YMSC Lectures

September 2024

1 Introduction

Hall conductance: robust quantity of insulating 2d electronic systems at near-zero temperature. Quantized to a rational number to an incredible accuracy.

Why is that?

Other quantities like this? How about thermal Hall conductance? It exists even when $U(1)$ symmetry is absent (say, in spin systems). Is it quantized as well?

2 Various explanations of robustness

1. Systems of non-interacting 2d fermions with a Fermi energy in a band gap. Here Hall conductance can be expressed as a Chern number of the valence bands. Thermal Hall conductance is proportional to Hall via the Wiedemann-Franz law.

2. Systems of non-interacting fermions with disorder. If the correlators of fermionic operators decay sufficiently fast, can define ”noncommutative Chern number” which determines both Hall and thermal Hall. Not a very intuitive definition.

3. Gapped interacting spin systems (on a lattice). One can either work on a torus or in infinite volume. On a torus, one can show that Hall is related to Berry curvature and is a rational number in the infinite-volume limit provided the ground-state degeneracy has a limit (as well as the integral of Berry curvature).

4. Laughlin's flux-insertion argument.

5. Field-theoretic explanations via bulk-boundary correspondence.

The last three approaches apply to interacting systems, but none of them can deal with thermal Hall conductance. The last two are also not mathematically rigoeous. In these lectures, I will develop a new formalism which in effect makes approaches 4 and 5 rigorous and extends them to a much larger class of systems and symmetries. But it is still not known how to explain the robustness of the thermal Hall conductance in the presence of interactions. This is an important open problem.

3 Gapped phases

We are interested in zero-temperature phases of quantum systems. Phases are assumed to be gapped so that correlation functions of local observables decay rapidly (faster than any power of distance). What is a gapped ground state? It is a unique ground state in infinite-volume and a nonzero minimal energy for all states created by acting with a local perturbation. Will give a more rigirous definition later.

4 Quantum Statistical Mechanics

Statistical Mechanics became a serious mathematical subject in 1950s, when it was realized that infinite systems with local interactions behave radically different from finite ones (thanks to Lars Onsager's exact solution of the 2d Ising model in 1944). Namely, even if the Hamiltonian is a smooth functions of parameters, the free energy is not, in general. Hence, phase transitions. Moreover, the notion of a "Gibbs state" becomes subtle for infinite systems, and there can be more than one Gibbs state for a fixed Hamiltonian and temperature. Hence, spontaneous symmetry breaking becomes possible.

We are dealing with systems at zero temperature, so QSM should be a useful tool for distinguishing such phases too. Instead of Gibbs states, we will be studying ground states. Also, we want to access invariant of phases which are not related to symmetry breaking. So we will assume that the ground states are invariant under all symmetries of the Hamiltonian.

5 Why Hilbert space is not the right thing to start with

The infinite tensor product

$$\otimes_{j\in\mathbb{Z}^d}\mathcal{H}_j$$

is rather pathological. First, it does not have a countable basis. Second, it does not have a natural inner product which would make it into a (non-separable) Hilbert space. This was first observed by John von Neumann who also proposed some ways out.

Simplest way out: consider the restricted tensor product where all vectors except a finite number are a particular $v_0 \in \mathcal{H}_0$. This is kind of like choosing a "fixed" boundary condition. Not clear how to choose v_0 though. The "correct" choice may depend on the Hamiltonian. Example: transverse-field Ising chain.

Better way: focus on observables. For a single site, the algebra of quantum observables is a matrix algebra $\mathscr{A}_j = L(\mathcal{H}_j, \mathcal{H}_j)$. Naively, the algebra of observables for $\mathbb{Z}^d$ should then be an infinite tensor product

$$\mathscr{A}_l = \otimes_{j\in\mathbb{Z}^d}\mathscr{A}_j.$$

This is a reasonably well-behaved algebra if we only consider observables which are 1 except for a finite number of sites. Then we have a countable basis and a good norm. The resulting normed algebra is called the algebra of local observables.

Many questions remain though.

- Is there a nice topology on $\mathscr{A}_l$?
- Observables should be operators on a Hilbert space. Where is the Hilbert space?
- How do we specify quantum dynamics? In other words, what sort of object is the quantum Hamiltonian?

Most of these issues are resolved if we enlarge a little our space of observables. This leads to Banach algebras and more specifically, to C^*-algebras.

6 Banach algebras and C^*-algebras

A normed algebra $\mathscr{A}$ is an algebra with a norm such that for all $a, b \in \mathscr{A}$ we have

$$\|a \cdot b\| \leq \|a\| \cdot \|b\|.$$

This "product inequality" is imposed because it makes all operations continuous in the topology arising from the norm.

Definition 1. A Banach algebra is a complete normed algebra.

We will work with algebras over complex numbers, so we need to add a bit more structure.

A $*$-algebra is an algebra with an involution (anti-linear anti-automorphism which squares to 1).

Definition 2. A Banach $*$-algebra is a Banach algebra over $\mathbb{C}$ with an involution such that $\|a^*\| = \|a\|$.

A complex $N \times N$ matrix algebra can be regarded as a Banach $*$-algebra, where $*$ is Hermitian conjugation and $\| \cdot \|$ is the "spectral norm":

$$\|a\| = \sup_{\|v\|=1} \|av\|,$$

where $\|v\|$ is the usual Euclidean norm on the N-dimensional vector space. Note that there are many (in fact, infinitely many) other norms on the algebra of complex matrices which make it into a Banach algebra, but the spectral norm is the most important one because it makes it into a C^*-algebra.

Definition 3. A C^*-algebra is a Banach $*$-algebra such that for all $a \in \mathscr{A}$ we have

$$\|a^* a\| = \|a\|^2.$$

Example 1. The algebra of bounded operators on a Hilbert space $\mathcal{H}$ is a C^*-algebra when equipped with the spectral norm. This algebra is denoted $B(\mathcal{H})$.

Example 2. The algebra of bounded functions on a set is a commutative C^*-algebra (when equipped with the usual $\|\cdot\|_\infty$ norm).

Example 3. The algebra of bounded continuous functions on a topological space is a commutative C^*-algebra (when equipped with the usual $\|\cdot\|_\infty$ norm).

Example 4. Let X be a locally compact Hausdorff topological space. The algebra of continuous functions on X which vanish "at infinity" is a commutative C^*-algebra (when equipped with the usual $\|\cdot\|_\infty$ norm). In fact, every commutative C^*-algebra is isomorphic to one of these (for some X).

In what follows, we will mostly work with Banach and C^*-algebras which have an identity element ("unital" algebras).

7 States on C^*-algebras

Let $\mathscr{A}$ be a C^*-algebra. Elements of the form $b = a^*a$ are called positive. One writes $b \geq 0$ in this case. The subset of positive elements in $\mathscr{A}$ is denoted $\mathscr{A}^+$. One can show that $\mathscr{A}^+$ is convex and invariant under $b \mapsto \lambda b$, $\lambda \in [0, +\infty)$. That is, $\mathscr{A}^+$ is a cone in $\mathscr{A}$.

Definition 4. Let $\mathscr{A}, \mathscr{B}$ be C^*-algebras. A linear map $f : \mathscr{A} \to \mathscr{B}$ is called positive if it commutes with $*$ and $f(\mathscr{A}^+) \subset \mathscr{B}^+$. In particular, a positive linear functional on $\mathscr{A}$ is a positive linear map $f : \mathscr{A} \to \mathbb{C}$. A state on $\mathscr{A}$ is a positive linear functional on $\mathscr{A}$ such that $f(1) = 1$.

Proposition 1. A linear functional ρ on a unital C^*-algebra is positive iff it is continuous and $\|\rho\| = \rho(1)$.

All states over $\mathscr{A}$ form a convex subset of the dual of $\mathscr{A}$ (i.e. of the space of continuous linear functionals $\mathscr{A} \to \mathbb{C}$). We have an obvious partial order on states and more general positive linear functionals: $\omega_1 \geq \omega_2$ iff $\omega_1 - \omega_2$ is positive.

Definition 5. A state ω is pure if it does not majorize a multiple of any other state except itself. Equivalently, it cannot be written as a convex linear combination of two distinct states.

8 Examples of states

For matrix algebras (i.e. linear operators on a f.d. Hilbert space $\mathcal{H}$), all states are described by density matrices:

$$a \mapsto \mathrm{Tr} \rho a,$$

where ρ must be positive (as a matrix) and normalized ($\text{Tr}\rho = 1$). Pure states are the ones for which ρ is a projector to a 1d space of $\mathcal{H}$.

For an infinite-dimensional $\mathcal{H}$, can also define a state using "density matrices" (more precisely, these are positive elements of $B(\mathcal{H})$ which are trace-class, i.e. $\text{Tr}\rho < \infty$. Such states are called normal. Surprisingly, not all states on $B(\mathcal{H})$ are normal! There exist pure states which vanish on all finite-rank projectors (Dixmier, 1969).

Remark 1. For states on a general C^*-algebra, we can define distance between states as

$$\|\omega - \omega'\| = \sup_{\|a\|=1} |\omega(a) - \omega'(a)|.$$

The distance between states is always 2 or less. This distance makes sense for arbitrary continuous linear functionals on $\mathscr{A}$, not just states. The state of continuous linear functionals on $\mathscr{A}$ thus becomes a Banach space. The corresponding topology on continuous linear functionals is called the strong topology (or "the topology of uniform convergence"). It does not come from a metric, in general. In the case of normal states on $B(\mathcal{H})$ the distance reduces to

$$\|\rho - \rho'\|_1 = \text{Tr}|\rho - \rho'|.$$

Note that this is different from the norm of the bounded operator $\rho - \rho'$ (in fact, for any bounded trace-class operator we have $\|\rho\|_1 \geq \|\rho\|$).

Remark 2. Another useful topology on continuous linear functionals is the topology of pointwise convergence (or weak-$*$ topology). It is defined by saying that a net of linear functionals ω_α converges to zero if for any $a \in \mathscr{A}$ the net of numbers $\omega_\alpha(a)$ converges to zero. If a net converges strongly, it also converges in weak-* topology, but the opposite is not true, in general.

We are interested in lattice systems. Here $\mathscr{A}$ is the completion of $\mathscr{A}_l = \bigotimes_{j\in\mathbb{Z}^d} \mathscr{A}_j$, where each $\mathscr{A}_j$ is a matrix algebra. We can first define a state on each $\mathscr{A}_j$ (given by a density matrix ρ_j), then define

$$\mathscr{A}_\Lambda = \bigotimes_{j\in\Lambda} \mathscr{A}_j$$

for every $\Lambda \in P_f(\mathbb{Z}^d)$ (i.e. for any finite subset of $\mathbb{Z}^d$), and define a state ω_Λ on $\mathscr{A}_\Lambda$ using a density matrix

$$\rho_\Lambda = \otimes_{j\in\Lambda} \rho_j.$$

States on different Λ are clearly compatible: if $\Lambda \subset \Lambda'$, then

$$\omega_\Lambda(a) = \omega_{\Lambda'}, \forall a \in \mathscr{A}_\Lambda.$$

Hence we get a positive normalized linear functional ω on $\mathscr{A}_l$. Clearly, we have

$$|\omega(a)| \leq \|a\|, \quad \forall a \in \mathscr{A}_l.$$

By the BLT theorem, ω extends to a continuous linear functional on $\mathscr{A}$. It is easy to check that it is positive. Hence we get a state on $\mathscr{A}$. Such states are called factorized states. If all ρ_j are pure, this state is also pure.

A special case of this construction is the "infinite-temperature state". This corresponds to taking all ρ_j to be proportional to the identity matrix. This state has the property

$$\omega(ab) = \omega(ba)$$

for all $a, b \in \mathscr{A}$. In other words, ω vanishes on commutators. A state which satisfies this condition is called a tracial state, for obvious reasons. Some algebras, like $B(\mathcal{H})$ for an infinite-dimensional $\mathcal{H}$, do not admit tracial states (one way to argue is to show that any bounded operator can be written as a sum of several commutators).

9 Representations

We need to learn how to "realize" abstract observables (elements of a C^*-algebra $\mathscr{A}$) as operators in a Hilbert space. I.e. we need to map each $a \in \mathscr{A}$ to $\pi(a) \in B(\mathcal{H})$ in a way compatible with algebra structure and $*$-operation (and maybe norm).

Definition 6. A representation of a $*$-algebra $\mathscr{A}$ on a Hilbert space $\mathcal{H}$ is a linear homomorphism $\pi : \mathscr{A} \to B(\mathcal{H})$ which "commutes with $*$": $\pi(a^*) = \pi(a)^\dagger$, $\forall a \in \mathscr{A}$.

This definition applies, in particular, to Banach $*$-algebras and C^*-algebras. For a C^*-algebra, a representation always preserves positivity. Indeed, if $a \in \mathscr{A}^+$, then we have $a = b^*b$, thus $\pi(a) = \pi(b)^\dagger \pi(b) \geq 0$.

Proposition 2. If $\mathscr{A}$ is a Banach $*$-algebra, then every representation is continuous, and in fact $\|\pi(a)\| \leq \|a\|$.

One says that a representation π is faithful if $\ker \pi = 0$. $\ker \pi$ is a two-sided ideal of $\mathscr{A}$. We will be dealing with algebras which are "simple" (i.e. have no non-trivial two-sided ideals), so all our representations will be faithful. One can show that for faithful representations $\|\pi(a)\| = \|a\|$ (basically, because π establishes an isomorphism between $\mathscr{A}$ and a C^*-sub-algebra of $B(\mathcal{H})$, and thus one can apply the preceding result to π^{-1} and deduce that $\|a\| \leq \|\pi(a)\|$).

Definition 7. Two representations π_1 and π_2 on Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$ are unitarily equivalent iff there exists a untary map $U : \mathcal{H}_1 \to \mathcal{H}_2$ such that for all $a \in \mathscr{A}$ we have

$$U\pi_1(a) = \pi_2(a)U.$$

Such a U is called an intertwiner.

Definition 8. A representation $\pi : \mathscr{A} \to B(\mathcal{H})$ is called irreducible iff the only closed $\pi(\mathscr{A})$-invariant subspaces of $\mathcal{H}$ are the zero subspace and the whole $\mathcal{H}$.

Definition 9. Let $\mathcal{B}$ be an algebra. For any $S \subset \mathscr{B}$ the commutant of S is the set of all elements of $\mathscr{B}$ which commute with all elements of S. The commutant of S is denoted S'.

Definition 10. A vector $v \in \mathcal{H}$ is called cyclic for a representation $\pi : \mathscr{A} \to B(\mathcal{H})$ iff the set of vectors of the form $\pi(a)v$, $a \in \mathscr{A}$, is dense in $\mathcal{H}$.

Proposition 3. The following conditions on $\pi : \mathscr{A} \to B(\mathcal{H})$ are equivalent:

- π is irreducible
- Any $v \in \mathcal{H}$ is cyclic for π
- The commutant of $\pi(\mathscr{A})$ consists of multiples of $1_{\mathcal{H}}$.

Proof. If $v \in \mathcal{H}$ is not cyclic, then the orthogonal complement of $\pi(\mathscr{A})v$ is a nontrivial closed linear subspace invariant under $\pi(\mathscr{A})$. Thus (1) implies (2). Now assume (2). Suppose $\pi(\mathscr{A})'$ contains a. We can choose a to be self-adjoint, and any spectral projection of a will also be in $\pi(\mathscr{A})'$. If the spectrum of a contains more than one point, then we can choose a spectral projection which does not contain some point $\lambda \in \sigma(a)$, and any vector in the image of this spectral projection will not be cyclic. Thus a must be a multiple of $1_{\mathcal{H}}$, and (2) implies (3). Now assume (3) and suppose that there is a closed invariant subspace for $\pi(\mathscr{A})$. Then the projector to this subspace will be in $\pi(\mathscr{A})'$, which means that (3) implies (1). □

Remark 3. Obviously, $\{1\}' = B(\mathcal{H})$. Therefore, for an irreducible $\pi : \mathscr{A} \to \mathcal{H}$, we have $\pi(\mathscr{A})'' = B(\mathcal{H})$. But this does not mean that $\pi(A) = B(\mathcal{H})$. That is, typically not every $A \in B(\mathcal{H})$ is of the form $\pi(a)$. In some sense, this is to be expected: for faithful representations (and in the cases of interest to us, $\mathscr{A}$ is simple, and thus all representations are faithful), π is an isomorphism between $\mathscr{A}$ and some C^*-sub-algebra of $B(\mathcal{H})$, and if $\pi(\mathscr{A})$ were the whole $B(\mathcal{H})$, it would mean that if a faithful irreducible representation of $\mathscr{A}$ exists, $\mathscr{A}$ is isomorphic to $B(\mathcal{H})$, and the whole theory would "trivialize". This does not happen precisely because $\pi(\mathscr{A})$ can be strictly smaller than $\pi(\mathscr{A})''$.

10 The GNS construction

Turns out one can construct a representation of $\mathscr{A}$ starting with a state ω on $\mathscr{A}$.

Start by defining an "scalar product" on $\mathscr{A}$:

$$(a, b) = \omega(a^* b).$$

It has all the right properties except it can happen that $(a, a) = 0$ but $a \neq 0$. (In other words, $\|a\|_2 = \sqrt{(a, a)}$ is not a norm, but a "seminorm"). Let's call such a isotropic.

Lemma 1. Let ω be a state on $\mathscr{A}$. The set of isotropic elements for ω is a left ideal of $\mathscr{A}$.

Proof. Suppose a is isotropic and b is arbitrary. By positivity of ω and Cauchy-Schwarz, we have

$$(ba, ba)^2 = \omega(a^* b^* ba)^2 \leq \omega(a^* a)\omega(b^* baa^* b^* b) = 0.$$

Thus $(ba, ba) = 0$. $\square$

Now we can form the quotient of the vector space $\mathscr{A}$ by the ideal J_ω of isotropic elements. We will denote by $[a]$ the equivalence class of a. The scalar product descends to the quotient $\mathscr{A}/J_\omega$ (again because of Cauchy-Schwarz). There are no nonzero isotropic elements in $\mathscr{A}/J_\omega$, because we declared them to be equivalent to zero. All that remains is to complete the space by adding the limits of all Cauchy sequences. This gives a Hilbert space $\mathcal{H}_\omega$.

For each $a \in \mathscr{A}$ we define a linear operator $\pi_\omega(a) : \mathcal{H}_\omega \to \mathcal{H}_\omega$ for by

$$\pi_\omega(a)[b] = [ab].$$

This operator is bounded because

$$(\pi_\omega(a)[b], \pi_\omega(a)[b]) = (ab, ab) = \omega(b^* a^* ab) \leq \omega(b^* b)\|a\|^2 = ([b], [b])\|a\|^2.$$

In fact, the above computation shows that $\|\pi_\omega(a)\| \leq \|a\|$.

We also have

$$\pi_\omega(a)\pi_\omega(b)[c] = [abc] = \pi_\omega(ab)[c],$$

as well as

$$([b], \pi_\omega(a^*)[c]) = \omega(b^* a^* c) = \omega((ab)^* c) = (\pi_\omega(a)[b], [c]).$$

Thus $\pi_\omega(a^*) = \pi_\omega(a)^\dagger$.

All this means that $a \mapsto \pi_\omega(a)$ is a representation of $\mathscr{A}$. This representation is called the GNS (Gelfand-Naimark-Segal) representation of $\mathscr{A}$ associated with ω. It has a cyclic vector $|0_\omega\rangle = \Omega_\omega = [1]$. Note that $(\Omega_\omega, \Omega_\omega) = \omega(1) = 1$.

Here is an important property of the GNS representation: every symmetry of $(\mathscr{A}, \omega)$ gives rise to a unitary operator U on $B(\mathcal{H}_\omega)$. More precisely, we have

Proposition 4. Let τ be a $*$-automorphism of $\mathscr{A}$ which preserves ω: $\omega \circ \tau = \omega$. Then there exists a unique unitary operator U such that $U\Omega_\omega = \Omega_\omega$ and

$$\pi_\omega(\tau(a)) = U\pi_\omega(a)U^{-1}.$$

One says that U implements symmetry τ.

Proof. U is defined on the image of π_ω by $U[a] = [\tau(a)]$. It is easily checked that this is a well-defined map (i.e. preserves J_ω) and that it preserves the scalar product. Then it extends to the whole $\mathcal{H}_\omega$ by the BLT theorem. $\square$

Lastly, we have

Theorem 1. π_ω is irreducible iff ω is pure.

Proof. 1. Suppose ω is pure. If $\pi_\omega(\mathscr{A})'$ contains elements other than multiples of identity, then we can pick a self-adjoint $A \in \pi_\omega(\mathscr{A})'$ such that $A \neq 1$. The spectrum of A has more than one point, hence we can choose a spectral projection of A which is less than 1. It is also in $\pi_\omega(\mathscr{A})'$. Now consider the following linear functional on $\mathscr{A}$:

$$\sigma(a) = (P\Omega_\omega, \pi_\omega(a)\Omega_\omega).$$

It is easy to check that $\sigma \geq 0$, and also $\omega - \sigma \geq 0$. Thus it must be that $\sigma = \lambda\omega$ (because ω is pure). In other words, for all $a \in \mathscr{A}$ we must have

$$(P\Omega_\omega, \pi_\omega(a)\Omega_\omega) = (P\Omega_\omega, \Omega_\omega)(\Omega_\omega, \pi_\omega(a)\Omega_\omega).$$

There are two ways it can happen. First, it may happen that Ω_ω is annihilated by P. That is, $\sigma = 0$. Since P was an arbitrary spectral projection of an element in $\pi_\omega(\mathscr{A})'$ which is not a multiple of identity, this means that any element b in $\pi_\omega(\mathscr{A})'$ which is not a multiple of identity annihilates all vectors of the form $\pi_\omega(a)\Omega_\omega$. But these vectors are dense in $\mathcal{H}_\omega$. So any such b is zero. So π_ω is irreducible.

The other option is that $\pi_\omega(a)\Omega_\omega$ is proportional to Ω_ω, for all $a \in \mathscr{A}$. But since Ω_ω is a cyclic vector, the whole $\mathcal{H}_\omega$ is one-dimensional, and π_ω is irreducible in a trivial way.

2. Now suppose π_ω is irreducible. Let σ be a positive linear functional majorized by ω. Then by Cauchy-Schwartz

$$|\sigma(a^*b)|^2 \leq \sigma(a^*a)\sigma(b^*b) \leq \omega(a^*a)\omega(b^*b).$$

This shows that the map $(a, b) \mapsto \sigma(a^*b)$ is bounded and vanishes when either a or b are in J_ω. Hence we can regard it as a bounded sesquilinear map $\mathcal{H}_\omega \times \mathcal{H}_\omega \to \mathbb{C}$. It is well known from basic functional analysis that all such maps arise from bounded operators. I.e. there is a bounded operator $T \in B(\mathcal{H}_\omega)$ such that

$$\sigma(a^*b) = ([a], T[b]).$$

We claim that T is in the commutant of $\pi_\omega(\mathscr{A})$. Indeed:

$$([a], \pi_\omega(b)T[c]) = ([b^*a], T[c]) = \sigma(a^*bc) = ([a], T[bc]) = ([a], T\pi_\omega(b)[c]).$$

Since this is true for arbitrary a, b, c, T commutes with $\pi_\omega(b)$ for all $b \in \mathscr{A}$. Hence $T = \lambda 1$ (because π_ω is irreducible), and thus $\sigma = \lambda\omega$. □

Remark 4. It is not true in general that for an irreducible $\pi : \mathscr{A} \to B(\mathcal{H})$ every operator $A \in B(\mathcal{H})$ can be approximated arbitrarily well by elements of the form $\pi(a)$ (in the norm topology). This might seem surprising. On the other hand, turns out there is a topology on $B(\mathcal{H})$ (weak topology) in which every element of $\pi(\mathscr{A})''$ is in the closure of $\pi(\mathscr{A})$. As a result, $\pi(\mathscr{A})''$ is always "weakly" closed and thus is a very special kind of sub-algebra of $B(\mathcal{H})$. Such sub-algebras are called von Neumann algebras. For an irreducible π, we have $\pi(\mathscr{A})'' = \{1\}' = B(\mathcal{H})$, but if π is reducible, $\pi(\mathscr{A})''$ can be something weird. For example, if we take ω to be the tracial state on a quasi-local $\mathscr{A}$ (see Section 1), the resulting von Neumann algebra $\pi_\omega(\mathscr{A})''$ is already weird (not isomorphic to $B(\mathcal{V})$ for any Hilbert space $\mathcal{V}$.) This was one of von Neumann's great discoveries.

11 Dynamics

Let $Aut(\mathscr{A})$ be the group of $*$-automorphisms of $\mathscr{A}$. Inner automorphisms are automorphisms of the form

$$a \mapsto uau^*,$$

where u is a unitary element of $\mathscr{A}$. Inner automorphisms from a normal subgroup of $Aut(\mathscr{A})$.

Definition 11. A one-parameter subgroup of automorphisms of a C^*-algebra $\mathscr{A}$ is a homomorphism $\mathbb{R} \to Aut(\mathscr{A})$.

That is, we have $t \mapsto \alpha_t$, and

$$\alpha_t \circ \alpha_s = \alpha_{t+s}.$$

"Hamiltonian" time-evolution is represented by one-parameter subgroups (in the Heisenberg picture). In the Schrödinger picture, we evolve states instead:

$$\omega(0) \mapsto \omega(t) = \omega(0) \circ \alpha_t.$$

Sometimes one also considers one-parameter semigroups of positive linear maps $\mathscr{A} \to \mathscr{A}$. They are used to represent irreversible evolution of open systems.

Naively, would like to write $\alpha_t(a) = e^{iHt}ae^{-iHt}$. But what is H? It is not an element of $\mathscr{A}$. To figure out what H means, let's try taking derivatives of α_t w. r. to t.

Definition 12. A one-parameter subgroup of automorphisms is strongly continuous if $t \mapsto \alpha_t(a)$ is a continuous function of t (for any $a \in \mathscr{A}$).

Actually, it is enough to require that $\lim_{t\to 0} \|\alpha_t(a) - a\| = 0$.

Definition 13. The generator of a strongly continuous 1-parameter family of automorphisms is a linear map $\delta : D(\delta) \to \mathscr{A}$, $D(\delta) \subset \mathscr{A}$ given by

$$\delta(a) = \lim_{t\to 0} \frac{1}{t}(\alpha_t(a) - a).$$

Definition 14. A (symmetric) derivation δ of a C^*-algebra $\mathscr{A}$ with a domain $D(\delta)$ is a linear map $D(\delta) \to \mathscr{A}$ such that $\delta(a^*) = \delta(a)^*$ and $\delta(ab) = \delta(a)b + a\delta(b)$ for all $a, b \in D(\delta)$.

Simplest derivations are inner derivations: derivations of the form

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

where $b^* = -b \in \mathscr{A}$. We will denote such a derivation ad_b. Inner derivations are defined everywhere, i.e. one can take $D(\delta) = \mathscr{A}$.

Proposition 5. The generator of a strongly continuous family of automorphisms is a derivation of $\mathscr{A}$ wwith a dense domain $D(\delta)$.

Proof. The first property ($\delta(a^*) = \delta(a)^*$) is obvious. For the second one, we write

$$\alpha_t(ab) - ab = \alpha_t(a)\alpha_t(b) - ab = (\alpha_t(a) - a)\alpha_t(b) + a(\alpha_t(b) - b).$$

Now divide by t and take the limit $t \to 0$.

As for the domain, consider elements of the form

$$a_f = \int_{-\infty}^{\infty} f(t)\alpha_t(a)dt,$$

where $f : \mathbb{R} \to \mathbb{C}$ is a smooth function with a compact support. It is easy to check that $\delta(a_f)$ is well-defined, namely

$$\delta(a_f) = a_{-f'}.$$

Checking that elements of the form a_f are dense in $\mathscr{A}$ is left as an exercise. □

One can show that a derivation which is defined everywhere is bounded. Sakai showed that for some C^*-algebras (for example, quasi-local algebras) all bounded derivations are inner. Every inner derivation ad_b is a generator of a 1-parameter subgroup of inner automorphisms: just set $\alpha_t(a) = e^{tb}ae^{-tb}$. This is well-defined because b is anti-self-adjoint. But these 1-parameter subgroups are not terribly interesting: the corresponding evolution is non-trivial only near some point (because b is quasi-local). That is, if we fix t and take a to be a local observable sufficiently far out, $\alpha_t(a) - a$ can be made arbitrarily small. This is not a physically sensible evolution.

Definition 15. A C^*-dynamical system is a pair $(\mathscr{A}, \alpha_t)$ where $\mathscr{A}$ is a C^*-algebra and α_t is a strongly continuous 1-parameter group of automorphisms of $\mathscr{A}$.

Continuous symmetries, such as a $U(1)$ symmetry, are also realized by 1-parameter groups of automorphisms, except that the parameter t is periodically identified (or more generally, takes value in the group manifold). More precisely, one says that G is a symmetry of a C^*-dynamical system if one is given a homomorphism $\tau : G \to Aut(\mathscr{A})$ such that $\tau(g)$ commutes with α_t for all $g \in G$ and all $t \in \mathbb{R}$. The corresponding generators live in the Lie algebra of the symmetry group and formally commute with δ. Again, interesting generators are not bounded. So strictly speaking the commutator of two derivations δ_1, δ_2 may not even be defined (if the domains of δ_1, δ_2 have no intersection). So for example the statement that generators of symmetries form a Lie algebra does not make sense, in general. It only makes sense if all symmetries (including time-translation symmetry) have a common dense domain in $\mathscr{A}$. We will see later how to deal with this difficulty.

A related issue is how to ”integrate” a derivation to a 1-parameter subgroup. Not every derivation with a dense domain can be integrated. Formally, one needs to solve the differential equation

$$\frac{d\alpha_t(a)}{dt} = \alpha_t(\delta(a))$$

or the equivalent integral equation. We will see later how this is done for special classes of derivations.

12 Evolution and symmetries in Hilbert space

In ordinary QM, generators of symmetries are (typically, unbounded) self-adjoint operators on the Hilbert space which commute with the Hamiltonian. Let's see how this comes about in the C^*-algebraic formalism.

First, we need to discuss/recall some facts about unbounded operators (see Reed-Simon, Vol. 1, Chapter VIII for details).

Definition 16. An unbounded operator A on $\mathcal{H}$ with a domain $D(A)$ (a linear subspace of $\mathcal{H}$) is a linear map $D(A) \to \mathcal{H}$.

One usually assumes that $D(A)$ is dense in $\mathcal{H}$, i.e. every element of $\mathcal{H}$ can be approximated arbitarily well by an element of $D(A)$.

Definition 17. An unbounded operator A is called symmetric if $(v', Av) = (Av', v)$ for all $v, v' \in D(A)$.

In the case when $\mathcal{H}$ is finite-dimensional, such operators are called self-adjoint. But in general, a self-adjoint unbounded operator is a more subtle notion:

Definition 18. Let A be a symmetric operator A with a dense domain $D(A)$. Suppose the following condition holds: if $v', w \in \mathcal{H}$ and for all $v \in D(A)$ one has $(v', Av) = (w, v)$, then $v' \in D(A)$ (and thus $w = Av'$). Equivalently, suppose for any $v' \in \mathcal{H}$ if there exists $C > 0$ such that $(v', Av) \leq C\|v\|$ for all $v \in D(A)$, then $v' \in D(A)$. Then A is called self-adjoint.

This notion is important because for self-adjoint unbounded operators there is a generalization of functional calculus:

Theorem 2. If A is self-adjoint, there is a $*$-homomorphism ϕ from the C^*-algebra of bounded measurable function on $\mathbb{R}$ to $B(\mathcal{H})$ which is positive (and thus norm-continuous,) and satisfies the following condition: if $\lim_{n\to\infty} f_n(x) \to x$ and $|f_n(x)| \leq |x|$ for all $x \in \mathbb{R}$, then $\lim_n (\phi \circ f_n)v = Av$ for all $v \in D(A)$.

So we can construct bounded operators from self-adjoint unbounded ones. In particular, if A is self-adjoint, then $U(t) = e^{iAt}$ is a well-defined unitary operator for all $t \in \mathbb{R}$ and $U(t)U(s) = U(t+s)$. Moreover, $U(t)v$ is a continuous function of t for all $v \in \mathcal{H}$. One says that $U(t)$ is a strongly continuous one-parameter group of unitaries. This is how we usually construct evolution in QM.

The operator A can be recovered from $U(t)$: for any $v \in D(A)$ one has

$$\lim_{t\to 0} \frac{1}{t}(U(t)v - v) = iAv,$$

Also, if for some $v \in \mathcal{H}$ this limit exists, then $v \in D(A)$.

In fact, all strongly-continuous one-parameter groups of unitaries arise in this way:

Theorem 3. (Stone's theorem) Let $U(t)$ be a strongly-continuous one-parameter group of unitaries in a Hilbert space $\mathcal{H}$. Then there exists a self-adjoint unbounded A such that $U(t) = e^{iAt}$.

Idea of the proof: for any infinitely-differentiable function $f : \mathbb{R} \to \mathcal{C}$ with compact support, we can consider vectors of the form $v_f = \int f(s)U(s)vds$. Then

$$U(t)v_f = \int f(s)U(t+s)vds = \int f(s-t)U(s)vds.$$

Then for such vectors we can define

$$iAv_f = v_{-f'}.$$

With some work, one can show that A can be extended to a self-adjoint operator, so e^{iAt} is well-defined. And then it is straightforward to check that $e^{iAt} = U(t)$ (basically, because they sastisfy the same differential equation). See Reed-Simon I, Chapter VIII for details.

Now suppose α_t is a strongly continuous one-parameter group of automorphisms of $\mathscr{A}$ which preserves a state ω:

$$\omega(\alpha_t(a)) = \omega(a), \quad \forall a \in \mathscr{A}.$$

Then we can define a unitary $U(t)$ in the GNS representation $(\pi_\omega, \mathcal{H}_\omega)$ such that $\pi_\omega(\alpha_t(a)) = U(t)\pi(a)U(t)^{-1}$. We just set $U(t)[a] = [\alpha_t(a)]$. Clearly, $U(t)U(s)[a] = U(t+s)[a]$. One can also show that $U(t)[a]$ is a continuous function of t. Thus by Stone's theorem we have $U(t) = e^{iHt}$ for some self-adjoint unbounded operator H on the GNS Hilbert space $\mathcal{H}_\omega$. We finally got a Hamiltonian on Hilbert space!

Note that $U(t)\Omega_\omega = U(t)[1] = [\alpha_t(1)] = [1] = \Omega_\omega$. By taking the derivative w.r. to t at $t = 0$, we see that $H\Omega_\omega = 0$. Thus the cyclic vector of the GNS representation has zero energy by definition. No need to subtract out "vacuum energy".

Let's discuss an important notion of a ground state. Suppose ω is pure, so that the GNS representation is irreducible, with the cyclic vector $\Omega_\omega = [1]$. Then usually the spectrum of H is bounded from below. The most interesting case is when Ω_ω has the lowest possible energy. Then we say that Ω_ω is the ground state of H. Of course, if this is the case, then $H \geq 0$.

Remark 5. It may happen that the ground state is not unique. For an infinite-volume system this is rare though.

Turns out one can define what a ground state directly in terms of $\mathscr{A}$ and ω.

Definition 19. ω is a ground state of a strongly-continuous one-parameter group of automorphisms of $\mathscr{A}$ iff for all $a \in D(\delta)$ we have $\omega(a^*(-i\delta(a))) \geq 0$.

Then we have

$$(\Omega_\omega, \pi_\omega(a)^\dagger[H, \pi_\omega(a)]\Omega_\omega) \geq 0,$$

or equivalently

$$([a], H[a]) \geq 0, \quad \forall a \in D(\delta).$$

Since $D(\delta)$ is dense in $\mathscr{A}$, and vectors of the form $[a]$ are dense in $\mathcal{H}$, this implies that $H \geq 0$, and thus Ω_ω is a ground state of H.

13 Approximately inner derivations

First, let's define a large class of derivations which are well-defined on $\mathscr{A}_l$. Let's fix an interaction Φ. For any finite $\Lambda \subset \mathbb{Z}^d$ and any $a \in \mathscr{A}_l$ we let

$$\delta_\Lambda^\Phi(a) = \sum_{X \subset \Lambda} i[\Phi(X), a]$$

This is obviously an inner derivation. We would like to define an "approximately inner" derivation $\mathscr{A}_l \to \mathscr{A}$ by

$$\delta_\Phi(a) = \lim_\Lambda \delta_\Lambda^\Phi(a).$$

Theorem 4. Let an interaction Φ satisfy

$$\sum_{X \ni 0} \|\Phi(X)\| < \infty.$$

Then the derivation δ_Φ is well-defined.

Proof. Suppose $a \in \mathscr{A}_Y$ for some finite Y and $\Lambda' \supset \Lambda$. Then

$$\delta_{\Lambda'}^\Phi(a) - \delta_\Lambda^\Phi(a) = \sum_{X \subset \Lambda', X \not\subset \Lambda} i[\Phi(X), a].$$

The norm of the above expression is upper-bounded by

$$2\|a\| \sum_{j \in Y} \sum_{X \ni j, X \subset \Lambda', X \not\subset \Lambda} \|\Phi(X)\|.$$

For any $\varepsilon > 0$ we can always choose R so that

$$\sum_{X \ni 0, X \not\subset B_0(R)} \|\Phi(X)\| < \varepsilon.$$

Then we for all Λ which contain $B_j(R)$ for all $j \in Y$ the above expression will be upper bounded by $2\varepsilon\|a\||Y|$. So $\delta_\Lambda(a)$ is a Cauchy net and converges to some element in $\mathscr{A}$. ☐

Now we want is to exponentiate δ_Φ and get a well-defined α_t. That is, we want to solve the equation

$$\frac{d\alpha_t^\Phi(a)}{dt} = \alpha_t^\Phi(\delta^\Phi(a)).$$

This requires further conditions on Φ.

Fix Φ. Note that δ_Λ is an inner derivation of $\mathscr{A}_\Lambda$, which is a matrix algebra. So we can define α_t^Λ by exponentiation without any problems. Namely, for $a \in \mathscr{A}_\Lambda$ we have

$$\alpha_t^\Lambda(a) = e^{it\sum_{X\subset\Lambda}\Phi(X)} a e^{-it\sum_{X\subset\Lambda}\Phi(X)}.$$

It clearly solves the equation we wanted to solve. α_t^Λ is defined everywhere on $\mathscr{A}_\Lambda$. If we want, we can extend it to the whole $\mathscr{A}_l$ using the fact that $\mathscr{A}_l$ is the tensor product of $\mathscr{A}_\Lambda$ and $\otimes_{j\notin\Lambda}\mathscr{A}_j$. α_t^Λ maps $\mathscr{A}_l$ to itself.

We want to ensure that

$$\alpha_t(a) = \lim_\Lambda \alpha_t^\Lambda(a)$$

exists for any $a \in \mathscr{A}_l$.

Definition 20. For any $r > 0$ let

$$\|\Phi\|^{(r)} = \sum_{X\ni 0} e^{r(|X|-1)}\|\Phi(X)\|$$

The set of interactions with $\|\Phi\|^{(r)} < \infty$ forms a Banach space $\mathcal{B}^{(r)}$.

Lemma 2. Let $\mathrm{ad}_{\Phi,\Lambda}(a) = \sum_{X\subset\Lambda}[\Phi(X), a]$. Then for $\Phi \in \mathcal{B}^{(r)}$ and $a \in \mathscr{A}_Y$ we have

$$\|\,\mathrm{ad}^n_{\Phi,\Lambda}(a)\| \le n!\left(\frac{2}{r}\|\Phi\|^{(r)}\right)^n e^{r|Y|}\|a\|.$$

Proof.

$$\mathrm{ad}^n_{\Phi,\Lambda}(a) = \sum_{X_1,\ldots,X_n\subset\Lambda}[\Phi(X_n),\ldots,[\Phi(X_1),a]\ldots].$$

The sets which contribute to this sum form a "chain" of overlapping sets: X_1 overlaps with Y, X_2 overlaps with either X_1 or Y, etc. Let

$$\phi_m = \sum_{X\ni 0,|X|=m}\|\Phi(X)\|.$$

Let's fix $X_1,\ldots,X_n$ of sizes $m_1,\ldots,m_n$. We call the corresponding contribution $f_n(X_1,\ldots,X_n)$. There are at most $|Y| + \sum_{j=1}^{n-1}(m_j-1)$ points in the union of Y and $X_1,\ldots,X_{n-1}$. Therefore summation over possible X_n with $|X_n| = m_n$ gives $f_n(X_1,\ldots,X_n)$ whose norm is upper-bounded by

$$2\phi_{m_n}(|Y| + \sum_{j=1}^{n-1}(m_j-1))\|f_{n-1}(X_1,\ldots,X_{n-1})\|.$$

Iterating this and summing over $m_1,\ldots,m_n$ we get an estimate for the multiple commutator:

$$\|a\|\sum_{m_i\ge 1}\prod_{k=1}^{n}(|Y| + \sum_{j=1}^{k-1} - (k-1))\prod_k(2\phi_{m_k})$$

Now, we can upper-bound the above expression by

$$\|a\| \sum_{m_i \geq 1} \left(|Y| + \sum_{j=1}^{n} (m_j - 1) \right)^n \prod_{k=1}^{n} (2\phi_{m_k}).$$

Then for any $r > 0$ we can write

$$y^n \leq n! e^{ry} r^{-n}.$$

Using this in the above expression, we can upper-bound it by

$$\|a\| n! (2/r)^n e^{r|Y|} \sum_{m_i \geq 1} \prod_{j=1}^{n} e^{r(m_j-1)} \phi_{m_j} = \|a\| n! (2/r)^n e^{r|Y|} \left(\|\Phi\|^r \right)^n$$

□

Theorem 5. Fix $r > 0$ and a $\Phi \in \mathcal{B}^{(r)}$. For any finite $Y \subset \mathbb{Z}^d$, any $a \in \mathscr{A}_Y$ there exists $T > 0$ such that for any $|t| < T_0$ the limit

$$\lim_{\Lambda} \alpha_t^{\Lambda}(a)$$

exists, and the convergence is uniform in t.

Proof. Just set $T = \frac{r}{k} (\|\Phi\|^{(r)})^{-1}$ for any $k > 2$ and use the estimate from the previous lemma. □

Theorem 6. The family of $*$-homomorhisms $\alpha_t : \mathscr{A}_l \to \mathscr{A}$, $|t| < T$, extends to a strongly continuous 1-parameter group of automorphisms of $\mathscr{A}$. These automorphisms commute with spatial translations.

Proof. First, since α_t^{Λ} is norm-preserving, so is α_t. Thus, by BLT theorem, we can uniquely extend it to $\mathscr{A}$ to a norm-presrving homomorphism. It is actually an automorphism, with the inverse of α_t being α_{-t}.

Second, let $s, r \in [-T, T]$. Then

$$\alpha_{s+r}^{\Lambda}(a) - \alpha_s \alpha_r(a) = \alpha_s^{\Lambda}(\alpha_r^{\Lambda}(a) - \alpha_r(a)) + (\alpha_s^{\Lambda} - \alpha_s)(\alpha_r(a)).$$

Now let's take the limit over Λ. Both terms converge on the r.h.s. converge to zero, so α_{s+r}^{Λ} converges to $\alpha_s \alpha_r(a)$. It follows from this that convergence is uniform on $[-2T, 2T]$, and by iteration for any bounded interval of t.

Third, since $\alpha_s^{\Lambda} \alpha_t^{\Lambda} = \alpha_{t+s}^{\Lambda}$, by passing to the limit we get that α_t is a one-parameter group of automorphisms.

Fourth, uniform convergence on any bounded subset for t and continuity of $\alpha_t^{\Lambda}(a)$ as a function of t imply continuity of $\alpha_t(a)$ as a function of t.

Fifth, translational invariance of α_t follows from the translational covariance of α_t^{Λ} and the convergence of α_t^{Λ} to α_t. □

14 Lieb-Robinson bounds

Turns out if $\|\Phi(X)\|$ decays rapidly with the size of X, dynamics defined by the interaction $X \mapsto \Phi(X)$ approximately preserves locality. That is, if $a \in \mathscr{A}_\ell$ is supported on some finite subset Γ, then for any fixed t $\|[b, a(t)]\|$ becomes very small if the support of $b \in \mathscr{A}_\ell$ is far from Γ. The first such estimate is due to Lieb and Robinson who assumed finite-range interactions. Then the decay with the distance to Γ is exponential. Many generalizations of this result are also called Lieb-Robinson bound. Here is one example.

Theorem 7. Let $F : [0, +\infty) \to (0, +\infty)$ be a monotonically decreasing function such that

$$\sum_{x \in \mathbb{Z}^d} F(|x|) < \infty$$

and

$$\sum_{y \in \mathbb{Z}_d} F(|x-y|)F(|y-z|) \leq C_F F(|x-z|)$$

for some $C > 0$. Let Φ satisfy

$$\|\Phi\|_F = \sup_{x,y} \frac{1}{F(|x-y|)} \sum_{Z \ni x,y} \|\Phi(Z)\| < \infty.$$

Then for any finite subsets $X, Y \subset \mathbb{Z}^d$ and any $a \in \mathscr{A}_X$ and $b \in \mathscr{A}_Y$ and any $t \in \mathbb{R}$ one has

$$\|[a(t), b]\| \leq \frac{2}{C_F} \|a\| \|b\| \left(e^{2\|\Phi\|_F C_F |t|} - \theta(d(X,Y)) \right) \sum_{x \in X, y \in Y} F(|x-y|).$$

For a proof, see https://arxiv.org/abs/1102.0842 .

These bounds have many applications. For example, they imply that if ω is a gapped state and interactions are finite-range, then $\omega(ab) - \omega(a)\omega(b)$ decays exponentially with the distance between supports of a and b (M. Hastings).

15 The noncommutative Chern number

As an aside, let's recall how the Hall conductance can be computed for systems of non-interacting electrons on $\mathbb{Z}^2$. There one works with 1-particle Hilbert space K which can be taken as $\oplus_{j \in \mathbb{Z}^2} L_j$, where L_j is the Hilbert space labelling fermionic creation and annihilation operators on site j. This can be used to construct a C^*-algebra (called CAR algebra, i.e. creation-annihilation algebras). A state of non-interacting electrons is entirely specified by the 2-point function, namely:

$$P_{ij} = \omega(c_i^\dagger c_j).$$

It can be viewed as a bounded operator on K satisfying $0 \leq P \leq 1$. Moreover, ω is pure if and only if P is a projector, $P^2 = P$.

For any region $X \subset \mathbb{Z}^d$ let $\Pi_X : K \to K$ be an operator which projects to X. That is, $\Pi_{X,ij} = 1$ if $i, j \in X$ and 0 otherwise.

If X is bounded, $\operatorname{Tr} \Pi_X$ is equal to $\sum_{j \in X} \dim V_j$. If X is not bounded, Π_X is not a trace-class operator. However, one can show that if $\mathbb{R}^2$ is decomposed into a union of three cones A_1, A_2, A_3, then

$$[P\Pi_{A_0}P, P\Pi_{A_1}P]$$

is trace-class, thus

$$\operatorname{Tr} P[P\Pi_{A_0}P, P\Pi_{A_1}P]$$

is well-defined. In fact, one can show that it is an integer times $4\pi i$ and is independent of the choice of cones once the orientation of $\mathbb{R}^2$ has been chosen. This integer is called the noncommutative Chern number of the state defined by the projector P (Avron, Seiler, Simon; Bellissard, van Elst, Schulz-Baldes).