a positive linear map $\theta_t : \mathcal{A} \to \mathbb{C}$ called the *state*. The evolution of the system is described by a 1-parameter group $\{u_t\}_{t \in \mathbb{R}}$ of $*$-algebra automorphisms of $\mathcal{A}$, so that $\theta_{t'+t} \circ u_t = \theta_{t'}$.

We can say more than this. In the classical case the flow $\{u_t\}$, which preserves the symplectic form, is generated by the Hamiltonian function $H : \mathcal{Y} \to \mathbb{R}$ according to the equation above. The corresponding fact in the quantum case is the evolution equation

$$da/dt = \mathrm{i}\hbar[H, a],$$

where $H \in \mathcal{A}$ is a self-adjoint element, again called the *Hamiltonian.* Here $[H, A]$ denotes the commutator $Ha - aH$ in $\mathcal{A}$, and $\hbar$ is Planck's constant.

The most obvious difference between the classical and quantum pictures is that a smooth manifold of states has been replaced by an algebra, and a point in the state-space has been replaced by a linear form on the algebra. But more fundamental is to realize that neither of the two parallel pictures just presented gives us any kind of description of the world. On the classical side we do get a picture of the world because we have the additional information that the manifold $\mathcal{Y}$ is $T\mathcal{X}$, where $\mathcal{X}$ is the manifold of configurations of particles and fields in space, which we are supposed to know about in advance, and thousands of years of physics have made us able to see the world as a collection of particles and fields. On the quantum side we interpret the system in two stages, first recognizing it as being close to a classical system, and then interpreting the classical system in terms of particles and fields.

Even to compare the classical and quantum pictures we need to know that a *commutative* algebra is the same kind of thing as a space, because a commutative algebra can be regarded as the algebra of functions on its spectrum[3]. When we say that the world is described by the quantum data $(\mathcal{A}, \theta_t)$ what we mean is that we have somehow fixed on a subset which we find especially relevant for us in the algebra $\mathcal{A}$ of observables, part of its attractiveness coming from two properties:

(i) it generates a subalgebra $\mathcal{C}_t$ of $\mathcal{A}$ which is to good enough approximation a commutative algebra $\mathcal{C}_t^{\text{class}}$, and

(ii) the state θ_t when restricted to $\mathcal{C}_t$ is very nearly a homomorphism of algebras $\theta_t^{\text{class}} : \mathcal{C}_t^{\text{class}} \to \mathbb{C}$.

We shall certainly look for many other properties as well; for example, that the chosen observables are not changing too quickly in time.

Once we have chosen $\mathcal{C}_t^{\text{class}}$ and θ_t^{class} we are back in the classical picture, for the commutative algebra $\mathcal{C}_t^{\text{class}}$ is a space of functions on its spectrum $\mathcal{Y}$, and θ_t^{class} is a point $y \in \mathcal{Y}$. We can then let the system evolve, and our description will work well for a time, but sooner or later we shall find that the set of observables we fixed on will no longer nearly commute, or perhaps that $\theta_{t'}|\mathcal{C}_{t'}$ is no longer nearly a homomorphism, and does not give us even approximately a point of $\mathcal{Y}$. The classical picture which we are using has gone out of focus. According to orthodox quantum mechanics there is then no substitute for looking at the world anew and choosing a new classical approximation $(\mathcal{C}_{t'}^{\text{class}}, \theta_{t'}^{\text{class}})$: we must "adjust our sets" to bring a classical picture back into focus. From this point of view the notorious "measurement problem" in quantum mechanics is to explain why there is a way — even a more-or-less canonical way — to do this. The measurement problem attracts endlessly ongoing discussion, and my formulation of it is not quite the usual one. Beyond that, I have nothing substantial to say about it; but fortunately that does not matter for these lectures. My own instinct is to think of the quantum picture as the real one, and to regard classical descriptions as provisional statements which relate only to a suitable selection of observables. But the quantum picture as I have described it can hardly be taken literally, if only because it is so resolutely non-relativistic. My perhaps heretical feeling is that the perspective of quantum field theory offers a more plausible foundation for physics than does traditional quantum mechanics, but, as it cannot accommodate gravity, it still cannot be pressed too far.

Let us at any rate accept that when we have a quantum system $(\mathcal{A}, \theta_t)$ we interpret it in terms of classical approximations. People often speak of the "classical limit", but we must remember that it is only a tiny subset of the whole algebra $\mathcal{A}$ that we take a limit of. Ignoring this point, a convenient mathematical formulation is as follows. We *recognize* $\mathcal{A}$ as a deformation $\mathcal{A} = \mathcal{A}_h$ of a

[3]The *spectrum* of a commutative algebra $\mathcal{C}$ over the complex numbers is the set of all algebra-homomorphisms $\mathcal{C} \to \mathbb{C}$. If the algebra $\mathcal{A}$ has a topology then we consider only *continuous* homomorphisms.

commutative algebra $\mathcal{A}_0$ by fitting it into a 1-parameter family of algebras $\{\mathcal{A}_h\}$ depending smoothly on the small parameter $\hbar$, with $\mathcal{A}_0$ commutative. We can then assume that the algebras have the same underlying vector space, so that the data is a 1-parameter family of multiplications $m_\hbar : \mathcal{A}_0 \times \mathcal{A}_0 \to \mathcal{A}_0$ deforming the given multiplication m_0 of $\mathcal{A}_0$. In this situation we can define a *Poisson bracket* $\{\ ,\ \}$ on $\mathcal{A}_0$ by

$$m_\hbar(f, g) = m_0(f, g) + \mathrm{i}\hbar\{f, g\} + O(\hbar^2).$$

If $\mathcal{A}_0$ is the algebra of smooth functions on a manifold $\mathcal{Y} = \mathrm{Spec}(\mathcal{A}_0)$ then the Poisson bracket, providing it is sufficiently nondegenerate, makes $\mathcal{Y}$ into a symplectic manifold (see Appendix x for more details).

However we look at the matter, we must remember that it makes no sense to say that a given algebra is a deformation of a commutative algebra without specifying a great deal more information. In particular, an algebra can perfectly well be a deformation of many quite different commutative algebras — some simple examples can be found in Appendix x — and so there can be quite different classical descriptions of the same quantum system, for example in terms of particles or in terms of waves.

An important feature of the traditional interpretation of quantum mechanics is that the result $\lambda \in \mathbb{R}$ of "measuring" an observable a — i.e. a self-adjoint element of the algebra $\mathcal{A}$ — is always an element of the spectrum of a, i.e. a generalized eigenvalue. Thus if a quantum system describes a collection of particles then there should be an observable whose eigenvalues are the positive integers, corresponding to the number of particles in a given region of space. It is to define the spectrum of an observable that we need the *topology* on the algebra $\mathcal{A}$ of observables. A number λ belongs to the spectrum of an operator A in Hilbert space precisely when $A - \lambda$ is not invertible, and it is common to define the spectrum of an element a of an arbitrary algebra $\mathcal{A}$ in terms of the invertibility of $a - \lambda$. But this is not satisfactory for our purposes, for it is very unstable when $\mathcal{A}$ is replaced by a dense subalgebra. (For example, a self-adjoint operator whose spectrum is the positive integers generates a subalgebra of operators isomorphic to the polynomial algebra $\mathbb{C}[x]$ in one variable, but so does an operator whose spectrum is the closed interval [1,2], or indeed any other infinite subset of $\mathbb{R}$.) It is better to define the spectrum of a self-adjoint element a of a topological $*$-algebra $\mathcal{A}$ as the set of all numbers $\theta(a)$ when θ runs through all the continuous linear $*$-maps[4] $\mathcal{A} \to \mathbb{C}$ whose restriction to the subalgebra of $\mathcal{A}$ generated by a is a homomorphism. [5] This gives us back the traditional Copenhagen lore about the

[4]This means that $\theta(b^*) = \overline{\theta(b)}$.

[5]The spectral theorem for a self-adjoint operator a in a Hilbert space $\mathcal{H}$ tells us that this coincides with the usual definition if a belongs to an algebra $\mathcal{A}$ of operators in $\mathcal{H}$ with the norm topology.

values taken by observables: if $(\mathcal{A}, \theta)$ is a small deformation of $(\mathcal{A}_0, \theta_0)$, with $\mathcal{A}_0$ commutative and θ_0 a homomorphism, then for any $a \in \mathcal{A}$ the number $\theta(a)$ will be close to $\theta_0(a)$, i.e. to the value of a continuous function on the space which is the spectrum of the algebra $\mathcal{A}_0$.

There is a kind of converse to the problem of finding a classical limit for a quantum system. If we are given a symplectic manifold $\mathcal{Y}$ then the algebra $\mathcal{A}_0 = C^\infty(\mathcal{Y})$ has a Poisson bracket[6], and we can ask whether there is a deformation $\{\mathcal{A}_h\}$ of the algebra which gives rise to the bracket. This is the problem of *deformation quantization*. For finite dimensional manifolds $\mathcal{Y}$ the definitive result, the culmination of a long history of work in this direction, is Kontsevich's theorem [K] which tells us that *formally* — i.e. if we work with formal power series in an indeterminate $\hbar$ rather than an actual real parameter — then such a deformation exists and is unique.

One final point before leaving the foundations of quantum theory: among the puzzling features of the conventional picture is that the observables are the self-adjoint elements of an *algebra*: for the multiplication in the algebra does not correspond to anything physically natural. (In particular, the product of two self-adjoint operators is usually not self-adjoint.) An attraction of the formulation of quantum field theory which I shall present in these lectures is that there is no "algebra" structure on the observables other than the obvious fact that one observation can be made *after* another. In fact when one has a quantum system $(\mathcal{A}, \theta_t)$ in the above sense it is easy to see — and will be explained in more detail later on — that all the information is retained if we simply know $\mathcal{A}$ as a vector space together with all the multilinear maps

$$\mathcal{A} \times \mathcal{A} \times \cdots \times \mathcal{A} \to \mathbb{C}$$

given by

$$(a_1, a_2, \cdots, a_n) \mapsto \theta_0(u_{-t_n}(a_n) u_{-t_{n-1}}(a_{n-1}) \cdots u_{-t_1}(a_1)),$$

for every finite sequence of times $t_1 < t_2 < \cdots < t_n$.

1.2 Variational principles and path integrals

The physics of a single massive particle moving freely in a Riemannian manifold M is the theory of geodesics. If $x : \mathbb{R} \to M$ is a smooth path in M then we define its *action* from time a to time b as

$$S_{a,b}(x) = \frac{1}{2} \int_a^b ||\dot{x}(t)||^2 dt,$$

[6] which encodes its symplectic structure completely — see Appendix x

where $||\dot{x}(t)||^2$ denotes the length-squared of the velocity-vector $\dot{x}(t)$ defined by the Riemannian metric of M. Thus $S_{a,b} : \mathcal{P}M \to \mathbb{R}$ is a smooth function on the manifold $\mathcal{P}M$ of smooth paths in M.

A *geodesic* or *classical trajectory* in M is a curve x for which, for all $a < b$, the action $S_{a,b}(x)$ is stationary under variations of x which vanish at the end-points a, b. By the usual integration-by-parts formula of the calculus of variations we can write the gradient of $S_{a,b}$ at x in the direction of the tangent vector ξ in the form

$$\mathrm{d}S_{a,b}(x,\xi) \;=\; \langle \dot{x}(b), \xi(b)\rangle \;-\; \langle \dot{x}(a), \xi(a)\rangle \;-\; \int_a^b \langle \ddot{x}(t), \xi(t)\rangle dt,$$

where $\ddot{x}(t)$ denotes the covariant derivative of $\dot{x}(t)$ defined using the Levi-Civita connection of M. Thus the submanifold $\mathcal{P}^{\mathrm{class}}M$ of classical trajectories in $\mathcal{P}M$ consists of the paths which satisfy the Euler-Lagrange equation $\ddot{x} = 0$, and on it we have

$$\mathrm{d}S_{a,b} \;= \varepsilon_b^*(\alpha) \;-\; \varepsilon_a^*(\alpha),$$

where $\varepsilon_t : \mathcal{P}M \to TM$ evaluates the path and its velocity at time t, and α is the obvious 1-form on TM. In particular, it follows that the exterior derivative $\omega = \mathrm{d}(\varepsilon_t^*\alpha)$ is a closed 2-form on $\mathcal{P}^{\mathrm{class}}M$ which does not depend on the time t. Furthermore, each evaluation-map ε_t is a diffeomorphism (at least if the manifold M is complete), and it is plain that the form ω is non-degenerate.

An exactly similar discussion applies to the generic situation of classical mechanics, where the instantaneous configuation of a system is a point of a finite-dimensional manifold X, and the system is described by a "Lagrangian" function $L : TX \to \mathbb{R}$ which is usually inhomogeneous quadratic on each tangent space to X. We define the action $S_{a,b} : \mathcal{P}X \to \mathbb{R}$ by

$$S_{a,b}(x) \;=\; \int_a^b L(x(t), \dot{x}(t))dt.$$

We obtain an Euler-Lagrange equation for the classical trajectories by integration by parts, and once again get the equation (*) on $\mathcal{P}^{\mathrm{class}}X$. As before, each $\varepsilon_t : \mathcal{P}X \to TX$ is a diffeomorphism, and we obtain a symplectic structure on $\mathcal{P}^{\mathrm{class}}X$.

When it was first discovered, the "principle of least action" had an attractive teleological aspect. The teleology becomes perhaps a little less persuasive when one realizes that the action is truly minimized only on sufficiently short segments of the trajectory. Whatever one thinks about that, however, the variational perspective gives a new and highly attractive approach to quantum mechanics, which was worked out by Dirac, Wheeler, and Feynmann. The idea is that, instead of just picking out the submanifold $\mathcal{P}^{\mathrm{class}}X$ in $\mathcal{P}X$, the action functions $S_{a,b}$ define a complex-valued *measure* μ_S on $\mathcal{P}X$, which is notionally written

$$\mathrm{d}\mu(x) \;=\; \mathrm{e}^{\mathrm{i}S(x)/\hbar}\mathcal{D}x$$

to suggest that each trajectory x has been "weighted" with $\mathrm{e}^{\mathrm{i}S(x)/\hbar}$. (Whether the action makes sense for a trajectory extending over all of time is probably best not asked at this point.) The measure μ_S defines a quantum system in the following way. For each time t, let $\mathcal{A}_t$ denote the vector space of smooth functions $f : \mathcal{P}X \to \mathbb{C}$ for which $f(x)$ depends only on the germ of the path x at time t. (Of course these spaces $\mathcal{A}_t$ are canonically isomorphic to each other by time-translation.) Then for each sequence of times $t_1 < t_2 < \cdots < t_n$ define

$$\mathcal{A}_{t_n} \times \cdots \times \mathcal{A}_{t_1} \ \to \ \mathbb{C}$$

by

$$(f_n, \cdots, f_1) \ \mapsto \ \int_{\mathcal{P}X} f_n(x) \cdots f_1(x) \mathrm{d}\mu(x).$$

As has been remarked above, this gives all the information we need to have a quantum system.

From the path-integral viewpoint the reason that classical mechanics approximates quantum mechanics is the occurrence of Planck's constant $\hbar$ — a very small quantity — in the weighting factor of the measure (*). This means that the measure is very highly oscillatory except on the critical set $\mathcal{P}^{\mathrm{class}}X$ in $\mathcal{P}X$. As $\hbar$ tends to 0 the measure μ_S tends to the delta-function along $\mathcal{P}^{\mathrm{class}}X$, and the classical theory is retrieved.

1.3 Waves and particles

Let us think about the classical systems formed by waves and particles in a space-manifold M, and how they might be approximations to the same quantum system.

To keep things as simple as possible, I shall suppose that the space M is a compact connected Riemannian manifold. The simplest kind of wave is described by a smooth real-valued function on M, so let us take the configuration-space of waves to be the vector space $\mathcal{X}^{\mathrm{wave}} = \mathrm{C}^\infty(M)$, and the state space to be $\mathcal{Y}^{\mathrm{wave}} = T\mathcal{X}^{\mathrm{wave}}$. The configuration space of an assembly of a finite number of particles in M looks quite different. It breaks up as the disjoint union of a sequence of connected components M_n indexed by the total number n of particles present, and M_n is the n-fold symmetric product

$$M_n = M^n/\mathrm{Symm}_n.$$

Because I have allowed more than one particle to occupy the same point of M — i.e. I have allowed multiplicities — M_n is not quite a manifold, but that is not a serious difficulty, and I shall overlook it for the present and write

$$\mathcal{Y}^{\mathrm{part}} = T\mathcal{X}^{\mathrm{part}} = \coprod_{n \geq 0} T(M_n).$$

According to the prescription for "quantization" in elementary quantum mechanics, the natural quantum system to put beside the linear classical phase-space $\mathcal{Y}^{\rm wave}$ is obtained by introducing the Hilbert space $\mathcal{H}$ of square-integrable functions[7] on the infinite dimensional vector space $\mathcal{X}^{\rm wave}$. The construction of $\mathcal{H}$ gives us a class of functions on $\mathcal{X}^{\rm wave}$ which are "measurable", and hence act on $\mathcal{H}$ as multiplication operators. These are the "position operators" of the waves. The linear function $\phi \mapsto \phi(x)$ on $\mathcal{X}^{\rm wave}$ which would tell us the value of the field at a point x of M is not measurable, but for each $f \in {\rm C}^\infty(M)$ we do have the "smeared" operator ϕ_f corresponding to the linear function $\phi \mapsto \int_M f(x)\phi(x)dx$. We expect the complete algebra $\mathcal{A}$ of the quantum system to be generated by the operators ϕ_f and their time-derivatives $\dot{\phi}_f$, which should obey standard commutation-relations

$$[\ \dot{\phi}_f\ ,\ \phi_g\ ]\ =\ {\rm i}\hbar \int_M fg\,{\rm d}x\ .$$

The position operators ϕ_f generate a commutative subalgebra $\mathcal{Q}$ of $\mathcal{A}$ which is isomorphic to the symmetric algebra of the vector space ${\rm C}^\infty(M) = \mathcal{X}^{\rm wave}$. It acts on $\mathcal{H}$ by unbounded operators. An algebra homomorphism $\mathcal{Q} \to \mathbb{R}$, i.e. a point of the spectrum of $\mathcal{Q}$, is given by an arbitrary continuous linear map $\mathcal{X}^{\rm wave} \to \mathbb{R}$, i.e. by a *distribution* on M. The class of distributions which arise depends on the dynamics defining the Hilbert space, but we certainly do not get only sums of delta functions at points of M, so this algebra of observables does not describe point excitations in any classical limit.

Nevertheless, the algebra $\mathcal{A}$ of operators on $\mathcal{H}$ contains another commutative subalgebra $\mathcal{C}$ which is also isomorphic to the symmetric algebra of ${\rm C}^\infty(M)$ but acts on the Hilbert space in a very different way. It is generated by unbounded operators which I shall denote by a_f for $f \in {\rm C}^\infty(M)$. The spectrum of $\mathcal{C}$ consists not of all distributions on M, but only of those which are finite sums of delta-functions at points of M. More precisely, $\mathcal{C}$ maps by a continuous algebra homomorphism to the algebra of smooth functions on $\mathcal{X}^{\rm part}$, and it is through this algebra that it acts on the Hilbert space. The map

$$\mathcal{C} \to {\rm C}^\infty(\mathcal{X}^{\rm part})$$

takes a_f to the sequence of functions $f_n : M_n \to \mathbb{R}$, where

$$f_n(x_1, \ldots, x_n) = \sum_i f(x_i).$$

[7]It is not obvious how to define this space: the appropriate definition depends on the dynamics of the system, i.e. on the Hamiltonian function $T\mathcal{X}^{\rm wave} \to \mathbb{R}$ expressing the total energy of the system. I shall assume for the moment that we have a free field theory, corresponding to non-interacting particles. This case will be discussed in detail in §xxx.

This homomorphism of algebras is injective, and its image is dense. The action of the polynomial algebra $\mathcal{C}$ on the quantum Hilbert space extends by continuity to an action of $\mathrm{C}^\infty(\mathcal{X}^{\mathrm{part}})$. This is the sense in which the quantum system describes collections of *particles* in M rather than smeared-out objects.

There is a basic algebraic mechanism in quantum field theory which ensures that an action of the symmetric algebra of $\mathrm{C}^\infty(M)$ has this property of describing point-particles in M. To understand it, let us begin by considering a single unbounded self-adjoint operator a in a Hilbert space.

To give an unbounded self-adjoint operator a in $\mathcal{H}$ is by definition the same as to give the bounded operator $f(a)$ for all f in the algebra $\mathrm{C}_0(\mathbb{R})$ of continuous complex-valued functions on the line $\mathbb{R}$ which tend to 0 at ∞, i.e. to give a is the same as to give a $*$-action of the algebra $\mathrm{C}_0(\mathbb{R})$ on $\mathcal{H}$. The *spectrum* of a is the closed subset σ_a of $\mathbb{R}$ such that the algebra of operators topologically generated by a is isomorphic to $\mathrm{C}_0(\sigma_a)$, with a itself corresponding to the inclusion $\sigma_a \to \mathbb{R}$. In other words, σ_a is the smallest subset of $\mathbb{R}$ such that the action of $\mathrm{C}_0(\mathbb{R})$ on $\mathcal{H}$ factorizes through $\mathrm{C}_0(\sigma)$. (As was mentioned above, in the usual language of quantum mechanics one says that "a measurement of a always yields a number in σ_a".)

The following simple property ensures that the spectrum of a is precisely $\mathbb{N} \subset \mathbb{R}$.

Proposition 1.3.1 *If an unbounded self-adjoint operator $a : \mathcal{H} \to \mathcal{H}$ in a Hilbert space $\mathcal{H}$ is of the form $a = bb^*$, where b is such that $[b^*, b] = b^*b - bb^* = 1$, then the spectrum σ_a is $\mathbb{N}$. More precisely, we have an orthogonal decomposition*

$$\mathcal{H} = \oplus_{k\geq 0}\mathcal{H}_k,$$

where $\mathcal{H}_k$ is the k-eigenspace of a, and b maps $\mathcal{H}_k$ isomorphically to $\mathcal{H}_{k+1}$, while b^* maps $\mathcal{H}_k$ isomorphically to $\mathcal{H}_{k-1}$ when $k > 1$, and maps $\mathcal{H}_0$ to 0.

This well-known result — a reformulation of the Stone-von Neumann theorem[8]— generalizes to give us a condition which ensures that an action of the symmetric algebra $\mathcal{C}$ of $\mathrm{C}^\infty(M)$ describes particles, i.e. that the action of $\mathcal{C}$ factorizes through the algebra $\mathrm{C}_0(\mathcal{X}^{\mathrm{part}})$. For to say that a system $(\mathcal{A}, \theta)$ describes a collection of particles in a space-manifold M we need to have commuting self-adjoint operators $a_U \in \mathcal{A}$ for each measurable subset U of M such that

(i) the spectrum of each a_U is $\mathbb{N} = \{0, 1, 2, \ldots\}$, and

(ii) $\sum a_{U_i} = a_U$ when U is the disjoint union of the sets U_i.

[8]If we write $b = Q + \mathrm{i}P$ with P and Q self-adjoint, then P and Q satisfy $[P, Q] = \mathrm{i}/2$, and the Stone-von Neumann theorem tells us that the algebra generated by P and Q has a unique irreducible representation, in which $bb^* = P^2 + Q^2 - 1/2$ acts with spectrum $\mathbb{N}$.

For a system of operators $\{a_U\}$, the condition (ii) simply amounts to saying that

$$a_U = \int_U a(x)dx,$$

where $a(x)$ is an operator-valued distribution on M. By smearing the operators $a(x)$ we obtain an action on $\mathcal{H}$ of the symmetric algebra $\mathcal{C}$ of the vector space $\mathrm{C}^\infty(M)$. Then each a_U has spectrum $\mathbb{N}$ if the action of $\mathcal{C}$ on $\mathcal{H}$ extends to an action of the commutative algebra $\mathrm{C}^\infty(\mathcal{X}^{\mathrm{part}})$. An easy generalization[9] of Proposition 1.1 shows that this is the case if there is another operator-valued distribution $b(x)$ such that $a(x) = b(x)b(x)^*$, and we have the commutation relations

$$[b(x)^*, b(x')] = \delta(x, x'),$$

$$[b(x), b(x')] \ = 0,$$

where $\delta(x, x')$ is the Dirac delta-function along the diagonal in $M \times M$. It is better, however, to state the result without introducing the ill-defined "product" $b(x)b(x)^*$, which involves a pointwise product of distributions. The essential property we want the product to have is

$$[b(x), b(x')b(x')^*] = \delta(x, x')b(x'),$$

so we can state the result as follows.

Proposition 1.3.2 *The action of a self-adjoint operator-valued distribution $\{a(x)\}$ extends to an action of the algebra $\mathrm{C}^\infty(\mathcal{X}^{\mathrm{part}})$ if there is another operator-valued distribution $\{b(x)\}$ which satisfies the commutation relations above, as well as*

$$[b(x), a(x')] = \delta(x, x')a(x'),$$

and if, furthermore, for any vector ξ, we have

$$a(x)\xi = 0 \text{ for all } x \iff b(x)^*\xi = 0 \text{ for all } x.$$

On a formal level this result is fairly plain. For the operator a_M is positive, and the commutation relations

$$[a_M, b(x)] = b(x) \qquad [a_M, b(x)^*] = -b(x)^*$$

show that the action of $b(x)$ (respectively of $b(x)^*$) takes an eigenvector of a_M to another eigenvector with the eigenvalue raised (respectively, lowered) by 1. There must therefore be vectors $\Omega \in \mathcal{H}$ which are annihilated by all the $b(x)^*$, and for each such Ω the vector

$$b(x_1)b(_2)\ldots b(x_k)\Omega$$

[9] I shall be cavalier here about the analytical details, as they will be treated carefully in §xxx below.

is an eigenvector of each a_f with the eigenvalue $f(x_1) + f(x_2) + \ldots + f(x_k)$.

The algebra generated by smearing operators $b(x)$ and $b(x)^*$ which satisfy the above commutation relations is isomorphic to the algebra $\mathcal{A}$ which was the natural quantization of the "wave" system $\mathcal{Y}^{\text{wave}}$. Indeed there is a wide choice of isomorphisms. For example, let us write

$$\psi(x) = (b(x) + b(x)^*)/\sqrt{2} \qquad \dot{\psi}(x) = \mathrm{i}\hbar(b(x) - b(x)^*)/\sqrt{2},$$

and define the self-adjoint operators

$$\psi_f = \int f(x)\psi(x)dx \qquad \dot{\psi}_f = \int f(x)\dot{\psi}(x)dx$$

for $f \in \mathrm{C}^\infty(M)$. Then the operators ψ_f commute among themselves, as do the $\dot{\psi}_f$, while we have

$$[\dot{\psi}_f, \psi_g] = \mathrm{i}\hbar \int fg dx.$$

These are precisely the relations which we used to define the quantum algebra $\mathcal{A}$ associated to the system $\mathcal{Y}^{\text{wave}}$. We get, however, the same relations if we replace ψ_f and $\dot{\psi}_f$ by $\phi_f = \psi_{Pf}$ and $\dot{\phi}_f = \dot{\psi}_{\check{P}f}$, where P is any linear automorphism of the vector space $\mathrm{C}^\infty(M)$ and $\check{P}$ is its contragredient, i.e. the transpose of its inverse. We shall see in §xxx that the correct choice for the usual theory of free particles of mass m is to take $P = (m^2 + \Delta)^{-1/2}$, where Δ is the (positive) Laplace operator of the Riemannian manifold M. This corresponds to "waves" obeying the Klein-Gordon equation

$$(\partial/\partial t)^2\phi + \Delta\phi = 0.$$

The quantized algebra $\mathcal{A}^{\text{part}}$ corresponding to the algebra of functions on $\mathcal{Y}^{\text{part}} = T\mathcal{X}^{\text{part}}$ is generated by the commutative algebra $\mathcal{C}$ of functions on $\mathcal{X}^{\text{part}}$ and their time derivatives. This algebra is contained in the algebra $\mathcal{A} = \mathcal{A}^{\text{wave}}$ of the wave theory generated by the ϕ_f and $\dot{\phi}_f$, but nevertheless $\mathcal{A}^{\text{wave}}$ is strictly larger than $\mathcal{A}^{\text{part}}$. For the Hilbert space $\mathcal{H}$ on which $\mathcal{A}^{\text{wave}}$ acts has a positive integer grading $\mathcal{H} = \oplus\mathcal{H}_n$ given by the "particle number", i.e. by the eigenspaces of a_M. The component $\mathcal{H}_n$ can be identified with the space of L^2 functions on M_n, and the operators a_f act on it through the algebra of functions on M_n. The algebra $\mathcal{A}^{\text{part}}$ preserves the grading of $\mathcal{H}$. The operators p_f and q_f, however, do not preserve the grading, and do not correspond to functions on $\mathcal{Y}^{\text{part}}$. They have no analogue in the particle description. In fact b_f and b_f^* respectively raise and lower the grading by 1, and are traditionally called *creation* and *annihilation* operators.

Chapter 2

Quantum field theory

Quantum field theory is part of quantum theory, but it introduces a quite new ontological component. Like classical physics it presupposes that we are given a definite space-time manifold, and that — oversimplifying a little —- the essential observables are *local* with respect to space-time. This means — roughly, again — that for each space-time point x there is given a vector space $\mathcal{A}_x$ of "local observables at x" contained in the algebra $\mathcal{A}$, and that the subspaces $\mathcal{A}_x$ together generate $\mathcal{A}$ as an algebra. This is a lot less than saying that the world is made up of particles, but it does allow us to ask what kinds of thing might or might not be present at a space-time point x. It is perfectly compatible with a particle picture, but it de-emphasizes the individuality of the particles, and leads one towards the idea that quantum field theory always describes *assemblies* of indistinguishable particles.