Quick (and informal) intro on the C^*-algebraic formulation of QM.

Alessandro

15 April 2014

The idea: to provide here only the general context, the basic definitions and properties (no spectral or functional analysis yet!), and an overview of the representation problem and of the conceptual path from an abstract C^*-algebra to the customary Hilbert-space formalism. This is meant both to answer a number of individual questions and to enlarge a bit the vision on the material from the first 2 lectures (of course only the material presented in class is requested for the exam!) I only state rigorously a number of definitions and theorems, the proof are often involved and the dedicated students can reconstruct them with the help of this note and using the reference books suggested in class.

Contents

1 Change of perspective: new axioms.

The Dirac-von Neumann axioms of Quantum Mechanics consist of cinematic rules

(1) The pure states of a quantum system are described by rays in a separable complex Hilbert space $\mathcal{H}$. More generally, mixed states are density matrices on $\mathcal{H}$.

(2) The observables are described by a subset of bounded, self-adjoint operators on $\mathcal{H}$.[1]

[1] In the original Dirac-von Neumann formulation any bounded and self-adjoint operator on $\mathcal{H}$ was an observable, so that in particular all orthogonal projections describe observables and therefore any ray in $\mathcal{H}$ describes a physically realizable state. Later it was realised that $\mathcal{H}$ may carry a reducible representation of the algebra of the observables, equivalently there exist bounded self-adjoint operators on $\mathcal{H}$ known as super-selection operators or "charges" (like the electric charge or the permutation operator of identical particles), which commute with all the observables. In the known cases the spectrum of the super-selection operators is discrete so that $\mathcal{H}$ decomposes as a direct sum of super-selection sectors, each carrying an irreducible representation of the observable algebra. Since observables have zero matrix elements between vectors belonging to different super-selection sectors,

the relative phase in the superposition of such states are not observable. The important physical implication is that not all projections are observables: it is impossible to prepare physical states which are coherent superposition of states belonging to different super-selection sectors.

(3) If a state ω is described by the vector $\Psi_\omega \in \mathcal{H}$, for any observable A the measure[2] (or "expectation") of A in the state ω is the real number $\omega(A) = \langle \Psi_\omega, A\Psi_\omega \rangle$ (or $\omega(A) = \mathrm{Tr}(\rho_\omega A)$ if a state is described by a density matrix ρ_ω).

one dynamical rule

(4) The dynamical evolution of the system is determined by the specification of a self-adjoint operator H through either of the algorithms $A \mapsto A_t = e^{itH} A e^{-itH}$ or $\psi \mapsto \psi_t = e^{-itH}\psi$. (In fact there is full complementarity of the two algorithms: $\langle \psi, A_t \psi \rangle = \langle \psi_t, A\psi_t \rangle$.)

and one operational constraint

(5) Whereas for any given observable A it is possible to prepare a state ω with no limitation on the smallness of the "dispersion" $\Delta_\omega(A) := \omega((A - \omega(A))^2) = \omega(A^2) - \omega(A)^2$, nevertheless there may be observables A, B for which $\Delta_\omega(A) + \Delta_\omega(B)$ is not arbitrarily small. Equivalently,[3] there may exist observables that do not commute.

The last point is in fact the axiomatization of the celebrated Heisenberg's indeterminacy principle.[4] It is therefore the translation of an experimental fact and *not* an axiom of mathematical nature.

In the C^*-algebraic formulation of Q.M. the axioms become:

(1*) The observables that define a quantum system are the self-adjoint elements of a non-commutative C^*-algebra $\mathcal{A}$.

(2*) The states of the given system are normalized positive linear functionals on $\mathcal{A}$.

(3*) The measure of the observable A in the state ω is the number $\omega(A)$.

(4*) The dynamics of the system is described by a one-parameter weakly continuous group $\{\alpha_t \,|\, t \in \mathbb{R}\}$ of $*$-automorphisms of $\mathcal{A}$ onto itself: the time evolution of $A \in \mathcal{A}$ is given by $A \mapsto A_t = \alpha_t(A)$, the time evolution of a state is given by $\omega \mapsto \omega_t$ where $\omega_t(A) = \omega(\alpha_t(A))$ $\forall A \in \mathcal{A}$, the map $t \mapsto \omega(\alpha_t(A))$ is continuous for every $A \in \mathcal{A}$ and every state ω on $\mathcal{A}$.

(5*) [*incorporated in* (1*)]

There is no Hilbert space in the algebraic axioms. There, the primal object is the C^*-algebra. The philosophical point is: I start from the physical properties I can *observe* in the system (i.e., the observables). The list of possible states in which I can prepare my system comes after.

A very clean and compact discussion (in fact, the best I know) of this change of perspective, from both the physical, mathematical, and philosophical point of view is in Chapters 1 and 2 of "*An Introduction to the Mathematical Structure of Quantum Mechanics. A Short Course for Mathematicians.*" by F. Strocchi, World Scientific, 2nd ed., 2008. I strongly recommend it.

The algebraic axioms were inferred from the Dirac-von Neumann's axioms by identifying in $\mathcal{B}(\mathcal{H})$, the bounded operators on a Hilbert space $\mathcal{H}$, those structural, algebraic, and analytical properties that constitute the definition of a C^*-algebra. This was quite a long process, in particular it was not clear a priori why the algebra containing observables should be closed (complete) *in norm*. Thus, when the Hilbert space operator theory began in the 1930s (Murray and von Neumann), the first physically meaningful structures to be identified were *weakly*-closed algebras of operators (nowadays referred to as von Neumann's algebras, or W^*-algebras – W = "weakly closed"). All in all what one measures in the lab is just expectations of observables, say $\langle \psi, A\psi \rangle$, and a limiting sequence of experimental apparatuses leads naturally to *weak* limits such as $\lim_n \langle \psi, A_n \psi \rangle$, at first a weakly closed algebra of observables seems the most natural. C^* (C = "Closed") came with Gelfand and Naimark in 1943, but its relevance to Q.M. was not fully appreciated for more than twenty years. Despite that, there has been a subsequent fruitful period of interplay between maths and physics which has instigated both interesting structural analysis of operator algebras and significant physical applications, notably to quantum statistical mechanics and relativistic quantum field theory.

[2] defined to be the average of the outcomes of replicated measurements of A performed on the system in the state ω

[3] The constraint itself is an axiom. The equivalence between its formulation for states and its formulation for observables is a theorem by von Neumann.

[4] It is easy to see that Heisenberg's commutation relation $[Q, P] = \mathrm{i}$ is not suited for a formulation within the algebra of bounded observables. But this is just technical: these relations can be re-formulated conveniently for bounded operators.

2 Def. of C^*-algebra

An ALGEBRA $\mathcal{A}$ over $\mathbb{C}$ is a vector space with respect to $+$ with an associative and distributive product $\cdot$, namely

$$\begin{aligned} A(BC) &= (AB)C \\ A(B+C) &= AB+AC \\ (A+B)C &= AC+BC \\ \alpha\beta(AB) &= (\alpha A)(\beta B) \qquad\qquad (\forall A,B,C \in \mathcal{A}\,, \quad \forall \alpha,\beta \in \mathbb{C}). \end{aligned}$$

$\mathcal{A}$ is said COMMUTATIVE or ABELIAN if $AB = BA\ \forall A, B \in \mathcal{A}$, otherwise it is NON-COMMUTATIVE.

A *-ALGEBRA is an algebra $\mathcal{A}$ with an involution $*$, namely, a map $* : \mathcal{A} \to \mathcal{A}$ such that

$$A^{**} = A, \qquad (AB)^* = B^*A^*, \qquad (\alpha A + \beta B)^* = \overline{\alpha}A^* + \overline{\beta}B^*$$

A NORMED ALGEBRA is an algebra $\mathcal{A}$ with a norm $\|\ \|$ compatible (sublinear) with $+$, namely

$$\begin{aligned} &\|A\| \geqslant 0 \\ &\|A\| = 0 \Leftrightarrow A = \mathbb{O} \\ &\|\alpha A\| = |\alpha|\,\|A\| \\ &\|A+B\| \leqslant \|A\| + \|B\|\,, \end{aligned}$$

and compatible (sub-multiplicative) with $\cdot$, namely

$$\|AB\| \leqslant \|A\|\,\|B\|.$$

A BANACH *-ALGEBRA $\mathcal{A}$ is a normed algebra and a $*$-algebra in which norm and involution satisfy the property

$$\|A^*\| \ = \ \|A\| \qquad \forall A \in \mathcal{A}\,,$$

and, additionally, as a vector space $\mathcal{A}$ is complete with respect to its norm (namely, it is a Banach space).

A C^*-algebra is a BANACH *-ALGEBRA $\mathcal{A}$ that satisfies the property

$$\|A^*A\| \ = \ \|A\|^2 \qquad \forall A \in \mathcal{A} \qquad\qquad (\text{“}C^*\text{-property”}).$$

If an algebra $\mathcal{A}$ contains a (multiplicative) UNIT $\mathbb{1}$, namely an element such that $\mathbb{1}A = A\mathbb{1} = A\ \forall A \in \mathcal{A}$, then $\mathcal{A}$ is said to be UNITAL, otherwise it is called NON-UNITAL. The unit element, if it exists, is clearly unique. In a C^*-algebra the unit can only have norm $\|\mathbb{1}\| = 1$ or $\|\mathbb{1}\| = 0$: in the second case necessarily the algebra is trivial, $\mathcal{A} = \{\mathbb{O}\}$, but we shall systematically ignore this case.

The presence of $\mathbb{1}$ is a strong structural constraint in the algebra. For instance, if $\mathcal{A}$ is a unital Banach algebra then no $A, B \in \mathcal{A}$ can satisfy an uncertainty relation of the form $AB - BA = \alpha\mathbb{1}$ (for some $\alpha \in \mathbb{C}\setminus\{0\}$).

Note, though: if a C^*-algebra $\mathcal{A}$ is non-unital, then there exists a unital C^*-algebra $\widetilde{\mathcal{A}}$ such that $\mathcal{A}$ is identifiable with a C^*-subalgebra of $\widetilde{\mathcal{A}}$. This fact allows one to circumvent many of the complications related with the absence of an identity.

3 The C^*-condition

The property $\|A^*A\| = \|A\|^2\ \forall A \in \mathcal{A}$ is crucial for the proof of all main facts of the theory. Despite its apparently technical nature, it is worth pointing out that the C^*-condition is the key ingredient for the proof of the functional calculus (Spectral Theorem) on $\mathcal{A}$.

To any element $A \in \mathcal{A}$ one associates a compact, non-empty set $\sigma(A)$ of complex numbers, called the SPECTRUM of A. In the case of a unital C^*-algebra, $\sigma(A)$ is defined as the set of numbers $\lambda \in \mathbb{C}$ for which

$A - \lambda\mathbb{1}$ is not invertible in $\mathcal{A}$. (If $\mathcal{A}$ is not unital, the spectrum is defined with respect to its unitalisation $\widetilde{A}$ mentioned above.) If $A \in \mathcal{A}$ is self-adjoint, i.e., $A = A^*$, then $\sigma(A) \subset \mathbb{R}$.

Details aside, the (continuous) functional calculus for a self-adjoint element $A = A^* \in \mathcal{A}$ is a map $C(\sigma(A)) \to \mathcal{A}$, $f \mapsto f(A)$, with $\|f(A)\| = \|f\|_{\sup}$, that generalises the construction of a polynomial p of A, where $p(A)$ is merely built by means of the algebraic operations $+$ and $\cdot$ of $\mathcal{A}$ (and with the convention $A^0 = \mathbb{1}$), and allows one to define consistently a function f of A for a generic $f \in C(\sigma(A))$. The map $f \mapsto f(A)$ is in fact constructed first on polynomials, in which case one proves the identity

$$\|p(A)\| = \|p\|_{C(\sigma(A))} \qquad \forall p \text{ polynomial over } \sigma(A)\,;$$

such an identity tells us that $p \mapsto p(A)$ has a unique linear extension to the closure of the polynomials in $C(\sigma(A))$, which by Stone-Weierstrass turns out to be the whole $C(\sigma(A))$.

It is precisely the identity above that requires the C^*-condition: indeed

$$\|p(A)\|^2 \overset{\downarrow}{=} \|p(A)^*p(A)\| \;=\; \|(\overline{p}\,p)(A)\| \;=\; \sup_{\lambda\in\sigma((\overline{p}\,p)(A))} |\lambda| \;=\; \sup_{\lambda\in\sigma(A)} |(\overline{p}\,p)(\lambda))| \;=\; \big(\sup_{\lambda\in\sigma(A)} |p(\lambda)|\,\big)^2 \;=\; \|p\|^2_{C(\sigma(A))}$$

(while the steps involving the sup over the spectrum follow from the self-adjointness of A). One may say that it is the C^*-condition that makes the Spectral Theorem valid for self-adjoint operators.[5]

4 States

Let $\mathcal{A}$ be a C^*-algebra.

A linear map $\omega : \mathcal{A} \to \mathbb{C}$ on $\mathcal{A}$ is said a FUNCTIONAL. A functional is POSITIVE if

$$\omega(A^*A) \geqslant 0 \qquad \forall A \in \mathcal{A}\,.$$

In the special case when the C^*-algebra $\mathcal{A}$ is unital, positive linear functionals are characterised as follows: if ω is a linear functional then

$$\omega \text{ is positive} \qquad \Leftrightarrow \qquad \omega \text{ is bounded } \textit{and } \; \|\omega\| := \sup_{\|A\|=1} |\omega(A)| \;=\; \omega(\mathbb{1})\,.$$

A positive linear functional ω on $\mathcal{A}$ has, among others, these properties for all $A, B \in \mathcal{A}$:

- $\omega(A^*B) = \overline{\omega(B^*A)}$
- $|\omega(A^*B)|^2 \leqslant \omega(A^*A)\omega(B^*B)$ (Cauchy-Schwarz)
- $\omega(A^*) = \overline{\omega(A)}$
- $|\omega(A)|^2 \leqslant \omega(A^*A)\|\omega\|$
- $|\omega(A^*BA)| \leqslant \omega(A^*A)\|B\|$
- $\|\omega\| = \sup_{\|A\|=1} \omega(A^*A)\,.$

A STATE ω on $\mathcal{A}$ is a linear functional that is positive (hence necessarily bounded), with $\|\omega\| = 1$. If $\mathcal{A}$ is unital one can equivalently define a state ω as a linear functional on $\mathcal{A}$ that is positive and with $\omega(\mathbb{1}) = 1$, or also as a linear functional on $\mathcal{A}$ that is bounded and with $\|\omega\| = 1 = \omega(\mathbb{1})$ (which makes ω necessarily positive).

A state ω is said PURE if the only positive linear functionals ρ majorized by ω, in the sense

$$0 \leqslant \rho \leqslant \omega \qquad \Leftrightarrow \qquad 0 \;\leqslant\; \rho(A^*A) \;\leqslant\; \omega(A^*A) \qquad \forall A \in \mathcal{A}\,,$$

are of the form $\rho = \lambda\omega$ with $\lambda \in [0,1]$. Equivalently, a state is pure if it cannot be written as a convex combination of other states. All this nomenclature is clearly borrowed from the usual Hilbert-space language

[5] All this is pretty much obvious, still, I never found it mentioned explicitly in textbooks: I myself was pointed out this fact in class by my M. Sc. supervisor (Gianni Morchio) some years ago when I was taking his “Algebraic methods in QM” course in Pisa.

where the action of a state $\Psi \in \mathcal{H}$ on $\mathcal{B}(\mathcal{H})$ is $A \mapsto \langle \Psi, A\Psi \rangle$. The meaning of axiom (4*) in is: for every state ω, $t \mapsto \omega(\alpha_t(A))$ is continuous.

If the C^*-algebra $\mathcal{A}$ is commutative then one can prove that a state ω on $\mathcal{A}$ is pure if and only if $\omega(AB) = \omega(A)\omega(B)$ $\forall A, B \in \mathcal{A}$.

As a consequence of the Hahn-Banach theorem (therefore: axiom of choice, namely, no constructive statement) there exist loads of states on a C^*-algebra in the following precise sense: if A is an arbitrary element of a C^*-algebra $\mathcal{A}$ then there exists a pure state ω_A on $\mathcal{A}$ such that

$$\omega_A(A^*A) \;=\; \|A\|^2\,.$$

Moreover (again from Hahn-Banach), the states separate the elements of $\mathcal{A}$, i.e., $\forall A, B \in \mathcal{A}$ with $A \neq B$ there exists a state ω such that $\omega(A) \neq \omega(B)$.

5 Main examples

$\mathcal{B}(\mathcal{H})$ (non-commutative) and $C(X)$ (commutative). (TO WRITE UP, DON'T HAVE TIME NOW)

6 Representations

Having recognised that the observables of a quantum system generate a non-commutative C^*-algebra and that the states of the system are normalised positive linear functionals on it, there comes the question of how such an abstract structure can be used for concrete physical problems, for calculations, predictions, etc. We have therefore to find concrete realisations of the above structure. It is not a priori obvious what the concrete realisations of C^*-algebras are. The tool is representation theory.

To this aim, recall: a $*$-HOMOMORPHISM π between two $*$-algebras $\mathcal{A}$ and $\mathcal{B}$ is a map $\pi : \mathcal{A} \to \mathcal{B}$ which preserves all the algebraic relations including the $*$, namely

$$\begin{aligned}
&\pi(aA + bB) = a\pi(A) + b\pi(B)\\
&\pi(AB) = \pi(A)\pi(B)\\
&\pi(A^*) = \pi(A)^*
\end{aligned}$$

$\forall A, B \in \mathcal{A}$ and $\forall a, b \in \mathbb{C}$. If the two $*$-algebras are unital, we also require

$$\pi(\mathbb{1}_{\mathcal{A}}) = \mathbb{1}_{\mathcal{B}}\,.$$

A REPRESENTATION π OF A C^*-ALGEBRA $\mathcal{A}$ ON A ("TARGET") HILBERT SPACE $\mathcal{H}$ is a $*$-homomorphism $\pi : \mathcal{A} \to \mathcal{B}(\mathcal{H})$. In symbols: $(\pi, \mathcal{H})$.

If $\mathcal{A}$ is a unital C^*-algebra and $(\pi, \mathcal{H})$ is a representation of $\mathcal{A}$, then

$$\|\pi(A)\| \leqslant \|A\| \qquad \forall A \in \mathcal{A}\,.$$

Thus, although the definition of a representation π of a C^*-algebra $\mathcal{A}$ does not require continuity (i.e., boundedness), continuity is automatically true for representations of unital C^*-algebras! This is a deep consequence of the intertwining between the algebraic and analytic properties of a C^*-algebra.

For a representation π of a C^*-algebra the following holds:

$$\pi \text{ is injective } (\ker\pi = \{\mathbb{O}\}) \qquad \Leftrightarrow \qquad \|\pi(A)\| = \|A\| \quad \forall A \in \mathcal{A}\,.$$

When any of the above two conditions holds true, π is said to be a FAITHFUL REPRESENTATION.

Among the most relevant C^*-algebras in the theory are the SIMPLE C^*-algebras, namely those for which the only norm-closed, two-sided, $*$-ideals are $\{\mathbb{O}\}$ and the algebra itself. Now, $\ker\pi$ is a norm-closed, two-sided, $*$-ideal in $\mathcal{A}$ for every representation of a C^*-algebra $\mathcal{A}$; thus, if $\mathcal{A}$ is simple then *every* representation of $\mathcal{A}$ is faithful.

Not only the $*$-homomorphisms of the form $\pi : \mathcal{A} \to \mathcal{B}(\mathcal{H})$ (namely: representations) are relevant in the theory. A $*$-ISOMORPHISM π between two C^*-algebras $\mathcal{A}$ and $\mathcal{B}$ is a $*$-homomorphism $\pi : \mathcal{A} \to \mathcal{B}$ which is bijective. A $*$-isomorphism between C^*-algebras turns out to preserve the norm. A $*$-isomorphism of a C^*-algebra $\mathcal{A}$ onto itself is called $*$-AUTOMORPHISM.

Back to representations of C^*-algebras, the most trivial one is the one such that $\pi(A) = 0\ \forall A \in \mathcal{A}$. A representation might be non-trivial but still have a "trivial part", namely the closed subspace[6]

$$\bigcap_{A\in\mathcal{A}} \ker \pi(A) \;=\; \{\psi \in \mathcal{H} \,|\, \pi(A)\psi = 0\ \forall A \in \mathcal{A}\}$$

of $\mathcal{H}$. A representation for which $\bigcap_{A\in\mathcal{A}} \ker \pi(A) = \{0\}$ is called NON-DEGENERATE. A representation of a unital C^*-algebra is always non-degenerate.

7 Irreps and cyclic representations

There are two most relevant types of representation of a C^*-algebra: cyclic and irreducible. Let π be a representation of a C^*-algebra $\mathcal{A}$ with target space $\mathcal{H}$.

π is called a CYCLIC REPRESENTATION if there exists a vector $\Omega \in \mathcal{H}$ which is cyclic for π in $\mathcal{H}$, which means that $\{\pi(A)\Omega \,|\, A \in \mathcal{A}\}$ is dense in $\mathcal{H}$. In symbols: $(\pi, \mathcal{H}, \Omega)$. Thus, the sole knowledge of Ω allows to reconstruct the whole $\mathcal{H}$ by application of representatives of $\mathcal{A}$. A cyclic representation is clearly non-degenerate.

π is said to be an IRREDUCIBLE REPRESENTATION ("irrep") if any of these *equivalent* properties holds true:

(1) $\{0\}$ and $\mathcal{H}$ are the only closed subspaces of $\mathcal{H}$ invariant under $\pi(\mathcal{A})$.

(2) *Every* non-zero vector $\psi \in \mathcal{H}$ is cyclic for π.

(3) The COMMUTANT
$$\pi(\mathcal{A})' := \{T \in \mathcal{B}(\mathcal{H}) \,|\, [\pi(A), T] = \mathbb{O}\ \forall A \in \mathcal{A}\}$$
consists of multiples of the identity operator.

Notice this obvious hierarchy of representations:

$$\text{irrep} \Rightarrow \text{cyclic} \Rightarrow \text{non-degenerate}$$

(the opposite implications are false in general).

Although cyclic representations are only special cases of representation of a C^*-algebra, they turn out to be a sort of "building blocks" of generic representations. In order to make this precise, let us first define the direct sum of representations.

Let $\{(\pi_\alpha, \mathcal{H}_\alpha)\}_{\alpha\in\mathcal{I}}$ be a collection of representations of a given C^*-algebra $\mathcal{A}$ (the index set $\mathcal{I}$ being countable or uncountable). The HILBERT DIRECT SUM of the target spaces is the space

$$\mathcal{H} \;:=\; \bigoplus_{\alpha\in\mathcal{I}} \mathcal{H}_\alpha$$

constructed as follows. The finite subsets F of $\mathcal{I}$ form a directed set[7] when ordered by inclusion. Consider the set of collections $\psi := \{\psi_\alpha\}_{\alpha\in\mathcal{I}}$, where each ψ_α belongs to $\mathcal{H}_\alpha$, such that $\lim_F \sum_{\alpha\in F} \|\psi_\alpha\|^2_{\mathcal{H}_\alpha} < \infty$ (the $\lim_F$ is well-defined precisely because the finite subsets F of $\mathcal{I}$ form a directed set). Equip this set with the obvious vector space structure inherited component-wise. On such a vector space define $\langle\psi, \phi\rangle \equiv \sum_{\alpha\in\mathcal{I}} \langle\psi_\alpha, \phi_\alpha\rangle_{\mathcal{H}_\alpha} := \lim_F \sum_{\alpha\in F} \langle\psi_\alpha, \phi_\alpha\rangle_{\mathcal{H}_\alpha}$. This turns out to be a scalar product which makes the space complete. This is $\oplus_{\alpha\in\mathcal{I}} \mathcal{H}_\alpha$.

[6]Note that in general $\mathcal{H} = \Big(\bigcap_{A\in\mathcal{A}} \ker \pi(A)\Big) \oplus \overline{\mathrm{Span}\{\pi(A)\psi \,|\, A \in \mathcal{A}, \psi \in \mathcal{H}\}}$.

[7]A directed set is an index set $\mathcal{J}$ together with an ordering $\prec$ that satisfies: (a) If $\alpha, \beta \in \mathcal{J}$ then there exists $\gamma \in \mathcal{J}$ such that $\alpha \prec \gamma$ and $\beta \prec \gamma$. (b) $\prec$ is a partial ordering in $\mathcal{J}$, i.e., $\prec$ is reflexive, transitive, and anti-symmetric.

The DIRECT SUM OF REPRESENTATIONS π_α is the map $\pi : \mathcal{A} \to \mathcal{B}(\oplus_{\alpha\in\mathcal{I}}\mathcal{H}_\alpha)$ defined by setting $\pi(A)$ equal to the operator $\pi_\alpha(A)$ on the component subspace $\mathcal{H}_\alpha$. In symbols: $\pi = \oplus_{\alpha\in\mathcal{I}}\pi_\alpha$. This definition yields bounded operators $\pi(A)$ on $\mathcal{H}$, because

$$\|\pi_\alpha(A)\| \leqslant \|A\| \qquad \forall A \in \mathcal{A},\, \forall\alpha \in \mathcal{I}\,.$$

One checks that $(\pi, \mathcal{H})$ is indeed a representation with

$$\|\pi(A)\| = \sup_{\alpha\in\mathcal{I}} \|\pi_\alpha(A)\|\,.$$

The notion of direct sum of representations allows one to state the following result:

Any *non-degenerate* representation $(\pi, \mathcal{H})$ of a C^*-algebra $\mathcal{A}$ is the direct sum of a family $\{(\pi_\alpha, \mathcal{H}_\alpha, \Omega_\alpha)\}_{\alpha\in\mathcal{I}}$ of cyclic (sub-) representations. That is, there exists $\{(\pi_\alpha, \mathcal{H}_\alpha, \Omega_\alpha)\}_{\alpha\in\mathcal{I}}$ such that $\mathcal{H} = \oplus_{\alpha\in\mathcal{I}}\mathcal{H}_\alpha$, $\pi = \oplus_{\alpha\in\mathcal{I}}\pi_\alpha$, and each Ω_α is cyclic for the corresponding π_α.

This fact follows by standard transfinite induction (Zorn's Lemma). Note that if the C^*-algebra $\mathcal{A}$ is unital, then *any* representation of $\mathcal{A}$ (no need to require non-degeneracy) is the direct sum of a family of sub-representations.

For practical purposes this result is quite useless: the decomposition $\pi = \oplus_{\alpha\in\mathcal{I}}\pi_\alpha$ can be highly not economical. Theoretically this result is of importance because it reduces the discussion of generic representations to that of cyclic representations and there is in fact a canonical manner for constructing representations (GNS).

8 GNS representation

So far the issue of the *existence* of representations was not touched. In fact, a C^*-algebra $\mathcal{A}$ has loads of representations. First of all, there is the following canonical one.

Let $\mathcal{A}$ be a C^*-algebra and let ω be a state on it (as seen above, ω exists always). Then there exists a cyclic representation $(\pi_\omega, \mathcal{H}_\omega, \Omega_\omega)$ such that

$$\omega(A) = \langle\Omega_\omega, \pi_\omega(A)\Omega_\omega\rangle \qquad \forall A \in \mathcal{A}$$

and, consequently,

$$\|\Omega_\omega\| = \|\omega\| = 1\,.$$

Such a $(\pi_\omega, \mathcal{H}_\omega, \Omega_\omega)$ is called the GNS-REPRESENTATION of $\mathcal{A}$ associated with ω. (GNS = Gelfand, Naimark, Segal.)

Any other cyclic representation $(\pi, \mathcal{H}, \Omega)$ such that $\omega(A) = \langle\Omega, \pi(A)\Omega\rangle$ $\forall A \in \mathcal{A}$ is UNITARILY EQUIVALENT to $(\pi_\omega, \mathcal{H}_\omega, \Omega_\omega)$, that is, there exists a unitary operator $U : \mathcal{H} \to \mathcal{H}_\omega$ such that

$$\begin{aligned} &U\Omega = \Omega_\omega \\ &U\pi(A)U^{-1} = \pi_\omega(A) \qquad \forall A \in \mathcal{A}\,. \end{aligned}$$

Such a U is nothing but the operator defined as $U\pi(A)\Omega := \pi_\omega(A)\Omega_\omega$ on the dense $\{\pi(A)\Omega \,|\, A \in \mathcal{A}\}$ of $\mathcal{H}$ and then extended by density (boundedness, unitarity, and the above properties for U are easily checked).

An important property of the GNS representation of a C^*-algebra $\mathcal{A}$ associated with a state ω on $\mathcal{A}$:

$$\pi_\omega \text{ is irreducible} \quad \Leftrightarrow \quad \omega \text{ is pure.}$$

This fact has two relevant consequences.

(1) Since (as seen above) for any element A of a C^*-algebra $\mathcal{A}$ there exists a pure state ω_A on $\mathcal{A}$ such that $\omega_A(A^*A) = \|A\|^2$, then for any $A \in \mathcal{A}$ there exists an irreducible representation $(\pi_A, \mathcal{H}_A)$ such that $\|\pi_A(A)\| = \|A\|$. ($A \to$ pure state $\omega_A \to$ irreducible GNS representation $\pi_A \equiv \pi_{\omega_A}$. The norm identity follows by $\|A\|^2 = \omega_A(A^*A) = \|\pi_{\omega_A}(A)\Omega_{\omega_A}\|^2 \leqslant \|\pi_{\omega_A}(A)\|^2 \leqslant \|A\|^2$.)

(2) For any $A \neq \mathbb{O}$ in $\mathcal{A}$ there exists a representation π such that $\pi(A) \neq \mathbb{O}$, i.e., the representations of a C^*-algebra separate points (this is usually referred to by saying that a C^*-algebra is REDUCED or SEMI-SIMPLE).

The GNS construction is mathematically important because it reduces the existence of Hilbert space representations of a C^*-algebra to the existence of states, which in turn is guaranteed by general (Hahn-Banach) existence theory. It is also physically relevant since it says that the (experimental) set of expectations of the observables given by a state have the customary Hilbert space interpretation. Thus, the basis of the mathematical description of Q.M. systems need not be postulated as in the Dirac-von Neumann axiomatic setting but it is merely a consequence of the C^*-algebra structure of the observables.

9 Structure theorem for C^*-algebras (Gelfand-Naimark)

Although every state ω on a C^*-algebra $\mathcal{A}$ defines a concrete realisation of a C^*-algebra as operators on a Hilbert space, this realisation may not be faithful (in the sense that π_ω might not be a faithful representation). Nevertheless, the following holds:

Every C^*-algebra $\mathcal{A}$ is $*$-isomorphic to a sub-algebra $\pi(\mathcal{A})$ of bounded operators on a Hilbert space which is closed in norm and under the adjoint: in other words, given $\mathcal{A}$ there exists a Hilbert space $\mathcal{H}$ and a representation $\pi : \mathcal{A} \to \mathcal{B}(\mathcal{H})$ which is faithful, i.e., norm-preserving.

This is the celebrated Gelfand-Naimark theorem. Mathematically it provides the basic structure theorem for C^*-algebras. Physically it completes the path from the C^*-algebraic formulation of Q.M. back to the Hilbert space formulation.

Such a faithful π is recovered as follows. For each state ω on $\mathcal{A}$ construct the associated GNS representation $(\pi_\omega, \mathcal{H}_\omega, \Omega_\omega)$ and form the direct sum representation $(\pi, \mathcal{H})$, $\mathcal{H} := \oplus_\omega \mathcal{H}_\omega$, $\pi := \oplus_\omega \pi_\omega$. As seen above, for each $A \in \mathcal{A}$ there exists a (pure) state ω_A such that $\|\pi_{\omega_A}(A)\| = \|A\|$. Since

$$\|\pi(A)\| \;=\; \sup_\omega \|\pi_\omega(A)\|$$

then $\|A\| = \|\pi_{\omega_A}(A)\| \leqslant \|\pi(A)\| \leqslant \|A\|$, whence $\|\pi(A)\| = \|A\|\ \forall A \in \mathcal{A}$. So π is faithful.

For us the true criterion for a "good" representation is faithfulness, i.e., the case where each element of the represented C^*-algebra is identified with one and only one Hilbert space operator. Cyclic representations are in some sense the building blocks ($\to$ Sec. 7) and the knowledge of the cyclic vector is enough to reconstruct the whole space. Also ($\to$ Sec. 8), given a cyclic representation $(\pi, \mathcal{H}, \Omega)$ of $\mathcal{A}$, this is (unitarily equivalent to) the GNS representation associated with the state $\langle \Omega, \pi(\cdot)\Omega \rangle$ on $\mathcal{A}$. Thus, any direct sum of cyclic representations is in fact a direct sum of GNS representations. As for irreps, they have the virtue of identifying invariant target spaces, but one single irrep in general brings only a very partial information on $\mathcal{A}$, the full information being encoded in the knowledge of all invariant subspaces (think of the irreps of the angular momentum, for example). We did not go in this direction, our interest was rather in the path from the C^*-algebraic axioms to the Dirac-von Neumann ones.

10 Commutative case

Let $\mathcal{A}$ be a commutative C^*-algebra.

A character ω on $\mathcal{A}$ is a non-zero linear map $\omega : \mathcal{A} \to \mathbb{C}$ such that

$$\omega(AB) \;=\; \omega(A)\omega(B)$$

for all $A, B \in \mathcal{A}$. In other words, a character is a non-zero multiplicative linear functional on $\mathcal{A}$. In fact (as already mentioned), this is not a new notion: one proves that a non-zero linear map $\omega : \mathcal{A} \to \mathbb{C}$ is a character on $\mathcal{A}$ if and only if it is a pure state on $\mathcal{A}$.

The set X of all characters (equivalently: of all pure states) on $\mathcal{A}$, equipped with the weak-$*$ topology inherited from the dual $\mathcal{A}^*$ of $\mathcal{A}$, is a locally compact Hausdorff space, which is compact if and only if $\mathcal{A}$ contains the identity.

Structural result (commutative Gelfand-Naimark theorem): $\mathcal{A}$ is $*$-isomorphic to the commutative C^*-algebra $C_0(X)$ of continuous functions on X which vanish at infinity, via the $*$-isomorphism (Gelfand transform) $\mathcal{A} \to C_0(X)$, $A \mapsto \widehat{A}$ defined by $\widehat{A}(\omega) := \omega(A)$.