Non-commutative graphs

Matthew Daws

Lancaster

Wrocław, May 2024

Graphs

A graph consists of a (finite) set of *vertices* V and a collection of *edges* $E \subseteq V \times V$.

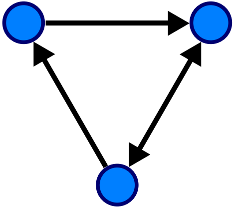

$V = \{A, B, C\}$ say, and $E = \{(A, B), (B, C), (C, B), (C, A)\}$.

So a graph $G = (V, E)$ is nothing but a *relation* on the set V.

- In general not reflexive (unless every vertex has a self-loop);
- Symmetric when G is undirected;
- Rarely transitive.

Channels

A channel sends an input message (element of a finite set A) to an output message (element of a finite set B) perhaps with *noise* so that there is a probability that $a \in A$ is mapped to various different $b \in B$.

$$p(b|a) = \text{probability that } b \text{ is received given that } a \text{ was sent}$$

Define a (simple, undirected) graph structure on A by

$$(a_1, a_2) \text{ an edge when } p(b|a_1)p(b|a_2) > 0 \text{ for some } b.$$

This is the *confusability graph* of the channel.
If we want to communicate with *zero error* then we seek a maximal *independent set* in A: a maximal subset of A which cannot be confused.

Physics notation

I will follow physics notation, so inner products $(\cdot|\cdot)$ are linear on the right.

- Use bra-ket notation: $|\psi\rangle$ is a vector in a Hilbert space H, and $\langle\psi|$ is a member of the dual space, identified with the conjugate $\overline{H}$.
- Then $\langle\psi|\phi\rangle = (\psi|\phi)$ the inner-product...
- and $|\phi\rangle\langle\psi|$ is the rank-one operator $H \to H;\ \alpha \mapsto (\psi|\alpha)\phi$.

Quantum Mechanics

Definition

A *state* is a unit vector $|\psi\rangle$ in a (finite dim) Hilbert space H.

Multiplying a state by a unit modulus complex number doesn't change the physics. One way to deal with this is to identify a state with the rank-one projection $|\psi\rangle\langle\psi|$.

Definition

A *density* is a positive, trace one operator $\rho \in \mathcal{B}(H)$.

- So a rank-one density is a state; we call a general density a *mixed* state.
- Mathematically, using trace-duality, a density is nothing but a (normal) state on the C^*-algebra $\mathcal{B}(H)$.

Quantum channels

Definition

A *(quantum) channel* is a trace-preserving, completely positive (CPTP) map $\mathcal{B}(H_A) \to \mathcal{B}(H_B)$.

- positive and trace-preserving so it maps densities to densities;
- completely positive so you can tensor with another system and still have positivity.

Theorem (Stinespring)

A linear map $\theta : A \to \mathcal{B}(H)$, from a C^*-algebra A, is completely positive if and only if it admits a dilation of the form

$$\theta(a) = V^*\pi(a)V \qquad (a \in A)$$

for $\pi : A \to \mathcal{B}(K)$ a $*$-homomorphism, and $V : H \to K$ a bounded linear map.

Stinespring and Kraus

Any CP map $\mathcal{E} : \mathcal{B}(H_A) \to \mathcal{B}(H_B)$ has the form

$$\mathcal{E}(x) = V^* \pi(x) V \qquad (x \in \mathcal{B}(H_A)),$$

where $V : H_B \to K$, and $\pi : \mathcal{B}(H_A) \to \mathcal{B}(K)$ is a $*$-representation.

- Any such π is of the form $\pi(x) = x \otimes 1$ where $K \cong H_A \otimes K'$.
- Take an o.n. basis (e_i) for K' so $V(\xi) = \sum_i K_i^*(\xi) \otimes e_i$ for some operators $K_i : H_A \to H_B$.

We arrive at the *Kraus form*:

$$\mathcal{E}(x) = \sum_i K_i x K_i^* \qquad (x \in \mathcal{B}(H_A)).$$

Trace-preserving if and only if $\sum_i K_i^* K_i = 1$.

Quantum zero-error

We turn $\mathcal{B}(H)$ into a Hilbert space using the trace: $(T|S) = \mathrm{tr}(T^*S)$.
A sensible notion of when densities ρ, σ are *distinguishable* is when they are orthogonal.
Let $\mathcal{E}(x) = \sum_i K_i x K_i^*$ be a quantum channel. We wish to consider when $\mathcal{E}(\rho) \perp \mathcal{E}(\sigma)$. As $\mathcal{E}$ is positive, this is equivalent to

$$\mathcal{E}(|\psi\rangle\langle\psi|) \perp \mathcal{E}(|\phi\rangle\langle\phi|) \qquad (\psi \in \mathrm{Image}\,\rho, \phi \in \mathrm{Image}\,\sigma).$$

Equivalently

$$\begin{aligned} 0 = \mathrm{tr}\left(\mathcal{E}(|\psi\rangle\langle\psi|)\mathcal{E}(|\phi\rangle\langle\phi|)\right) &= \sum_{i,j} \mathrm{tr}\left(K_i|\psi\rangle\langle\psi|K_i^* K_j|\phi\rangle\langle\phi|K_j^*\right) \\ &= \sum_{i,j} |\langle\psi|K_i^* K_j|\phi\rangle|^2 \end{aligned}$$

which is equivalent to $\langle\psi|K_i^* K_j|\phi\rangle = 0$ for each i, j.

To operator systems

So ψ, ϕ are distinguishable after applying $\mathcal{E}$ when

$$\langle\psi|T|\phi\rangle = 0 \quad \text{for each} \quad T \in \operatorname{lin}\{K_i^* K_j\}.$$

Set $\mathcal{S} = \operatorname{lin}\{K_i^* K_j\}$ which has the properties:

- $\mathcal{S}$ is a linear subspace;
- $T \in \mathcal{S}$ if and only if $T^* \in \mathcal{S}$;
- $1 \in \mathcal{S}$ (as $\sum_i K_i^* K_i = 1$ as $\mathcal{E}$ is CPTP).

That is, $\mathcal{S}$ is an *operator system*, which depends only on $\mathcal{E}$ and not the choice of (K_i).

Theorem (Duan)

For any operator system $\mathcal{S} \subseteq \mathcal{B}(H_A)$ there is some quantum channel $\mathcal{E} : \mathcal{B}(H_A) \to \mathcal{B}(H_B)$ giving rise to $\mathcal{S}$.

In the classical case

Given a classical channel from A to B with probabilities $p(b|a)$, we encode this as follows:

- Let $H_A = \ell^2(A)$ with o.n. basis $\{|a\rangle : a \in A\}$; and the same for B.
- Define Kraus operators

$$K_{ab} = p(b|a)^{1/2}|b\rangle\langle a| : H_A \to H_B.$$

Then $\mathcal{E} : \rho \mapsto \sum_{a,b} K_{ab}\rho K_{ab}^*$ sends a pure state $|c\rangle\langle c|$ to

$$\sum_{ab} K_{ab}|c\rangle\langle c|K_{ab}^* = \sum_{ab} p(b|a)|b\rangle\langle a|c\rangle\langle c|a\rangle\langle b| = \sum_{b} p(b|c)|b\rangle\langle b|.$$

That is, the combination of pure states which can be received, given that message c was sent.

The associated operator system

The Kraus operators are

$$K_{ab} = p(b|a)^{1/2}|b\rangle\langle a| : H_A \to H_B.$$

Hence

$$\begin{aligned}\mathcal{S} &= \mathrm{lin}\{K_{ab}^* K_{cd}\} = \mathrm{lin}\{p(b|a)^{1/2} p(d|c)^{1/2} |a\rangle\langle b|d\rangle\langle c|\} \\ &= \mathrm{lin}\{p(b|a)^{1/2} p(b|c)^{1/2} |a\rangle\langle c|\} \\ &= \mathrm{lin}\{|a\rangle\langle c| : a \sim c\},\end{aligned}$$

where $a \sim c$ exactly when $p(b|a)p(b|c) > 0$ for some b.
Thus $\mathcal{S}$ is directly linked to the confusability graph of the channel: it is the span of the matrix units e_{ac} for each edge (a, c) in the graph.
(Notice here our "graphs" are finite, simple, but we allow (single, unoriented) loops at vertices.)

Quantum relations

Simultaneously, and motivated more by "noncommutative geometry":

Definition (Weaver)

Let $M \subseteq \mathcal{B}(H)$ be a von Neumann algebra. A *quantum relation* on M is a weak*-closed subspace $S \subseteq \mathcal{B}(H)$ with $M'SM' \subseteq S$. We say that the relation is:

1. *reflexive* if $M' \subseteq S$;
2. *symmetric* if $S^* = S$ where $S^* = \{x^* : x \in S\}$;
3. *transitive* if $S^2 \subseteq S$ where $S^2 = \overline{\operatorname{lin}}^{w^*}\{xy : x, y \in S\}$.

When $M = \ell^\infty(X) \subseteq \mathcal{B}(\ell^2(X))$ there is a bijection between the usual meaning of "relation" on X and quantum relations on M, given by

$$x \sim y \text{ when } e_{x,y} \in S, \qquad S = \overline{\operatorname{lin}}^{w^*}\{e_{x,y} : x \sim y\}.$$

Operator bimodules

The condition that $M'SM' \subseteq S$ means that S is an *operator bimodule* over M'.
(Not to be confused with Hilbert C^*-modules!)

- We assume $M \subseteq \mathcal{B}(H)$ and $S \subseteq \mathcal{B}(H)$.
- If $M \subseteq \mathcal{B}(K)$ as well, we of course want a $T \subseteq \mathcal{B}(K)$ corresponding to S.
- This can be found by using the structure theory for normal $*$-homomorphisms $\theta : M \to \mathcal{B}(K)$. Essentially θ is a dilation followed by a cut-down in the commutant.
- That S is a bimodule over M' is needed to get this correspondence with T.

So this notion is really independent of the choice of embedding $M \subseteq \mathcal{B}(H)$. [Weaver] gives an intrinsic notion just using M.

Quantum graphs

As a graph on a (finite) vertex set V is simply a relation, and as:

- undirected graphs correspond to symmetric relations;
- a reflexive relation corresponds to having a "loop" at every vertex.

Definition (Weaver)

A *quantum graph* on a von Neumann algebra $M \subseteq \mathcal{B}(H)$ is a reflexive, symmetric quantum relation. That is, a unital, self-adjoint, weak*-closed subspace $S \subseteq \mathcal{B}(H)$, which is an M'-bimodule ($M'SM' \subseteq S$).

If $M = \mathcal{B}(H)$ with H finite-dimensional, then as $M' = \mathbb{C}$, a quantum graph is just an operator system: that is, exactly what we had before! [Duan, Severini, Winter; Stahlke]

Adjacency matrices

Given a graph $G = (V, E)$ consider the $\{0, 1\}$-valued matrix A with

$$A_{i,j} = \begin{cases} 1 & : (i, j) \in E, \\ 0 & : \text{otherwise}, \end{cases}$$

the *adjacency matrix* of G.

- A is idempotent for the *Schur product*;
- G is undirected if and only if A is self-adjoint;
- A has 1s down the diagonal when G has a loop at every vertex.

We can think of A as an operator on $\ell^2(V)$. This is the GNS space for the C^*-algebra $\ell^\infty(V)$ for the state induced by the uniform measure.

General C^*-algebras

Let B be a finite-dimensional C^*-algebra, and let φ be a faithful state on B, with GNS space $L^2(B)$. Thus B bijects with $L^2(B)$ as a vector space, and so we get:

- The multiplication on B induces a map $m : L^2(B) \otimes L^2(B) \to L^2(B)$;
- the Hilbert space structure now allows us to define $m^* : L^2(B) \to L^2(B) \otimes L^2(B)$.
- The unit in B induces a map $\eta : \mathbb{C} \to L^2(B)$;
- similarly we obtain $\eta^* : L^2(B) \to \mathbb{C}$, which is just φ.

We get an analogue of the Schur product:

$$x \bullet y = m(x \otimes y)m^* \qquad (x, y \in \mathcal{B}(L^2(B))).$$

Quantum adjacency matrix

Definition (Many authors)

A *quantum adjacency matrix* is a self-adjoint $A \in \mathcal{B}(L^2(B))$ with:

- $m(A \otimes A)m^* = A$ (so Schur product idempotent);
- $(1 \otimes \eta^* m)(1 \otimes A \otimes 1)(m^* \eta \otimes 1) = A$;
- $m(A \otimes 1)m^* = \mathrm{id}$ (a "loop at every vertex");

The middle axiom is a little mysterious: it roughly corresponds to "undirected".

Subspaces to projections

Fix a finite-dimensional C^*-algebra (von Neumann algebra) M. A "quantum graph" is either:

- A subspace of $\mathcal{B}(H)$ (where $M \subseteq \mathcal{B}(H)$) with some properties; or
- An operator on $L^2(M)$ with some properties.

How do we move between these?

$S \subseteq \mathcal{B}(H)$ is a bimodule over M'. As H is finite-dimensional, $\mathcal{B}(H)$ is a Hilbert space for

$$(x|y) = \mathsf{tr}(x^* y).$$

Then $M \otimes M^{\mathsf{op}}$ is represented on $\mathcal{B}(H)$ via

$$\pi : M \otimes M^{\mathsf{op}} \to \mathcal{B}(\mathcal{B}(H)); \quad \pi(x \otimes y) : T \mapsto xTy.$$

- The commutant of $\pi(M \otimes M^{\mathsf{op}})$ is $M' \otimes (M')^{\mathsf{op}}$.
- An M'-bimodule of $\mathcal{B}(H)$ corresponds to an $M' \otimes (M')^{\mathsf{op}}$-invariant subspace of the Hilbert space $\mathcal{B}(H)$;
- which corresponds to a *projection* in $M \otimes M^{\mathsf{op}}$.

Operators to algebras

So how can we relate:

- Operators $A \in \mathcal{B}(L^2(M))$;
- Projections in $M \otimes M^{\mathrm{op}}$?

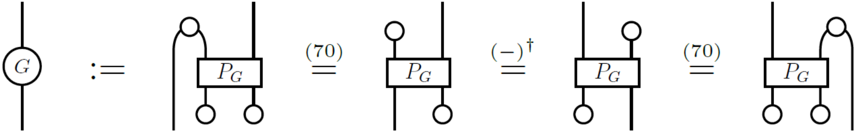

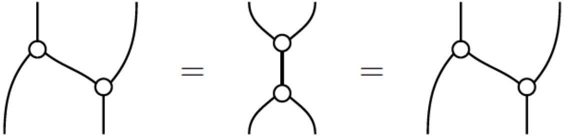

[Musto, Reutter, Verdon]

Operators to algebras 2

Recall the GNS construction for a *tracial* state ψ on M:

$$\Lambda : M \to L^2(M); \quad (\Lambda(x)|\Lambda(y)) = \psi(x^*y).$$

As $L^2(M)$ is finite-dimensional, every operator on $L^2(M)$ is a linear combination of rank-one operators. So we may define a bijection

$$\Psi : \mathcal{B}(L^2(M)) \to M \otimes M^{\mathsf{op}}; \quad |\Lambda(b)\rangle\langle\Lambda(a)| \mapsto b \otimes a^*,$$

and extend by linearity!

Operators to algebras 3

$$\Psi : \mathcal{B}(L^2(M)) \to M \otimes M^{\mathrm{op}}; \quad |\Lambda(b)\rangle\langle\Lambda(a)| \mapsto b \otimes a^*,$$

- Ψ is a homomorphism for the "Schur product" $A_1 \bullet A_2 = m(A_1 \otimes A_2)m^*$;
- $A \mapsto (1 \otimes \eta^* m)(1 \otimes A \otimes 1)(m^*\eta \otimes 1)$ corresponds to the anti-homomorphism $\sigma : a \otimes b \mapsto b \otimes a$ on $M \otimes M^{\mathrm{op}}$;
- $A \mapsto A^*$ corresponds to $e \mapsto \sigma(e)^*$.

Conclude: A quantum adjacency matrix corresponds to an idempotent $e \in M \otimes M^{\mathrm{op}}$ with $\sigma(e) = e$ and $\sigma(e)^* = e$.
That is, a projection e with $\sigma(e) = e$.
But: There is no clean one-to-one correspondence between the axioms.

KMS States

Any faithful state ψ is KMS: there is an automorphism σ' of M with

$$\psi(ab) = \psi(b\sigma'(a)) \qquad (a, b \in M).$$

Indeed, there is $Q \in M$ positive and invertible with

$$\psi(a) = \mathrm{tr}(Qa) \qquad \sigma'(a) = QaQ^{-1}.$$

Theorem (D.)

Twisting our bijection Ψ using σ' allows us to establish a bijection between:

- *Quantum adjacency operators $A \in \mathcal{B}(L^2(M))$;*
- *projections $e \in M \otimes M^{\mathrm{op}}$ with $e = \sigma(e)$ and $(\sigma' \otimes \sigma')(e) = e$;*
- *self-adjoint M'-bimodules $S \subseteq \mathcal{B}(H)$ with $QSQ^{-1} = S$.*

So this is *more restrictive* than the tracial case.

Invariance under the modular automorphism

Why do we end up with $(\sigma' \otimes \sigma')(e) = e$?

- The "middle axiom" is a bit mysterious: we already assume that A is self-adjoint, and shouldn't this alone correspond to the graph being undirected? (Both conditions together is a bit strong.)
- [Matsuda] looked at a different condition, that of A being "real" which means that $A : L^2(B) \to L^2(B)$, thought of as a map $B \to B$, is $*$-preserving.
- [D.] showed that replacing "self-adjoint and axiom (2)" with "real" gives a simple bijection with projections.
- [Wasilewski] has recently shown that looking at "KMS inner-products" not "GNS inner-products" is a nice framework to view this in.

(However, we are stuck with the existing literature.)

Towards homomorphisms

Let B_1, B_2 be finite-dimensional C^*-algebras (maybe just $B_i = \mathcal{B}(H_i)$), and let $\theta : B_1 \to B_2$ be a CPTP map with Kraus form

$$\theta(x) = \sum_{i=1}^{n} b_i x b_i^*.$$

For $i = 1, 2$ let $B_i \subseteq \mathcal{B}(H_i)$ and let $S_i \subseteq \mathcal{B}(H_i)$ be a quantum graph/relation over B_i.

Definition (Weaver)

The *pushforward* of S_1 is

$$\overrightarrow{S_1} = B_2'\text{-bimodule } \{b_i x b_j^* : x \in S_1, 1 \leqslant i, j \leqslant n\}.$$

The *pullback* of S_2 is

$$\overleftarrow{S_2} = B_1'\text{-bimodule } \{b_i^* y b_j : y \in S_2, 1 \leqslant i, j \leqslant n\}.$$

Motivation

Let $G = (V_G, E_G)$, $H = (V_H, E_H)$ be graphs.

- For $f : V_G \to V_H$ a map, define

$$\theta : C(V_H) \to C(V_G),$$
$$\theta(a)(u) = a(f(u)) \qquad (u \in V_G, a \in C(V_H)).$$

- So θ is a $*$-homomorphism, in particular, a UCP map.

We find a Kraus form for θ. Given $x \in V_H$ there might be many (or none!) $u \in V_G$ with $f(u) = x$; enumerate the u in some way. Define $b_i : \ell^2(V_G) \to \ell^2(V_H)$ by

$$b_i(\delta_u) = \delta_x \qquad \text{if } u \text{ is the } i\text{th vertex with } f(u) = x.$$

Then indeed

$$\sum_i b_i^* a b_i(\delta_u) = a(f(u)) = \theta(a)(\delta_u) \qquad (a \in C(V_H), u \in V_G).$$

CPTP maps

These (b_i) satisfy the pleasing fact that

$$\sum_i b_i e_u b_i^* = e_{f(u)} \qquad (u \in V_G),$$

where $e_u \in \ell^\infty(V_G)$ is the minimal projection. So we also obtain a TPCP map $\widehat{\theta} : C(V_G) \to C(V_H)$.

The operator system associated to G is

$$S_G = \operatorname{lin}\{e_{u,v} : (u,v) \in E_G\} \subseteq \mathbb{M}_{V_G}.$$

Then, using $\widehat{\theta}$,

$$\overrightarrow{S_G} = \operatorname{lin}\{e_{f(u),f(v)} : (u,v) \in E_G\}.$$

Similarly, given S_H, we find that

$$\overleftarrow{S_H} = \operatorname{lin}\{e_{u,v} : (f(u),f(v)) \in E_H\}.$$

Homomorphisms

$$\overrightarrow{S_G} = \mathrm{lin}\{e_{f(u),f(v)} : (u,v) \in E_G\}.$$

So $\overrightarrow{S_G} \subseteq S_H$ means exactly that

$$(u,v) \in E_G \quad \implies \quad (f(u), f(v)) \in E_H.$$

That is, $f : V_G \to V_H$ induces a graph homomorphism.

- So we've captured the concept of a graph homomorphism using $\overrightarrow{S_G}$.
- For general quantum graphs, and general TPCP maps, Stahlke takes this as the definition of a *homomorphism*.
- Weaver calls these *CP morphisms*; tentatively suggests we should start with a $*$-homomorphism if we want a "homomorphism".

Pullbacks

[Time?] [We "reverse the arrows" and use UCP maps not TPCP maps.]
Let $\theta : M \to N$ be a normal CP map between von Neumann algebras $M \subseteq \mathcal{B}(H_M)$ and $N \subseteq \mathcal{B}(H_N)$. The Stinespring dilation takes a special form:

- there is a Hilbert space K and $U : H_N \to H_M \otimes K$;
- $\theta(x) = U^*(x \otimes 1)U$ for $x \in M \subseteq \mathcal{B}(H_M)$;
- there is a normal $*$-homomorphism $\rho : N' \to H_M \otimes K$ with $Ux' = \rho(x')U$ for $x' \in N'$.

Proposition (D.)

The pullback satisfies

$$\overleftarrow{S} = \textit{weak}^*\textit{-closure}\{U^*xU : x \in S\overline{\otimes}\mathcal{B}(K)\},$$

independent of choice of U. In particular, this is already an N'-bimodule.

Duality

Let B_1, B_2 be finite-dimensional with faithful traces φ_i. Given a UCP map $\theta : B_2 \to B_1$ there is a TPCP map $\hat{\theta} : B_1 \to B_2$ satisfying/defined by

$$\varphi_1(a\theta(b)) = \varphi_2(\hat{\theta}(a)b) \qquad (a \in B_1, b \in B_2).$$

(“Accardi–Cecchini adjoint”.)

Proposition (D.)

Let φ_i be the “Markov Traces”, and given θ form $\hat{\theta}$. Then a pushforward of a quantum relation using θ is the same as the pullback using $\hat{\theta}$.

(We saw this for our maps on ℓ^∞ and ℓ^1. The general case is more complicated, but follows roughly the same idea.)

Homomorphisms

Recall that $\theta : M \to N$ is a *homomorphism* / *CP-morphism* $S_1 \to S_2$ when $\overrightarrow{S_2} \subseteq S_1$.

Theorem (Stahlke)

Let $\theta : C(V_H) \to C(V_G)$ be a UCP map giving a homomorphism G to H (that is, with $\overrightarrow{S_G} \subseteq S_H$). Then there is some map $f : V_G \to V_H$ which is a (classical) graph homomorphism.

- In general θ need not be directly related to f.
- However, often we just care about the *existence* of a homomorphism.
- E.g. a k-colouring of G corresponds to some homomorphism $G \to K_k$, the complete graph. (This requires our graphs not to have loops!)

Automorphisms

An *automorphism* of a graph $G = (V, E)$ is a bijection $\theta : V \to V$ which satisfies that $(i, j) \in E \Leftrightarrow (\theta(i), \theta(j)) \in E$.
Set $V = \{1, \cdots, n\}$ for ease, so the adjacency matrix A is in $\mathbb{M}_n$.

Lemma

Let $P_\theta \in \mathbb{M}_n$ be permutation matrix associated with a bijection θ. Then θ is an automorphism of G if and only if $P_\theta A = AP_\theta$.

Proof.

$P_\theta A = AP_\theta$ is equivalent to $(\theta^{-1}(i), j) \in E \Leftrightarrow (i, \theta(j)) \in E$, which in turn is equivalent to $(i, j) \in E \Leftrightarrow (\theta(i), \theta(j)) \in E$. □

Non-commutative topology

I am under obligation to provide this table:

Spaces	Algebras
Locally compact Hausdorff space	Commutative C*-algebra
Compact	Unital
(Proper) continuous map	$*$-Homomorphism
Cartesian Product	Tensor product

Remember that this relationship is *contravariant*.

How might we deal with (Compact) *groups*?

As the product $G \times G \to G$ and the inverse $G \to G$ and the identity $* \to G$ are continuous maps, we could specify a commutative C^*-algebra A, and $*$-homomorphisms

$$A \to A \otimes A, \qquad A \to A, \qquad A \to \mathbb{C},$$

satisfying appropriate axioms.

What are groups?

Definition

A group is a set G with an associative product $G \times G \to G$ such that:

- There is $e \in G$ with $eg = ge = g$ for each $g \in G$;
- For each $g \in G$ there are $h, k \in G$ with $gh = kg = e$.

So really the identity and inverse are "properties" of the semigroup G, not "structure".

It turns out that we get a (much) more interesting theory if we similarly focus on the product, and think about an extra property.

Compact Quantum groups

Definition (Woronowicz)

A *compact quantum group* is a unital C^*-algebra A together with a unital $*$-homomorphism, the *coproduct*, $\Delta : A \to A \otimes A$, which is coassociative, $(\Delta \otimes \mathrm{id})\Delta = (\mathrm{id} \otimes \Delta)\Delta$, and such that:

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

both have dense linear span in $A \otimes A$.

Theorem

Let (A, Δ) be a compact quantum group with A commutative. There is a compact group G with $A = C(G)$ and $\Delta : C(G) \to C(G) \otimes C(G) = C(G \times G)$ given by

$$\Delta(f)(s,t) = f(st) \qquad (f \in C(G), s, t \in G).$$

Quantum group (co)actions

An (right) action of a group G on a space/set X is a map

$$X \times G \to X.$$

So we get a $*$-homomorphism

$$\alpha : C(X) \to C(X) \otimes C(G),$$

- $(\mathrm{id} \otimes \Delta)\alpha = (\alpha \otimes \mathrm{id})\alpha$ corresponds to $x \cdot st = (x \cdot s) \cdot t$;
- $\mathrm{lin}\{\alpha(b)(1 \otimes a) : a \in C(G), b \in C(X)\}$ is dense in $C(X) \otimes C(G)$ corresponds to $x \cdot e = x$.

Definition (Podleś)

A (right) coaction of a compact quantum group (A, Δ) on a C^*-algebra B is a unital $*$-homomorphism $\alpha : B \to B \otimes A$ with these two conditions.

Coactions on ℓ^∞_n

Fix a compact quantum group (A,Δ).

- The algebra ℓ^∞_n is spanned by projections $(e_i)_{i=1}^n$.
- So $\alpha : \ell^\infty_n \to \ell^\infty_n \otimes A$ is determined by (u_{ij}) in A with

$$\alpha(e_i) = \sum_{j=1}^n e_j \otimes u_{ji}.$$

- α is a $*$-homomorphism $\Leftrightarrow$ each u_{ji} a projection and $u_{ji}u_{jk} = \delta_{ik}u_{ji}$;
- α is unital $\Leftrightarrow$ $\sum_i u_{ji} = 1$;
- α satisfies the coaction equation $\Leftrightarrow$ $\Delta(u_{ji}) = \sum_k u_{jk} \otimes u_{ki}$;
- α satisfies the Podleś density condition $\Leftrightarrow$ $\sum_i u_{ji} = 1$.
- General Theory $\implies$ $\sum_j u_{ji} = 1$.
- So $u = (u_{ij})$ is a matrix of projections, each row and column sums to 1. A *quantum permutation matrix* or *magic unitary*.

Quantum symmetry group of the space of n points

For $\ell^\infty_n = C(\{1, 2, \cdots, n\})$,

$$\alpha(e_i) = \sum_{j=1}^{n} e_j \otimes u_{ji},$$

with $u = (u_{ij})$ a magic unitary.

Theorem (Wang)

Let S_n^+ be the "universal" C^*-algebra generated by a magic unitary. Then S_n^+ is the "largest" compact quantum group which acts on $\mathbb{C}^n$ is a "non-degenerate" way.

We think of S_n^+ as the "quantum symmetry group" of $\{1, 2, \cdots, n\}$.

(Co)actions on graphs

Recall that a permutation θ gives an automorphism of G when

$$P_\theta A_G = A_G P_\theta.$$

Here A_G is the adjacency matrix of G, which we can think of as also a linear map $\ell^\infty_n \to \ell^\infty_n$.

So $\mathrm{Aut}(G)$ acts in a way which preserves A_G:

$$\alpha : \ell^\infty_n \to \ell^\infty_n \otimes C(\mathrm{Aut}(G)); \quad \alpha A_G = (A_G \otimes \mathrm{id})\alpha.$$

Definition (Banica)

The *quantum automorphism group* of G is the maximal compact quantum group $\mathrm{QAut}(G)$ with a coaction satisfying

$$\alpha : \ell^\infty_n \to \ell^\infty_n \otimes \mathrm{QAut}(G); \quad \alpha A_G = (A_G \otimes \mathrm{id})\alpha.$$

Equivalently, the underlying magic unitary $U = (u_{ij})$ has to commute with the adjacency matrix A_G. This allows us to construct $\mathrm{QAut}(G)$ as a quotient of S_n^+.

Examples

We say that a graph *has quantum symmetry* if $\mathrm{Aut}(G) \neq \mathrm{QAut}(G)$.

- By now, we have many examples.
- For example, the Petersen graph has no quantum symmetry [Schmidt].

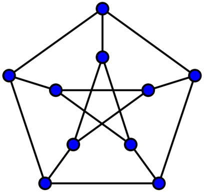

[CC-BY-SA, Leshabirukov, Wikipedia]

- [Roberson, Schmidt] have constructed G with $\mathrm{Aut}(G) \neq \mathrm{QAut}(G)$ and yet $\mathrm{QAut}(G)$ is finite.
- [Dobben de Bruyn, Roberson, Schmidt] have constructed G with $\mathrm{Aut}(G)$ trivial and $\mathrm{QAut}(G)$ non-trivial.

(Co)actions on operator bimodules

What is an "automorphism" of $\mathcal{S} \subseteq \mathcal{B}(\ell^2(V))$?

- Start with a bijection $\theta : V \to V$, hence giving $P_\theta \in \mathcal{B}(\ell^2(V))$.
- Then get an action on $\mathcal{B}(\ell^2(V))$ as $\hat{\theta} : x \mapsto P_\theta x P_\theta^*$ (as $P_\theta^* = P_\theta^{-1}$).
- When is $\mathcal{S}$ left invariant: $P_\theta \mathcal{S} P_\theta^* = \mathcal{S}$?

Notice that

$$P_\theta e_{ij} P_\theta^* = e_{\theta(i),\theta(j)}$$

- So if G is a graph, and $\mathcal{S} = \mathcal{S}_G$ the canonical operator system;
- then $P_\theta \mathcal{S}_G P_\theta^* = \mathcal{S}$ exactly when $(i,j) \in E \Leftrightarrow (\theta(i), \theta(j)) \in E$;
- that is, θ is an automorphism of G.

How to phrase this in terms of coactions?

Unitary implementations

Given a coaction $\alpha : \ell^\infty(V) \to \ell^\infty(V) \otimes A$ of (A, Δ) on $\ell^\infty(V)$, we saw before that α gives rise to a magic unitary $u = (u_{ij})_{i,j \in V}$,

$$\alpha(e_i) = \sum_{j \in V} e_j \otimes u_{ji} \qquad (i \in V).$$

Lemma

Let $\ell^\infty(V) \subseteq \mathcal{B}(\ell^2(V))$. Then

$$\alpha(x) = u(x \otimes 1)u^* \qquad (x \in \ell^\infty(V)).$$

Coactions on operator bimodules

$$\alpha(x) = u(x \otimes 1)u^* \qquad (x \in \ell^\infty(V) \subseteq \mathcal{B}(\ell^2(V))).$$

It hence make sense...

Definition

α is a coaction on $\mathcal{S} \subseteq \mathcal{B}(\ell^2(V))$ exactly when $u(x \otimes 1)u^* \in \mathcal{S} \otimes A$ for each $x \in \mathcal{S}$.

One can check (non-trivially) that we then get the following.

Theorem (Eifler)

If a graph G is associated to the $\ell^\infty(V)$-operator bimodule $\mathcal{S}$, then a coaction of (A, Δ) on $\ell^\infty(V)$ gives a coaction on G if and only if it gives a coaction on $\mathcal{S}$.

Coactions on C^*-algebras

A coaction of (A, Δ) on B is, as before,

$$\alpha : B \to B \otimes A; \quad (\mathrm{id} \otimes \Delta)\alpha = (\alpha \otimes \mathrm{id})\alpha,$$

and satisfying the Podleś density condition.
(So simply replace ℓ^∞_n by an arbitrary B.)

Theorem (Wang)

There is no maximal compact quantum group coacting on B.
If ψ is a faithful state on B, there is a maximal compact quantum group coacting on B and preserving ψ, meaning:
$(\psi \otimes \mathrm{id})\alpha(x) = \psi(x)1$ *for* $x \in B$. *Write* $\mathrm{QAut}(B, \psi)$ *for this.*

Coactions on quantum adjacency matrices

There is now a clear definition:

Definition (Brannan et al.)

Let A_G be a quantum adjacency matrix on (B, ψ). We say that (A, Δ) coacts on A_G when $\alpha : B \to B \otimes A$ is a coaction, which preserves ψ, and with $(A_G \otimes \mathrm{id})\alpha = \alpha A_G$.

- Here we regard A_G as a linear map on B.
- That α preserves ψ allows us to define a unitary $U \in \mathcal{B}(L^2(B)) \otimes A$ which implements α, as $\alpha(x) = U(x \otimes 1)U^*$. Indeed, one way to prove Wang's theorem is to start with such a U and impose certain conditions on it (compare Compact Quantum Matrix Groups).
- Then, equivalently, we require that U and $A_G \otimes 1$ commute.

Coactions on operator bimodules

A coaction α which preserves ψ gives a unitary U (which is a *corepresentation*) and it is then easy to see that

$$\alpha_U : \mathcal{B}(L^2(B)) \to \mathcal{B}(L^2(B)) \otimes A; \quad x \mapsto U(x \otimes 1)U^*$$

is a coaction (which extends α).

Might this leave $\mathcal{S} \subseteq \mathcal{B}(L^2(B))$ invariant if and only if U commutes with A_G?

- No, as the "trivial quantum graph" is $\mathcal{S} = B'$, which should always be invariant, but α_U leaves B invariant, not B'.
- Instead, we can use the *modular conjugation* J and *antipode* to form a "commutant" coaction α'_U; or equivalently, look at α_U but work with

$$\mathcal{S}' := \{JTJ : T \in \mathcal{S}\}.$$

Theorem (D.)

α leaves A_G invariant if and only if α_U leaves $\mathcal{S}'$ invariant.

Further

- For a “homomorphism” do we really want our UCP map to be a $*$-homomorphism?
- It turns out some ideas from “quantum games” [Brannan et al.] naturally separate out the conditions on a “CP-morphism”, and these actually force a $*$-homomorphism.
- Also related to trying to “ignore loops”.

Possible future things:

- What are the “correct axioms”? E.g. self-adjointness or “reality”? Applications which might motivate this?
- Is there some sort of infinite-dimensional theory?