THE SYMBIOSIS OF C*- AND W*-ALGEBRAS

NATHANIAL P. BROWN

Abstract. These days it is common for young operator algebraists to know a lot about C*-algebras, or a lot about von Neumann algebras – but not both. Though a natural consequence of the breadth and depth of each subject, this is unfortunate as the interplay between the two theories has deep historical roots and has led to many beautiful results. We review some of these connections, in the context of amenability, with the hope of convincing (younger) readers that tribalism impedes progress.

1. Introduction

I was raised a hardcore C*-algebraist. My thesis focused on C*-dynamical systems, and never once required a weak topology. As a fresh PhD my knowledge of von Neumann algebras was superficial, at best. I didn't really like von Neumann algebras, didn't understand them, and certainly didn't need them to prove theorems. Conversations with new W*-PhDs make it clear that this goes both ways; they often know little about C*-algebras, and care less.

Today, I still know relatively little about von Neumann algebras, but I have grown to love them. And they are an indispensable tool in my C*-research. Conversely, recent work of Narutaka Ozawa (some in collaboration with Sorin Popa) has shown that C*-techniques can have deep applications to the structure theory of certain von Neumann algebras. In other words, there are very good reasons for C*-algebraists and von Neumann algebraists to learn something about each other's craft. In the "old" days (say 30 or more years ago), the previous sentence would have been silly (indeed, some "old" timers may still find it silly) as the field of operator algebras was small enough for students to become well acquainted with most of it. That's no longer the case. Hence, I hope these notes will help my generation, and those that follow, to see the delightfully intertwined theories of C*- and W*-algebras as an indivisible unit.

I do not intend to write an encyclopedia of C*- and W*-interactions. Amenability (for groups, actions and operator algebras) is a perfect context for illustrating some of the most important interactions, so these notes are organized around that theme.[1]

2. C*-algebras vs. W*-algebras

This expository section will be written later – it is irrelevant to the math contained in these notes. However, we'll need the following basic fact.

Theorem 2.1. *Let X be a Banach space, M be a von Neumann algebra and $T_\lambda\colon X \to M$ be a bounded net of linear maps. Then $\{T_\lambda\}_{\lambda\in\Lambda}$ has a cluster point in the point-ultraweak topology.*

3. Five classical theorems

Here are some general tools that facilitate the passage between norm-closed and weakly-closed algebras. The first result, one of the oldest in the subject, is still used daily.

[1]Large portions of these notes are cut-n-pasted directly from [2]. This preliminary manuscript still requires much polishing, so please don't distribute beyond the participants of the summer school in Santander.

Theorem 3.1 (Bicommutant Theorem). *Let $A \subset \mathbb{B}(\mathcal{H})$ be a nondegenerate C*-algebra. Then the weak-operator-topology closure of A is equal to the double commutant A''.*

Theorem 3.2 (Kaplansky's Density Theorem). *Let $A \subset \mathbb{B}(\mathcal{H})$ be a nondegenerate C*-algebra. Then the unit ball of A is weakly dense in the unit ball of A''.*

Theorem 3.3 (Up-Down Theorem). *Let $A \subset \mathbb{B}(\mathcal{H})$ be a nondegenerate C*-algebra on a separable Hilbert space $\mathcal{H}$. For each self-adjoint $x \in A''$, there exists a decreasing sequence of self-adjoints $x_n \geq x_{n+1} \geq \cdots$ in A'' such that*

(1) *$x_n \to x$ in the strong operator topology, and*
(2) *for each $n \in \mathbb{N}$, there exists an increasing sequence of self-adjoint $y_k^{(n)} \leq y_{k+1}^{(n)}$ in A, such that $y_k^{(n)} \to x_n$ (as $k \to \infty$) in the strong operator topology.*

Theorem 3.4 (Lusin's Theorem). *Let $A \subset \mathbb{B}(\mathcal{H})$ be a nondegenerate C*-algebra. For every finite set of vectors $\mathfrak{F} \subset \mathcal{H}$, $\varepsilon > 0$, projection $p_0 \in A''$ and self-adjoint $y \in A''$, there exist a self-adjoint $x \in A$ and a projection $p \in A''$ such that $p \leq p_0$, $\|p(h) - p_0(h)\| < \varepsilon$ for all $h \in \mathfrak{F}$, $\|x\| \leq \min\{2\|yp_0\|, \|y\|\} + \varepsilon$ and $xp = yp$.*

Theorem 3.5 (Double Dual Theorem). *The (Banach space) double dual A^{**} of a C*-algebra A is a von Neumann algebra. Moreover, the ultraweak topology on A^{**} (coming from its von Neumann algebra structure) agrees with the weak-$*$ topology (coming from A^*), and hence restricts to the weak topology on A (coming from A^*).*

From a C*-algebraist's point of view, the double dual theorem is probably the most important as it allows one to come back from the world of von Neumann algebras. That is, suppose one wants to prove a theorem about a C*-algebra, exploiting the fact that von Neumann algebras have far more structure (projections, traces, etc.) and vastly more powerful tools (e.g. Borel functional calculus, polar decompositions, trivial representation theory). Well, in any representation of the given C*-algebra one could take the weak closure and switch to von Neumann algebra techniques. But how to come back to the C*-algebra of interest? Answer: The Hahn-Banach Theorem, of course!! That is, rather than work in any old weak closure, one should work in the double dual von Neumann algebra, where the Hahn-Banach theorem implies that convex sets in $A \subset A^{**}$ have the same norm and weak closures.

That probably makes little sense, so here's an illustrative example (that we'll use later on). Let's show that if A^{**} is semidiscrete, then A is nuclear.

Semidiscreteness and Nuclearity. First we have to recall a few facts about an important class of morphisms. We say a linear map $\varphi\colon A \to B$ is *completely positive* (c.p.) if $\varphi_n\colon \mathbb{M}_n(A) \to \mathbb{M}_n(B)$, defined by

$$\varphi_n([a_{i,j}]) = [\varphi(a_{i,j})],$$

is positive (i.e., maps positive matrices to positive matrices) for every $n \in \mathbb{N}$. When the domain or range is a matrix algebra, complete positivity can be reformulated.

Proposition 3.6. *Let A be a* C**-algebra and $\{e_{i,j}\}$ be matrix units of $\mathbb{M}_n(\mathbb{C})$. A map $\varphi\colon \mathbb{M}_n(\mathbb{C}) \to A$ is c.p. if and only if $[\varphi(e_{i,j})]$ is positive in $\mathbb{M}_n(A)$.*

Given a linear map $\varphi\colon A \to \mathbb{M}_n(\mathbb{C})$, define a functional $\widehat{\varphi}$ on $\mathbb{M}_n(A)$ by

$$\widehat{\varphi}([a_{i,j}]) = \sum_{i,j=1}^{n} \varphi(a_{i,j})_{i,j},$$

where $\varphi(a_{i,j})_{i,j}$ means the $(i,j)^{\text{th}}$ entry of the matrix $\varphi(a_{i,j})$.

Proposition 3.7. *Let A be a unital* C*-*algebra. A map* $\varphi\colon A \to \mathbb{M}_n(\mathbb{C})$ *is c.p. if and only if* $\widehat{\varphi}$ *is positive on* $\mathbb{M}_n(A)$.

Definition 3.8. A (contractive c.p.) map $\theta\colon A \to B$ is *nuclear* if there exist contractive c.p. maps $\varphi_n\colon A \to \mathbb{M}_{k(n)}(\mathbb{C})$ and $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to B$ such that $\psi_n \circ \varphi_n(a) \to \theta(a)$ in norm, for all $a \in A$.

A C*-algebra A is *nuclear* if id: $A \to A$ is nuclear – i.e., if there exist contractive c.p. maps $\varphi_n\colon A \to \mathbb{M}_{k(n)}(\mathbb{C})$ and $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to A$ such that $\psi_n \circ \varphi_n(a) \to a$ in norm, for all $a \in A$.

A W*-algebra M is *semidiscrete* if there exist contractive c.p. maps $\varphi_n\colon M \to \mathbb{M}_{k(n)}(\mathbb{C})$, $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to M$ such that $\psi_n \circ \varphi_n(x) \to x$ ultraweakly, for all $x \in M$.

Proposition 3.9. *If A^{**} is semidiscrete, then A is nuclear.*

Proof. I'll sketch the argument, highlighting the use of the double dual theorem and neglecting some nontrivial details (see [2, Proposition 2.3.8]).

Let $\varphi_n\colon A^{**} \to \mathbb{M}_{k(n)}(\mathbb{C})$, $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to A^{**}$ be such that $\psi_n \circ \varphi_n(x) \to x$ ultraweakly, for all $x \in A^{**}$. An approximation argument, using Proposition 3.6 and the fact that $\mathbb{M}_k(A)$ is ultraweakly dense in $\mathbb{M}_k(A^{**})$ for all $k \in \mathbb{N}$, allows us to assume that $\psi_n(\mathbb{M}_{k(n)}(\mathbb{C})) \subset A$ for all n.

Here's the punchline: for each $a \in A$, since the ultraweak topology on A^{**} restricts to the weak topology on A, the Hahn-Banach theorem implies that a belongs to the *norm*-closed convex hull of $\{\psi_n(\varphi_n(a))\}$!! It's not too hard to see that one can replace a convex combination of maps into different matrix algebras with a single map into a single matrix algebra, and a standard direct-sum trick allows us to replace individual operators $a \in A$ with finite sets, thereby completing the proof. $\square$

4. Reduced Group C*-algebras

For a discrete group Γ we let $\lambda\colon \Gamma \to \mathbb{B}(\ell^2(\Gamma))$ denote the *left regular representation*: $\lambda_s(\delta_t) = \delta_{st}$ for all $s, t \in \Gamma$, where $\{\delta_t : t \in \Gamma\} \subset \ell^2(\Gamma)$ is the canonical orthonormal basis. There is also a *right regular representation* $\rho\colon \Gamma \to \mathbb{B}(\ell^2(\Gamma))$, defined by $\rho_s(\delta_t) = \delta_{ts^{-1}}$. Note that λ and ρ are unitarily equivalent; the intertwining unitary is defined by $U\delta_t = \delta_{t^{-1}}$.

We denote the *group ring* of Γ by $\mathbb{C}[\Gamma]$. By definition, it is the set of formal sums

$$\sum_{s \in \Gamma} a_s s,$$

where only finitely many of the scalar coefficients $a_s \in \mathbb{C}$ are nonzero, and multiplication is defined by

$$(\sum_{s \in \Gamma} a_s s)(\sum_{t \in \Gamma} a_t t) = \sum_{s,t \in \Gamma} a_s a_t st.$$

The group ring $\mathbb{C}[\Gamma]$ acquires an involution by declaring $(\sum_{s\in\Gamma} a_s s)^* = \sum_{s\in\Gamma} \overline{a_s} s^{-1}$. Note that the left regular representation can be extended to an injective $*$-homomorphism $\mathbb{C}[\Gamma] \to \mathbb{B}(\ell^2(\Gamma))$, which we also denote by λ. Evidently, there is a one-to-one correspondence between unitary representations of Γ and $*$-representations of $\mathbb{C}[\Gamma]$.

Both amenable and exact groups are defined in terms of their canonical actions on $\ell^\infty(\Gamma)$. For $f \in \ell^\infty(\Gamma)$ and $s \in \Gamma$ we let $s.f \in \ell^\infty(\Gamma)$ be the function $s.f(t) = f(s^{-1}t)$; simple calculations show that $f \mapsto s.f$ defines a group action of Γ on $\ell^\infty(\Gamma)$. An important fact is that this action is *spatially implemented by the left regular representation.* That is, if we regard $\ell^\infty(\Gamma) \subset \mathbb{B}(\ell^2(\Gamma))$ as multiplication operators (i.e., $f\delta_t = f(t)\delta_t$), then a calculation shows

$$\lambda_s f \lambda_s^* = s.f$$

for all $f \in \ell^\infty(\Gamma)$ and $s \in \Gamma$.

The *reduced* C*-algebra of Γ, denoted $C^*_\lambda(\Gamma)$,[2] is the completion of $\mathbb{C}[\Gamma]$ with respect to the norm

$$\|x\|_r = \|\lambda(x)\|_{\mathbb{B}(\ell^2(\Gamma))}.$$

Though isomorphic to $C^*_\lambda(\Gamma)$, it is sometimes useful to consider $C^*_\rho(\Gamma)$, which is just the closure of $\mathbb{C}[\Gamma]$ in the right regular representation.

Example 4.1. If $\Gamma = \mathbb{Z}$, then $C^*_\lambda(\Gamma) = C(\mathbb{T})$, the continuous functions on the circle. Indeed, the Fourier transform identifies $\ell^2(\mathbb{Z})$ with $L^2(\mathbb{T}, \text{Lebesgue})$ and one checks that this unitary takes $C^*_\lambda(\mathbb{Z})$ to (continuous) multiplication operators. More generally, for every abelian group Γ, Pontryagin duality gives an identification of $C^*_\lambda(\Gamma)$ with $C(\hat{\Gamma})$, the continuous functions on the dual group.

Proposition 4.2. *The vector state $x \mapsto \langle x\delta_e, \delta_e\rangle$ defines a faithful tracial state on $C^*_\lambda(\Gamma)$.*

Proof. A simple calculation shows this state to be tracial.

Clearly ρ_s commutes with every operator in $C^*_\lambda(\Gamma)$ (since this is easily seen on the generators $\lambda_g \in C^*_\lambda(\Gamma)$). It follows that δ_e is a *separating* vector, meaning that $x\delta_e = y\delta_e$ if and only if $x = y$ (for all $x, y \in C^*_\lambda(\Gamma)$). Indeed, if $x\delta_e = y\delta_e$, then

$$x\delta_s = \rho_s^* x\delta_e = \rho_s^* y\delta_e = y\delta_s,$$

for all $s \in \Gamma$. Since such vectors span $\ell^2(\Gamma)$, it follows that $x = y$. With this observation, faithfulness is simple: If $0 \le x \in C^*_\lambda(\Gamma)$ and $0 = \langle x\delta_e, \delta_e\rangle$, then $0 = \|x^{1/2}\delta_e\|$ and this implies $x^{1/2} = 0$. Thus $x = 0$ too. $\square$

The *group von Neumann algebra of* Γ is defined to be

$$L(\Gamma) = C^*_\lambda(\Gamma)'' \subset \mathbb{B}(\ell^2(\Gamma)).$$

A fundamental theorem of Murray and von Neumann states that $L(\Gamma)$ is the commutant of the right regular representation $\rho\colon \Gamma \to \mathbb{B}(\ell^2(\Gamma))$ – i.e., $L(\Gamma) = \rho(\mathbb{C}[\Gamma])'$ and $L(\Gamma)' = \rho(\mathbb{C}[\Gamma])''$.

Another way of describing $L(\Gamma)$ is as the set of $T \in \mathbb{B}(\ell^2(\Gamma))$ such that T *is constant down the diagonals* – meaning that for every $s, t, x, y \in \Gamma$ such that $ts^{-1} = yx^{-1}$, we have $\langle T\delta_s, \delta_t\rangle = \langle T\delta_x, \delta_y\rangle$.[3] A simple calculation shows that every unitary $\lambda_s \in \mathbb{B}(\ell^2(\Gamma))$ is constant down all diagonals; hence any finite linear combination has this property; thus anything in the weak closure $C^*_\lambda(\Gamma)'' = L(\Gamma)$ does too. The converse, that every such operator is a weak limit of something in $C^*_\lambda(\Gamma)$, uses the bicommutant theorem. Indeed, assume $T \in \mathbb{B}(\ell^2(\Gamma))$ and assume there exist scalars $\{\alpha_s\}_{s\in\Gamma} \subset \mathbb{C}$ such that $\langle T\delta_g, \delta_h\rangle = \alpha_{hg^{-1}}$ for all $g, h \in \Gamma$. A simple calculation shows $\langle T\rho_s\delta_g, \delta_h\rangle = \langle \rho_s T\delta_g, \delta_h\rangle$, for all $s \in \Gamma$, and hence $T \in \rho(\mathbb{C}[\Gamma])' = L(\Gamma)$.

Definition 4.3. A function $\varphi\colon \Gamma \to \mathbb{C}$ is said to be *positive definite* if the matrix

$$[\varphi(s^{-1}t)]_{s,t\in F} \in \mathbb{M}_F(\mathbb{C})$$

is positive for every finite set $F \subset \Gamma$.

[2] You will also see $C^*_r(\Gamma)$ in the literature.

[3] "What? Why is that 'constant down the diagonals'?" you may wonder. Well, if $\Gamma = \mathbb{Z}$ and you write down the matrix of such an operator (with respect to the canonical basis), you'll see that it really is constant down the diagonals.

Fix a positive definite function φ and let $C_c(\Gamma)$ be the finitely supported functions on Γ. Define a sesquilinear form $C_c(\Gamma) \times C_c(\Gamma) \to \mathbb{C}$ by

$$\langle f, g\rangle_\varphi = \sum_{s,t\in\Gamma} \varphi(s^{-1}t)f(t)\overline{g(s)}.$$

This form is positive semidefinite. Indeed, if $f \in C_c(\Gamma)$ has support F, then

$$\langle f, f\rangle_\varphi = \sum_{s,t\in\Gamma} \varphi(s^{-1}t)f(t)\overline{f(s)} = \big\langle [\varphi(s^{-1}t)]_{s,t\in F}(f), (f)\big\rangle,$$

where the inner product on the right is the standard one on $\ell^2(F)$. Since φ is positive definite, $\langle f, f\rangle_\varphi \geq 0$ as asserted. Hence we can mod out by the zero elements and complete to get a Hilbert space $\ell^2_\varphi(\Gamma)$. For $f \in C_c(\Gamma)$ we let $\hat{f} \in \ell^2_\varphi(\Gamma)$ denote its natural image. Here's a GNS construction for the present context.

Definition 4.4. If φ is a positive definite function on Γ, then $\lambda^\varphi\colon \Gamma \to \mathbb{B}(\ell^2_\varphi(\Gamma))$ is the unitary representation given by $\lambda^\varphi_s(\hat{f}) = \widehat{s.f}$, where $s.f(t) = f(s^{-1}t)$, for all $t \in \Gamma$.[4]

Note that

$$\langle \lambda^\varphi_s \hat{\delta}_e, \hat{\delta}_e\rangle_\varphi = \langle \hat{\delta}_s, \hat{\delta}_e\rangle_\varphi = \varphi(s),$$

for all $s \in \Gamma$, and hence we recover φ from the vector functional $\langle \cdot\ \hat{\delta}_e, \hat{\delta}_e\rangle$.

Perhaps the construction of $\ell^2_\varphi(\Gamma)$ seems familiar? It should. Suppose φ is a positive linear functional on $C^*(\Gamma)$. Then $s \mapsto \varphi(s)$ is a positive definite function on Γ: for $s_1, \ldots, s_n \in \Gamma$ we have

$$[\varphi(s_i^{-1}s_j)]_{i,j} = (\mathrm{id}_n \otimes \varphi)\left(\begin{bmatrix} s_1 & s_2 & \cdots & s_n \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}^* \begin{bmatrix} s_1 & s_2 & \cdots & s_n \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \right),$$

which is positive since φ is a c.p. map. It is a simple exercise to show that the GNS space of $C^*(\Gamma)$ with respect to φ is nothing but $\ell^2_\varphi(\Gamma)$.

We will need one more important fact about positive definite functions: they naturally give rise to completely positive maps at the C*-level. First a bit more notation.

Definition 4.5. Let $\varphi\colon \Gamma \to \mathbb{C}$ be a function. We define a corresponding linear functional $\omega_\varphi\colon \mathbb{C}[\Gamma] \to \mathbb{C}$ by

$$\omega_\varphi(\sum_{t\in\Gamma} \alpha_t t) = \sum_{t\in\Gamma} \varphi(t)\alpha_t$$

and *multiplier* $m_\varphi\colon \mathbb{C}[\Gamma] \to \mathbb{C}[\Gamma]$ by

$$m_\varphi(\sum_{t\in\Gamma} \alpha_t t) = \sum_{t\in\Gamma} \varphi(t)\alpha_t t.$$

Theorem 4.6. *Let $\varphi\colon \Gamma \to \mathbb{C}$ be a function with $\varphi(e) = 1$. The following are equivalent:*

1. *the function φ is positive definite;*
2. *there exists a unitary representation λ_φ of Γ on a Hilbert space $\mathcal{H}_\varphi$ and a unit vector ξ_φ such that*
$$\varphi(s) = \langle \lambda_\varphi(s)\xi_\varphi, \xi_\varphi\rangle;$$
3. *the functional ω_φ extends to a state on $C^*(\Gamma)$;*

[4]It isn't hard, just tedious, to check that this is really a unitary representation.

(4) *the multiplier* m_φ *extends to a u.c.p. map on either* $C^*(\Gamma)$ *or* $C^*_\lambda(\Gamma)$*, or extends to a normal u.c.p. map on* $L(\Gamma)$.

Proof. (1) $\Rightarrow$ (2): This follows from Definition 4.4.

(2) $\Rightarrow$ (3): Trivial.

(3) $\Rightarrow$ (4): First we handle the von Neumann algebra case. We can identify ω_φ with a vector state in the universal representation $C^*(\Gamma) \subset \mathbb{B}(\mathcal{H})$ (i.e., the direct sum of all GNS representations). By Fell's absorption principle, there is a unitary operator which conjugates $C^*_\lambda(\Gamma) \otimes 1$ onto the "diagonal" subalgebra of $C^*_\lambda(\Gamma) \otimes C^*(\Gamma)$; that is, the mapping $\sigma\colon C^*_\lambda(\Gamma) \to C^*_\lambda(\Gamma) \otimes C^*(\Gamma)$ defined by

$$\sum_t \alpha_t \lambda_t \mapsto \sum_t \alpha_t (\lambda_t \otimes t)$$

is a $*$-homomorphism and it extends to a normal $*$-homomorphism (also denoted σ) from $L(\Gamma)$ into $L(\Gamma) \bar{\otimes} \mathbb{B}(\mathcal{H})$ (since Fell's principle is spatially implemented). Notice that m_φ coincides with the continuous u.c.p. map

$$(\mathrm{id}_{L(\Gamma)} \otimes \omega_\varphi) \circ \sigma\colon L(\Gamma) \to L(\Gamma),$$

and this completes the von Neumann case (which evidently implies the reduced C*-algebra case as well).

For $C^*(\Gamma)$ we consider the diagonal map $C^*(\Gamma) \to C^*(\Gamma) \otimes C^*(\Gamma)$, $s \mapsto s \otimes s$, and repeat the argument above.

(4) $\Rightarrow$ (1): If m_φ is u.c.p., then for any finite sequence $s_1, \ldots, s_n \in \Gamma$,

$$[\varphi(s_i^{-1} s_j)]_{ij} = \mathrm{diag}(s_1, \ldots, s_n)[m_\varphi(s_i^{-1} s_j)]_{ij}\, \mathrm{diag}(s_1^{-1}, \ldots, s_n^{-1})$$

is positive since $[s_i^{-1} s_j]_{ij} \in \mathbb{M}_n(\mathbb{C}(\Gamma))$ is positive. $\square$

5. Amenable groups

Definition 5.1. A group Γ is *amenable* if there exists a state μ on $\ell^\infty(\Gamma)$ which is invariant under the left translation action: for all $s \in \Gamma$ and $f \in \ell^\infty(\Gamma)$, $\mu(s.f) = \mu(f)$.

Such a state μ is called an *invariant mean.*

Definition 5.2. For a discrete group Γ, we let $\mathrm{Prob}(\Gamma)$ be the space of all probability measures on Γ:

$$\mathrm{Prob}(\Gamma) = \{\mu \in \ell^1(\Gamma) : \mu \geq 0 \text{ and } \sum_{t \in \Gamma} \mu(t) = 1\}.$$

Note that the left translation action of Γ on $\ell^\infty(\Gamma)$ leaves the subspace $\mathrm{Prob}(\Gamma)$ invariant; hence we can also use $\mu \mapsto s.\mu$ to denote the canonical action of Γ on $\mathrm{Prob}(\Gamma)$.

Definition 5.3. We say Γ has an *approximate invariant mean* if for any finite subset $E \subset \Gamma$ and $\varepsilon > 0$, there exists $\mu \in \mathrm{Prob}(\Gamma)$ such that

$$\max_{s \in E} \|s.\mu - \mu\|_1 < \varepsilon.$$

Recall that the *symmetric difference* of two sets E and F, denoted $E \triangle F$, is $E \cup F \setminus E \cap F$.

Definition 5.4. We say Γ satisfies the *Følner condition* if for any finite subset $E \subset \Gamma$ and $\varepsilon > 0$, there exists a finite subset $F \subset \Gamma$ such that

$$\max_{s \in E} \frac{|sF \triangle F|}{|F|} < \varepsilon,$$

where $sF = \{st : t \in F\}$.[5] A sequence of finite sets $F_n \subset \Gamma$ such that

$$\frac{|sF_n \triangle F_n|}{|F_n|} \to 0$$

for every $s \in \Gamma$ is called a *Følner sequence*.

Note that this implies the existence of an approximate invariant mean given by normalized characteristic functions. Indeed, if χ_F is the characteristic function over F, then $\frac{1}{|F|}\chi_F \in \mathrm{Prob}(\Gamma)$ and a computation confirms that

$$\|s.(\frac{1}{|F|}\chi_F) - \frac{1}{|F|}\chi_F\|_1 = \frac{|sF \triangle F|}{|F|}.$$

It turn out that all the definitions above give rise to the same class of groups. Before the proof, however, a few examples might be nice.

Example 5.5 (Elementary amenable groups)**.** It is not hard to see that finite groups are amenable (take the state which maps $\chi_{\{s\}}$ to $1/|\Gamma|$, for each group element). So are abelian groups, as the Markov-Kakutani fixed point theorem easily implies. (There is an alternate proof below.) It is also true that the class of amenable groups is closed under taking subgroups, extensions, quotients and inductive limits. (These all make excellent exercises.) Hence anything built out of finite or abelian groups, using the four operations above, is also amenable; by definition, these are the *elementary amenable groups*. In particular, all solvable (hence all nilpotent) groups are amenable.

Example 5.6 (Groups with subexponential growth)**.** A group Γ is said to have *subexponential growth* if $\limsup |E^n|^{1/n} = 1$ for every finite subset $E \subset \Gamma$. ($E^n = \{g_1 g_2 \cdots g_n : g_i \in E\}$.) It is clear that if a particular finite set E satisfies the above condition, then every subset $F \subset E^n$ will too. Hence if Γ is generated by a finite subset $E \subset \Gamma$ as a semigroup, then it suffices to check the growth condition only for E.

Such groups are amenable. To see this, we construct an increasing sequence $E_0 \subset E_1 \subset E_2 \subset \cdots$ of finite subsets of G, whose union equals Γ, such that $E_n^{-1} = E_n$, $E_m E_n \subset E_{m+n}$, and $\liminf |E_n|^{1/n} = 1$. (Start with any finite set, keep throwing in group elements, and then take higher powers as in the definition of subexponential growth.) It turns out that some subsequence of $\{E_n\}$ must be a Følner sequence. Indeed, for any $g \in E_k$, we have $gE_{n-k} \subset E_n$, and thus $|gE_n \cap E_n| \geq |gE_{n-k}| = |E_{n-k}|$. The proof of the ratio test, from elementary calculus, contains the following general fact:

$$\liminf_{n\to\infty} \frac{a_n}{a_{n-k}} \leq \liminf_{n\to\infty} a_n^{k/n},$$

for $a_n \geq 0$ and any fixed $k \in \mathbb{N}$. Applying the reciprocal of this inequality, we have

$$\limsup_{n\to\infty} \frac{|gE_n \cap E_n|}{|E_n|} \geq \limsup_{n\to\infty} \frac{|E_{n-k}|}{|E_n|} \geq \limsup_{n\to\infty} \frac{1}{|E_n|^{k/n}} = 1.$$

It is a fun combinatorial exercise to show all abelian groups have subexponential growth.

Here is the simplest example of something nonamenable.

Example 5.7 (Nonabelian free groups)**.** The free group $\mathbb{F}_2$ of rank two is not amenable. Let $a, b \in \mathbb{F}_2$ be the free generators and set

$$A^+ = \{\text{all reduced words starting with } a\} \subset \mathbb{F}_2.$$

[5] Since $sF \triangle F = [sF \setminus (sF \cap F)] \cup [F \setminus (sF \cap F)]$, it follows that $\frac{|sF \triangle F|}{|F|} = 2 - 2\frac{|F \cap sF|}{|F|}$. Hence the Følner condition is equivalent to requiring $\max_{s\in E} \frac{|sF\cap F|}{|F|} > 1 - \varepsilon/2$, which is often how it gets used in our context.

Similarly, let A^- be the reduced words beginning with a^{-1} and likewise define B^+ and B^-. Then, for $C = \{1, b, b^2, \dots\} \subset \mathbb{F}_2$, we have

$$\begin{aligned}\mathbb{F}_2 &= A^+ \sqcup A^- \sqcup (B^+ \setminus C) \sqcup (B^- \cup C) \\ &= A^+ \sqcup aA^- \\ &= b^{-1}(B^+ \setminus C) \sqcup (B^- \cup C).\end{aligned}$$

This kind of decomposition is said to be *paradoxical*.[6] Note that the existence of an invariant mean μ on $\ell^\infty(\Gamma)$ would lead to a contradiction:

$$\begin{aligned}1 = \mu(1) &= \mu(\chi_{A^+}) + \mu(\chi_{A^-}) + \mu(\chi_{B^+\setminus C}) + \mu(\chi_{B^-\cup C}) \\ &= \mu(\chi_{A^+}) + \mu(a.\chi_{A^-}) + \mu(b^{-1}.\chi_{B^+\setminus C}) + \mu(\chi_{B^-\cup C}) \\ &= \mu(\chi_{A^+} + a.\chi_{A^-} + b^{-1}.\chi_{B^+\setminus C} + \chi_{B^-\cup C}) \\ &= 2\mu(1) = 2.\end{aligned}$$

Since amenability passes to subgroups, it follows that all nonabelian free groups (on any number of generators) are nonamenable.

Here is a small sample of the known characterizations of amenable groups.

Theorem 5.8. *Let Γ be a discrete group. The following are equivalent:*

(1) *Γ is amenable;*
(2) *Γ has an approximate invariant mean;*
(3) *Γ satisfies the Følner condition;*
(4) *the trivial representation τ_0 is weakly contained in the regular representation λ (i.e., there exist unit vectors $\xi_i \in \ell^2(\Gamma)$ such that $\|\lambda_s(\xi_i) - \xi_i\| \to 0$ for all $s \in \Gamma$);*
(5) *there exists a net (φ_i) of finitely supported positive definite functions on Γ such that $\varphi_i \to 1$ pointwise;*
(6) *$C^*(\Gamma) = C^*_\lambda(\Gamma)$;*
(7) *$C^*_\lambda(\Gamma)$ has a character (i.e., one-dimensional representation);*
(8) *for any finite subset $E \subset \Gamma$, we have*
$$\|\frac{1}{|E|}\sum_{s\in E}\lambda_s\| = 1;$$
(9) *$C^*_\lambda(\Gamma)$ has the CPAP;*
(10) *$L(\Gamma)$ is semidiscrete.*

Proof. (1) $\Rightarrow$ (2): Take an invariant mean μ on $\ell^\infty(\Gamma)$. Being the predual of $\ell^\infty(\Gamma)$, $\ell^1(\Gamma)$ is dense in $\ell^\infty(\Gamma)^*$ and thus we can find a net (μ_i) in $\mathrm{Prob}(\Gamma)$ which converges to μ in the $\sigma(\ell^\infty(\Gamma)^*, \ell^\infty(\Gamma))$-topology. Note that for each $s \in \Gamma$, the net $(s.\mu_i - \mu_i)$ converges to zero weakly in $\ell^1(\Gamma)$ (not just weak* in $\ell^\infty(\Gamma)^*$). Hence, for any finite subset $E \subset \Gamma$, the weak closure of the convex subset $\bigoplus_{s\in E}\{s.\mu - \mu : \mu \in \mathrm{Prob}(\Gamma)\}$ contains zero. Since the weak and norm closures coincide, by the Hahn-Banach Theorem, assertion (2) follows.

(2) $\Rightarrow$ (3): Let a finite subset $E \subset \Gamma$ and $\varepsilon > 0$ be given. Choose $\mu \in \mathrm{Prob}(\Gamma)$ such that
$$\sum_{s\in E}\|s.\mu - \mu\|_1 < \varepsilon.$$
Given a positive function $f \in \ell^1(\Gamma)$ and $r \geq 0$, we define a set $F(f, r) = \{t \in \Gamma : f(t) > r\}$ and let $\chi_{F(f,r)}$ be the characteristic function of this set. For a pair of positive functions

[6] This paradoxical decomposition leads to the famous Banach-Tarski paradox. See Eric Weisstein's website *Mathworld* (mathworld.wolfram.com) for more.

$f, h \in \ell^1(\Gamma)$ and $t \in \Gamma$, observe that $|\chi_{F(f,r)}(t) - \chi_{F(h,r)}(t)| = 1$ if and only if r lies between the numbers $f(t)$ and $h(t)$. If both f and h are bounded above by 1, it follows that

$$|f(t) - h(t)| = \int_0^1 |\chi_{F(f,r)}(t) - \chi_{F(h,r)}(t)| dr.$$

Applying this observation to μ and $s.\mu$, we get

$$\begin{aligned}
\|s.\mu - \mu\|_1 &= \sum_{t\in\Gamma} |s.\mu(t) - \mu(t)| \\
&= \sum_{t\in\Gamma} \int_0^1 |\chi_{F(s.\mu,r)}(t) - \chi_{F(\mu,r)}(t)| dr \\
&= \int_0^1 \big(\sum_{t\in\Gamma} |\chi_{sF(\mu,r)}(t) - \chi_{F(\mu,r)}(t)| \big) dr \\
&= \int_0^1 |sF(\mu, r) \triangle F(\mu, r)| dr.
\end{aligned}$$

Hence we have

$$\varepsilon \int_0^1 |F(\mu, r)| dr = \varepsilon > \sum_{s\in E} \|s\mu - \mu\|_1 = \int_0^1 \sum_{s\in E} |sF(\mu, r) \triangle F(\mu, r)| dr.$$

Thus for some r we must have

$$\sum_{s\in E} |sF(\mu, r) \triangle F(\mu, r)| < \varepsilon |F(\mu, r)|,$$

which shows that $F(\mu, r)$ is almost invariant under translation by the elements in E.

(3) $\Rightarrow$ (4): Let (F_i) be a Følner sequence and $\xi_i = |F_i|^{-1/2}\chi_{F_i}$ be the normalized characteristic functions of the F_i's (viewed as unit vectors in $\ell^2(\Gamma)$). The same calculation used in the ℓ^1 context (see the paragraph after Definition 5.4) shows that $\|\lambda_s(\xi_i) - \xi_i\|_{\ell^2(\Gamma)} \to 0$ for every $s \in \Gamma$.

(4) $\Rightarrow$ (5): Consider the vector states $x \mapsto \langle x\xi_i, \xi_i\rangle$. As noted in the previous section, these restrict to positive definite functions on Γ and obviously tend to 1 pointwise. To make them finitely supported, one simply forces each ξ_i to be a finitely supported ℓ^2 function.

(5) $\Rightarrow$ (6): Take a net (φ_i) as in condition (5). By Theorem 4.6, the multipliers m_{φ_i} (resp. $\tilde{m}_{\varphi_i}$) are u.c.p. on $C^*(\Gamma)$ (resp. $C^*_\lambda(\Gamma)$). We note that $\lambda \circ m_{\varphi_i} = \tilde{m}_{\varphi_i} \circ \lambda$ on $C^*(\Gamma)$ since the two maps are continuous and coincide on the dense subspace $\mathbb{C}[\Gamma]$. Observe that $m_{\varphi_i}(x) \to x$ for every $x \in C^*(\Gamma)$ since this is true for $x \in \mathbb{C}[\Gamma]$. Now suppose $x \in C^*(\Gamma)$ and $\lambda(x) = 0$. Then, we have

$$\lambda(m_{\varphi_i}(x)) = \tilde{m}_{\varphi_i}(\lambda(x)) = 0$$

for every i. But since φ_i is finitely supported, we have $m_{\varphi_i}(x) \in \mathbb{C}[\Gamma]$, and hence $\lambda(m_{\varphi_i}(x)) = 0$ implies $m_{\varphi_i}(x) = 0$. Therefore, $x = \lim_i m_{\varphi_i}(x) = 0$ and the $*$-homomorphism $\lambda\colon C^*(\Gamma) \to C^*_\lambda(\Gamma)$ is injective.

(6) $\Rightarrow$ (7): The trivial representation $\Gamma \to \mathbb{C}$ extends to $C^*(\Gamma) = C^*_\lambda(\Gamma)$.

(7) $\Rightarrow$ (1): Let $\tau\colon C^*_\lambda(\Gamma) \to \mathbb{C}$ be any $*$-homomorphism, but regard it as a state. Extending to $\mathbb{B}(\ell^2(\Gamma))$, we may assume that τ is also defined on $\ell^\infty(\Gamma) \subset \mathbb{B}(\ell^2(\Gamma))$. Since the left translation action is spatially implemented,

$$\tau(s.f) = \tau(\lambda_s f \lambda_s^*) = \tau(\lambda_s)\tau(f)\overline{\tau(\lambda_s)} = \tau(f)$$

for all $s \in \Gamma$ and $f \in \ell^\infty(\Gamma)$ (the unitaries λ_s belong to the multiplicative domain of τ). Hence, τ is an invariant mean as desired.

At this point we have shown the first seven conditions to be equivalent.

(4) $\Leftrightarrow$ (8): The $\Rightarrow$ direction is easy. For the converse, it suffices to show that if E is a finite *symmetric* set (meaning $E = E^{-1}$) satisfying condition (8), then E generates an amenable group. In this situation, the norm-one operator $S = \frac{1}{|E|}\sum_{s\in E}\lambda_s$ is self-adjoint. Thus, for any $\varepsilon > 0$, we can find a unit vector $\xi \in \ell^2(\Gamma)$ such that $|\langle S\xi, \xi\rangle| > 1-\epsilon$. Letting $|\xi|$ be the pointwise absolute value of ξ, a straightforward calculation confirms

$$1-\varepsilon < |\langle S\xi,\xi\rangle| \le \langle S|\xi|, |\xi|\rangle = \frac{1}{|E|}\sum_{s\in E}\langle \lambda_s|\xi|, |\xi|\rangle.$$

Since the cardinality of E is fixed, by taking ε sufficiently small, we deduce that all the numbers $\langle \lambda_s|\xi|, |\xi|\rangle$ must be close to 1; hence the norms $\|\lambda_s|\xi| - |\xi|\|$ are small, for all $s \in E$.

(1) $\Rightarrow$ (9): Let $F_k \subset \Gamma$ be a sequence of Følner sets. For each k let $P_k \in \mathbb{B}(\ell^2(\Gamma))$ be the orthogonal projection onto the finite-dimensional subspace spanned by $\{\delta_g : g \in F_k\}$. Identify $P_k\mathbb{B}(\ell^2(\Gamma))P_k$ with the matrix algebra $\mathbb{M}_{F_k}(\mathbb{C})$ and let $\{e_{p,q}\}_{p,q\in F_k}$ be the canonical matrix units of $\mathbb{M}_{F_k}(\mathbb{C})$. One can check that for each $s \in \Gamma$ we have $e_{p,p}\lambda_s e_{q,q} = 0$ unless $sq = p$, and $e_{p,p}\lambda_s e_{q,q} = e_{p,q}$ if $sq = p$. Since $P_k = \sum_{p\in F_k} e_{p,p}$, we have

$$P_k\lambda_s P_k = \sum_{p,q\in F_k} e_{p,p}\lambda_s e_{q,q} = \sum_{p\in F_k\cap sF_k} e_{p,s^{-1}p}.$$

Let $\varphi_k\colon C^*_\lambda(\Gamma) \to \mathbb{M}_{F_k}(\mathbb{C})$ be the u.c.p. map defined by $x \mapsto P_k x P_k$. Now define a map $\psi_k\colon \mathbb{M}_{F_k}(\mathbb{C}) \to C^*_\lambda(\Gamma)$ by sending

$$e_{p,q} \mapsto \frac{1}{|F_k|}\lambda_p\lambda_{q^{-1}}.$$

Evidently this map is unital; it is also completely positive, as one can check.

The φ_k's and ψ_k's do the trick. Since the linear span of $\{\lambda_s : s \in \Gamma\}$ is norm dense in $C^*_\lambda(\Gamma)$, it suffices to check that $\|\lambda_s - \psi_k\circ\varphi_k(\lambda_s)\| \to 0$ for all $s \in \Gamma$. This follows from the definition of Følner sets together with the following computation:

$$\psi_k\circ\varphi_k(\lambda_s) = \psi_k(\sum_{p\in F_k\cap sF_k} e_{p,s^{-1}p}) = \sum_{p\in F_k\cap sF_k}\frac{1}{|F_k|}\lambda_s = \frac{|F_k\cap sF_k|}{|F_k|}\lambda_s.$$

Hence the reduced group C*-algebra is nuclear.

(1) $\Rightarrow$ (10): The maps constructed above also prove semidiscreteness of $L(\Gamma)$. It suffices to show that for every $x \in L(\Gamma)$ and $g, h \in \Gamma$,

$$\langle \psi_k\circ\varphi_k(x)\delta_g, \delta_h\rangle \to \langle x\delta_g, \delta_h\rangle.$$

If $x \in L(\Gamma)$ is given, then we can find unique scalars $\{\alpha_s\}_{s\in\Gamma}$ such that $\langle x\delta_g, \delta_h\rangle = \alpha_s$, whenever $hg^{-1} = s$. A computation shows

$$\psi_k\circ\varphi_k(x) = \sum_{s\in\Gamma}\alpha_s\frac{|F_k\cap sF_k|}{|F_k|}\lambda_s.$$

It follows that for each fixed pair $g, h \in \Gamma$,

$$\langle\psi_k\circ\varphi_k(x)\delta_g, \delta_h\rangle = \langle\sum_{s\in\Gamma}\alpha_s\frac{|F_k\cap sF_k|}{|F_k|}\lambda_s\delta_g, \delta_h\rangle = \alpha_{hg^{-1}}\frac{|F_k\cap hg^{-1}F_k|}{|F_k|}$$

converges to $\langle x\delta_g, \delta_h\rangle = \alpha_{hg^{-1}}$ as $k \to \infty$.

(9) $\Rightarrow$ (1): Let $\varphi_n\colon C^*_\lambda(\Gamma) \to \mathbb{M}_{k(n)}(\mathbb{C})$ and $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to C^*_\lambda(\Gamma)$ be u.c.p. maps converging to $\mathrm{id}_{C^*_\lambda(\Gamma)}$ in the point-norm topology. By Arveson's Extension Theorem we may assume that the φ_n's are actually defined on all of $\mathbb{B}(\ell^2(\Gamma))$. In other words, letting

$\Phi_n = \psi_n \circ \varphi_n$, we have u.c.p. maps $\Phi_n\colon \mathbb{B}(\ell^2(\Gamma)) \to C^*_\lambda(\Gamma)$ such that $\Phi_n(x) \to x$ for all $x \in C^*_\lambda(\Gamma)$. Taking a point-ultraweak limit point of $\{\Phi_n\}$ (see Theorem 2.1), we get a u.c.p. map $\Phi\colon \mathbb{B}(\ell^2(\Gamma)) \to L(\Gamma)$ which restricts to the identity on $C^*_\lambda(\Gamma)$. This is all we need to show amenability of Γ.

Let τ be the canonical vector trace on $L(\Gamma)$ and consider the state

$$\eta = \tau \circ \Phi$$

on $\mathbb{B}(\ell^2(\Gamma))$. Restricting to $\ell^\infty(\Gamma) \subset \mathbb{B}(\ell^2(\Gamma))$, we get an invariant mean. Indeed, for any $T \in \mathbb{B}(\ell^2(\Gamma))$ and $s \in \Gamma$ we have

$$\eta(\lambda_s T \lambda_s^*) = \tau\big(\lambda_s \Phi(T) \lambda_s^*\big) = \tau(\Phi(T)) = \eta(T),$$

where the first equality uses the fact that Φ restricts to the identity on $C^*_\lambda(\Gamma)$ (hence $C^*_\lambda(\Gamma)$ falls in the multiplicative domain of Φ) and the second uses the fact that τ is a trace. Thus if $T \in \ell^\infty(\Gamma)$, we have $\eta(s.T) = \eta(\lambda_s T \lambda_s^*) = \eta(T)$, since left translation is spatially implemented.

(10) $\Rightarrow$ (1): Semidiscreteness allows one to construct a u.c.p. map $\Phi\colon \mathbb{B}(\ell^2(\Gamma)) \to L(\Gamma)$ which restricts to the identity on $L(\Gamma)$ (use Arveson's Extension Theorem and Theorem 2.1). This is more than enough to imply amenability of Γ, as we saw above. □

Remark 5.9. This theorem not only shows that amenable groups give rise to a very natural class of nuclear C*-algebras, but it also gives our first examples of nonnuclear C*-algebras (since there are plenty of nonamenable groups).

Remark 5.10. A similar theorem holds in the locally compact case, but not everything generalizes. For example, nuclearity or semidiscreteness need not imply amenability in general; Connes proved in [3] that if G_0 denotes the connected component of G and if G/G_0 is amenable, then $C^*_\lambda(G)$ is nuclear and $L(G)$ is semidiscrete. In particular, all connected Lie groups have nuclear reduced C*-algebras (though they need not be amenable).

6. Tensor Products and The Trick

To a von Neumann algebraist there is only one tensor product. This is a problem. Indeed, a wonderful feature of the C*-theory is its complexity. This exposes new ideas and sometimes, in the right hands, provides insight that would otherwise remain out of sight.

The spatial and maximal C*-norms

When A and B are C*-algebras, it can happen that numerous different norms make $A \odot B$ (the algebraic tensor product) into a pre-C*-algebra. In other words, $A \odot B$ may carry more than one C*-norm.

Definition 6.1. A C*-*norm* $\|\cdot\|_\alpha$ on $A \odot B$ is a norm such that $\|xy\|_\alpha \leq \|x\|_\alpha \|y\|_\alpha$, $\|x^*\|_\alpha = \|x\|_\alpha$ and $\|x^*x\|_\alpha = \|x\|_\alpha^2$ for all $x, y \in A \odot B$. We will let $A \otimes_\alpha B$ denote the completion of $A \odot B$ with respect to $\|\cdot\|_\alpha$.

The following example is both of fundamental importance and also illustrates the fact that even "trivial" examples in this subject can have subtleties which require care.

Proposition 6.2. *For each* C*-*algebra* A *there is a* C*-*norm on the algebraic tensor product* $\mathbb{M}_n(\mathbb{C}) \odot A$ *and it is unique.*

Proof. We assume the reader knows how to make $\mathbb{M}_n(A)$ into a C*-algebra and hence the existence of a C*-norm follows from the existence of an algebraic $*$-isomorphism

$$\mathbb{M}_n(\mathbb{C}) \odot A \cong \mathbb{M}_n(A).$$

Uniqueness is then a consequence of the fact that C*-algebras have unique norms since $\mathbb{M}_n(\mathbb{C}) \odot A$ is a C*-algebra with respect to the norm it gets from $\mathbb{M}_n(A)$.[7] □

It requires a little work, but it's a fact that C*-norms on algebraic tensor products always exist. Here are the two most natural candidates.

Definition 6.3. (Maximal norm) Given A and B, we define the *maximal* C*-*norm* on $A \odot B$ to be

$$\|x\|_{\max} = \sup\{\|\pi(x)\| : \pi\colon A \odot B \to \mathbb{B}(\mathcal{H}) \text{ a (cyclic) } *\text{-homomorphism}\}$$

for $x \in A \odot B$. We let $A \otimes_{\max} B$ denote the completion of $A \odot B$ with respect to $\|\cdot\|_{\max}$.

Definition 6.4. (Spatial norm) Let $\pi\colon A \to \mathbb{B}(\mathcal{H})$ and $\sigma\colon B \to \mathbb{B}(\mathcal{K})$ be *faithful* representations. Then the *spatial* (or *minimal*) C*-norm on $A \odot B$ is

$$\|\sum a_i \otimes b_i\|_{\min} = \|\sum \pi(a_i) \otimes \sigma(b_i)\|_{\mathbb{B}(\mathcal{H}\otimes\mathcal{K})}.$$

The completion of $A \odot B$ with respect to $\|\cdot\|_{\min}$ is denoted $A \otimes B$.[8]

Remark 6.5 (Von Neumann algebra tensor products). If $M \subset \mathbb{B}(\mathcal{H})$ and $N \subset \mathbb{B}(\mathcal{K})$ are von Neumann algebras, then there are a number of C*-norms that one can put on $M \odot N$. However, the norm completions won't be von Neumann algebras and researchers have virtually forgotten about the subject of C*-tensor products of von Neumann algebras. On the other hand, the *von Neumann algebraic* tensor product is still very important and is denoted by $M \bar{\otimes} N$. By definition, this is the von Neumann algebra generated by $M \otimes \mathbb{C}1_{\mathcal{K}} \subset \mathbb{B}(\mathcal{H}\otimes\mathcal{K})$ and $\mathbb{C}1_{\mathcal{H}} \otimes N \subset \mathbb{B}(\mathcal{H}\otimes\mathcal{K})$ – i.e., the weak closure of $M \otimes N \subset \mathbb{B}(\mathcal{H}\otimes\mathcal{K})$.

Remark 6.6 (Operator space tensor products). For completeness we also mention that one defines the spatial tensor product norm on operator systems (or spaces) in exactly the same way. Given X and Y, we take embeddings $X \subset \mathbb{B}(\mathcal{H})$ and $Y \subset \mathbb{B}(\mathcal{K})$ which induce the given operator space structures and then define $X \otimes Y$ to be the norm closure of the span of $\{x \otimes y \in \mathbb{B}(\mathcal{H}\otimes\mathcal{K}) : x \in X, y \in Y\}$. As we'll soon see for C*-algebras, $X \otimes Y$ is independent of the embeddings (so long as they induce the proper operator space structures, of course).

There are numerous technical points which one should worry about. The first is whether or not $\|\cdot\|_{\max}$ is even finite. This is the case thanks to the existence of restrictions – i.e., given a $*$-representation $\pi : A \odot B \to \mathbb{B}(\mathcal{H})$, there are $*$-representations $\pi_A\colon A \to \mathbb{B}(\mathcal{H})$ and $\pi_B\colon B \to \mathbb{B}(\mathcal{H})$ such that $\pi(a \otimes b) = \pi_A(a)\pi_B(b)$ for all $a \in A$ and $b \in B$ (see [2, Theorem 3.2.6]). This easily implies that $\|x\|_{\max} < \infty$ for all $x \in A \odot B$.

The remainder of this section is devoted to resolving the following technical issues.

(1) Are $\|\cdot\|_{\max}$ and $\|\cdot\|_{\min}$ norms (as opposed to seminorms)?[9]
(2) Is $\|\cdot\|_{\min}$ independent of the choice of faithful representations?
(3) Can one usually reduce the nonunital case to the unital case?

All three questions have affirmative answers, though none are completely obvious.

Let us first tackle the norm vs. seminorm question. The following *universal property of* $\|\cdot\|_{\max}$ implies that it suffices to show $\|\cdot\|_{\min}$ is a norm.

[7] Depending on what the phrase "C*-algebras have unique norms" means to you, there may or may not be a subtlety here. If this statement only means, "Whenever an algebra B is a C*-algebra with respect to two norms $\|\cdot\|$ and $\|\cdot\|'$, then those norms agree," then the proof of uniqueness has a gap. Luckily, the more general statement, "If $(B, \|\cdot\|)$ is a C*-algebra and $(B, \|\cdot\|')$ is a pre-C*-algebra (i.e., not necessarily complete), then $\|\cdot\| = \|\cdot\|'$," is true and this is what we are using above.

[8] You will also see $A \otimes_{\min} B$ in the literature.

[9] Since both of these (semi)norms are defined via $*$-representations and honest C*-norms, an affirmative answer to this question will imply that both $\|\cdot\|_{\max}$ and $\|\cdot\|_{\min}$ are C*-norms.

Proposition 6.7 (Universality). *If $\pi\colon A \odot B \to C$ is a $*$-homomorphism, then there exists a unique $*$-homomorphism $A \otimes_{\max} B \to C$ which extends π. In particular, any pair of $*$-homomorphisms with commuting ranges $\pi_A\colon A \to C$ and $\pi_B\colon B \to C$ induces a unique $*$-homomorphism*

$$\pi_A \times \pi_B\colon A \otimes_{\max} B \to C.$$

Proof. Faithfully representing C on some Hilbert space, this fact follows from the definition of $\|\cdot\|_{\max}$. □

Corollary 6.8. *The norm $\|\cdot\|_{\max}$ is the largest possible* C*-*norm on $A \odot B$.*

Proof. If $\|\cdot\|_\alpha$ is any other C*-norm on $A \odot B$, then, by universality, there is a (surjective) $*$-homomorphism $A \otimes_{\max} B \to A \otimes_\alpha B$. Hence, $\|x\|_\alpha \leq \|x\|_{\max}$ for every $x \in A \odot B$. □

In particular, $\|\cdot\|_{\max}$ dominates $\|\cdot\|_{\min}$ and thus, if $\|x\|_{\min} = 0 \Rightarrow x = 0$, then it will follow that both $\|\cdot\|_{\max}$ and $\|\cdot\|_{\min}$ are honest norms.

Lemma 6.9. *The product $*$-homomorphism $\mathbb{B}(\mathcal{H}) \odot \mathbb{B}(\mathcal{K}) \to \mathbb{B}(\mathcal{H} \otimes \mathcal{K})$, induced by the commuting $*$-representations $\mathbb{B}(\mathcal{H}) \cong \mathbb{B}(\mathcal{H}) \otimes \mathbb{C}1_{\mathcal{K}} \subset \mathbb{B}(\mathcal{H} \otimes \mathcal{K})$ and $\mathbb{B}(\mathcal{K}) \cong \mathbb{C}1_{\mathcal{H}} \otimes \mathbb{B}(\mathcal{K}) \subset \mathbb{B}(\mathcal{H} \otimes \mathcal{K})$, is injective.*

Proof. We must show that if a finite sum of tensor product operators $\sum_i S_i \otimes T_i \in \mathbb{B}(\mathcal{H} \otimes \mathcal{K})$ is zero, then the corresponding sum of elementary tensors $\sum_i S_i \otimes T_i \in \mathbb{B}(\mathcal{H}) \odot \mathbb{B}(\mathcal{K})$ is also zero. We may assume that the operators $\{S_i\} \subset \mathbb{B}(\mathcal{H})$ are linearly independent.

If $0 = \sum_i S_i \otimes T_i \in \mathbb{B}(\mathcal{H} \otimes \mathcal{K})$, then for all vectors $v, w \in \mathcal{H}$ and $\xi, \eta \in \mathcal{K}$ we have

$$\langle (\sum_i S_i \otimes T_i) v \otimes \xi, w \otimes \eta \rangle = 0.$$

Rearranging terms, we get

$$\begin{aligned} \langle (\sum_i S_i \otimes T_i) v \otimes \xi, w \otimes \eta \rangle &= \sum_i \langle S_i \otimes T_i (v \otimes \xi), w \otimes \eta \rangle \\ &= \sum_i \langle S_i v, w \rangle \langle T_i \xi, \eta \rangle \\ &= \langle (\sum_i \langle T_i \xi, \eta \rangle S_i) v, w \rangle. \end{aligned}$$

Since this holds for all $v, w \in \mathcal{H}$, it follows that the operator $\sum_i \langle T_i \xi, \eta \rangle S_i \in \mathbb{B}(\mathcal{H})$ is zero and hence, by linear independence, that each of the coefficients $\langle T_i \xi, \eta \rangle$ is zero. Since this holds for all $\xi, \eta \in \mathcal{K}$, it follows that $0 = T_i \in \mathbb{B}(\mathcal{K})$ for all i, and the proof is complete. □

Corollary 6.10. *For each $x \in A \odot B$, if $\|x\|_{\min} = 0$, then $x = 0$.*

Proof. If $\pi\colon A \to \mathbb{B}(\mathcal{H})$ and $\sigma\colon B \to \mathbb{B}(\mathcal{K})$ are faithful representations, then the tensor product map

$$\pi \odot \sigma\colon A \odot B \to \mathbb{B}(\mathcal{H}) \odot \mathbb{B}(\mathcal{K})$$

is also injective. Together with the previous lemma this implies the result. □

We now resolve the second technical question.

Proposition 6.11. *The spatial tensor product norm is independent of the choices of* faithful *representations $\pi\colon A \to \mathbb{B}(\mathcal{H})$ and $\sigma\colon B \to \mathbb{B}(\mathcal{K})$.*

Proof. For the moment we will let $\|\cdot\|_{\min}^{(\pi,\sigma)}$ denote the minimal norm with respect to π and σ. Evidently it suffices to prove that if $\sigma'\colon B \to \mathbb{B}(\mathcal{K}')$ is another faithful representation, then $\|\cdot\|_{\min}^{(\pi,\sigma)} = \|\cdot\|_{\min}^{(\pi,\sigma')}$.

For notational reasons it is slightly more convenient to give the proof in the separable setting. It is a simple exercise to net-ify the argument and deduce the general case. Let $P_1 \leq P_2 \leq \cdots$ be finite-rank projections in $\mathbb{B}(\mathcal{H})$ such that P_n has rank n and $\|P_n(h)-h\| \to 0$ for all $h \in \mathcal{H}$. Then it is not hard to show that for every $X \in \mathbb{B}(\mathcal{H}\otimes\mathcal{K})$ we have

$$\|X\| = \sup_n\{\|(P_n \otimes 1_{\mathcal{K}})X(P_n \otimes 1_{\mathcal{K}})\|\}.$$

Thus, if $\sum a_i \otimes b_i \in A \odot B$ is arbitrary, we have

$$\|\sum a_i \otimes b_i\|_{\min}^{(\pi,\sigma)} = \sup_n\{\|\sum(P_n\pi(a_i)P_n)\otimes\sigma(b_i)\|\}$$

and

$$\|\sum a_i \otimes b_i\|_{\min}^{(\pi,\sigma')} = \sup_n\{\|\sum(P_n\pi(a_i)P_n)\otimes\sigma'(b_i)\|\}.$$

But since $P_n\mathbb{B}(\mathcal{H})P_n$ is naturally isomorphic to $\mathbb{M}_n(\mathbb{C})$, we have

$$\|\sum(P_n\pi(a_i)P_n)\otimes\sigma(b_i)\| = \|\sum(P_n\pi(a_i)P_n)\otimes\sigma'(b_i)\|$$

for each n, since there is a unique C*-norm on $\mathbb{M}_n(\mathbb{C})\odot B$ (Proposition 6.2). □

Finally we present a result which allows many nonunital questions to be handled (relatively) painlessly. For a nonunital C*-algebra A we will let $\tilde{A}$ denote the unitization.

Corollary 6.12. *If A is nonunital, then any* C**-norm $\|\cdot\|_\alpha$ on $A\odot B$ can be extended to a* C**-norm on $\tilde{A}\odot B$. Hence, when both A and B are nonunital, any* C**-norm $\|\cdot\|_\alpha$ on $A\odot B$ can be extended to a* C**-norm on $\tilde{A}\odot\tilde{B}$.*

Proof. See [2, Corollary 3.3.12]. □

Exercises

Exercise 6.1. Show that both $\|\cdot\|_{\min}$ and $\|\cdot\|_{\max}$ are *commutative* tensor product norms – i.e., there are canonical isomorphisms $A\otimes B \cong B\otimes A$ and $A\otimes_{\max} B \cong B\otimes_{\max} A$.

Exercise 6.2. Show that both $\|\cdot\|_{\min}$ and $\|\cdot\|_{\max}$ are *associative* – i.e., there are canonical isomorphisms $(A\otimes B)\otimes C \cong A\otimes(B\otimes C)$ and $(A\otimes_{\max} B)\otimes_{\max} C \cong A\otimes_{\max}(B\otimes_{\max} C)$. How would you define the maximal or minimal tensor product of n algebras?

Exercise 6.3. Give an example of a $*$-representation $\pi\colon A\odot B \to \mathbb{B}(\mathcal{H})$ such that both π_A and π_B are injective but π is not. (Hint: Think finite dimensional and abelian.)

Exercise 6.4. Prove that if $\pi\colon A \to \mathbb{B}(\mathcal{H})$ and $\sigma\colon B \to \mathbb{B}(\mathcal{K})$ are arbitrary (not necessarily faithful) representations, then there exists a unique extending $*$-homomorphism $\pi\otimes\sigma\colon A\otimes B \to \mathbb{B}(\mathcal{H}\otimes\mathcal{K})$ such that $\pi\otimes\sigma(a\otimes b) = \pi(a)\otimes\sigma(b)$. (Hint: Dilate π and σ to faithful representations and then cut back down.)

Exercise 6.5. If $\pi\colon A \to C$ and $\sigma\colon B \to D$ are $*$-homomorphisms, prove that there is a unique $*$-homomorphism $\pi\otimes\sigma\colon A\otimes B \to C\otimes D$ such that $\pi\otimes\sigma(a\otimes b) = \pi(a)\otimes\sigma(b)$.

Exercise 6.6. Prove that $\mathbb{B}(\ell^2)\odot\mathbb{B}(\ell^2) \subset \mathbb{B}(\ell^2\otimes\ell^2)$ (see Lemma 6.9) is not dense in norm.

Takesaki's Theorem

It's a nontrivial fact that $\|\cdot\|_{\min}$ is really the smallest possible C*-norm on $A \odot B$. For a proof see [2].

Theorem 6.13 (Takesaki)**.** *For arbitrary* C*-*algebras A and B, $\|\cdot\|_{min}$ is the smallest* C*-*norm on $A \odot B$.*

The following corollary gets used, both explicitly and implicitly, all of the time. For example, in the literature it is often written that $A \odot B$ has a unique C*-norm if and only if $A \otimes_{\max} B = A \otimes B$.

Corollary 6.14. *For any A and B and any* C*-*norm $\|\cdot\|_\alpha$ on $A \odot B$ we have natural surjective $*$-homomorphisms*

$$A \otimes_{\max} B \to A \otimes_\alpha B \to A \otimes B.$$

Continuity of tensor product maps

Continuity of maps on tensor products requires some care. It turns out that nothing funny happens so long as one sticks to c.p. maps, but this is the largest class of maps which always behave well. To get a feel for what can go wrong, let's consider a finite-dimensional example.

Proposition 6.15. *Let $\varphi\colon \mathbb{M}_n(\mathbb{C}) \to \mathbb{M}_n(\mathbb{C})$ be the usual transpose map on the $n \times n$ matrices. Then φ is a unital, positive isometry but the norm of*

$$\varphi \otimes \mathrm{id}_{\mathbb{M}_n(\mathbb{C})}\colon \mathbb{M}_n(\mathbb{C}) \otimes \mathbb{M}_n(\mathbb{C}) \to \mathbb{M}_n(\mathbb{C}) \otimes \mathbb{M}_n(\mathbb{C})$$

is greater than or equal to n.[10]

Proof. It's well known, and easily verified, that the transpose map is a positive isometry. Let $\{e_{i,j}\}_{1\le i,j\le n}$ be a system of matrix units for $\mathbb{M}_n(\mathbb{C})$ and consider

$$S := \sum_{i,j=1}^{n} e_{i,j} \otimes e_{j,i}.$$

Evidently S is a permutation matrix – hence unitary – and has norm one. (If $\{\delta_k\}$ is an orthonormal basis, then $S(\delta_k \otimes \delta_l) = \delta_l \otimes \delta_k$.) On the other hand

$$\varphi \otimes \mathrm{id}_{\mathbb{M}_n(\mathbb{C})}(S) = \sum_{i,j=1}^{n} e_{j,i} \otimes e_{j,i}$$

and a straightforward computation shows that this matrix is equal to nP where P is the one-dimensional projection onto the span of the vector

$$v = \sum_{k=1}^{n} \delta_k \otimes \delta_k.$$

□

Unlike the case of states, this shows that the tensor product of norm-one maps need not have norm one – even on the 2×2 matrices when one map is the identity and the other is a positive unital isometry! The next result follows easily from the previous one.

Proposition 6.16. *Let $\varphi\colon A \to A$ be a positive, unital isometry. It can happen that $\varphi \odot \mathrm{id}_{\mathrm{A}}\colon A \odot A \to A \odot A$ is unbounded. For example, let φ be the transpose map on the unitization of the compact operators.*

[10] Actually, it is equal to n, but we won't need this fact.

Not wanting to dwell on the problems that occur for more general maps, let's treat the case of c.p. maps and move on. We will need the following result approximately $\aleph_0$ times (maybe more).

Theorem 6.17 (Continuity of tensor product maps)**.** *Let $\varphi\colon A \to C$ and $\psi\colon B \to D$ be c.p. maps. Then the algebraic tensor product map*

$$\varphi \odot \psi\colon A \odot B \to C \odot D$$

extends to a c.p. (hence continuous) map on both the minimal and maximal tensor products. Moreover, letting $\varphi \otimes_{\max} \psi\colon A \otimes_{\max} B \to C \otimes_{\max} D$ and $\varphi \otimes \psi\colon A \otimes B \to C \otimes D$ denote the extensions, we have

$$\|\varphi \otimes_{\max} \psi\| = \|\varphi \otimes \psi\| = \|\varphi\|\|\psi\|.$$

Proof. We first handle the spatial tensor product case. Assume $C \subset \mathbb{B}(\mathcal{H})$ and $D \subset \mathbb{B}(\mathcal{K})$. Let $\pi_A\colon A \to \mathbb{B}(\tilde{\mathcal{H}})$, $\pi_B\colon B \to \mathbb{B}(\tilde{\mathcal{K}})$ be the Stinespring dilations of φ and ψ, respectively, and $V_A\colon \mathcal{H} \to \tilde{\mathcal{H}}$, $V_B\colon \mathcal{K} \to \tilde{\mathcal{K}}$ the associated bounded linear operators. By Exercise 6.4 there is a natural $*$-homomorphism $\pi_A \otimes \pi_B\colon A \otimes B \to \mathbb{B}(\tilde{\mathcal{H}} \otimes \tilde{\mathcal{K}})$. Hence we may define $\varphi \otimes \psi\colon A \otimes B \to C \otimes D$ by the formula

$$\varphi \otimes \psi(x) = (V_A \otimes V_B)^* \pi_A \otimes \pi_B(x)(V_A \otimes V_B).$$

Note that on elementary tensors we have

$$\varphi \otimes \psi(a \otimes b) = (V_A^* \pi_A(a) V_A) \otimes (V_B^* \pi_B(b) V_B) = \varphi(a) \otimes \psi(b).$$

Hence $\varphi \otimes \psi$ really is a c.p. extension of $\varphi \odot \psi$ which takes values in $C \otimes D \subset \mathbb{B}(\mathcal{H} \otimes \mathcal{K})$. Finally note that the completely bounded norm satisfies

$$\|\varphi \otimes \psi\|_{\mathrm{cb}} \le \|V_A \otimes V_B\|^2 = \|V_A\|^2 \|V_B\|^2 = \|\varphi\|\|\psi\|.$$

The other inequality is easy and will be left to the reader (consider elementary tensors).

For the maximal tensor product, let's first tackle the case that $B = D$ and $\psi = \mathrm{id}_B$. Fix a faithful representation $C \otimes_{\max} B \subset \mathbb{B}(\mathcal{H})$. By the existence of restrictions, we may assume that $C \subset \mathbb{B}(\mathcal{H})$ and $B \subset \mathbb{B}(\mathcal{H})$ commute (and generate $C \otimes_{\max} B$) thus allowing us to regard φ as a c.p. map into $\mathbb{B}(\mathcal{H})$ with $B \subset \varphi(A)'$. Applying Stinespring to φ – also lifting B with the commutant $\varphi(A)'$ (see [2, Proposition 1.5.6]) – we get a $*$-representation of $A \otimes_{\max} B$ (by universality) which we can cut to recover the original map $\varphi \odot \mathrm{id}_B\colon A \odot B \to C \otimes_{\max} B \subset \mathbb{B}(\mathcal{H})$ (just as in the spatial tensor product case above).

Since an arbitrary map $\varphi \odot \psi\colon A \odot B \to C \odot D$ can be decomposed as $(\varphi \odot \mathrm{id}_D) \circ (\mathrm{id}_A \odot \psi)$, the proof is complete. $\square$

The next result is a trivial consequence (that we will use frequently and without reference).

Corollary 6.18. *Assume $\theta\colon A \to C$ and $\sigma\colon B \to D$ are c.c.p. maps and $\theta_n\colon A \to C$ are c.c.p. maps converging to θ in the point-norm topology. Then*

$$\theta_n \otimes_{\max} \sigma \to \theta \otimes_{\max} \sigma$$

and

$$\theta_n \otimes \sigma \to \theta \otimes \sigma$$

in the point-norm topology as well.

Inclusions and The Trick

In this section we discuss one of the important subtleties of C*-tensor products. We also introduce one of the great tensor product tricks, a technique so important that it should not be regarded as a trick, but rather The Trick.

The issue at hand is whether or not inclusions of C*-algebras give inclusions of tensor products. So long as one stays at the algebraic (i.e., pre-C*-algebra) level, nothing funny happens. This simple fact implies that *spatial* tensor products are also well behaved in this regard.

Proposition 6.19. *If $A \subset B$ and C are* C*-*algebras, then there is a natural inclusion*

$$A \otimes C \subset B \otimes C.$$

Proof. Perhaps we should first point out what this proposition is really asserting. Since we have a natural algebraic inclusion

$$A \odot C \subset B \odot C,$$

one can ask what sort of norm we would get on $A \odot C$ if we took the spatial norm on $B \odot C$ and restricted it. This proposition asserts that we just get the spatial norm on $A \odot C$.

Having understood the meaning of the result, the proof is now an immediate consequence of Proposition 6.11. □

Applying this fact again on the right hand side implies that a pair of inclusions $A \subset B$ and $C \subset D$ induces a natural inclusion $A \otimes C \subset B \otimes D$.

For maximal tensor products the question then becomes: If $A \subset B$ and C are given, do we have a natural inclusion $A \otimes_{\max} C \subset B \otimes_{\max} C$? In general this turns out to be false and may seem a little puzzling at first. However, when reformulated at the algebraic level, it becomes clear what can go wrong. Indeed, what we are really asking is whether or not the maximal norm on $B \odot C$ restricts to the maximal norm on $A \odot C \subset B \odot C$. But the maximal norm is defined via a supremum over representations and since every representation of $B \odot C$ gives a representation of the smaller algebra $A \odot C$, it is clear that the supremum only over representations of $B \odot C$ will always be less than or equal to the supremum over all representations of $A \odot C$.

Having seen what the problem could be, it's not too hard to formulate a condition which ensures that inclusions behave nicely for maximal tensor products.

Proposition 6.20. *Let $A \subset B$ be an inclusion of* C*-*algebras and assume that for every nondegenerate $*$-homomorphism $\pi\colon A \to \mathbb{B}(\mathcal{H})$ there exists a c.c.p. map $\varphi\colon B \to \pi(A)''$ such that $\varphi(a) = \pi(a)$ for all $a \in A$. Then for every* C*-*algebra C there is a natural inclusion*

$$A \otimes_{\max} C \subset B \otimes_{\max} C.$$

Proof. By universality, we have a canonical $*$-homomorphism $A \otimes_{\max} C \to B \otimes_{\max} C$. Our goal is to show that if $x \in A \otimes_{\max} C$ is in the kernel of this map, then $x = 0$.

Let $\pi\colon A \otimes_{\max} C \to \mathbb{B}(\mathcal{H})$ be a faithful representation and $\pi_A\colon A \to \mathbb{B}(\mathcal{H})$, $\pi_C\colon C \to \mathbb{B}(\mathcal{H})$ be the restrictions. Note that $\pi_C(C) \subset \pi_A(A)'$ and hence the commuting inclusions $\pi_A(A)'' \hookrightarrow \mathbb{B}(\mathcal{H})$, $\pi_C(C) \hookrightarrow \mathbb{B}(\mathcal{H})$ induce, by universality, a product $*$-homomorphism

$$\pi_A(A)'' \otimes_{\max} \pi_C(C) \longrightarrow \mathbb{B}(\mathcal{H}).$$

Extend π_A to a c.c.p. map $\varphi\colon B \to \pi(A)''$ such that $\varphi(a) = \pi(a)$ for all $a \in A$. By Theorem 6.17 we have the following commutative diagram:

$$\begin{tikzcd} B \otimes_{\max} C \arrow[r, "\varphi\otimes_{\max}\pi_C"] & \pi_A(A)'' \otimes_{\max} \pi_C(C) \arrow[d] \\ A \otimes_{\max} C \arrow[u] \arrow[r, "\pi"] & \mathbb{B}(\mathcal{H}). \end{tikzcd}$$

The fact that π is faithful implies that the map on the left is also injective. □

We will soon introduce The Trick and provide the converse to the previous result, but first we consider two nice corollaries.

Corollary 6.21. *If $A \subset B$, A is nuclear and C is arbitrary, then we have a natural inclusion*

$$A \otimes_{\max} C \subset B \otimes_{\max} C.$$

Proof. Let $\pi\colon A \to \mathbb{B}(\mathcal{H})$ be a representation and $\varphi_n\colon A \to \mathbb{M}_{k(n)}(\mathbb{C})$, $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to \pi(A)$ be c.c.p. maps converging to π in the point-norm topology.[11] By Arveson's Extension Theorem we may assume that each of the φ_n's is actually defined on all of B. Letting $\Phi\colon B \to \pi_A(A)''$ be any point-ultraweak cluster point of the maps $\psi_n \circ \varphi_n\colon B \to A \subset \pi_A(A)''$, we get the c.c.p. extension of π required to invoke Proposition 6.20. □

Corollary 6.22. *If $A \subset B$ is a hereditary subalgebra, then for every C we have a natural inclusion*

$$A \otimes_{\max} C \subset B \otimes_{\max} C.$$

Proof. If $\{e_n\} \subset A$ is an approximate unit, then the c.c.p. maps $\varphi_n\colon B \to A$, $\varphi_n(b) = e_n b e_n$ have the property that $\varphi_n(a) \to a$ for all $a \in A$. With this observation, the proof is similar to the previous corollary, so we leave the details to the reader. □

Proposition 6.23 (The Trick)**.** *Let $A \subset B$ and C be C*-algebras, $\|\cdot\|_\alpha$ be a C*-norm on $B \odot C$ and $\|\cdot\|_\beta$ be the C*-norm on $A \odot C$ obtained by restricting $\|\cdot\|_\alpha$ to $A \odot C \subset B \odot C$. If $\pi_A\colon A \to \mathbb{B}(\mathcal{H})$, $\pi_C\colon C \to \mathbb{B}(\mathcal{H})$ are representations with commuting ranges and if the product $*$-homomorphism*

$$\pi_A \times \pi_C\colon A \odot C \to \mathbb{B}(\mathcal{H})$$

is continuous with respect to $\|\cdot\|_\beta$, then there exists a c.c.p. map $\varphi\colon B \to \pi_C(C)'$ which extends π_A.

Proof. Assume first that A, B and C are all unital and, moreover, that $1_A = 1_B$. Let

$$\pi_A \times_\beta \pi_C\colon A \otimes_\beta C \to \mathbb{B}(\mathcal{H})$$

be the extension of the product map to $A \otimes_\beta C$. Since $A \otimes_\beta C \subset B \otimes_\alpha C$, we apply Arveson's Extension Theorem to get a u.c.p. extension $\Phi\colon B \otimes_\alpha C \to \mathbb{B}(\mathcal{H})$. The desired map is just $\varphi(b) = \Phi(b \otimes 1_C)$.

To see that φ takes values in $\pi_C(C)'$ is a simple multiplicative domain argument. Indeed, $\mathbb{C}1_B \otimes C$ lives in the multiplicative domain of Φ since $\Phi|_{\mathbb{C}1_B \otimes C} = \pi_C$ is a $*$-homomorphism. Since $B \otimes \mathbb{C}1_C$ commutes with $\mathbb{C}1_B \otimes C$ and u.c.p. maps are bimodule maps over their multiplicative domains, a simple calculation completes the proof.

The nonunital case is a bit more irritating but can be deduced from the unital case as follows. For a C*-algebra D, let $\tilde{D}$ be the unitization if D is nonunital and $\tilde{D} = D$ if D is already unital. For an arbitrary inclusion $A \subset B$ and auxiliary algebra C we may extend any C*-norm $\|\cdot\|_\alpha$ on $B \odot C$ to unitizations (Corollary 6.12) and get an inclusion

[11] It's crucial that $\pi\colon A \to \pi(A)$ be nuclear; the result need not hold if $\pi(A)$ is replaced by $\mathbb{B}(\mathcal{H})$.

$B \otimes_\alpha C \subset \tilde{B} \otimes_\alpha \tilde{C}$. Let $A_1 = A + \mathbb{C}1_{\tilde{B}}$ (which may or may not be the same as $\tilde{A}$) and note that $A_1 \odot C \subset \tilde{B} \otimes_\alpha \tilde{C}$. Hence $\|\cdot\|_\beta$ extends to a norm which yields an inclusion $A \otimes_\beta C \subset A_1 \otimes_\beta \tilde{C}$. The key observation is that $A \otimes_\beta C$ is an ideal in $A_1 \otimes_\beta \tilde{C}$ and hence any representation of $A \otimes_\beta C$ extends to a representation of $A_1 \otimes_\beta \tilde{C}$. Given this fact, it is easy to deduce the general case from the unital one proved above. □

At first glance, the utility of The Trick is far from obvious, but please be patient as the mileage one can get out of this simple observation is remarkable. Let us briefly explain what the point is and then we will give an application.

Given an inclusion $A \subset B$ and a representation $\pi\colon A \to \mathbb{B}(\mathcal{H})$, Arveson's Extension Theorem always allows one to extend π to a c.c.p. map $\varphi\colon B \to \mathbb{B}(\mathcal{H})$. When The Trick is applicable, it gives one the ability to better control the range of φ and this is how it gets used. As our first example we provide the converse of Proposition 6.20, promised earlier. An inclusion satisfying one of the following equivalent conditions is called *relatively weakly injective*.

Proposition 6.24. *Let $A \subset B$ be an inclusion. Then the following are equivalent:*

(1) *there exists a c.c.p. map $\varphi\colon B \to A^{**}$ such that $\varphi(a) = a$ for all $a \in A$;*
(2) *for every $*$-homomorphism $\pi\colon A \to \mathbb{B}(\mathcal{H})$ there exists a c.c.p. map $\varphi\colon B \to \pi(A)''$ such that $\varphi(a) = \pi(a)$ for all $a \in A$;*
(3) *for every C*-algebra C there is a natural inclusion*

$$A \otimes_{\max} C \subset B \otimes_{\max} C.$$

Proof. Since every representation of A extends to a normal representation of A^{**}, the equivalence of the first two statements is easy.

Assume condition (3) and let $\pi\colon A \to \mathbb{B}(\mathcal{H})$ be a representation. Let $C = \pi(A)'$ and, by universality, we can apply The Trick to the product map induced by the commuting representations $\pi\colon A \to \mathbb{B}(\mathcal{H})$ and $\pi(A)' \hookrightarrow \mathbb{B}(\mathcal{H})$. That's it. □

Definition 6.25. A C*-algebra $A \subset \mathbb{B}(\mathcal{H})$ is said to have Lance's *weak expectation property* (WEP) if there exists a u.c.p. map $\Phi\colon \mathbb{B}(\mathcal{H}) \to A^{**}$ such that $\Phi(a) = a$ for all $a \in A$.

A simple application of Arveson's Extension Theorem shows that the WEP is independent of the choice of faithful representation.

Corollary 6.26. *A C*-algebra A has the WEP if and only if for every inclusion $A \subset B$ and arbitrary C we have a natural inclusion $A \otimes_{\max} C \subset B \otimes_{\max} C$.*

Proof. Assume first that $A \subset \mathbb{B}(\mathcal{H})$ has the WEP and $A \subset B$. The inclusion $A \hookrightarrow \mathbb{B}(\mathcal{H})$ extends to a c.c.p. map $\Psi\colon B \to \mathbb{B}(\mathcal{H})$ by Arveson's Extension Theorem. Composing with Φ gives a map $B \to A^{**}$ which restricts to the identity on A and then Proposition 6.24 applies. The converse uses The Trick just as in the previous proposition. This time take $B = \mathbb{B}(\mathcal{H}_\mathcal{U})$, the universal representation of A, and $C = (A^{**})'$. □

We opened this section by claiming that inclusions of tensor products can be tricky and then proceeded to give several instances where they behave well. Here is an example where inclusions misbehave.

Proposition 6.27. *Let Γ be a discrete group. Then the following are equivalent:*

(1) *Γ is amenable;*
(2) *$C^*_\lambda(\Gamma)$ has the WEP;*
(3) *the natural inclusion $\iota\colon C^*_\lambda(\Gamma) \hookrightarrow \mathbb{B}(\ell^2(\Gamma))$ induces an injective tensor product map*

$$\iota \otimes_{\max} \mathrm{id}\colon C^*_\lambda(\Gamma) \otimes_{\max} C^*_\lambda(\Gamma) \hookrightarrow \mathbb{B}(\ell^2(\Gamma)) \otimes_{\max} C^*_\lambda(\Gamma).$$

In particular, the natural map

$$\iota \otimes_{\max} \mathrm{id}\colon C^*_\lambda(\Gamma) \otimes_{\max} C^*_\lambda(\Gamma) \to \mathbb{B}(\ell^2(\Gamma)) \otimes_{\max} C^*_\lambda(\Gamma)$$

has a nontrivial kernel for every nonamenable group.

Proof. (1) $\Rightarrow$ (2) follows from Theorem 5.8, since Theorem 2.1 easily implies that every nuclear C*-algebra has the WEP. The implication (2) $\Rightarrow$ (3) is immediate from Corollary 6.26.

For the final implication we use The Trick to produce a u.c.p. map $\Phi\colon \mathbb{B}(\ell^2(\Gamma)) \to L(\Gamma)$ such that $\Phi(x) = x$ for all $x \in C^*_\lambda(\Gamma)$. We already saw in the proof of Theorem 5.8 that this is enough to imply amenability of Γ.

So let $B = \mathbb{B}(\ell^2(\Gamma))$, $C = C^*_\lambda(\Gamma)$ and recall that the commutant of the right regular representation is $L(\Gamma)$. In other words, if $C^*_\lambda(\Gamma) \otimes_{\max} C^*_\lambda(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))$ is the product of the left and right regular representations, then the extension $\varphi\colon \mathbb{B}(\ell^2(\Gamma)) \to \mathbb{B}(\ell^2(\Gamma))$ given by The Trick takes values in $L(\Gamma)$. $\square$

Exercises

Exercise 6.7. Let Γ be a discrete group and let

$$\lambda \times \rho\colon C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))$$

be the product of the left and right regular representations. Prove that Γ is amenable if and only if $\lambda \times \rho$ is continuous with respect to the minimal tensor product norm.

Exercise 6.8. Let X be a locally compact Hausdorff space and $C_0(X)$ be the continuous functions vanishing at ∞. For a C*-algebra A we let

$$C_0(X, A) = \{f\colon X \to A : f \text{ is continuous and } f(\infty) = 0\}.$$

Show that there is a natural isomorphism

$$C_0(X) \otimes_{\max} A \cong C_0(X) \otimes A \cong C_0(X, A)$$

such that $h \otimes a$ maps to the function $x \mapsto h(x)a$. (Hint: A partition of unity argument will show density.)

7. Nuclearity, Injectivity and Semidiscreteness

This section is what we've been after, one of the greatest applications of W*-theory to C*-theory that I'm aware of. The results below are due to Connes, Choi-Effros and Kirchberg.

Lemma 7.1. [12] *Assume that $\theta\colon A \to B$ is a nuclear map. Then for every* C*-*algebra C the map $\theta \otimes_{\max} \mathrm{id}_C\colon A \otimes_{\max} C \to B \otimes_{\max} C$ factors through $A \otimes C$. That is, there exists a c.c.p. map $\Psi\colon A \otimes C \to B \otimes_{\max} C$ such that the diagram*

$$\begin{tikzcd} A \otimes_{\max} C \arrow[rr, "\theta \otimes_{\max} \mathrm{id}_C"] \arrow[dr] & & B \otimes_{\max} C \\ & A \otimes C \arrow[ur, dashed, "\Psi"'] & \end{tikzcd}$$

commutes, where $A \otimes_{\max} C \to A \otimes C$ is the canonical quotient map.

[12]The converse of this lemma also holds, but we won't need it.

Proof. Let $\varphi_n\colon A \to \mathbb{M}_{k(n)}(\mathbb{C})$ and $\psi_n\colon \mathbb{M}_{k(n)}(\mathbb{C}) \to B$ be c.c.p. maps converging to θ in the point-norm topology. Due to the fact that there is a unique C*-norm on $\mathbb{M}_{k(n)}(\mathbb{C}) \odot C$, we get an approximately commuting diagram

$$\begin{tikzcd} A \otimes_{\max} C \arrow[r, "\theta\otimes_{\max}\mathrm{id}_C"] \arrow[d] \arrow[dr, "\varphi_n\otimes_{\max}\mathrm{id}_C"] & B \otimes_{\max} C \\ A \otimes C \arrow[r, "\varphi_n\otimes\mathrm{id}_C"'] & \mathbb{M}_{k(n)}(\mathbb{C}) \otimes C. \arrow[u, "\psi_n\otimes_{\max}\mathrm{id}_C"'] \end{tikzcd}$$

Hence we can define a sequence of c.c.p. maps $\Psi_n\colon A \otimes C \to B \otimes_{\max} C$ by

$$\Psi_n = (\psi_n \otimes_{\max} \mathrm{id}_C) \circ (\varphi_n \otimes \mathrm{id}_C).$$

It follows that the algebraic tensor product map $\theta \odot \mathrm{id}_C\colon A \odot C \to B \odot C$ is contractive from the spatial norm on $A \odot C$ to the maximal norm on $B \odot C$ (since $\Psi_n(x) \to \theta \odot \mathrm{id}_C(x)$ for all $x \in A \odot C$) and hence it extends to a contractive linear map $\Psi\colon A \otimes C \to B \otimes_{\max} C$. Finally, one checks that Ψ is the point-norm limit of the Ψ_n's, hence is completely positive. □

Proposition 7.2. *If A is nuclear, then for every C*-algebra C there is a unique C*-norm on $A \odot C$. In other words, the canonical quotient mapping*

$$A \otimes_{\max} C \to A \otimes C$$

is injective.

Proof. Apply Lemma 7.1 to $\theta = \mathrm{id}_A\colon A \to A$. □

The conclusion of the previous result is the historical definition of nuclearity. For now, let us say that A is $\otimes$-*nuclear* if $A \otimes_{\max} B = A \otimes B$ for every C*-algebra B – thus nuclearity implies $\otimes$-nuclearity. But what is $\otimes$-nuclearity good for? Well, here's something.

Proposition 7.3. *If A is $\otimes$-nuclear and $\pi\colon A \to \mathbb{B}(\mathcal{H})$ is a nondegenerate representation, then $\pi(A)''$ is injective – i.e., there is a u.c.p. map $\Phi\colon \mathbb{B}(\mathcal{H}) \to \pi(A)''$ such that $\Phi(x) = x$ for all $x \in \pi(A)''$.*[13]

Proof. A simple application of The Trick, applied to

$$\pi(A)' \otimes_{\max} A = \pi(A)' \otimes A \subset \mathbb{B}(\mathcal{H}) \otimes A,$$

shows that $\pi(A)'$ is injective. The result now follows from the triviality of W*-representation theory, together with the existence of Haagerup standard form. □

Theorem 7.4. *For a C*-algebra A, the following are equivalent:*

(1) *A is $\otimes$-nuclear;*
(2) *A is nuclear;*
(3) *A^{**} is semidiscrete;*
(4) *A^{**} is injective.*

Proof. We've already seen (3) $\Longrightarrow$ (2) (Proposition 3.9), (2) $\Longrightarrow$ (1) (Proposition 7.2) and (1) $\Longrightarrow$ (4) (Proposition 7.3). The remaining implication is very hard and played no small role in the awarding of a Fields Medal to Alain Connes (cf. [3]).[14] For details see [2]. □

[13] It follows from Arveson's Extension Theorem that injectivity defined this way is equivalent to the usual definition found in homological algebra books (where the category has operator systems as objects and u.c.p. maps as morphisms).

[14] Today there are simpler proofs, but there is still no simple proof.

If $J \triangleleft A$ is a closed, 2-sided ideal, then $A^{**} = J^{**} \oplus (A/J)^{**}$. Hence the next corollary is easily deduced from the previous theorem.

Corollary 7.5. *Nuclearity passes to quotients.*

To any C*-algebraist that still isn't convinced of the usefulness of von Neumann algebras, I have a challenge: give a C*-proof of the previous corollary. Good luck...

Knowing the equivalence of (2) and (3) in Theorem 7.4 would be enough to deduce Corollary 7.5. Which begs the following question: is there a simple proof of (2) $\Longrightarrow$ (3)?

8. Reduced Crossed Products

Definition 8.1. Let Γ be a discrete group and A be a C*-algebra. An *action* of Γ on A is a group homomorphism α from Γ into the group of $*$-automorphisms on A. A C*-algebra equipped with a Γ-action is called a Γ-*C*-algebra*.[15]

Our goal is to construct a single C*-algebra which encodes the action of Γ on A. In group theory, this procedure is well known and is called the *semidirect product*. We will adapt this idea and create an algebra $A \rtimes_\alpha \Gamma$ with the property that there is a copy of Γ inside the unitary group of $A \rtimes_\alpha \Gamma$ (at least when A is unital) and there is a natural inclusion $A \subset A \rtimes_\alpha \Gamma$ such that (a) $A \rtimes_\alpha \Gamma$ is generated by A and Γ and (b) $\alpha_g(a) = gag^*$ for all $a \in A$ and $g \in \Gamma$ (i.e., the action of Γ becomes inner).

For a Γ-C*-algebra A, we denote by $C_c(\Gamma, A)$ the linear space of finitely supported functions on Γ with values in A. A typical element S in $C_c(\Gamma, A)$ is written as a finite sum $S = \sum_{s\in\Gamma} a_s s$. We equip $C_c(\Gamma, A)$ with an α-twisted convolution product and $*$-operation as follows: for $S = \sum_{s\in\Gamma} a_s s, T = \sum_{t\in\Gamma} b_t t \in C_c(\Gamma, A)$ we declare

$$S *_\alpha T = \sum_{s,t\in\Gamma} a_s \alpha_s(b_t) st \text{ and } S^* = \sum_{s\in\Gamma} \alpha_{s^{-1}}(a_s^*) s^{-1}.$$

The twisted convolution is a generalization of the classical convolution of two $\ell^2(\mathbb{Z})$ functions, but the algebraic explanation of these formulas is perhaps more enlightening. Indeed, we are trying to turn $C_c(\Gamma, A)$ into a $*$-algebra where the action becomes inner and hence the definition above comes from the formal calculation

$$(\sum_{s\in\Gamma} a_s s)(\sum_{t\in\Gamma} b_t t) = \sum_{s,t\in\Gamma} a_s(s b_t s^*)st = \sum_{s,t\in\Gamma} a_s \alpha_s(b_t) st.$$

However you care to think about it, $C_c(\Gamma, A)$ is the smallest $*$-algebra which encodes the action of Γ on A. Note that when $A = \mathbb{C}$ and the action α is trivial, we simply recover the group ring $\mathbb{C}[\Gamma]$. Now the question is, "How shall we complete $C_c(\Gamma, A)$?" Just as for group C*-algebras, there are two natural choices, a universal and a reduced completion.

A *covariant representation* $(u, \pi, \mathcal{H})$ of the Γ-C*-algebra A consists of a unitary representation $(u, \mathcal{H})$ of Γ and a $*$-representation $(\pi, \mathcal{H})$ of A such that $u_s \pi(a) u_s^* = \pi(\alpha_s(a))$ for every $s \in \Gamma$ and $a \in A$. It is not hard to see that every covariant representation gives rise to a $*$-representation of $C_c(\Gamma, A)$ and, conversely, every (nondegenerate) $*$-representation of $C_c(\Gamma, A)$ arises this way. For a covariant representation $(u, \pi, \mathcal{H})$, we denote by $u \times \pi$ the associated $*$-representation of $C_c(\Gamma, A)$.

[15] Of course, there could be many different actions of Γ on A, giving rise to different Γ-C*-algebra structures.

Definition 8.2. The *full crossed product* (sometimes called the "universal" crossed product) of a C*-dynamical system (A, α, Γ), denoted $A \rtimes_\alpha \Gamma$, is the completion of $C_c(\Gamma, A)$ with respect to the norm

$$\|x\|_u = \sup \|\pi(x)\|,$$

where the supremum is over all (cyclic) $*$-homomorphisms $\pi \colon C_c(\Gamma, A) \to \mathbb{B}(\mathcal{H})$.

Though it isn't completely obvious, we will soon see that there are lots of representations $C_c(\Gamma, A) \to \mathbb{B}(\mathcal{H})$. (In particular, $\|\cdot\|_u$ really is a norm, as opposed to seminorm, on $C_c(\Gamma, A)$ and hence we have a natural inclusion $C_c(\Gamma, A) \subset A \rtimes_\alpha \Gamma$.) Evidently our definition implies the following *universal property*.

Proposition 8.3 (Universal property)**.** *For every covariant representation $(u, \pi, \mathcal{H})$ of a Γ-C*-algebra A, there is a $*$-homomorphism $\sigma \colon A \rtimes_\alpha \Gamma \to \mathbb{B}(\mathcal{H})$ such that*

$$\sigma(\sum_{s \in \Gamma} a_s s) = \sum_{s \in \Gamma} \pi(a_s) u_s,$$

for all $\sum_{s \in \Gamma} a_s s \in C_c(\Gamma, A)$.

To define the reduced crossed product, we begin with a *faithful* representation $A \subset \mathbb{B}(\mathcal{H})$. Define a new representation of A on $\mathcal{H} \otimes \ell^2(\Gamma)$ by

$$\pi(a)(v \otimes \delta_g) = (\alpha_{g^{-1}}(a)(v)) \otimes \delta_g,$$

where $\{\delta_g\}_{g \in G}$ is the canonical orthonormal basis of $\ell^2(\Gamma)$. Under the identification $\mathcal{H} \otimes \ell^2(\Gamma) \cong \bigoplus_{g \in \Gamma} \mathcal{H}$ we have simply taken the direct sum representation

$$\pi(a) = \bigoplus_{g \in \Gamma} \alpha_g^{-1}(a) \in \mathbb{B}(\bigoplus_{g \in \Gamma} \mathcal{H}).$$

The point of doing this is that now the left regular representation of Γ spatially implements the action α: for all elementary tensors we have

$$\begin{aligned}(1 \otimes \lambda_s)\pi(a)(1 \otimes \lambda_s^*)(v \otimes \delta_g) &= (1 \otimes \lambda_s)\pi(a)(v \otimes \delta_{s^{-1}g}) \\ &= (1 \otimes \lambda_s)((\alpha_{g^{-1}s}(a)(v)) \otimes \delta_{s^{-1}g}) \\ &= (\alpha_{g^{-1}s}(a)(v)) \otimes \delta_g \\ &= (\alpha_{g^{-1}}(\alpha_s(a))(v)) \otimes \delta_g \\ &= \pi(\alpha_s(a))(v \otimes \delta_g).\end{aligned}$$

Hence we get an induced covariant representation $(1 \otimes \lambda) \times \pi$, called a *regular representation*.[16]

Definition 8.4. The *reduced crossed product* of a C*-dynamical system (A, Γ, α), denoted $A \rtimes_{\alpha,r} \Gamma$, is defined to be the norm closure of the image of a regular representation $C_c(\Gamma, A) \to \mathbb{B}(\mathcal{H} \otimes \ell^2(\Gamma))$.

For notational simplicity, we will usually forget about the representation π and the fact that we had to inflate the left regular representation of Γ – i.e., we often (abuse notation slightly and) denote a typical element $x \in C_c(\Gamma, A) \subset A \rtimes_{\alpha,r} \Gamma$ as a finite sum $x = \sum_{s \in \Gamma} a_s \lambda_s$.

Though the following proposition should come as no surprise, the proof contains some important calculations.

Proposition 8.5. *The reduced crossed product $A \rtimes_{\alpha,r} \Gamma$ does not depend on the choice of the faithful representation $A \subset \mathbb{B}(\mathcal{H})$.*

[16] Regular representations are easily seen to be injective on $C_c(\Gamma, A)$; hence the universal norm really is a norm.

Proof. The proof boils down to the fact that there is a unique C*-norm on $\mathbb{M}_n(A)$, just as in the proof of Proposition 6.11. For a finite set $F \subset \Gamma$, let $P \in \mathbb{B}(\ell^2(\Gamma))$ be the finite-rank projection onto the span of $\{\delta_g : g \in F\}$. Rather than compute the norm of $x \in \mathbb{B}(\mathcal{H} \otimes \ell^2(\Gamma))$, we will cut by the (infinite-rank) projections $1 \otimes P$ and show that the norm of the compression is independent of the representation $A \subset \mathbb{B}(\mathcal{H})$ – taking a limit over finite sets in Γ, we conclude the same for x.

Let $\{e_{p,q}\}_{p,q \in F}$ be the canonical matrix units of $P\mathbb{B}(\ell^2(\Gamma))P \cong \mathbb{M}_F(\mathbb{C})$ and fix some arbitrary elements $a \in A$ and $s \in \Gamma$. Let $\pi \colon A \to \mathbb{B}(\mathcal{H} \otimes \ell^2(\Gamma))$ be a regular representation. Our first claim is that

$$(1 \otimes P)\pi(a) = (1 \otimes P)\pi(a)(1 \otimes P) = \sum_{q \in F} \alpha_q^{-1}(a) \otimes e_{q,q}.$$

This is clear if one thinks of $\pi(a)$ as a diagonal matrix in $\mathbb{B}(\bigoplus_{g \in \Gamma} \mathcal{H})$; in the tensor product picture we have

$$\pi(a) = \sum_{q \in \Gamma} \alpha_q^{-1}(a) \otimes e_{q,q},$$

where convergence is in the strong operator topology.

Thus we see that

$$\begin{aligned}(1 \otimes P)\pi(a)(1 \otimes \lambda_s)(1 \otimes P) &= \Big(\sum_{q \in F} \alpha_q^{-1}(a) \otimes e_{q,q}\Big)(1 \otimes P\lambda_s P) \\ &= \Big(\sum_{q \in F} \alpha_q^{-1}(a) \otimes e_{q,q}\Big)\Big(\sum_{p \in F \cap sF} 1 \otimes e_{p,s^{-1}p}\Big) \\ &= \sum_{p \in F \cap sF} \alpha_p^{-1}(a) \otimes e_{p,s^{-1}p} \in A \otimes \mathbb{M}_F(\mathbb{C}).\end{aligned}$$

Now if $x = \sum a_s \lambda_s \in C_c(\Gamma, A) \subset \mathbb{B}(\mathcal{H} \otimes \ell^2(\Gamma))$, then we have

$$(1 \otimes P)x(1 \otimes P) = \sum_{s \in \Gamma} \sum_{p \in F \cap sF} \alpha_p^{-1}(a_s) \otimes e_{p,s^{-1}p} \in A \otimes \mathbb{M}_F(\mathbb{C})$$

and thus the norm of $(1 \otimes P)x(1 \otimes P)$ does not depend on the embedding $A \subset \mathbb{B}(\mathcal{H})$. □

The following description of positive elements is sometimes handy.

Corollary 8.6. *An element $x = \sum_{s \in \Gamma} a_s \lambda_s \in C_c(\Gamma, A)$ is positive in $A \rtimes_{\alpha,r} \Gamma$ if and only if for any finite sequence $s_1, \dots, s_n \in \Gamma$, the operator matrix $[\alpha_{s_i}^{-1}(a_{s_i s_j^{-1}})]_{i,j} \in \mathbb{M}_n(A)$ is positive.*

Proof. Since an operator is positive if and only if its compression by any projection is positive, the result follows from a calculation above:

$$(1 \otimes P)x(1 \otimes P) = \sum_{s \in \Gamma} \sum_{p \in F \cap sF} \alpha_p^{-1}(a_s) \otimes e_{p,s^{-1}p} \in A \otimes \mathbb{M}_F(\mathbb{C}).$$

Indeed, if $F = \{s_1, \dots, s_n\}$, then we can identify this double sum with the operator matrix in the statement of the corollary. (Let $p = s_i$ and $s_j = s^{-1}p$.) □

Here is a C*-dynamical version of Fell's absorption principle, with identical proof.

Proposition 8.7 (Fell's absorbtion principle II)**.** *If $(u, \mathrm{id}_A, \mathcal{H})$ is a covariant representation (i.e., $A \subset \mathbb{B}(\mathcal{H})$ and the action α is spatially implemented in this representation), then the covariant representation*

$$(u \otimes \lambda, \mathrm{id}_A \otimes 1, \mathcal{H} \otimes \ell^2(\Gamma))$$

is unitarily equivalent to a regular representation. In particular, we have a natural $*$-isomorphism

$$C^*\big((u\otimes\lambda)(\Gamma), A\otimes 1\big) \cong A\rtimes_{\alpha,r}\Gamma.$$

Proof. Let $(u, \mathrm{id}_A, \mathcal{H})$ be a covariant representation and define a unitary U on $\mathcal{H}\otimes\ell^2(\Gamma)$ by $U(\xi\otimes\delta_t) = u_t\xi\otimes\delta_t$. One checks that

$$U(1\otimes\lambda_s)U^* = (u_s\otimes\lambda_s) \quad \text{and} \quad U(\sum_t \alpha_t^{-1}(a)\otimes e_{t,t})U^* = a\otimes 1$$

for every $s\in\Gamma$ and $a\in A$. □

We close this section with the existence of conditional expectations. First, a lemma.

Lemma 8.8. *Let ψ be a faithful state on B. Then $\mathrm{id}_A\otimes\psi\colon A\otimes B\to A$ is faithful.*

Proof. Observe that $\{f\otimes g : f\in A^*,\ g\in B^*\}\subset (A\otimes B)^*$ separates the points of $A\otimes B$. Indeed, $A\otimes B\subset\mathbb{B}(\mathcal{H}\otimes\mathcal{K})$ and vector states arising from elementary tensors $h\otimes k\in\mathcal{H}\otimes\mathcal{K}$ separate all of $\mathbb{B}(\mathcal{H}\otimes\mathcal{K})$.

So, if $x\in(A\otimes B)_+$ is nonzero, we can find a state φ on A such that $(\varphi\otimes\mathrm{id}_B)(x)\in B$ is nonzero (and positive). Since ψ is faithful, we have

$$0<\psi((\varphi\otimes\mathrm{id}_B)(x)) = \varphi((\mathrm{id}_A\otimes\psi)(x)),$$

which implies $(\mathrm{id}_A\otimes\psi)(x)$ is nonzero. □

Proposition 8.9. *The map $E\colon C_c(\Gamma, A)\to A$, $E(\sum_s a_s\lambda_s) = a_e$, extends to a faithful conditional expectation from $A\rtimes_{\alpha,r}\Gamma$ onto A. In particular,*

$$\max_{s\in\Gamma}\|a_s\|_A \le \|\sum_{s\in\Gamma} a_s\lambda_s\|_{A\rtimes_{\alpha,r}\Gamma}.$$

Proof. Let $(u, \mathrm{id}_A, \mathcal{H})$ be a covariant representation. By Fell's absorption principle, we may view $A\rtimes_{\alpha,r}\Gamma$ as the C*-algebra generated by $A\otimes 1$ and $(u\otimes\lambda)(\Gamma)$ – in particular, it is a subalgebra of $\mathbb{B}(\mathcal{H})\otimes C^*_r(\Gamma)$. The key observation is that in this representation our map E is nothing but the restriction of $\mathrm{id}_{\mathbb{B}(\mathcal{H})}\otimes\tau$, where τ is the canonical faithful tracial state on $C^*_r(\Gamma)$ (which is clear since $\tau(\lambda_s) = 0$, whenever $s\neq e$). Thus the previous lemma implies that E is faithful.

Finally, note that $a_s = E(z\lambda_s^*)$ for $z = \sum_s a_s\lambda_s$. This implies the asserted inequality, so the proof is complete. □

Remark 8.10. Note that $E\colon A\rtimes_{\alpha,r}\Gamma\to A$ is Γ-*equivariant*: $E(\lambda_s z\lambda_s^{-1}) = \alpha_s(E(z))$ for every $s\in\Gamma$ and $z\in A\rtimes_{\alpha,r}\Gamma$.

More generally, if α and β are actions of Γ on sets X and, respectively, Y, we will say a map $\Phi\colon X\to Y$ is Γ-*equivariant* if $\Phi\circ\alpha_g = \beta_g\circ\Phi$ for all $g\in\Gamma$.

9. Crossed Products by Amenable Groups

Let's construct explicit approximating maps on crossed products by amenable groups. The analysis is a little boorish, but it has been very important for other purposes (e.g. noncommutative entropy theory or calculating Haagerup invariants).

Suppose $A\subset\mathbb{B}(\mathcal{H})$ and $\alpha\colon\Gamma\to\mathrm{Aut}(A)$ is a homomorphism. If Γ is amenable, it turns out that we can easily construct approximating maps by cutting to Følner sets and then mapping back to the crossed product. We will need a few simple lemmas.

Lemma 9.1. *Let A be a C*-algebra and let $n\in\mathbb{N}$. Every positive element in $\mathbb{M}_n(A)$ is a sum of n elements of the form $[a_i^*a_j]_{i,j=1}^n$.*

Proof. Take an arbitrary positive element $x \in \mathbb{M}_n(A)$ and decompose it as a product $x = [b_{ij}]^*[b_{ij}]$. Now one writes

$$[b_{ij}] = A_1 + A_2 + \cdots + A_n$$

and

$$[b_{ij}]^* = A_1^* + A_2^* + \cdots + A_n^*,$$

where

$$A_1 = \begin{bmatrix} b_{11} & b_{12} & \cdots & b_{1n} \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}, A_2 = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ b_{21} & b_{22} & \cdots & b_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix},$$

and so on. Since $A_j^* A_i = 0$ whenever $i \neq j$, a straightforward calculation completes the proof. □

Lemma 9.2. *If A is a Γ-C*-algebra and $F \subset \Gamma$ is a finite set, then for each set $\{a_p\}_{p\in F} \subset A$, the element*

$$\sum_{p,q\in F} \alpha_p(a_p^* a_q)\lambda_{pq^{-1}} \in C_c(\Gamma, A)$$

is positive as an element in $A \rtimes_\alpha \Gamma$ (or $A \rtimes_{\alpha,r} \Gamma$).

Proof. The element in question is equal to $(\sum_{p\in F} a_p\lambda_{p^{-1}})^*(\sum_{p\in F} a_p\lambda_{p^{-1}})$. □

Here are the approximating maps we're after.

Lemma 9.3. *If A is a Γ-C*-algebra and $F \subset \Gamma$ is a finite set, then there exist c.c.p. maps $\varphi\colon A \rtimes_{\alpha,r} \Gamma \to A \otimes \mathbb{M}_F(\mathbb{C})$ and $\psi\colon A \otimes \mathbb{M}_F(\mathbb{C}) \to C_c(\Gamma, A) \subset A \rtimes_{\alpha,r} \Gamma$ such that for all $a \in A$ and $s \in \Gamma$ we have*

$$\psi \circ \varphi(a\lambda_s) = \frac{|F \cap sF|}{|F|} a\lambda_s.$$

Proof. In the proof of Proposition 8.5 we saw that there is a c.c.p. map $\varphi\colon A \rtimes_{\alpha,r} \Gamma \to A \otimes \mathbb{M}_F(\mathbb{C})$ such that

$$\varphi(a\lambda_s) = \sum_{p\in F\cap sF} \alpha_p^{-1}(a) \otimes e_{p,s^{-1}p}.$$

It suffices to prove that $\psi\colon A \otimes \mathbb{M}_F(\mathbb{C}) \to C_c(\Gamma, A) \subset A \rtimes_{\alpha,r} \Gamma$ defined by

$$\psi(a \otimes e_{p,q}) = \frac{1}{|F|}\alpha_p(a)\lambda_{pq^{-1}}$$

is a c.c.p. map, as a simple calculation confirms the asserted formula.

In fact, it suffices to prove that ψ is positive since there is a natural commutative diagram

$$\begin{array}{ccc} \mathbb{M}_n(\mathbb{C}) \otimes (A \otimes \mathbb{M}_F(\mathbb{C})) & \xrightarrow{\cong} & (\mathbb{M}_n(\mathbb{C}) \otimes A) \otimes \mathbb{M}_F(\mathbb{C}) \\ \downarrow & & \downarrow \\ \mathbb{M}_n(\mathbb{C}) \otimes (A \rtimes_{\alpha,r} \Gamma) & \xrightarrow{\cong} & (\mathbb{M}_n(\mathbb{C}) \otimes A) \rtimes_{\tau\otimes\alpha,r} \Gamma. \end{array}$$

By Lemma 9.1, we only need to check that for every set $\{a_p\}_{p\in F} \subset A$, $\psi(\sum a_p^* a_q \otimes e_{p,q}) \geq 0$. But

$$\psi(\sum_{p,q\in F} a_p^* a_q \otimes e_{p,q}) = \sum_{p,q\in F} \frac{1}{|F|}\alpha_p(a_p^* a_q)\lambda_{pq^{-1}},$$

so the previous lemma completes the proof. □

Theorem 9.4. *For any amenable group* Γ *and action* $\alpha\colon \Gamma \to \mathrm{Aut}(A)$*, the following statements hold:*

(1) $A \rtimes_\alpha \Gamma = A \rtimes_{\alpha,r} \Gamma$*;*
(2) A *is nuclear if only if* $A \rtimes_\alpha \Gamma$ *is nuclear.*

Proof. Proof of (1): It suffices to show that there exist c.c.p. maps

$$\Psi_n\colon A \rtimes_{\alpha,r} \Gamma \to A \rtimes_\alpha \Gamma$$

such that $\|x - \Psi_n \circ \pi(x)\|_u \to 0$ for all $x \in C_c(\Gamma, A) \subset A \rtimes_\alpha \Gamma$, where

$$\pi\colon A \rtimes_\alpha \Gamma \to A \rtimes_{\alpha,r} \Gamma$$

is the canonical quotient map (coming from universality).

The key observation is that the proof of Lemma 9.3 is algebraic. In other words, if $F_n \subset \Gamma$ is a Følner sequence and φ_n, ψ_n are the corresponding maps constructed in Lemma 9.3, then we can define c.c.p. maps $\Psi_n\colon A \rtimes_{\alpha,r} \Gamma \to A \rtimes_\alpha \Gamma$ by $\Psi_n = \psi_n \circ \varphi_n$, but simply regarding the ψ_n's as taking values in the universal crossed product (as opposed to the reduced one, since the range of ψ_n is contained in $C_c(\Gamma, A)$). The formula in Lemma 9.3 still holds, and hence for $x = \sum_{k\in\Gamma} a_k k \in C_c(\Gamma, A)$ we have

$$\|x - \Psi_n(\pi(x))\|_{A\rtimes_\alpha\Gamma} = \|\sum_{k\in\Gamma}(1 - \frac{|F_n \cap (k + F_n)|}{|F_n|})a_k k\|_{A\rtimes_\alpha\Gamma} \to 0$$

since only finitely many a_k's are nonzero.

Proof of (2): The "if" direction is trivial since there is a conditional expectation $A \rtimes_\alpha \Gamma \to A$. For the converse, note that another way of stating Lemma 9.3 is this: there exist c.c.p. maps $\varphi_n\colon A \rtimes_\alpha \Gamma \to A \otimes \mathbb{M}_{k(n)}(\mathbb{C})$ and $\psi_n\colon A \otimes \mathbb{M}_{k(n)}(\mathbb{C}) \to A \rtimes_\alpha \Gamma$ such that $\psi_n \circ \varphi_n \to \mathrm{id}$ in the point-norm topology. Since $A \otimes \mathbb{M}_{k(n)}(\mathbb{C})$ is nuclear whenever A is nuclear, the remainder of the proof is trivial. □

10. Amenable Actions

We now step up the generality ladder and consider crossed products by amenable actions – i.e., the group involved need not be amenable, but we require it to act nicely. When defined "appropriately" (not the definition usually found in the literature, but an equivalent one that makes our present work easier) and done abstractly, finding approximating maps on a crossed product by an amenable action is only slightly harder than the last section.

Given a Γ-C*-algebra A, we put a third norm on the (α-twisted) convolution algebra $C_c(\Gamma, A)$: for finitely supported functions $S, T\colon \Gamma \to A$ we define

$$\langle S, T\rangle = \sum S(g)^* T(g) \in A$$

and

$$\|S\|_2 = \|\langle S, S\rangle\|^{1/2}.$$

The informed reader will notice that we have made a Hilbert C*-module. The Cauchy-Schwarz inequality holds in this context: $\|\langle S, T\rangle\|_A \leq \|S\|_2 \|T\|_2$, for all $S, T \in C_c(\Gamma, A)$.[17]

Definition 10.1. An action $\alpha\colon \Gamma \to \mathrm{Aut}(A)$ on a unital C*-algebra A is *amenable* if there exist finitely supported functions $T_i\colon \Gamma \to A$ with the following properties:

(1) $0 \leq T_i(g) \in \mathcal{Z}(A)$ (the center of A) for all $i \in \mathbb{N}$ and $g \in \Gamma$;

[17]This is a general fact about Hilbert modules, but here we only need the case that A is abelian. If $A = C(X)$, the asserted inequality follows from the usual Cauchy-Schwarz inequality, applied pointwise in X.

(2) $\langle T_i, T_i\rangle = \sum_{g\in\Gamma} T_i(g)^2 = 1_A$;
(3) $\|s *_\alpha T_i - T_i\|_2 \to 0$ for all $s \in \Gamma$, where $s \in C_c(\Gamma, A)$ is the function which sends $s \mapsto 1_A$ and all other group elements to zero.[18]

The functions T_i will replace the Følner sets we used in the previous section.

Lemma 10.2. *Let A be a Γ-C*-algebra and $T\colon \Gamma \to A$ be a finitely supported function such that $0 \le T(g) \in \mathcal{Z}(A)$ for all $g \in \Gamma$ and $\sum_g T(g)^2 = 1_A$. Then,*

(1) $T *_\alpha T^*(s) = \sum\limits_{p\in F\cap sF} T(p)\alpha_s(T(s^{-1}p))$, *where F is the support of T, and*
(2) $\|1_A - T *_\alpha T^*(s)\| \le \|T - s *_\alpha T\|_2$, *for all $s \in \Gamma$.*

Proof. Statement (1) is a trivial calculation, using the fact that $T(g)^* = T(g)$ for all $g \in \Gamma$.

To prove the second statement, we first note that $s *_\alpha T(p) = \alpha_s(T(s^{-1}p))$ for all $p \in \Gamma$. Now we compute

$$\begin{aligned} 1_A - T *_\alpha T^*(s) &= \sum_{p\in\Gamma} T(p)^2 - \sum_{p\in\Gamma} T(p)\alpha_s(T(s^{-1}p)) \\ &= \sum_{p\in\Gamma} T(p)\Big(T(p) - \alpha_s(T(s^{-1}p))\Big) \\ &= \langle T, T - s *_\alpha T\rangle. \end{aligned}$$

Hence the desired inequality follows from the Cauchy-Schwarz inequality, since $\|T\|_2 = 1$. $\square$

Here is the analogue of Lemma 9.3 for crossed products by amenable actions.

Lemma 10.3. *Let A be a unital Γ-C*-algebra and $T\colon \Gamma \to A$ be a finitely supported function with support F, such that $0 \le T(g) \in \mathcal{Z}(A)$ for all $g \in \Gamma$ and $\sum_g T(g)^2 = 1_A$. Then, there exist u.c.p. maps $\varphi\colon A \rtimes_{\alpha,r} \Gamma \to A \otimes \mathbb{M}_F(\mathbb{C})$ and $\psi\colon A \otimes \mathbb{M}_F(\mathbb{C}) \to A \rtimes_{\alpha,r} \Gamma$ such that for all $s \in \Gamma$ and $a \in A$,*

$$\psi \circ \varphi(a\lambda_s) = \big(T *_\alpha T^*(s)\big)a\lambda_s.$$

Proof. We already have a u.c.p. compression map $\varphi\colon A \rtimes_{\alpha,r} \Gamma \to A \otimes \mathbb{M}_F(\mathbb{C})$ such that

$$\varphi(a\lambda_s) = \sum_{p\in F\cap sF} \alpha_p^{-1}(a) \otimes e_{p,s^{-1}p} \in A \otimes \mathbb{M}_F(\mathbb{C}).$$

Define

$$X = \sum_{p\in F} \alpha_p^{-1}(T(p)) \otimes e_{p,p}$$

and note that $X = X^*$. Hence compression by X is a c.p. map and a computation confirms that

$$X\varphi(a\lambda_s)X = \sum_{p\in F\cap sF} \alpha_p^{-1}(T(p)a)\alpha_{s^{-1}p}^{-1}(T(s^{-1}p)) \otimes e_{p,s^{-1}p}.$$

We know the map $A \otimes M_F(\mathbb{C}) \to A \rtimes_{\alpha,r} \Gamma$ defined by

$$b \otimes e_{x,y} \mapsto \alpha_x(b)\lambda_{xy^{-1}}$$

is u.c.p., so we get another u.c.p. map $\psi\colon A \otimes \mathbb{M}_F(\mathbb{C}) \to A \rtimes_{\alpha,r} \Gamma$ by composing it with compression by X:

$$\psi\colon A \otimes \mathbb{M}_F(\mathbb{C}) \stackrel{X\cdot X}{\to} A \otimes \mathbb{M}_F(\mathbb{C}) \stackrel{b\otimes e_{x,y}\mapsto\alpha_x(b)\lambda_{xy^{-1}}}{\longrightarrow} A \rtimes_{\alpha,r} \Gamma.$$

[18]This definition comes from the characterization of (classical) amenability in terms of weak containment of the trivial representation in the left regular representation.

Finally, since $T(g) \in \mathcal{Z}(A)$ for all $g \in \Gamma$, we have

$$\begin{aligned}\psi \circ \varphi(a\lambda_s) &= \sum_{p \in F \cap sF} \alpha_p\Big(\alpha_p^{-1}(T(p)a)\alpha_{s^{-1}p}^{-1}(T(s^{-1}p))\Big)\lambda_s \\ &= \Big(\sum_{p \in F \cap sF} T(p)\alpha_s(T(s^{-1}p))\Big)a\lambda_s \\ &= \big(T *_\alpha T^*(s)\big)a\lambda_s.\end{aligned}$$

□

The proof of the next theorem is but a tiny perturbation of that given for Theorem 9.4. We leave the details to the reader.

Theorem 10.4. *For any amenable action of Γ on A, the following statements hold:*

(1) $A \rtimes_\alpha \Gamma = A \rtimes_{\alpha,r} \Gamma$;
(2) *A is nuclear if only if $A \rtimes_\alpha \Gamma$ is nuclear.*

We call a compact[19] space X a Γ-*space* if it is equipped with an action of Γ (by homeomorphisms). Let $x \mapsto s.x$ denote the action of $s \in \Gamma$ on $x \in X$. To help distinguish, we let $\alpha_s\colon C(X) \to C(X)$ denote the induced automorphism of $C(X)$ (i.e., $\alpha_s(f)(x) = f(s^{-1}.x)$). The notion of an amenable action comes from classical (i.e., abelian) dynamical systems. As already mentioned, our definition at the C*-level is not very common in the literature. Here is a more popular version.

Definition 10.5. An action of Γ on a compact space X is called (topologically) *amenable* (or, equivalently, X is an *amenable* Γ-*space*) if there exists a net of continuous maps $m_i\colon X \to \mathrm{Prob}(\Gamma)$, such that for each $s \in \Gamma$,

$$\lim_{i\to\infty}\Big(\sup_{x\in X}\|s.m_i^x - m_i^{s.x}\|_1\Big) = 0,$$

where $s.m_i^x(g) = m_i^x(s^{-1}g)$.[20]

Remark 10.6. Let $\mathrm{Prob}(X)$ be the set of all regular Borel probability measures on X. In Proposition 14.1 we will show for a countable group Γ that amenability can be reformulated as: For any finite subset $E \subset \Gamma$, $\varepsilon > 0$ and any $m \in \mathrm{Prob}(X)$, there exists a Borel map $\mu\colon X \to \mathrm{Prob}(\Gamma)$ (i.e., the function $X \to \mathbb{R}$, $x \mapsto \mu^x(t)$, is Borel for every $t \in \Gamma$) such that

$$\max_{s\in E}\int_X \|s.\mu^x - \mu^{s.x}\|_1\, dm(x) < \varepsilon.$$

Lemma 10.7. *An action $\alpha\colon \Gamma \to \mathrm{Homeo}(X)$ is amenable if and only if the induced action on $C(X)$ is amenable in the sense of Definition 10.1.*

Proof. The proofs of both directions are similar. First assume the action is amenable in the sense of Definition 10.5. Let $m_i\colon X \to \mathrm{Prob}(\Gamma)$ be a sequence of continuous maps such that for each $s \in \Gamma$,

$$\lim_{i\to\infty}\big(\sup_{x\in X}\|s.m_i^x - m_i^{s.x}\|_1\big) = 0.$$

Define $S_i\colon \Gamma \to C(X)$ by

$$S_i(g)(x) = m_i^x(g).$$

[19]As usual, compactness includes the Hausdorff axiom.

[20]By definition, $\mathrm{Prob}(\Gamma)$ is the set of probability measures on Γ – which we identify with the set of positive, norm-one elements in $\ell^1(\Gamma)$. Continuity means with respect to the restriction of the weak-$*$ topology on $\ell^1(\Gamma)$. In other words, $m\colon X \to \mathrm{Prob}(\Gamma)$ is continuous if and only if for each convergent net $x_i \to x \in X$ we have $m^{x_i}(g) \to m^x(g)$ for all $g \in \Gamma$.

Then for each $x \in X$ we have

$$\sum_g S_i(g)(x) = \sum_g m_i^x(g) = 1,$$

since m_i^x is a probability measure. Defining $\tilde{T}_i(g) = \sqrt{S_i(g)}$, it follows that for each i,

$$\langle \tilde{T}_i, \tilde{T}_i \rangle = \sum_g \tilde{T}_i(g)^2 = 1_{C(X)}.^{21}$$

Of course, the $\tilde{T}_i$'s are not finitely supported (we will fix that later) but note that for each $x \in X$,

$$(s *_\alpha \tilde{T}_i)(g)(x) = \alpha_s\big(\tilde{T}_i(s^{-1}g)\big)(x) = \tilde{T}_i(s^{-1}g)(s^{-1}.x) = \sqrt{s.m_i^{s^{-1}.x}(g)}.$$

Using the inequality $(a-b)^2 \le |a^2 - b^2|$ for all positive numbers a, b, we then get

$$\begin{aligned}
\|s *_\alpha \tilde{T}_i - \tilde{T}_i\|_2^2 &= \sup_{x \in X} \Big(\sum_{g \in \Gamma} |\sqrt{s.m_i^{s^{-1}.x}(g)} - \sqrt{m_i^x(g)}|^2 \Big) \\
&\le \sup_{x \in X} \Big(\sum_{g \in \Gamma} |s.m_i^{s^{-1}.x}(g) - m_i^x(g)| \Big) \\
&\overset{x=s.y}{=} \sup_{y \in X} \Big(\sum_{g \in \Gamma} |s.m_i^y(g) - m_i^{s.y}(g)| \Big) \\
&= \sup_{y \in X} \|s.m_i^y - m_i^{s.y}\|_1 \to 0.
\end{aligned}$$

Hence the $\tilde{T}_i$'s have the right properties, except for finite support. Fixing this problem is easy once we prove the following claim.

Claim. If $T \colon \Gamma \to C(X)$ is a positive function such that $\langle T, T \rangle = 1_{C(X)}$, then there exists a sequence of *finitely supported* positive functions $T_n \colon \Gamma \to C(X)$ such that $\langle T_n, T_n \rangle = 1_{C(X)}$ for all n and

$$\|s *_\alpha T_n - T_n\|_2 \to \|s *_\alpha T - T\|_2,$$

for all $s \in \Gamma$.

To prove this claim, we let $F_n \subset F_{n+1}$ be a sequence of finite subsets of Γ such that $\bigcup F_n = \Gamma$. Since

$$\sum_{g \in \Gamma} T(g)^2 = 1_{C(X)}$$

and convergence is uniform, it follows that for all sufficiently large n,

$$\sum_{g \in F_n} T(g)^2 > 0,$$

meaning bounded uniformly away from 0. Hence we can define T_n by declaring

$$T_n(g) = \sqrt{\frac{1}{\sum_{g \in F_n} T(g)^2}} T(g),$$

for all $g \in F_n$ and $T_n(g) = 0$ if $g \notin F_n$. Tedious and unenlightening calculations (left to the diligent few) show that these functions do the trick.

[21]Note that Dini's Theorem implies this sum converges uniformly, since everything is positive.

To prove the opposite direction of Lemma 10.7, one basically reverses the procedure above. That is, define $m_i^x(g) = T_i(g)^2(x)$ and calculate away. It should be noted that the Cauchy-Schwarz inequality gets used in the following way:

$$\sum |a_i^2 - b_i^2| = \sum |a_i - b_i|(a_i + b_i) \leq \|(a_i) - (b_i)\|_2 \|(a_i) + (b_i)\|_2.$$

□

The proof of this lemma shows that if X is an amenable Γ-space, then we can assume each map $m_i \colon X \to \mathrm{Prob}(\Gamma)$ has the property that there exists a finite set $F_i \subset \Gamma$ with $\mathrm{supp}(m_i^x) \subset F_i$, for every $x \in X$. Here is a direct proof of this fact.

Lemma 10.8. *Let $m \colon X \to \mathrm{Prob}(\Gamma)$ be a continuous map. Then, for any $\varepsilon > 0$, there exist $\tilde{m} \colon X \to \mathrm{Prob}(\Gamma)$ and a finite subset $F \subset \Gamma$ such that $\mathrm{supp}\, \tilde{m}^x \subset F$ and $\|m^x - \tilde{m}^x\|_1 < \varepsilon$ for all $x \in X$.*

Proof. For every finite subset $F \subset \Gamma$, let $U(F) = \{x \in X : \|m^x \chi_F\|_1 > 1 - \varepsilon/2\} \subset X$, where χ_F is the characteristic function of F. It is easily seen that $\{U(F)\}_F$ is an open cover of X which is upward directed. Since X is compact, there exists F such that $X = U(F)$. It follows that $\tilde{m}^x = m^x \chi_F + \|m^x \chi_{\Gamma \setminus F}\|_1 \delta_e$ has the desired property. □

Exercises

Exercise 10.1. Assume $\Gamma \times \Gamma$ acts on $C(X)$ and there exist two $\Gamma \times \Gamma$-invariant subalgebras $A, B \subset C(X)$ such that (a) $\Gamma \times \{e\}|_A$ is amenable while $\{e\} \times \Gamma|_A$ is trivial and (b) $\Gamma \times \{e\}|_B$ is trivial while $\{e\} \times \Gamma|_B$ is amenable. Prove that the action of $\Gamma \times \Gamma$ on $C(X)$ is amenable. (In addition to helping cement Definition 10.5 in your mind, this exercise will be needed later; see Corollary 16.4.)

11. The Uniform Roe Algebra

Let Γ be a discrete group and $E \subset \Gamma$ be a finite subset. The *tube of width* E is the subset $\mathrm{Tube}(E)$ in $\Gamma \times \Gamma$ given by

$$\mathrm{Tube}(E) = \{(s,t) \in \Gamma \times \Gamma : st^{-1} \in E\}.^{22}$$

By the generic term *tube* we mean a tube of width E for some finite subset $E \subset \Gamma$. The *uniform algebra* (or the *uniform Roe algebra*) $C_u^*(\Gamma)$ of Γ is the C*-subalgebra of $\mathbb{B}(\ell^2(\Gamma))$ generated by $C_\lambda^*(\Gamma)$ and $\ell^\infty(\Gamma)$. Thinking of operators in $\mathbb{B}(\ell^2(\Gamma))$ as infinite matrices indexed by Γ, it is instructive to convince yourself of the following fact: $x = [x_{s,t}]_{s,t \in \Gamma} \in \mathbb{B}(\ell^2(\Gamma))$ belongs to the $*$-algebra generated by $\lambda(\mathbb{C}[\Gamma])$ and $\ell^\infty(\Gamma)$ if and only if x is supported in a tube (i.e., there exists a finite set $E \subset \Gamma$ such that $x_{s,t} = 0$ whenever $(s,t) \notin \mathrm{Tube}(E)$).[23]

It turns out that the uniform Roe algebra is an old friend incognito.

Proposition 11.1. *Let $\alpha \colon \Gamma \to \mathrm{Aut}(\ell^\infty(\Gamma))$ be the left translation action. Then*

$$C_u^*(\Gamma) \cong \ell^\infty(\Gamma) \rtimes_{\alpha,r} \Gamma.$$

[22] One should be careful about st^{-1} and $s^{-1}t$. We use here the *right* invariant tube so that $\lambda(s)$ is supported on a tube. However, when we deal with the Cayley graph later, we use the *left* invariant metric to make the left multiplication action isometric.

[23] If you aren't familiar with this point of view, it is good to start with $\ell^\infty(\Gamma)$; all of these operators are supported in $\mathrm{Tube}(\{e\})$. Next consider an element from the group ring $\lambda(\mathbb{C}[\Gamma])$. Such an operator is "constant down the diagonals," so which tube is it supported in?

Proof. We may apply the construction of $\ell^\infty(\Gamma) \rtimes_{\alpha,r} \Gamma$ to any faithful representation of $\ell^\infty(\Gamma)$, so we start with the canonical inclusion $\ell^\infty(\Gamma) \subset \mathbb{B}(\ell^2(\Gamma))$.

Define a unitary $U \colon \ell^2(\Gamma) \otimes \ell^2(\Gamma) \to \ell^2(\Gamma) \otimes \ell^2(\Gamma)$ by $U(\delta_x \otimes \delta_y) = \delta_x \otimes \delta_{yx}$. Now we compute

$$\begin{aligned} U\pi(f)(\delta_s \otimes \delta_t) &= U((\alpha_t^{-1}(f)\delta_s) \otimes \delta_t) \\ &= U((f(ts)\delta_s) \otimes \delta_t) \\ &= f(ts)\delta_s \otimes \delta_{ts} \\ &= \delta_s \otimes (f(ts)\delta_{ts}) \\ &= (1 \otimes f)(U(\delta_s \otimes \delta_t)). \end{aligned}$$

It follows that $U\pi(f)U^* = 1 \otimes f$ for all $f \in \ell^\infty(\Gamma)$. A similar calculation shows that U commutes with $1 \otimes \lambda_g$ for all $g \in \Gamma$ and hence

$$U(\ell^\infty(\Gamma) \rtimes_{\alpha,r} \Gamma)U^* = \mathbb{C}1 \otimes C^*(\ell^\infty(\Gamma), C^*_\lambda(\Gamma)) \cong C^*_u(\Gamma).$$

□

Note that $\ell^\infty(\Gamma) \rtimes_{\alpha,r} \Gamma$ is universal in the following sense: if X is a compact Hausdorff space with homeomorphic Γ-action β, then there is a covariant homomorphism $C(X) \to \ell^\infty(\Gamma)$. To see this simply pick a point $x \in X$, consider the orbit $\{\beta_g(x) : g \in \Gamma\}$, and define a homomorphism $C(X) \to \ell^\infty(\Gamma)$ by restriction: $f \mapsto (f(\beta_g(x)))_{g \in \Gamma}$. It is an easy exercise to check that if β is an amenable action, then translation on $\ell^\infty(\Gamma)$ is also amenable – simply map the functions $T \colon \Gamma \to C(X)$ over to $\ell^\infty(\Gamma)$ using the covariant map $C(X) \to \ell^\infty(\Gamma)$. It follows that Γ *admits an amenable action on some compact space if and only if the left translation action on* $\ell^\infty(\Gamma)$ *is amenable.*

Corollary 11.2. *If* Γ *admits an amenable action on some compact Hausdorff space, then* $C^*_u(\Gamma)$ *is nuclear.*

12. Exact C*-algebras and Solid von Neumann Algebras

Though I won't get into the details, there is an important connection with exactness that must be mentioned.

Definition 12.1. A C*-algebra A is *exact* if there exists a faithful, nuclear $*$-representation $\pi \colon A \to \mathbb{B}(\mathcal{H})$. A discrete group Γ is exact if $C^*_\lambda(\Gamma)$ is exact.

As with nuclearity, this is not the historically correct definition. A deep theorem of Kirchberg states that the original definition (which involved tensor products and short exact sequences) is equivalent to the definition above.

The following result connects exactness with amenable actions. It is due to Ozawa (and, independently, Anantharaman-Delaroche), following an important contribution by Guentner and Kaminker. See [2, Chapter 5] for details.

Theorem 12.2. *For a discrete group* Γ*, the following are equivalent:*

1. Γ *is exact;*
2. $C^*_u(\Gamma)$ *is nuclear;*
3. Γ *acts amenably on some compact topological space.*

Moreover, if Γ *is countable and exact, then* Γ *acts amenably on a compact metrizable space.*

It turns out that exact C*-algebras enjoy an important approximation property that isn't as well known as it should be.

Definition 12.3. A C*-algebra A is *locally reflexive* if for every finite-dimensional operator system[24] $E \subset A^{**}$, there exists a net of c.c.p. maps $\varphi_i \colon E \to A$ which converges to id_E in the point-ultraweak topology.

The following result is due to Kirchberg, but depends on some older tensor product work of Archbold and Batty. See [2, Chapter 9] for details.

Theorem 12.4. *Exact* C*-*algebras are locally reflexive.*

Though it's technical and hard (at this point) to see the significance of, let's prove an important lemma which illustrates the usefulness of local reflexivity.

Lemma 12.5. *Let $M \subset \mathbb{B}(\mathcal{H})$ be a von Neumann algebra which contains a weakly dense exact C^*-algebra $B \subset M$. Assume $N \subset M$ is a von Neumann subalgebra with a weakly continuous conditional expectation $\Phi \colon M \to N$, such that there exists a u.c.p. map $\Psi \colon \mathbb{B}(\mathcal{H}) \to M$ with the property that $\Phi|_B = \Psi|_B$. Then N is injective.*

Proof. Let $E \subset N$ be a finite-dimensional operator system and $\varphi_n \colon E \to B$ be contractive c.p. maps converging to id_E in the point-ultraweak topology. By Arveson's Extension Theorem, we may assume each φ_n is defined on all of $\mathbb{B}(\mathcal{H})$ (and now takes values in $\mathbb{B}(\mathcal{H})$). Then one readily checks that $\Phi \circ \Psi \circ \varphi_n \colon \mathbb{B}(\mathcal{H}) \to N$ are u.c.p. maps with the property that $\Phi \circ \Psi \circ \varphi_n(x) \to x$ ultraweakly for all $x \in E$ (since $\Phi \circ \Psi \circ \varphi_n(x) = \Phi(\varphi_n(x))$ for all $x \in E$). Taking a cluster point in the point-ultraweak topology we get a u.c.p. map $\theta_E \colon \mathbb{B}(\mathcal{H}) \to N$ which restricts to the identity on E. Taking another cluster point of the maps θ_E (over all finite-dimensional operator systems $E \subset N$) we get the desired conditional expectation $\mathbb{B}(\mathcal{H}) \to N$. □

Specializing to group von Neumann algebras and applying The Trick we get:

Lemma 12.6. *Assume Γ is exact and let $N \subset L(\Gamma)$ be a von Neumann sublagebra with trace-preserving conditional expectation $\Phi \colon L(\Gamma) \to N$. If*

$$\Phi|_{C^*_\lambda(\Gamma)} \times \mathrm{id}_{C^*_\rho(\Gamma)} \colon C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))$$

is $\otimes$-continuous, then N is injective.

Proof. Since $C^*_\lambda(\Gamma) \otimes C^*_\rho(\Gamma) \subset \mathbb{B}(\ell^2(\Gamma)) \otimes C^*_\rho(\Gamma)$, The Trick applied to $\Phi|_{C^*_\lambda(\Gamma)} \times \mathrm{id}_{C^*_\rho(\Gamma)}$ yields a u.c.p. map $\Psi \colon \mathbb{B}(\ell^2(\Gamma)) \to L(\Gamma)$ such that $\Psi|_{C^*_\lambda(\Gamma)} = \Phi|_{C^*_\lambda(\Gamma)}$. Hence Lemma 12.5 completes the proof. □

Definition 12.7. A von Neuman algebra M is called *solid* if the relative commutant of every diffuse von Neumann subalgebra is injective.

The consequences of this definition will be discussed later, but now is an appropriate time to prove a celebrated theorem of Ozawa.

Theorem 12.8 (Ozawa, 2003)**.** *Assume Γ is exact and the canonical map $C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma)) \to \mathbb{B}(\ell^2(\Gamma))/\mathbb{K}(\ell^2(\Gamma))$ is $\otimes$-continuous. Then $L(\Gamma)$ is solid.*

Proof. Let $M \subset L(\Gamma)$ be diffuse and $A \subset M$ be a masa (in M). Since $A' \cap L(\Gamma) \supset M' \cap L(\Gamma)$ and there is a conditional expectation $A' \cap L(\Gamma) \to M' \cap L(\Gamma)$, it suffices to prove $A' \cap L(\Gamma)$ is injective. Since M is diffuse, A is non-atomic – i.e., $A \cong L^\infty[0,1]$. Hence we can find a generating unitary $u \in A$ such that $u^n \to 0$ ultraweakly.

[24]That is, E is a self-adjoint linear subspace containing the unit of A^{**}.

Let $N = A' \cap L(\Gamma)$ and $\Phi\colon L(\Gamma) \to N$ be the unique trace-preserving conditional expectation. As is well known, we can define a conditional expectation $\Psi\colon \mathbb{B}(\ell^2(\Gamma)) \to A' \cap \mathbb{B}(\ell^2(\Gamma))$ by taking a cluster point of the maps

$$\varphi_n(T) := \frac{1}{n}\sum_{i=1}^{n} u^i T u^{-i}.$$

Evidently $\Psi|_{L(\Gamma)}$ is a trace-preserving conditional expectation of $L(\Gamma)$ onto N; hence, by uniqueness, $\Psi|_{L(\Gamma)} = \Phi$. Moreover, since $C^*_\rho(\Gamma) \subset A' \cap \mathbb{B}(\ell^2(\Gamma))$ we have that $C^*_\rho(\Gamma)$ lies in the multiplicative domain of Ψ. Thus,

$$\Psi(\sum_{j=1}^{k} x_j y_j) = \sum_{j=1}^{k} \Phi(x_j) y_j,$$

for all $x_j \in L(\Gamma)$ and $y_j \in C^*_\rho(\Gamma)$. In particular,

$$\Phi|_{C^*_\lambda(\Gamma)} \times \mathrm{id}_{C^*_\rho(\Gamma)}(x \otimes y) = \Phi(x)y = \Psi(xy),$$

for all $x \in C^*_\lambda(\Gamma)$ and $y \in C^*_\rho(\Gamma)$.

By Lemma 12.6, to prove N is injective, it suffices to show that

$$\Phi|_{C^*_\lambda(\Gamma)} \times \mathrm{id}_{C^*_\rho(\Gamma)}\colon C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))$$

is $\otimes$-continuous. However, by the previous paragraph, and our assumption that

$$C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma)) \to \mathbb{B}(\ell^2(\Gamma))/\mathbb{K}(\ell^2(\Gamma))$$

is $\otimes$-continuous, it suffices to show that Ψ contains $\mathbb{K}(\ell^2(\Gamma))$ in its kernel (since this means it factors through the Calkin algebra). But this is a routine exercise, so the proof is complete. □

Of course, Ozawa's theorem would be of little interest if it didn't apply to any examples. So our next task is to provide such examples. Though it isn't easy, we'll eventually see that all hyperbolic groups satisfy the hypotheses of Theorem 12.8. Before that, however, let's give a simple consequence of the definition of solidity. A II_1-factor M is said to be *prime* if it can't be written as the tensor product of two II_1-factors.

Proposition 12.9. *If M is a solid II_1-factor and $N \subset M$ is a non-injective subfactor, then N is prime.*

Proof. Assume N is not prime. Then we can write $N = N_1 \bar{\otimes} N_2$ where N_1 is diffuse and N_2 is not injective. By definition, the relative commutant of N_1 (in M) is injective. But this commutant contains N_2, which is a contradiction (since there is a conditional expectation $N_1' \cap M \to N_2$). □

13. The Case of Free Groups

There are several different ways of proving exactness for free groups. Here is a straightforward proof based on measures.

Proposition 13.1. *Free groups are exact.*

Proof. It suffices to show $\mathbb{F}_2$ is exact. For $t \in \mathbb{F}_2$ of length l with reduced form $t = t_1 \cdots t_l$, we denote $t(k) = t_1 \cdots t_k$ for $k \le l$ and $t(k) = t$ for $k > l$. Fix $N \in \mathbb{N}$ and define $m\colon \mathbb{F}_2 \to \mathrm{Prob}(\mathbb{F}_2)$ by

$$m_t = \frac{1}{N}\sum_{k=0}^{N-1} \delta_{t(k)}.$$

Clearly, all m_t are supported on the finite set of elements in $\mathbb{F}_2$ with length $< N$. An instructive calculation left to the reader confirms that $\|s.m_t - m_{st}\| \leq 2|s|/N$ for every $s, t \in \mathbb{F}_2$. Letting $N \to \infty$ completes the proof. $\square$

The proof above is more transparent when viewed geometrically in the Cayley graph of $\mathbb{F}_2$. Moreover, a geometric point of view will be crucial in the next two sections, so let's develop our intuition by proving that free groups act amenably on their ideal boundaries.

Fix $r \in \mathbb{N}$ and let $\mathbb{F}_r = \langle g_1, \ldots, g_r \rangle$ be the rank-r free group (think of the $r = 2$ case for now). Then, its *ideal boundary* is the set

$$\partial\mathbb{F}_r = \{(a_i) \in \prod_{\mathbb{N}} \{g_1, g_1^{-1}, \ldots, g_r, g_r^{-1}\} : \forall i \in \mathbb{N}, a_{i+1} \neq a_i^{-1}\}.$$

The complement of $\partial\mathbb{F}_r$ (in $\prod_{\mathbb{N}}\{g_1, g_1^{-1}, \ldots, g_r, g_r^{-1}\}$) is easily seen to be open in the product topology; hence $\partial\mathbb{F}_r$ is compact. For geometric intuition, it is better to identify $\partial\mathbb{F}_r$ with the set of infinite paths in the Cayley graph of $\mathbb{F}_r$ which start at the neutral element. Indeed, given $x = (x_i) \in \partial\mathbb{F}_r$, we first think of x as the infinite word $x_1x_2x_3\cdots$ (note that this is in reduced form, since no cancellation occurs); then we identify this word with the path determined by the sequence of vertices $\{x_1, x_1x_2, x_1x_2x_3, \ldots\}$ in the Cayley graph of $\mathbb{F}_r$. Thinking of $\partial\mathbb{F}_r$ as infinite reduced words, it is easy to see that $\mathbb{F}_r$ acts continuously on $\partial\mathbb{F}_r$ by left multiplication (and rectifying possible cancellation).

For $x \in \partial\mathbb{F}_r$ with reduced word form $x = x_1x_2\cdots$, we set $x(0) = e$ and $x(k) = x_1\cdots x_k$ for all $k > 0$. Fix $N \in \mathbb{N}$ and define a continuous map $\mu\colon \partial\mathbb{F}_r \to \mathrm{Prob}(\mathbb{F}_r)$ by

$$\mu^x = \frac{1}{N}\sum_{k=0}^{N-1} \delta_{x(k)}.$$

Looking at the Cayley graph in Figure 1, μ^x is just the normalized characteristic function of the first N steps along the infinite path determined by x. Here's an important observation/exercise: For each $s \in \mathbb{F}_r$ and $x \in \partial\mathbb{F}_r$, there exists a unique geodesic path starting at s and eventually merging with the path determined by $s.x \in \partial\mathbb{F}_r$ (see Figure 2); moreover, $s.\mu^x$ is just the normalized characteristic function of the first N steps along this geodesic.

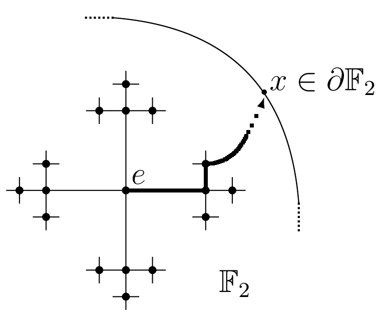

FIGURE 13.1. The Cayley graph of $\mathbb{F}_2$ and the boundary $\partial\mathbb{F}_2$

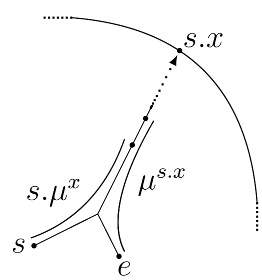

FIGURE 13.2. Amenability of $\mathbb{F}_2$ acting on $\partial\mathbb{F}_2$

With this geometric picture in mind, one checks that $\|s.\mu^x - \mu^{s.x}\| \leq 2|s|/N$ for all $s \in \mathbb{F}_r$ and $x \in \partial\mathbb{F}_r$. Letting $N \to \infty$, this shows that $\mathbb{F}_r$ acts amenably on its ideal boundary.

14. Groups Acting on Trees

For any Γ-space K, we denote the stabilizer subgroup of $a \in K$ by $\Gamma^a = \{s \in \Gamma : s.a = a\}$. Our goal here is to show that a group acting on a tree is exact whenever all the vertex stabilizers are exact. This gives an alternate proof of the fact that amalgamated free products of exact groups are exact, since $\Gamma = \Gamma_1 *_\Lambda \Gamma_2$ acts on a tree in such a way that the vertex stabilizers are conjugates of Γ_1 or Γ_2.

The flexibility provided by Borel maps makes the following result very useful.

Proposition 14.1. *Let Γ be a countable group, X a compact Γ-space and K a countable Γ-space. Assume that the stabilizer subgroups Γ^a are exact, for all $a \in K$, and that there exists a net of Borel maps $\zeta_n \colon X \to \mathrm{Prob}(K)$ (i.e., the function $X \ni x \mapsto \zeta_n^x(a) \in \mathbb{R}$ is Borel for every $a \in K$) such that*

$$\lim_n \int_X \|s.\zeta_n^x - \zeta_n^{s.x}\| \, dm(x) = 0$$

for every $s \in \Gamma$ and every regular Borel probability measure m on X. Then Γ is exact.

Moreover, if X is amenable as a Γ^a-space for every $a \in K$, then X is an amenable Γ-space.

Proof. We first claim that for every $\varepsilon > 0$ and finite subset $E \subset \Gamma$, there exists a continuous map $\eta \colon X \to \mathrm{Prob}(K)$ such that

$$\max_{s \in E} \sup_{x \in X} \|s.\eta^x - \eta^{s.x}\| < \varepsilon.$$

Let $E \subset \Gamma$ be a fixed finite symmetric subset containing e. For every continuous map $\eta \colon X \to \mathrm{Prob}(K)$, we define $f_\eta \in C(X)$ by

$$f_\eta(x) = \sum_{s \in E} \|s.\eta^x - \eta^{s.x}\| = \sum_{s \in E} \sum_{a \in K} |\eta^x(s^{-1}.a) - \eta^{s.x}(a)|.$$

Observe that $f_{\sum_k \alpha_k \zeta_k} \le \sum_k \alpha_k f_{\zeta_k}$ for every $\alpha_k \ge 0$ with $\sum_k \alpha_k = 1$. Hence, it suffices to show that 0 is in the norm-closed convex hull of $\{f_\eta : \eta \colon X \to \mathrm{Prob}(K) \text{ is continuous}\}$. By the Hahn-Banach separation theorem, it actually suffices to show 0 is in the weak closure of this set. That is, by the Riesz representation theorem, we must show that for every finite set of regular Borel probability measures $\mu_1, \ldots, \mu_n$ on X, there exists a continuous function $\eta \colon X \to \mathrm{Prob}(K)$ such that $\int f_\eta \, d\mu_i < \varepsilon$, for $i = 1, \ldots, n$.

Letting $m = \frac{1}{n} \sum \mu_i$, a little thought reveals that we really only have to find η such that $\int f_\eta \, dm < \varepsilon$ (for a smaller ε than that above). So, let $\varepsilon > 0$ be arbitrary. By assumption, there exists a Borel map $\zeta \colon X \to \mathrm{Prob}(K)$ such that

$$\sum_{s \in E} \int_X \|s.\zeta^x - \zeta^{s.x}\| \, dm(x) < \frac{\varepsilon}{9}.$$

Fubini's Theorem and the fact that $\zeta^{s.x}$ is a probability measure implies

$$1 = \int_X \big(\sum_{a \in K} \zeta^{s.x}(a) \big) \, dm(x) = \sum_{a \in K} \int_X \zeta^{s.x}(a) \, dm(x),$$

for every s. Hence we can find a finite subset $F \subset K$ such that

$$\sum_{s \in E} \int_X \sum_{a \in \Gamma \setminus F} \zeta^{s.x}(a) \, dm(x) < \frac{\varepsilon}{9}.$$

By Lusin's Theorem (applied to the measure $\sum_{s\in E} s.m$) we can approximate, for each $a \in F$, the Borel function $x \mapsto \zeta^x(a)$ by a continuous function $x \mapsto \eta^x(a)$ so that

$$\sum_{s\in E}\sum_{a\in F}\int_X |\eta^{s.x}(a) - \zeta^{s.x}(a)|\, dm(x) < \frac{\varepsilon}{9}.$$

Now fix $a_0 \in K \setminus F$ and define $\eta^x(a_0) = 1 - \sum_{a\in F} \eta^x(a)$, for every $x \in X$. For $b \notin F \cup \{a_0\}$ we define $\eta^x(b) = 0$ for all $x \in X$. We may assume that $\eta^x(a_0) \geq 0$ and regard η as a continuous map $\eta\colon X \to \mathrm{Prob}(K)$ such that $\operatorname{supp}\eta^x \subset F \cup \{a_0\}$ for all $x \in X$. It follows that

$$\sum_{s\in E}\int_X \|\eta^{s.x} - \zeta^{s.x}\|\, dm(x) < \frac{4\varepsilon}{9}.$$

This implies

$$\int_X f_\eta(x)\, dm(x) < \int_X \sum_{s\in E} \|s.\zeta^x - \zeta^{s.x}\|\, dm(x) + \frac{8\varepsilon}{9} < \varepsilon$$

and we obtain the claim.

Now, let a finite subset $E \subset \Gamma$ and $\varepsilon > 0$ be given. By our work above, there exists a continuous map η such that $\sup_{x\in X} \|s.\eta^x - \eta^{s.x}\| < \varepsilon$ for every $s \in E$. We may assume that there exists a finite subset $F \subset \Gamma$ such that $\operatorname{supp}\eta^x \subset F$ for all $x \in X$. Picking one point out of every orbit, we can find a Γ-fundamental domain $V \subset K$ – i.e., K decomposes into the disjoint union $\bigsqcup_{v\in V} \Gamma v$ – and let $v\colon K \to V$ be the corresponding projection (v takes every element in an orbit to its representative in V). Next we fix a cross section $\sigma\colon K \to \Gamma$ such that $a = \sigma(a).v(a)$ for every $a \in K$. Since the map v is constant along orbits, $\sigma(s.a)^{-1}s\sigma(a) \in \Gamma^{v(a)}$ for every $s \in \Gamma$ and $a \in K$. For each $v \in V$, we set

$$E^v = \{\sigma(s.a)^{-1}s\sigma(a) : a \in F \cap \Gamma v \text{ and } s \in E\} \subset \Gamma^v.$$

Let Y be a compact Γ-space which is amenable as a Γ^v-space for every $v \in V$. (Such Y always exists when each Γ^a is exact – take $Y = \beta\Gamma$). The proof of our proposition will be complete once we see that $X \times Y$ is an amenable Γ-space (with the diagonal action). Indeed, this will imply Γ is exact; moreover, if we can take $Y = X$, then the diagonal embedding $X \hookrightarrow X \times X$ is Γ-equivariant and continuous – hence amenability of X will follow from amenability of the diagonal action on $X \times X$. Since Y is Γ^v-amenable and E^v is finite, there exists a continuous map $\nu_v\colon Y \to \mathrm{Prob}(\Gamma)$ such that

$$\max_{s\in E^v}\sup_{y\in Y} \|s.\nu_v^y - \nu_v^{s.y}\| < \varepsilon.$$

Now, we define $\mu\colon X \times Y \to \mathrm{Prob}(\Gamma)$ by

$$\mu^{x,y} = \sum_{a\in F} \eta^x(a)\, \sigma(a).\nu_{v(a)}^{\sigma(a)^{-1}.y}.$$

The map μ is clearly continuous. Moreover, we have

$$\begin{aligned}
s.\mu^{x,y} &= \sum_{a\in K} \eta^x(a)\, s\sigma(a).\nu_{v(a)}^{\sigma(a)^{-1}.y} \\
&= \sum_{a\in F} \eta^x(a)\, \sigma(s.a).\big(\sigma(s.a)^{-1} s\sigma(a).\nu_{v(a)}^{\sigma(a)^{-1}.y}\big) \\
&\approx_\varepsilon \sum_{a\in K} \eta^x(a)\, \sigma(s.a).\nu_{v(s.a)}^{\sigma(s.a)^{-1}s.y} \\
&\approx_\varepsilon \sum_{a\in K} \eta^{s.x}(s.a)\, \sigma(s.a).\nu_{v(s.a)}^{\sigma(s.a)^{-1}s.y} \\
&= \mu^{s.x,s.y}
\end{aligned}$$

for every $s \in E$ and $(x, y) \in X \times Y$. □

Remark 14.2. This result really does generalize the fact that extensions of exact groups are exact. Let Γ be a group and $\Lambda \triangleleft \Gamma$ be a normal subgroup such that Λ and Γ/Λ are exact. The hypotheses of Proposition 14.1 are satisfied with $X = \beta(\Gamma/\Lambda)$ and $K = \Gamma/\Lambda$.

The hard work is essentially over. We now recall lots of definitions and prove a few well-known facts about trees, compactifications and groups acting on these objects.

Let $\mathbf{T}$ be a tree, which we identify (as a metric space) with its vertex set. A finite or infinite sequence $x(0), x(1), \ldots$ in $\mathbf{T}$ is called a *geodesic path* if $d(x(n), x(m)) = |n - m|$ for every n and m.[25] For convenience, if $(x(n))_{n=0}^N$ is a finite geodesic path, then we extend it to an infinite sequence by setting $x(m) = x(N)$ for every $m \geq N$; we still call this sequence a (finite) geodesic, even though it isn't, strictly speaking. Two geodesic paths x and x' are *equivalent* if they eventually flow together, i.e., if there exist $m_0, n_0 \in \mathbb{N}$ such that $x(m_0 + n) = x'(n_0 + n)$ for every $n \geq 0$. We can (and will) identify $\mathbf{T}$ with a subset of the equivalence classes of geodesics: every point x in $\mathbf{T}$ is identified with the equivalence class of geodesic paths ending at x. The *ideal boundary* $\partial\mathbf{T}$ of $\mathbf{T}$ is defined as the set of all equivalence classes of infinite geodesic paths. We define the *compactification of the tree* $\mathbf{T}$ to be $\bar{\mathbf{T}} = \mathbf{T} \sqcup \partial\mathbf{T}$ (a topology will be described shortly). If $(x(n))_n$ is a geodesic path with equivalence class $x \in \bar{\mathbf{T}}$, then we say the geodesic path $(x(n))_n$ *connects* $x(0)$ with x. For a bi-infinite geodesic path $(x(n))_{n=-\infty}^{\infty}$, we let $x(\infty) \in \partial\mathbf{T}$ (resp. $x(-\infty) \in \partial\mathbf{T}$) be the equivalence class of the geodesic path $(x(n))_{n\geq 0}$ (resp. $(x(-n))_{n\geq 0}$), and we say $(x(n))_n$ *connects* $x(-\infty)$ with $x(\infty)$.

Lemma 14.3. *Let $x \in \mathbf{T}$ and $y \in \partial\mathbf{T}$. Then, there exists a unique geodesic connecting x with y.*

Proof. Pictorially, the proof is totally transparent. Here's the recipe in words: Let $(y(n))$ be a representative of y and let $(w(j))_{j=1}^N$ be a finite geodesic connecting x with $y(0)$ (which exists, since $\mathbf{T}$ is connected). Let $N_0 \leq N$ be the first integer such that there exists n_0 with $w(N_0) = y(n_0)$ – i.e., find the first point of intersection of the two geodesics. Define a new geodesic $(z(k))$ by $z(k) = w(k)$ for $1 \leq k \leq N_0$ and $z(k) = y(n_0 + (k - N_0))$ for $k > N_0$. Evidently $(z(k))$ is a geodesic connecting x with y.

Uniqueness of $(z(k))$ follows from the fact that $\mathbf{T}$ is a tree – any other geodesic connecting x with y would yield a loop in $\mathbf{T}$. □

The lemma above is really a special case. Indeed, essentially the same proof yields the following fundamental fact (left to the reader): If $x, y \in \bar{\mathbf{T}}$, then there exists a geodesic

[25]In a tree, a path is geodesic if and only if it never backtracks.

connecting x with y (since $\mathbf{T}$ is connected), and it is unique (since $\mathbf{T}$ is a tree); this path will be denoted $[x, y]$.

If $[x, w_0]$ is a finite geodesic path and $[w_0, y]$ is any other geodesic, then we let $[x, w_0] \cup [w_0, y]$ denote the concatenation of these two paths (which is equal to $[x, y]$, of course). The following important lemma will be used repeatedly.

Lemma 14.4. *Given $x, y, z \in \bar{\mathbf{T}}$, $[x, y] \cap [y, z] \cap [z, x]$ is a singleton (i.e., there exists a unique point $w_0 \in \mathbf{T}$ such that $[x, y] = [x, w_0] \cup [w_0, y]$, $[x, z] = [x, w_0] \cup [w_0, z]$ and $[y, z] = [y, w_0] \cup [w_0, z]$).*

Proof. Again, the proof is trivial pictorially, so we only state the main idea. First note that $[x, y]$ and $[z, y]$ are equivalent geodesics. Letting w_0 be the first point of intersection of $[x, y]$ and $[z, y]$, the remainder of the proof is routine. □

We're now ready to topologize $\bar{\mathbf{T}}$. For $x \in \bar{\mathbf{T}}$ and a finite subset $F \subset \mathbf{T}$, we define

$$U(x; F) = \{x\} \cup \{y \in \bar{\mathbf{T}} : [x, y] \cap F = \emptyset\}.$$

One checks that $\{U(x; F)\}_{x,F}$ forms a basis for a topology (if $x \in U(x_1, F_1) \cap U(x_2, F_2)$, then Lemma 14.4 implies that $U(x, F_1 \cup F_2) \subset U(x_1, F_1) \cap U(x_2, F_2)$) and that the resulting topology is Hausdorff (given x, y, and any point $z_0 \neq x, y$ on the geodesic $[x, y]$, Lemma 14.4 implies $U(x, \{z_0\}) \cap U(y, \{z_0\}) = \emptyset$; and if x and y are adjacent, then $U(x, \{y\}) \cap U(y, \{x\}) = \emptyset$). This topology is very visual: cut finitely many edges in $\mathbf{T}$ and the connected components of $\bar{\mathbf{T}}$ are open (first verify this when only one edge is cut). Finally, it is worth checking that for a point $x \in \mathbf{T}$, the set $\{x\}$ is open if and only if x has finite degree.

Proposition 14.5. *The topological space $\bar{\mathbf{T}}$ is compact and any automorphism (i.e., isometric bijection) of the tree $\mathbf{T}$ extends to a homeomorphism of $\bar{\mathbf{T}}$.*

Proof. We must show that any net $(x_\alpha)_{\alpha \in A}$ in $\bar{\mathbf{T}}$ has an accumulation point.

Fix a base point $o \in \mathbf{T}$ and identify every $x_\alpha \in \bar{\mathbf{T}}$ with the unique geodesic path connecting o to x_α. Let N be the largest integer (possibly 0 or ∞) such that there exist $x(0), \ldots, x(N)$ satisfying

$$\bigcap_{n=0}^{N} \{\alpha \in A : x_\alpha(n) = x(n)\} \in \mathcal{U}.$$

We observe that for each n there exists at most one $x(n)$ such that $\{\alpha : x_\alpha(n) = x(n)\} \in \mathcal{U}$ and such that $(x(n))_{n=0}^N$ is a geodesic path. Thus, if $N = \infty$, then the boundary point represented by $(x(n))_{n=0}^\infty$ is an accumulation point. On the other hand, if $N < \infty$, then $x(N)$ is an accumulation point. Thus $\bar{\mathbf{T}}$ is compact.

If s is an automorphism of $\mathbf{T}$, then it naturally acts on $\bar{\mathbf{T}} = \mathbf{T} \sqcup \partial\mathbf{T}$. Namely, for every $x \in \bar{\mathbf{T}}$, we define $s.x \in \bar{\mathbf{T}}$ to be the equivalence class of the geodesic path $(s.x(n))_n$, where $(x(n))_n$ is a representative of x. It is routine to check that this is a homeomorphism. □

Lemma 14.6. *Let $\mathbf{T}$ be a countable tree with fixed base point o. There exists a sequence of Borel maps*

$$\zeta_n \colon \bar{\mathbf{T}} \to \mathrm{Prob}(\mathbf{T})$$

such that

$$\sup_{x \in \bar{\mathbf{T}}} \|s.\zeta_n^x - \zeta_n^{s.x}\| \leq \frac{2d(s.o, o)}{n}$$

for every automorphism s on $\mathbf{T}$.

Proof. As before, we identify every $x \in \bar{\mathbf{T}}$ with the unique geodesic path $(x(n))_n$ connecting o to x. (Recall our convention that $x(k) = x$ when $x \in \mathbf{T}$ and $k \geq \mathrm{dist}(x, o)$.) The maps ζ_n, defined by

$$\zeta_n^x = \frac{1}{n}\sum_{k=0}^{n-1} \delta_{x(k)} \in \mathrm{Prob}(\mathbf{T}),$$

satisfy the desired inequality. Indeed, $s.\zeta_n^x(p) > 0$ if and only if $s^{-1}.p$ is one of the first n points in the geodesic from o to x; equivalently, p is one of the first n points in the geodesic from $s.o$ to $s.x$. Similarly, $\zeta_n^{s.x}(q) > 0$ if and only if q is one of the first n points in the geodesic from o to $s.x$. Hence, for $n > d(s.o, o)$ we have cancellation in the difference $s.\zeta_n^x - \zeta_n^{s.x}$, because the geodesics from $s.o$ to $s.x$ and o to $s.x$ are equivalent.

Finally, it is easy to see that the ζ_n's are Borel since the set $\{x \in \bar{\mathbf{T}} : x(k) = z\}$ is clopen in $\bar{\mathbf{T}}$ for every $k \geq 0$ and $z \in \mathbf{T}$. □

Theorem 14.7. *Let Γ be a countable group and $\mathbf{T}$ be a countable tree on which Γ acts. If every vertex stabilizer Γ^x of $x \in \mathbf{T}$ is exact, then Γ is exact. In particular, an amalgamated free product of exact groups is exact.*

Proof. This follows from Proposition 14.1 and Lemma 14.6. □

We close this section with a result that won't be needed until later.

Lemma 14.8. *Let Γ be a group and $\mathbf{T}$ be a tree on which Γ acts. Let (s_n) be a net in Γ such that $s_n \notin s\Lambda$ eventually*[26] *for every $s \in \Gamma$ and every edge stabilizer Λ. If $s_n.x \to z$ for some $x \in \mathbf{T}$ and $z \in \bar{\mathbf{T}}$, then $s_n.y \to z$ for every $y \in \mathbf{T}$.*

Proof. We consider the open neighborhood of z given by a finite set F of edges in $\mathbf{T}$. It suffices to show that the geodesic paths $[s_n.x, s_n.y]$ between $s_n.x$ and $s_n.y$ do not cross F eventually. Since $[s_n.x, s_n.y] = s_n.[x, y]$, this reduces to showing that $s_n.\mathbf{e} \neq \mathbf{e}'$ eventually for any edges $\mathbf{e}, \mathbf{e}'$ in $\mathbf{T}$. Take $s \in \Gamma$ such that $s\mathbf{e} = \mathbf{e}'$. (If there is no such s, then we are already done.) Then $s_n.\mathbf{e} = \mathbf{e}'$ if and only if $s_n \in s\Gamma^{\mathbf{e}}$, where $\Gamma^{\mathbf{e}}$ is the edge stabilizer of $\mathbf{e}$. Hence $s_n.\mathbf{e} \neq \mathbf{e}'$ eventually, by assumption. □

Exercises

Exercise 14.1. Let X be a compact space which has no isolated points. Prove that the cardinality of X is at least c (cardinality of the continuum).

Exercise 14.2. Let X be a compact Γ-space whose cardinality is countable. Prove that there is a point $x \in X$ whose stabilizer subgroup Γ^x has finite index in Γ. (This explains the fact that Γ is exact if each stabilizer subgroup Γ^x is exact, which follows from Proposition 14.1 with $K = X$ and $\zeta\colon X \ni x \mapsto \delta_x \in \mathrm{Prob}(K)$.)

Exercise 14.3. Let $\mathbf{T}$ be a countable tree which has no infinite geodesic path. Prove that $\mathrm{Aut}(\mathbf{T})$ fixes either a point or an unoriented edge (i.e., a pair of points).

15. Hyperbolic groups

In this section we study an important class of graphs, namely those which are hyperbolic in the sense of Gromov. The main result is that groups which act properly on such graphs are exact.

Let K be a connected graph. We view K as a discrete metric space with the graph metric d. As in the previous section, a geodesic path α is a sequence of vertices such that

[26] That is, $\forall s, \forall\Lambda$ there exists n_0 such that $\forall n$, $n \geq n_0 \Rightarrow s_n \notin s\Lambda$.

$d(\alpha(m), \alpha(n)) = |m - n|$ for every m and n. Since K is connected, for every pair $x, y \in K$, there exists a (not necessarily unique) geodesic path connecting x to y. Though not exactly well-defined, we often use $[x, y]$ to denote a geodesic path from x to y (multiple such geodesics may exist). For every subset $A \subset K$ and $r > 0$, we define

$$d(x, A) = \inf\{d(x, a) : a \in A\} \quad \text{and} \quad N_r(A) = \{x \in K : d(x, A) < r\}.$$

The set $N_r(A)$ is called the r-tubular neighborhood of A in K. For subsets $A, B \subset K$, the *Hausdorff distance* between A and B is defined by

$$d_H(A, B) = \inf\{r : A \subset N_r(B) \text{ and } B \subset N_r(A)\}.$$

Definition 15.1. Let K be a connected graph. A *geodesic triangle* $\triangle$ in K consists of three points x, y, z in K and three geodesic paths $[x, y], [y, z], [z, x]$ connecting them.

Definition 15.2 (Hyperbolic graph)**.** For $\delta > 0$, we say a geodesic triangle $\triangle$ is *δ-slim* if each of its sides is contained in the open δ-tubular neighborhood of the union of the other two – i.e., $[x, y] \subset N_\delta([y, z] \cup [z, x])$ and similarly for the other two sides. We say that the graph K is *hyperbolic* if there exists $\delta > 0$ such that every geodesic triangle in K is δ-slim.

Note that hyperbolicity makes sense for any geodesic metric space (i.e., metric space where geodesics always exist). To get a feel for this concept, one should check that a tree is ε-hyperbolic (i.e., every geodesic triangle is ε-slim) for every $\varepsilon > 0$.

A *comparison tripod* is a geodesic triangle in a tree. It is not too hard to see that for every geodesic triangle $\triangle$ in a graph K there exist a unique tripod and a unique map f from $\triangle$ into the tripod that is isometric on all edges. Indeed, the lengths of the legs of the comparison tripod are determined by *the Gromov product*

$$\langle y, z\rangle_x = \frac{1}{2}(d(y, x) + d(z, x) - d(y, z)).$$ [27]

Definition 15.3. For $\delta > 0$, we say that a geodesic triangle $\triangle$ is *δ-thin* if $u, v \in \triangle$ and $f(u) = f(v)$ imply that $d(u, v) < \delta$, where f is the unique map to $\triangle$'s comparison tripod.

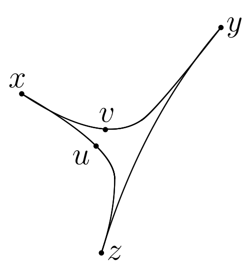

FIGURE 15.1. Thin geodesic triangle

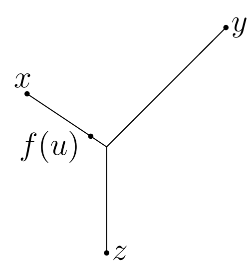

FIGURE 15.2. Comparison tripod

It is clear that any δ-thin geodesic triangle is δ-slim. The converse almost holds.

Proposition 15.4. *Let K be a hyperbolic graph. Then there exists $\delta > 0$ such that every geodesic triangle $\triangle$ is δ-thin.*

Proof. It will be convenient to think of K with the edges thrown in and each having length 1; that is, view K as a (continuous, rather than discrete) connected geodesic metric space. We will show that if every geodesic triangle in K is δ-slim, then they are all 4δ-thin.

[27] This number is the distance from x to the intersection point in Figure 15.2. It is not an integer, in general, of course.

Let $\triangle = [x,y]\cup[y,z]\cup[z,x]$ be a geodesic triangle and choose points u on $[x,y]$ and v on $[z,x]$ such that

$$d := d(u,x) = d(v,x) \leq \langle y,z\rangle_x.$$

By the intermediate value theorem, there is y' on $[x,y]$ such that $\langle y',z\rangle_x = d$. We note that u is on the subpath $[x,y']$ of $[x,y]$. Let $[y',z]$ be any geodesic connecting y' to z and let w be the point on $[y',z]$ such that $d(w,y') = d(u,y')$. It follows that $f(u) = f(v) = f(w)$ for the unique comparison map f from the geodesic triangle $\triangle' = [x,y']\cup[y',z]\cup[z,x]$ onto its comparison tripod. Since $\triangle'$ is δ-slim, $u \in N_\delta([y'z]\cup[z,x])$ and $v \in N_\delta([x,y']\cup[y'z])$. If $u \in N_\delta([z,x])$ or $v \in N_\delta([x,y'])$, then we must have $d(u,v) < 2\delta$ by the triangle inequality. Otherwise, we have $d(u,w) < 2\delta$ and $d(v,w) < 2\delta$. Therefore, we have $d(u,v) < 4\delta$ in either case. □

Now let Γ be a finitely generated group and $\mathcal{S}$ be a finite symmetric set of generators. We always equip the Cayley graph $\mathbf{X}(\Gamma,\mathcal{S})$ with the graph metric d (which is left invariant). Suppose that $\mathcal{S}'$ is another finite symmetric set of generators and let d' be the graph metric on $\mathbf{X}(\Gamma,\mathcal{S}')$. The vertex sets of $\mathbf{X}(\Gamma,\mathcal{S})$ and $\mathbf{X}(\Gamma,\mathcal{S}')$ are the same, of course; however their metric structures are different. But not that different. Indeed, if we choose $n\in\mathbb{N}$ so that $\mathcal{S}\subset(\mathcal{S}')^n = \{s_1s_2\cdots s_n : s_i\in\mathcal{S}'\}$ and $\mathcal{S}'\subset\mathcal{S}^n$, then it is readily seen that

$$n^{-1}d(x,y) \leq d'(x,y) \leq nd(x,y)$$

for every $x,y\in\Gamma$. Thus the formal identity from $\mathbf{X}(\Gamma,\mathcal{S})$ to $\mathbf{X}(\Gamma,\mathcal{S}')$ is *quasi-isometric.* More generally, we say that a map $f\colon (K,d)\to(K',d')$ between metric spaces is a *quasi-isometric embedding* if there exist $C\geq 1$ and $r>0$ such that

$$C^{-1}d(x,y) - r \leq d'(f(x),f(y)) \leq Cd(x,y) + r$$

for every $x,y\in K$. Thus, if Γ is finitely generated, its Cayley graph (with respect to a finite generating set) is unique up to quasi-isometry. Hence it is natural to look for properties which are quasi-isometry invariants, as they will provide invariants of groups.

Hyperbolicity turns out to be just such an invariant. This follows from the important fact that hyperbolic metric spaces enjoy *geodesic stability* – i.e., if a path is "close to being geodesic," then it is close (in Hausdorff distance) to an honest geodesic. To make this precise, we must define "close to being geodesic." For $C\geq 1$ and $r>0$, we say that a finite sequence[28] α in K is (C,r)*-quasigeodesic* if

$$C^{-1}d(\alpha(m),\alpha(n)) - r \leq |m-n| \leq Cd(\alpha(m),\alpha(n)) + r$$

for every m,n.

Proposition 15.5. *Let K be a hyperbolic graph, $C\geq 1$ and $r>0$. Then, there exists $D>0$ with the following property: For any (C,r)-quasigeodesic sequence α and any geodesic path β having the same origin and terminal point as α, one has $d_H(\alpha,\beta)<D$.*

In particular, a graph is hyperbolic if it quasi-isometrically embeds into some hyperbolic graph.[29]

Proof. Let α and β be given. We set $D_0 = \max\{d(p,\alpha) : p \text{ on } \beta\}$ (hence β is contained in the D_0 tubular neighborhood of α).

Now suppose that q_0 is a point on α such that $d(q_0,\beta)\geq D_0$. By maximality of D_0, for every point u on β, there exists a point u' on α such that $d(u,u')\leq D_0$. Since the endpoints of α and β are the same, "the intermediate value theorem" implies the existence

[28]It is important that we don't require α to be a path in this definition, because quasi-isometric embeddings don't always map paths to paths – i.e., neighbors need not map to neighbors.

[29]Think about the δ-slim condition and this is easily deduced.

of consecutive u_0, u_1 on β such that u_0' is on the origin (or "left") side and u_1' is on the terminus (or "right") side of q_0. (That is, u' is on the left of q_0 when u is the starting point of β, and it is on the right when u is the endpoint – hence there is a place where u' jumps over q_0.) Since $d(u_0', u_1') \le 2D_0 + 1$, the length of the subsequence of α from u_0' to u_1' is at most $C(2D_0+1)+r$. It follows that

$$d(u_0, q_0) \le d(u_0, u_0') + d(u_0', q_0) \le D_0 + C(C(2D_0+1)+2r).$$

Therefore, we have $d_H(\alpha, \beta) \le D$ for $D = D_0 + C(C(2D_0+1)+2r)$. Hence, we must show that D_0 is bounded above by a function depending only on δ, C and r, where δ is the constant satisfying Definition 15.2.

Choose a point p_0 on β such that $d(p_0, \alpha) = D_0$. Choose two points b_0 and b_1 on β, one coming before p_0 and one after, such that $d(b_0, p_0) = 2D_0 = d(b_1, p_0)$ – or, if this isn't possible, take an endpoint of β. Let a_k, $k = 1, 2$, be points on α such that $d(b_k, a_k) = d(b_k, \alpha)$ and choose geodesic paths γ_0 and γ_1 connecting b_0 to a_0 and b_1 to a_1, respectively. (Note that if b_k is an endpoint, then $a_k = b_k$, so we take γ_k to be the single point $a_k = b_k$ in this case.) Maximality of D_0 implies that $d(b_k, a_k) \le D_0$, and hence $d(p_0, \gamma_k) \ge D_0$. Let α' be the subsequence of α connecting a_0 to a_1. (It may flow backward.) Since

$$d(a_0, a_1) \le d(a_0, b_0) + d(b_0, b_1) + d(b_1, a_1) \le 6D_0$$

and α is a (C, r)-quasigeodesic path, the length of α' is at most $6CD_0 + r$. By joining γ_0, α' and γ_1, we obtain a sequence γ connecting b_0 to b_1. For the reader's convenience, we list the properties of γ: it connects b_0 to b_1; $d(p_0, \gamma) \ge D_0$; the length $|\gamma|$ of the sequence γ is at most $(6C+2)D_0 + r$; and $d(\gamma(k), \gamma(k+1)) \le C(1+r)$ for every k.

Now we apply the Weierstrass bisection process. Set $b_k^0 = b_k$, $p_0^0 = p_0$ and $\gamma^0 = \gamma$. Let c^0 be (one of) the midpoint(s) of γ^0 and consider a geodesic triangle $[b_0^0, c^0] \cup [c^0, b_1^0] \cup [b_1^0, b_0^0]$. Since K is hyperbolic, there exists p_0^1 in $[b_0^0, c^0] \cup [c^0, b_1^0]$ such that $d(p_0^0, p_0^1) \le \delta$. If p_0^1 is on $[b_0^0, c^0]$, then we set $b_0^1 = b_0^0$ and $b_1^1 = c^0$ – otherwise let $b_0^1 = c^0$ and $b_1^1 = b_1^0$. Let γ^1 be the subsequence of γ connecting b_0^1 to b_1^1. We note that $|\gamma^1| \le (2/3)|\gamma|$. Now, we continue this process by letting c^1 be (one of) the midpoint(s) of γ^1, and so on. This process terminates in l steps, with $l \le \log|\gamma| / \log(3/2)$, and gives rise to p_0^l on $[b_0^l, b_1^l]$ such that b_0^l and b_1^l are consecutive points on γ. It follows that

$$D_0 \le d(p_0, \gamma) \le l\delta + d(b_0^l, b_1^l) \le \delta \log(3/2)^{-1} \log\big((6C+2)D_0 + r\big) + C(1+r).$$

Since linear functions grow faster than logarithms, it follows that for each fixed C, r and δ, the numbers D_0 must be uniformly bounded (independent of α and β). □

Having established geodesic stability, the following definition is independent of the choice of generating set.

Definition 15.6 (Hyperbolic group)**.** Let Γ be a finitely generated group. We say that Γ is *hyperbolic* if its Cayley graph is hyperbolic.

Remark 15.7. Since the Cayley graph of $\mathbb{F}_n$ is a tree, free groups are hyperbolic. Other examples include co-compact lattices in simple Lie groups of real rank one and the fundamental groups of compact Riemannian manifolds of negative sectional curvature (cf. [5]).

Our next goal is to define the Gromov boundary; drawing lots of pictures will help.

Let K be a hyperbolic graph. We say that two infinite geodesic paths α and β in K are *equivalent* if

$$\liminf_{m,n\to\infty} \langle \alpha(m), \beta(n) \rangle_o = \infty$$

for some $o \in K$. Geometrically, this means α and β are pointing in the same direction. It is clear that the definition is independent of the choice of $o \in K$. It's not so clear that we have an equivalence relation, hence a lemma.

Lemma 15.8. *There exists a constant $C = C(K) > 0$ with the following property: For any two equivalent infinite geodesic paths α and β in K and any $m \geq d(\alpha(0), \beta(0))$, there exists n with $|m - n| \leq d(\alpha(0), \beta(0))$ such that $d(\alpha(m), \beta(n)) < C$.*

In particular, α and β are equivalent if and only if $\sup_m d(\alpha(m), \beta(m)) < \infty$ (and this is clearly an equivalence relation).

Proof. Choose $\delta > 0$ so that every geodesic triangle is δ-thin. Let $m \geq d(\alpha(0), \beta(0))$ be given. Let $o = \alpha(0)$ and find $m_1, n_1 \in \mathbb{N}$ such that $\langle \alpha(m_1), \beta(n_1)\rangle_o > m$. Choose any geodesic path $[o, \beta(n_1)]$ connecting o to $\beta(n_1)$. Let x be the vertex on $[o, \beta(n_1)]$ such that $d(o, x) = d(o, \alpha(m)) = m$. Since $\langle \alpha(m_1), \beta(n_1)\rangle_o > m$, we have $d(x, \alpha(m)) < \delta$. Let n be such that $n < n_1$ and $d(x, \beta(n_1)) = d(\beta(n), \beta(n_1))$. Since $d(o, x) \geq d(\alpha(0), \beta(0))$, such an n exists. Moreover $|m - n| \leq d(\alpha(0), \beta(0))$ and $d(x, \beta(n)) < \delta$. It follows that $d(\alpha(m), \beta(n)) < 2\delta$. This proves the first assertion.

For the second assertion, let equivalent geodesic paths α and β and $m \geq d(\alpha(0), \beta(0))$ be given. Then, by the first assertion, there is n with $|m - n| \leq d(\alpha(0), \beta(0))$ such that $d(\alpha(m), \beta(n)) < C$. Hence, $d(\alpha(m), \beta(m)) \leq d(\alpha(m), \beta(n)) + |m - n| \leq C + d(\alpha(0), \beta(0))$. Conversely, suppose $d_H(\alpha, \beta) < \infty$ and take $m_0, n_0 \geq 0$ such that $d(\alpha(m_0), \beta(n_0)) \leq d_H(\alpha, \beta) + 1$. Then, for any $m \geq m_0$ and $n \geq n_0$, one has

$$\begin{aligned} 2\langle \alpha(m), \beta(n)\rangle_o &= d(\alpha(m), o) + d(\beta(n), o) - d(\alpha(m), \beta(n)) \\ &\geq m - d(\alpha(0), o) + n - d(\beta(0), o) \\ &\qquad - ((m - m_0) + d(\alpha(m_0), \beta(n_0)) + (n - n_0)) \\ &\geq m_0 + n_0 - (d(\alpha(0), o) + d(\beta(0), o) + d_H(\alpha, \beta) + 1). \end{aligned}$$

This proves $\liminf_{m,n\to\infty} \langle \alpha(m), \beta(n)\rangle_o = \infty$. ☐

Definition 15.9. We define the *Gromov boundary* ∂K of a hyperbolic graph K to be the set of all equivalence classes of infinite geodesic paths. We call $\bar{K} = K \cup \partial K$ the *Gromov compactification* of K (we soon describe the topology). For a hyperbolic group Γ, $\bar{\Gamma}$ denotes the Gromov compactification of its Cayley graph.

Definition 15.10. For a finite or infinite geodesic path $\alpha = x_0 x_1 \cdots$ in K, we denote by $\alpha_- = x_0$ its starting point and α_+ its terminal point (i.e., the boundary point α represents, in the infinite case). As before, we say that α *connects* α_- with α_+.

The Cayley graph of a finitely generated group is always *uniformly locally finite* (or has *bounded geometry*) – i.e., there is a uniform bound on the degree of vertices. This is a nice property for a graph to have, so *from now on, we assume that the hyperbolic graph K is uniformly locally finite* and, in particular, countable. We leave it as an exercise to check that every $x \in K$ can be connected to every $z \in \partial K$ via a geodesic path.

Although the topology on $\bar{K}$ can be defined like that of a tree, we give a different description. Fix a base point $o \in K$. For $z \in \partial K$ and $R > 0$, we set

$$\begin{aligned} U(z, R) = \{x \in \bar{K} : \exists \text{ geodesic paths } \alpha, \beta \text{ with } \alpha_+ = x,\ \beta_+ = z \\ \text{such that } \liminf_{m,n\to\infty} \langle \alpha(m), \beta(n)\rangle_o > R\}, \end{aligned}$$

where in the case $x \in K$, we choose the "geodesic" $\alpha(m) = x$ for all m. We also define

$$U'(z, R) = \{x \in \bar{K} : \forall \text{ geodesic paths } \alpha, \beta \text{ with } \alpha_+ = x,\ \beta_+ = z,$$
$$\text{we have } \liminf_{m,n\to\infty} \langle \alpha(m), \beta(n)\rangle_o > R\}.$$

It turns out these sets satisfy the axioms for a neighborhood basis. The resulting topology on $\bar{K}$ is as expected: $\bar{K}$ is compact and K is a dense open discrete subset. Let's prove this.

It is clear that $U'(z, R) \subset U(z, R)$. On the other hand, we have

Lemma 15.11. *There exists $C = C(K) > 0$ with the following property: If α, α' and β, β' are geodesic paths such that $\alpha_+ = \alpha'_+$ and $\beta_+ = \beta'_+$, then*

$$\liminf_{m,n\to\infty} \langle \alpha'(m), \beta'(n)\rangle_o \geq \liminf_{m,n\to\infty} \langle \alpha(m), \beta(n)\rangle_o - C.$$

In particular, $U'(z, R) \supset U(z, R + C)$ for every $z \in \partial K$ and $R > 0$.

Proof. This follows from Lemma 15.8 and the inequality

$$\langle \alpha'(m'), \beta'(n')\rangle_o \geq \langle \alpha(m), \beta(n)\rangle_o - (d(\alpha'(m'), \alpha(m)) + d(\beta'(n'), \beta(n))).$$

□

Lemma 15.12. *For any $R > 0$, there exists $S > 0$ with the following property: For any $y, z \in \partial K$ with $y \in U(z, S)$, we have $U(y, S) \subset U(z, R)$.*

Proof. Choose some $\delta > 0$ such that every geodesic triangle is δ-thin. By Lemma 15.11, it suffices to show that if $y \in U'(z, N)$ for $y, z \in \partial K$ and $N \in \mathbb{N}$, then $U'(y, N) \subset U(z, N - \delta)$. Let $x \in U'(y, N)$ and take geodesic paths α, β and γ connecting o to x, y and z, respectively. Since $\liminf \langle \gamma(n), \beta(m)\rangle_o > N$, we have $d(\gamma(N), \beta(N)) < \delta$. Similarly, $d(\beta(N), \alpha(N)) < \delta$ and hence $d(\gamma(N), \alpha(N)) < 2\delta$. It follows that for every $m, n \geq N$,

$$\begin{aligned} 2\langle \alpha(m), \gamma(n)\rangle_o &= m + n - d(\alpha(m), \gamma(n)) \\ &\geq m + n - (m - N + d(\alpha(N), \gamma(N)) + n - N) \\ &\geq 2N - 2\delta. \end{aligned}$$

This shows $x \in U(z, N - \delta)$. □

Now, it is easy to check that $\{U(z, R)\}_{R>0}$ defines a (not necessarily open) neighborhood basis and the resulting topology is Hausdorff.

Definition 15.13. We equip $\bar{K} = K \cup \partial K$ with a topology by declaring that a subset $O \subset \bar{K}$ is open if and only if for every $z \in \partial K \cap O$, there exists $R > 0$ such that $U(z, R) \subset O$. We note that for every $x \in K$, the singleton set $\{x\}$ is open in $\bar{K}$.

It is clear that this topology is independent of the choice of the base point o. Moreover, for a hyperbolic group Γ, the Gromov compactification $\bar{\Gamma}$ is independent of the choice of finite generating subset (thanks to Proposition 15.5).

Theorem 15.14. *Let K be a locally finite hyperbolic graph. Then the topological space $\bar{K}$ defined above is compact and contains K as a dense open subset. Every automorphism (i.e., isometric bijection) on K extends uniquely to a homeomorphism on $\bar{K}$.*

Proof. The proof is similar to that of Proposition 14.5. We only prove compactness; the rest is trivial. It suffices to show that an arbitrary net $(x_i)_{i\in I}$ in $\bar{K}$ has an accumulation point. For every i, choose a geodesic path α_i connecting o to x_i. For convenience, we set $\alpha_i(n) = x_i$ when $n \geq d(o, x_i)$. Let $\mathcal{U}$ be a cofinal ultrafilter on the directed set I. Since K is locally finite, for every n, there exists a unique point $\alpha(n) \in K$ such that $\{i : \alpha_i(n) = \alpha(n)\} \in \mathcal{U}$.

Since each α_i is a geodesic path, α is also a geodesic path (or perhaps a path which is eventually constant). It is not too hard to see that $\alpha_+ \in \bar{K}$ is an accumulation point. $\square$

Here is the exactness result we have been after.

Theorem 15.15. *Let K be a uniformly locally finite hyperbolic graph and Γ be a group acting properly*[30] *on it. Then the action of Γ on the Gromov compactification $\bar{K}$ is amenable. In particular, every hyperbolic group is exact (since it acts properly on its Cayley graph).*

Proof. For $x, y \in K$, we denote by $T(x, y)$ the set of $z \in \partial K$ such that there exists a geodesic path connecting x to z which passes through y. It is not hard to see that $T(x, y)$ is a closed subset of ∂K. For every $x \in K$, $z \in \partial K$ and integers l, k with $l \geq k$, we define a subset $S(x, z, l, k) \subset K$ by declaring

$$S(x, z, l, k) = \{\alpha(l) : \alpha \text{ a geodesic path in } K \text{ such that } d(\alpha_-, x) \leq k \text{ and } \alpha_+ = z\}.$$

Note that for every $x, y \in K$ and k, l, we have

$$\begin{aligned}&\{z \in \partial K : y \in S(x, z, l, k)\}\\ &\quad= \bigcup \{T(x', y) : x' \in K \text{ with } d(x', x) \leq k \text{ and } d(x', y) = l\}\end{aligned}$$

and hence the former set is Borel in $\bar{K}$. Also, note the inclusion $S(x, z, l, k) \subset S(x, z, l, k')$ whenever $k \leq k'$.

Let $C = C(K) > 0$ be the constant appearing in Lemma 15.8. Since K is uniformly locally finite, there exists $D > 0$ such that every ball in K of radius C contains at most $D/3$ points. By Lemma 15.8, the subset $S(x, z, l, k)$ is contained in the C-tubular neighborhood of a subpath $\alpha([l-k, l+k])$ of any geodesic path α connecting x to z. This implies that $|S(x, z, l, k)| \leq Dk$ for all x, z, k, l with $l \geq k$. For a finite subset $S \subset K$, we denote by $\chi_S \in \mathrm{Prob}(K)$ the normalized characteristic function on S. Define a sequence of Borel functions $\eta_n \colon K \times \partial K \to \mathrm{Prob}(K)$ by

$$\eta_n(x, z) = \frac{1}{n} \sum_{k=n+1}^{2n} \chi_{S(x,z,3n,k)}.$$

We claim that, for each $x, x' \in K$, we have

$$\lim_{n \to \infty} \sup_{z \in \partial K} \|\eta_n(x, z) - \eta_n(x', z)\| = 0.$$

Let $d = d(x, x')$. Fix $z \in \partial K$ and $n \geq d$ and set $S_k = S(x, z, 3n, k)$ and $S'_k = S(x', z, 3n, k)$. Then, we have $S_k \cup S'_k \subset S_{k+d}$ and $S_k \cap S'_k \supset S_{k-d}$ for every $n < k \leq 2n$. It follows that

$$\|\chi_{S_k} - \chi_{S'_k}\| = 2\left(1 - \frac{|S_k \cap S'_k|}{\max\{|S_k|, |S'_k|\}}\right) \leq 2\left(1 - \frac{|S_{k-d}|}{|S_{k+d}|}\right)$$

[30] In this case, being proper is equivalent to saying every vertex stabilizer is finite.

for $n < k \leq 2n$. Since $|S_k| \leq Dk$, we have

$$\begin{aligned}
\|\eta_n(x,z) - \eta_n(x',z)\| &\leq \frac{1}{n}\sum_{k=n+1}^{2n} \|\chi_{S_k} - \chi_{S'_k}\| \\
&\leq 2\left(1 - \frac{1}{n}\sum_{k=n+1}^{2n} \frac{|S_{k-d}|}{|S_{k+d}|}\right) \\
&\leq 2\left(1 - \left(\prod_{k=n+1}^{2n} \frac{|S_{k-d}|}{|S_{k+d}|}\right)^{1/n}\right) \\
&= 2\left(1 - \left(\frac{\prod_{k=n+1-d}^{n+d} |S_k|}{\prod_{k=2n+1-d}^{2n+d} |S_k|}\right)^{1/n}\right) \\
&\leq 2\left(1 - (D(2n+d))^{-2d/n}\right).
\end{aligned}$$

Since $(D(2n+d))^{-2d/n} \to 1$ as $n \to \infty$, this proves the claim.

Now we fix a base point $o \in K$, set $\zeta_n^z = \eta_n(o,z)$ and observe that the maps $\zeta_n \colon \partial K \to \mathrm{Prob}(K)$ are Borel. Since we have $(s.\eta_n)(x,z) = \eta_n(s.x, s.z)$ for every $s \in \Gamma$ and $(x,z) \in K \times \partial K$, it follows that

$$\lim_{n\to\infty} \sup_{z \in \partial K} \|s.\zeta_z^n - \zeta_{s.z}^n\| = \lim_{n\to\infty} \sup_{z\in\partial K} \|\eta_n(s.o, s.z) - \eta(o, s.z)\| = 0.$$

Finally, for $x \in K$ we set $\zeta_n^x = \delta_x$, and one can check that the Borel maps ζ_n satisfy the hypotheses of Proposition 14.1 – hence the action of Γ on $\bar{K}$ is amenable. □

16. The Akemann-Ostrand Property

At this point we've established that all hyperbolic groups are exact. The other hypothesis required by Ozawa's theorem (Theorem 12.8) is $\otimes$-continuity of the canonical map $C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))/\mathbb{K}(\ell^2(\Gamma))$. We now show hyperbolic groups have this property, too. For free groups this fact was established by Akemann and Ostrand in [1].

Definition 16.1. A *compactification* of a group Γ is a compact topological space $\bar{\Gamma} = \Gamma \cup \partial\Gamma$ containing Γ as an open dense subset. We assume that a compactification is (left) *equivariant* in the sense that the left translation action of Γ on Γ extends to a continuous action on $\bar{\Gamma}$. The compactification $\bar{\Gamma}$ is said to be *small at infinity* if for every net $\{s_n\} \subset \Gamma$ converging to a boundary point $x \in \partial\Gamma$ and every $t \in \Gamma$, one has that $s_n t \to x$.

By Gelfand duality, there is a one-to-one correspondence between compactifications $\bar{\Gamma}$ and C*-algebras $C(\bar{\Gamma})$, where $c_0(\Gamma) \subset C(\bar{\Gamma}) \subset \ell^\infty(\Gamma)$ is left-translation invariant. The proof of the following lemma is a good exercise.

Lemma 16.2. *Let Γ be a group and $\bar{\Gamma} = \Gamma \cup \partial\Gamma$ be a compactification. The following are equivalent:*

(1) *the compactification $\bar{\Gamma}$ is small at infinity;*
(2) *the right translation action extends to a continuous action on $\bar{\Gamma}$ in such a way that it is trivial on $\partial\Gamma$;*
(3) *one has $f^t - f \in c_0(\Gamma)$ for every $f \in C(\bar{\Gamma})$ and $t \in \Gamma$, where $f^t(s) = f(st^{-1})$ for $f \in \ell^\infty(\Gamma)$.*

Proposition 16.3. *For any hyperbolic group Γ, the Gromov compactification $\bar{\Gamma}$ is small at infinity.*

Proof. Let a sequence $\{s_n\}$ converging to a boundary point $x \in \partial\Gamma$ and $t \in \Gamma$ be given. Let β be a geodesic path converging to x. Since $d(s_n t, s_n) = d(t, e)$ for every n, we have

$$\liminf_{m,n\to\infty} \langle s_m t, \beta(n)\rangle_e \geq \liminf_{m,n\to\infty} \langle s_m, \beta(n)\rangle_e - d(t,e) = \infty.$$

This means that $s_n t \to x$. □

We've seen that the left-translation action of a hyperbolic group Γ on $\ell^\infty(\Gamma)$ is amenable, but much more is true: the action of $\Gamma \times \Gamma$ on $\ell^\infty(\Gamma)$ (given by the left and right translations) is amenable mod $c_0(\Gamma)$.

Corollary 16.4. *If Γ is hyperbolic, then $\Gamma \times \Gamma$ acts amenably on the quotient algebra $\ell^\infty(\Gamma)/c_0(\Gamma)$.*

Proof. The previous proposition ensures that we can find a $(\Gamma \times \Gamma)$-invariant subalgebra $A \subset \ell^\infty(\Gamma)/c_0(\Gamma)$ such that the restriction of the $\Gamma \times \Gamma$ action to A is amenable on $\Gamma \times \{e\}$ and trivial on $\{e\} \times \Gamma$ (just let A be the image of $C(\bar{\Gamma})$ under the quotient map). By symmetry, we can also find $B \subset \ell^\infty(\Gamma)/c_0(\Gamma)$ such that the restriction of the $\Gamma \times \Gamma$ action to B is trivial on $\Gamma \times \{e\}$ and amenable on $\{e\} \times \Gamma$. The result now follows from Exercise 10.1. □

We're finally ready to show that Theorem 12.8 is not vacuous.

Corollary 16.5. *Let Γ be hyperbolic, λ and ρ be the left and, respectively, right regular representations and $\pi\colon \mathbb{B}(\ell^2(\Gamma)) \to \mathbb{B}(\ell^2(\Gamma))/\mathbb{K}(\ell^2(\Gamma))$ be the quotient map. Then, the $*$-homomorphism*

$$C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma) \ni \sum_k a_k \otimes x_k \mapsto \pi(\sum_k a_k x_k) \in \mathbb{B}(\ell^2(\Gamma))/\mathbb{K}(\ell^2(\Gamma))$$

is continuous with respect to the minimal tensor norm.

Proof. This is in fact an immediate corollary of Corollary 16.4 and Theorem 10.4, but we give a different proof here. It suffices to show that there exists a nuclear C*-algebra $A \subset \mathbb{B}(\ell^2(\Gamma))$ such that $C^*_\lambda(\Gamma) \subset A$ and $\pi(A)$ commutes with $\pi(C^*_\rho(\Gamma))$. Indeed, if such A exists, then we have an inclusion $C^*_\lambda(\Gamma) \otimes C^*_\rho(\Gamma) \subset A \otimes C^*_\rho(\Gamma) = A \otimes_{\max} C^*_\rho(\Gamma)$ and a natural $*$-homomorphism $A \otimes_{\max} C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))/\mathbb{K}(\ell^2(\Gamma))$.

So, let $\bar{\Gamma} = \Gamma \cup \partial\Gamma$ be the Gromov compactification and embed $C(\bar{\Gamma}) \subset \ell^\infty(\Gamma)$ as above. By Theorems 15.15 and 10.4, the C*-subalgebra A of the uniform Roe algebra generated by $C(\bar{\Gamma})$ and $C^*_\lambda(\Gamma)$ is nuclear. (It is $*$-isomorphic to $C(\bar{\Gamma}) \rtimes_r \Gamma$ by Proposition 11.1.) By Proposition 16.3, we have

$$\rho_t^* f \rho_t - f = f^t - f \in c_0(\Gamma) \subset \mathbb{K}(\ell^2(\Gamma))$$

for any $f \in C(\bar{\Gamma})$ and any $t \in \Gamma$, which implies that $\pi(A)$ commutes with $\pi(C^*_\rho(\Gamma))$, as desired. □

For free groups, the following fact is a celebrated result of Liming Ge [4].

Corollary 16.6. *For every hyperbolic group Γ, $L(\Gamma)$ is prime.*

Proof. Since Theorem 12.8 applies to hyperbolic groups, the result follows from Proposition 12.9. In fact, the result even holds for noninjective subfactors of $L(\Gamma)$. □

17. More Applications to von Neumann algebras

Shortly after Ozawa proved Theorem 12.8, he teamed up with Popa and proved several more remarkable results. This sections contains a unified and succinct approach to much of their work.

We make the blanket assumption that all groups are countable and von Neumann algebras have separable predual.

Requisite Results

Unfortunately, we require some nontrivial von Neumann algebra results (mostly due to Popa) that won't be proved here – see [2, Appendix F] for details.

Theorem 17.1. *Let $A \subset M$ be finite von Neumann algebras with separable predual and let $p \in M$ be a nonzero projection. Then, for a von Neumann subalgebra $B \subset pMp$, the following are equivalent:*

1. *there is no sequence (w_n) of unitary elements*[31] *in B such that $\|E_A(b^*w_na)\|_2 \to 0$ for every $a, b \in M$;*
2. *there exists a positive element $d \in \langle M, A\rangle$ with $\mathrm{Tr}(d) < \infty$ such that the ultraweakly closed convex hull of $\{w^*dw : w \in B \text{ unitary}\}$ does not contain 0;*
3. *there exists a B-A-submodule $\mathcal{H}$ of $pL^2(M)$ with $\dim_A \mathcal{H} < \infty$;*
4. *there exist nonzero projections $e \in A$ and $f \in B$, a unital normal $*$-homomorphism $\theta\colon fBf \to eAe$ and a nonzero partial isometry $v \in M$ such that*

$$\forall x \in fBf, \quad xv = v\theta(x)$$

and such that $v^*v \in \theta(fBf)' \cap eMe$ and $vv^* \in (fBf)' \cap fMf$.

Definition 17.2. Let $A \subset M$ and $B \subset pMp$ be finite von Neumann algebras. We say B *embeds in A inside M* if one of (and hence all of) the conditions in Theorem 17.1 holds.

Note that if there is a nonzero projection $p_0 \in B$ such that p_0Bp_0 embeds in A inside M, then B embeds in A inside M (as condition (4) in Theorem 17.1 evidently implies). Recall that a (nonzero) projection $f \in B$ is minimal if and only if $fBf = \mathbb{C}f$, and a von Neumann algebra B is *diffuse* if it has no minimal projections.

Corollary 17.3. *Let M be a finite von Neumann algebra with separable predual and (A_n) be a sequence of von Neumann subalgebras. Let $N \subset pMp$ be a von Neumann subalgebra such that N does not embed in A_n inside M for any n. Then, there exists a diffuse abelian von Neumann subalgebra $B \subset N$ such that B does not embed in A_n inside M for any n.*

Lemma 17.4. *Let $A, B \subset M$ be diffuse finite von Neumann algebras such that A and $B' \cap M$ are factors. (This implies that M and B are also factors.) Assume that $A_0' \cap M \subset A$ for any diffuse von Neumann subalgebra $A_0 \subset A$. If B embeds in A inside M, then there exists a unitary element $u \in M$ such that $uBu^* \subset A$.*

A subgroup $\Lambda \subset \Gamma$ is called *malnormal* if for every $s \in \Gamma \setminus \Lambda$ one has $s\Lambda s^{-1} \cap \Lambda = \{e\}$.

Theorem 17.5. *Let $\Lambda \subset \Gamma$ be a malnormal subgroup and $A_0 \subset L(\Lambda)$ be a diffuse von Neumann subalgebra. Then, $A_0' \cap L(\Gamma) \subset L(\Lambda)$. More generally, if $u \in L(\Gamma)$ is a unitary element such that $uA_0u^* \subset L(\Lambda)$, then $u \in L(\Lambda)$.*

[31] A unitary element w in B is a partial isometry in M such that $w^*w = p = ww^*$.

Subalgebras with noninjective relative commutants

Definition 17.6. Let Γ be a group and $\mathcal{G}$ be a family of subgroups of Γ. We say a subset Ω of Γ is *small relative to* $\mathcal{G}$ if it is contained in a finite union of $s\Lambda t$'s, where $s, t \in \Gamma$ and $\Lambda \in \mathcal{G}$. (Here $s\Lambda t = \{sat : a \in \Lambda\} \subset \Gamma$.)

Let $c_0(\Gamma; \mathcal{G}) \subset \ell^\infty(\Gamma)$ be the C*-subalgebra generated by functions whose supports are small relative to $\mathcal{G}$. More intuitively, for a net (s_i) in Γ, we write $s_i \to \infty/\mathcal{G}$ if $s_i \notin s\Lambda t$ eventually[32] for every $s, t \in \Gamma$ and $\Lambda \in \mathcal{G}$. Hence, for $f \in \ell^\infty(\Gamma)$ we have that $f \in c_0(\Gamma; \mathcal{G})$ $\Leftrightarrow$ $\{x \in \Gamma : |f(x)| > \varepsilon\}$ is small relative to $\mathcal{G}$ for every $\varepsilon > 0 \Leftrightarrow \lim_{s \to \infty/\mathcal{G}} f(s) = 0$.

Definition 17.7. We say the group Γ is *bi-exact relative to* $\mathcal{G}$ if it is exact and there exists a map

$$\mu \colon \Gamma \to \mathrm{Prob}(\Gamma)$$

such that for every $s, t \in \Gamma$, one has

$$\lim_{x \to \infty/\mathcal{G}} \|\mu(sxt) - s.\mu(x)\| = 0.$$

Remark 17.8. A countable exact group Γ is bi-exact relative to $\mathcal{G}$ if for every finite subset $E \subset \Gamma$ and $\varepsilon > 0$, there exists $\mu \colon \Gamma \to \mathrm{Prob}(\Gamma)$ such that for every $s, t \in E$ the subset $\{x : \|\mu(sxt) - s.\mu(x)\| \geq \varepsilon\}$ is small relative to $\mathcal{G}$. Since we won't need this fact, we won't prove it. However, the main points are laid out in Exercise 17.1.

Let $\mathbb{K}(\Gamma; \mathcal{G})$ be the hereditary C*-subalgebra of $\mathbb{B}(\ell^2(\Gamma))$ generated by $c_0(\Gamma; \mathcal{G})$:

$$\mathbb{K}(\Gamma; \mathcal{G}) = \text{the norm closure of } c_0(\Gamma; \mathcal{G})\mathbb{B}(\ell^2(\Gamma))c_0(\Gamma; \mathcal{G}).$$

Since the left and right regular representations λ and, respectively, ρ normalize $c_0(\Gamma; \mathcal{G})$, the reduced group C*-algebras $C^*_\lambda(\Gamma)$ and $C^*_\rho(\Gamma)$ are in the multipliers of $\mathbb{K}(\Gamma; \mathcal{G})$.

Lemma 17.9. *Let Γ be an exact group and $\mathcal{G}$ be a nonempty family of subgroups of Γ. Then Γ is bi-exact relative to $\mathcal{G}$ if and only if there exists a u.c.p. map*

$$\theta \colon C^*_\lambda(\Gamma) \otimes C^*_\rho(\Gamma) \to \mathbb{B}(\ell^2(\Gamma))$$

such that $\theta(a \otimes b) - ab \in \mathbb{K}(\Gamma; \mathcal{G})$ for every $a \in C^*_\lambda(\Gamma)$ and $b \in C^*_\rho(\Gamma)$.

Proof. We first prove the "if" direction. Let θ be a u.c.p. map such that $\theta(a \otimes b) - ab \in \mathbb{K}(\Gamma; \mathcal{G})$. By Voiculescu's Theorem, there is an isometry $V \colon \ell^2(\Gamma) \to \ell^2(\Gamma \times \Gamma)$ such that $\theta(a \otimes b) - V^*(a \otimes b)V \in \mathbb{K}(\ell^2(\Gamma))$ for every a and b. It follows that

$$V^*(\lambda(s) \otimes \rho(t))V - \lambda(s)\rho(t) \in \mathbb{K}(\Gamma; \mathcal{G})$$

for every $s, t \in \Gamma$. Define a map $\mu \colon \Gamma \to \mathrm{Prob}(\Gamma)$ by

$$\mu(x)(y) = \sum_{z \in \Gamma} |(V\delta_x)(y, z)|^2.$$

It follows that

$$\begin{aligned} \|\mu(sxt) - s.\mu(x)\|_1 &\leq \|\, |V\delta_{sxt}|^2 - |(\lambda(s) \otimes \rho(t))^{-1}V\delta_x|^2\|_1 \\ &\leq 2\|V\delta_{sxt} - (\lambda(s) \otimes \rho(t))^{-1}V\delta_x\|_2 \to 0 \end{aligned}$$

as $x \to \infty/\mathcal{G}$.

Now we prove the "only if" direction. Define a unitary operator U on $\ell^2(\Gamma) \otimes \ell^2(\Gamma)$ by $U(\delta_x \otimes \delta_y) = \delta_x \otimes \delta_{x^{-1}y}$, so that $U^*(\lambda(s) \otimes \rho(t))U = (\lambda \otimes \lambda)(s)(1 \otimes \rho)(t)$ (cf. Fell's absorption principle). Let $\mu \colon \Gamma \to \mathrm{Prob}(\Gamma)$ be a map as in the definition of bi-exactness and define an

[32] $\forall s, t, \forall \Lambda, \exists i_0$ such that $\forall i$ we have the implication $i \geq i_0 \Rightarrow s_i \notin s\Lambda t$.

isometry $V\colon \ell^2(\Gamma) \to \ell^2(\Gamma)\otimes\ell^2(\Gamma)$ by $V\delta_x = U(\mu(x)^{1/2}\otimes\delta_x)$. Then, it is routine to check that

$$V^*(\lambda(s)\otimes\rho(t))V\delta_x = \langle\lambda(s)(\mu(x)^{1/2}), \mu(sxt^{-1})^{1/2}\rangle\lambda(s)\rho(t)\delta_x.$$

Since

$$\lim_{x\to\infty/\mathcal{G}} \|\lambda(s)(\mu(x)^{1/2}) - \mu(sxt^{-1})^{1/2}\|_2^2 \le \lim_{x\to\infty/\mathcal{G}} \|s.\mu(x) - \mu(sxt^{-1})\|_1 = 0,$$

we have $V^*(\lambda(s)\otimes\rho(t))V - \lambda(s)\rho(t) \in \mathbb{K}(\Gamma;\mathcal{G})$ for every $s,t\in\Gamma$. $\square$

Here is the main theorem of this section. (See Definition 17.2 for the terminology "N embeds in $L(\Lambda)$ inside $L(\Gamma)$".)

Theorem 17.10. *Let Γ be a countable group and $\mathcal{G}$ be a countable family of subgroups of Γ. Assume that the group Γ is bi-exact relative to $\mathcal{G}$. Let $p \in L(\Gamma)$ be a projection and $N \subset pL(\Gamma)p$ be a von Neumann subalgebra. If the relative commutant $N'\cap pL(\Gamma)p$ is noninjective, then there exists $\Lambda\in\mathcal{G}$ such that N embeds in $L(\Lambda)$ inside $L(\Gamma)$.*

The proof of this result requires some preparation. Let $M\subset\mathbb{B}(\mathcal{H})$ be a von Neumann algebra and consider the $*$-homomorphism

$$\Phi_M\colon M\odot M' \ni \sum a_k\otimes b_k \mapsto \sum a_kb_k \in \mathbb{B}(\mathcal{H}).$$

We note that Φ_M is min-continuous if and only if M is injective. We will need a refinement of this result for von Neumann subalgebras contained in corners $P\subset pMp$.

Proposition 17.11. *Let $M\subset\mathbb{B}(\mathcal{H})$ be a finite von Neumann algebra and $p\in M$ be a projection. Let $P\subset pMp$ be a von Neumann subalgebra and $E_P\colon pMp\to P$ be the trace-preserving conditional expectation. Consider the bi-normal u.c.p. map*

$$\Phi_P\colon M\odot M' \ni \sum_k a_k\otimes b_k \mapsto \sum_k E_P(pa_kp)b_kp \in \mathbb{B}(p\mathcal{H}).$$

Suppose that there are weakly dense C*-*subalgebras $C_l\subset M$ and $C_r\subset M'$ such that C_l is exact and Φ_P is* min-*continuous on $C_l\odot C_r$. Then P is injective.*

Proof. It can be shown that our assumptions imply that Φ_P is min-continuous on $M\odot M'$ (cf. [2, Lemma 9.2.9]). By The Trick, $\Phi_P|_M$ extends to a u.c.p. map ψ from $\mathbb{B}(\mathcal{H})$ into $(pM')' = pMp$. (Note that the argument for The Trick only requires $\Phi_P|_{\mathbb{C}1\otimes M'}$ to be $*$-homomorphic.) It follows that $E_P\circ\psi|_{\mathbb{B}(p\mathcal{H})}$ is a conditional expectation from $\mathbb{B}(p\mathcal{H})$ onto P. $\square$

We primarily consider Φ_P in the case where $P = B'\cap pMp$ for a projection $p\in M$ and a *diffuse* abelian von Neumann subalgebra $B\subset pMp$ (meaning B has no nonzero minimal projections). Every diffuse abelian von Neumann algebra B with separable predual is $*$-isomorphic to $L^\infty[0,1]$ and hence is generated by a single unitary element $u_0\in B$ (e.g., $u_0(t) = e^{2\pi it}$). Fixing such a generator, we define a c.p. map Ψ_B from $\mathbb{B}(\mathcal{H})$ into $\mathbb{B}(p\mathcal{H})$ by

$$\Psi_B(x) = \text{ultraweak-}\lim_n \frac{1}{n}\sum_{k=1}^n u_0^kxu_0^{-k},$$

where the limit is taken along some fixed ultrafilter. It is not hard to see that Ψ_B is a (nonunital) conditional expectation onto $B'\cap\mathbb{B}(p\mathcal{H})$ and that $\Psi_B|_{pMp}$ is a trace-preserving conditional expectation from pMp onto $B'\cap pMp$. By uniqueness of the trace-preserving conditional expectation, one has $\Psi_B(a) = E_P(pap)$ for every $a\in M$. It follows that

$$\Psi_B(\sum_k a_kb_k) = \sum_k E_P(pa_kp)b_kp = \Phi_P(\sum_k a_k\otimes b_k)$$

for $a_k \in M$ and $b_k \in M'$.

Proof of Theorem 17.10. By contradiction, suppose that the conclusion of the theorem is not true. Then, by Corollary 17.3, there is a diffuse abelian von Neumann subalgebra $B \subset N$ such that B does not embed in $L(\Lambda)$ inside $M = L(\Gamma)$ for any Λ. We will use Theorem 17.1 with $A = L(\Lambda)$. For this, observe that $\chi_\Lambda \in \ell^\infty(\Gamma) \subset \mathbb{B}(L^2(M))$ is nothing but the orthogonal projection e_A onto $L^2(A)$ and hence $\chi_{s\Lambda} = \lambda(s)e_A\lambda(s)^* \in \langle M, A\rangle_+$ with $\mathrm{Tr}(\chi_{s\Lambda}) = 1$. It follows that $\Psi_B(\chi_{s\Lambda})$ is a positive element in $p\langle M, A\rangle p \cap B'$ such that $\mathrm{Tr}(\Psi_B(\chi_{s\Lambda})) \leq 1$. By assumption and Theorem 17.1, $\Psi_B(\chi_{s\Lambda}) = 0$. Since $\rho(\Gamma)$ is in the multiplicative domain of Ψ_B, this implies that $\Psi_B(\chi_{s\Lambda t}) = 0$ for every $s, t \in \Gamma$ and $\Lambda \in \mathcal{G}$, or equivalently, $\mathbb{K}(\Gamma; \mathcal{G}) \subset \ker \Psi_B$. Hence, for the u.c.p. map θ given in Lemma 17.9, one has $\Phi_P = \Psi_B \circ \theta$ and Φ_P is min-continuous on $C^*_\lambda(\Gamma) \odot C^*_\rho(\Gamma)$. Injectivity of $P = B' \cap pMp$ now follows from Proposition 17.11. $\square$

Exercise

Exercise 17.1. Prove the claim made in Remark 17.8. Here is a hint: Let $\{e\} = E_0 \subset E_1 \subset E_2 \cdots$ be an increasing sequence of finite symmetric subsets of Γ with $\bigcup E_n = \Gamma$. Find μ_n for E_n and $\varepsilon = 1/n$. Define relatively small sets Ω_n inductively by

$$\Omega_n = \bigcup_{s,t \in E_n} \{x : \|\mu_n(sxt) - s.\mu_n(x)\| \geq 1/n\} \cup E_n \Omega_{n-1} E_n.$$

Set $|x| = \min\{n : x \in \Omega_n\}$ and $\mu(x) = |x|^{-1} \sum_{n=1}^{|x|} \mu_n(x)$.

On bi-exactness

Definition 17.12. Let Γ be a group and $\mathcal{G}$ be a family of subgroups of Γ. For $f \in \ell^\infty(\Gamma)$ and $t \in \Gamma$, we define the right translation $f^t \in \ell^\infty(\Gamma)$ by $f^t(s) = f(st^{-1})$. Note that $(f^t)^{t'} = f^{tt'}$. Now define a compact space $\bar{\Gamma}^{\mathcal{G}}$ by

$$C(\bar{\Gamma}^{\mathcal{G}}) = \{f \in \ell^\infty(\Gamma) : f - f^t \in c_0(\Gamma; \mathcal{G}) \text{ for every } t \in \Gamma\}$$

and view it as a Γ-space, where Γ acts by *left* translation. We define another compact Γ-space $\Delta^{\mathcal{G}}\Gamma \subset \bar{\Gamma}^{\mathcal{G}}$ by

$$C(\Delta^{\mathcal{G}}\Gamma) = C(\bar{\Gamma}^{\mathcal{G}})/c_0(\Gamma; \mathcal{G})$$

and we call it the $\mathcal{G}$-boundary of Γ.

Remark 17.13. It is not hard to see that $x \in \bar{\Gamma}^{\mathcal{G}}$ belongs to $\Delta^{\mathcal{G}}$ if and only if there is a net (s_n) in Γ such that $s_n \to x$ and $s_n \to \infty/\mathcal{G}$.

It is possible that $\mathcal{G} = \emptyset$ and $c_0(\Gamma; \mathcal{G}) = \{0\}$, but otherwise we have $c_0(\Gamma) \subset c_0(\Gamma; \mathcal{G}) \subset C(\bar{\Gamma}^{\mathcal{G}})$ and $\bar{\Gamma}^{\mathcal{G}}$ is an equivariant compactification of Γ.[33] By Gelfand duality, there is a one-to-one correspondence between equivariant compactifications $\bar{\Gamma}$ of Γ and intermediate C*-subalgebras $c_0(\Gamma) \subset C(\bar{\Gamma}) \subset \ell^\infty(\Gamma)$ which are left translation invariant. It is possible that $\Gamma \in \mathcal{G}$ and $c_0(\Gamma; \mathcal{G}) = \ell^\infty(\Gamma)$ and $\Delta^{\mathcal{G}}\Gamma = \emptyset$.

Note that $f \in C(\bar{\Gamma}^{\mathcal{G}})$ if $f - f^t \in c_0(\Gamma; \mathcal{G})$ for all t in some generating subset of Γ, since $f - f^{tt'} = (f - f^{t'}) + (f - f^t)^{t'}$.

Proposition 17.14. *Let Γ be a countable group and $\mathcal{G}$ be a nonempty family of subgroups of Γ. Then the following are equivalent:*

(1) *Γ is bi-exact relative to $\mathcal{G}$;*

[33] A *compactification* is a compact space $\bar{\Gamma}$ containing Γ as an open dense subset; it is *equivariant* if the left translation action of Γ on Γ extends continuously to $\bar{\Gamma}$. (This is the same as Definition 16.1, where equivariance was assumed.)

(2) *the* $\mathcal{G}$*-boundary* $\Delta^{\mathcal{G}}\Gamma$ *is amenable;*[34]
(3) *the Gelfand spectrum of* $\ell^\infty(\Gamma)/c_0(\Gamma;\mathcal{G})$ *is amenable as a* $\Gamma\times\Gamma$*-space (with the left-times-right translation action).*

Proof. Assume condition (1) and let $\mu\colon \Gamma\to \mathrm{Prob}(\Gamma)$ be a map as in Definition 17.7. Then, the u.c.p. map $\mu_*\colon \ell^\infty(\Gamma)\to\ell^\infty(\Gamma)$ defined by $\mu_*(f)(x)=\langle f,\mu(x)\rangle$ has the property that $\mu_*(s.f)-s.\mu_*(f)^t\in c_0(\Gamma;\mathcal{G})$. In particular, $\mu_*(f)\in C(\bar\Gamma^{\mathcal{G}})$. Let $Q\colon C(\bar\Gamma^{\mathcal{G}})\to C(\Delta^{\mathcal{G}}\Gamma)$ be the quotient map. Then, $Q\circ\mu_*$ is a Γ-equivariant u.c.p. map from $\ell^\infty(\Gamma)$ into $C(\Delta^{\mathcal{G}}\Gamma)$. One can now deduce the amenability of $\Delta^{\mathcal{G}}\Gamma$ from Exercise 17.3.

Next, we assume condition (2) and let X denote the Gelfand spectrum of $\ell^\infty(\Gamma)/c_0(\Gamma;\mathcal{G})$. The inclusion $C(\bar\Gamma^{\mathcal{G}})\subset\ell^\infty(\Gamma)$ induces a continuous map $\varphi_l\colon X\to\Delta^{\mathcal{G}}\Gamma$ which is $\Gamma\times\Gamma$-equivariant, where the right action of Γ on $\Delta^{\mathcal{G}}\Gamma$ is trivial. By symmetry, there exists a continuous $\Gamma\times\Gamma$-equivariant map $\varphi_r\colon X\to\Delta^{\mathcal{G}}_r\Gamma$, where $\Delta^{\mathcal{G}}_r\Gamma$ is amenable as a $1\times\Gamma$-space and is trivial as a $\Gamma\times 1$-space. Thus, the $\Gamma\times\Gamma$-space $\Delta^{\mathcal{G}}\Gamma\times\Delta^{\mathcal{G}}_r\Gamma$ is amenable and $\varphi_l\times\varphi_r$ is a $\Gamma\times\Gamma$-equivariant continuous map from X into it. Therefore, X is amenable.

Finally, assume condition (3) and define a C*-algebra D by

$$D=C^*(\lambda(\Gamma),\rho(\Gamma),\ell^\infty(\Gamma))+\mathbb{K}(\Gamma;\mathcal{G})\subset\mathbb{B}(\ell^2(\Gamma)).$$

It is not hard to see that $\mathbb{K}(\Gamma;\mathcal{G})$ is an ideal in D and $D/\mathbb{K}(\Gamma;\mathcal{G})$ is a quotient of the crossed product of $\ell^\infty(\Gamma)/c_0(\Gamma;\mathcal{G})$ by $\Gamma\times\Gamma$ (actually, it's isomorphic to this crossed product). By assumption, the canonical $*$-homomorphism $C^*_\lambda(\Gamma)\odot C^*_\rho(\Gamma)\to D/\mathbb{K}(\Gamma;\mathcal{G})$ is min-continuous and $D/\mathbb{K}(\Gamma;\mathcal{G})$ is nuclear. Hence, the quotient map from D to $D/\mathbb{K}(\Gamma;\mathcal{G})$ has a u.c.p. splitting on any separable C*-subalgebra, by the Choi-Effros Lifting Theorem. Thanks to Lemma 17.9, we are done. □

It will be more convenient to work with $\Delta^{\mathcal{G}}\Gamma$ than the original definition of bi-exactness. This allows us to exploit the technology developed in previous chapters.

Definition 17.15. Let $\bar\Gamma$ be an equivariant compactification of Γ. We say $\bar\Gamma$ is *small at infinity relative to* $\mathcal{G}$ if the following holds: If (s_n) is a net in Γ such that $s_n\to x\in\bar\Gamma$ and $s_n\to\infty/\mathcal{G}$, then $s_nt\to x$ for every $t\in\Gamma$.

One should check that an equivariant compactification $\bar\Gamma$ of Γ is small at infinity relative to $\mathcal{G}$ if and only if the identity map on Γ extends to a continuous map from $\bar\Gamma^{\mathcal{G}}$ onto $\bar\Gamma$. The image of $\Delta^{\mathcal{G}}\Gamma$ under this map is the set of $x\in\bar\Gamma$ such that there is a net (s_n) in Γ with the property that $s_n\to x$ and $s_n\to\infty/\mathcal{G}$.

Example 17.16. In the examples below, amenability of $\Delta^{\mathcal{G}}\Gamma$ follows from that of $\bar\Gamma^{\mathcal{G}}$.

(1) Let $\mathcal{G}$ be the empty family. Then, $c_0(\Gamma;\mathcal{G})=\{0\}$ and $\bar\Gamma^{\mathcal{G}}$ is a one-point set. Hence $\bar\Gamma^{\mathcal{G}}$ is amenable if and only if Γ is amenable.
(2) Let $\mathcal{G}=\{\mathbf{1}\}$, where $\mathbf{1}$ is the trivial subgroup consisting of the neutral element. Then, $c_0(\Gamma;\mathcal{G})=c_0(\Gamma)$ and $\bar\Gamma^{\mathcal{G}}$ is the universal compactification which is small at infinity. Recall from Section 15 that if Γ is a hyperbolic group, then there is a Γ-equivariant continuous map from $\bar\Gamma^{\mathcal{G}}$ onto the Gromov compactification – hence $\bar\Gamma^{\mathcal{G}}$ is amenable for a hyperbolic group (Corollary 16.4).
(3) Suppose $\Gamma\in\mathcal{G}$. Then $c_0(\Gamma;\mathcal{G})=C(\bar\Gamma^{\mathcal{G}})=\ell^\infty(\Gamma)$ and $\bar\Gamma^{\mathcal{G}}=\beta\Gamma$. Hence $\bar\Gamma^{\mathcal{G}}$ is amenable if and only if Γ is exact.

It is often useful to ignore amenable subgroups.

[34] By convention, we say that the empty Γ-space $\emptyset$ is amenable if Γ is exact.

Lemma 17.17. *Let Γ be an exact group, $\Upsilon \subset \Gamma$ be an amenable subgroup and $\mathcal{G}$ be a family of subgroups of Γ. If there is a map*

$$\zeta\colon \Gamma \to \ell^1(\Gamma/\Upsilon)$$

such that

$$\lim_{x\to\infty/\mathcal{G}} \frac{\|\zeta(sxt) - s.\zeta(x)\|}{\|\zeta(x)\|} = 0$$

for every $s, t \in \Gamma$, then Γ is bi-exact relative to $\mathcal{G}$.

Proof. We define $\mu\colon \Gamma \to \mathrm{Prob}(\Gamma/\Upsilon)$ by $\mu(x) = \|\zeta(x)\|^{-1}|\zeta(x)|$. Then,

$$\begin{aligned}\|\mu(sxt) - s.\mu(x)\| &\le \left|1 - \frac{\|\zeta(sxt)\|}{\|\zeta(x)\|}\right| + \frac{\|\zeta(sxt) - s.\zeta(x)\|}{\|\zeta(x)\|} \\ &\le 2\frac{\|\zeta(sxt) - s.\zeta(x)\|}{\|\zeta(x)\|} \to 0 \quad \text{ as } x \to \infty/\mathcal{G}\end{aligned}$$

for every $s, t \in \Gamma$. Let $\mu_*\colon \ell^\infty(\Gamma/\Upsilon) \to \ell^\infty(\Gamma)$ be the u.c.p. map defined by $\mu_*(f)(x) = \langle \mu(x), f\rangle$. It is not hard to see that $\mu_*(f) \in C(\bar{\Gamma}^{\mathcal{G}})$ and composed with the quotient map, it gives rise to a Γ-equivariant u.c.p. map from $\ell^\infty(\Gamma/\Upsilon)$ into $C(\Delta^{\mathcal{G}}\Gamma)$. We view $\ell^\infty(\Gamma/\Upsilon)$ as the C*-subalgebra of right Υ-invariant functions in $\ell^\infty(\Gamma)$. Since Υ is amenable, by taking an "average" over the right Υ-action, one can find a (left) Γ-equivariant conditional expectation from $\ell^\infty(\Gamma)$ onto $\ell^\infty(\Gamma/\Upsilon)$. Combining these two Γ-equivariant u.c.p. maps, we obtain a Γ-equivariant u.c.p. map from $\ell^\infty(\Gamma)$ into $C(\Delta^{\mathcal{G}}\Gamma)$. Amenability of $\Delta^{\mathcal{G}}\Gamma$ now follows from Exercise 17.3. □

Proposition 17.18. *Let Γ be a group. For families $\mathcal{G}$ and $\mathcal{G}'$ of subgroups of Γ, define*

$$\mathcal{G} \wedge \mathcal{G}' = \{\Lambda \cap s\Lambda' s^{-1} : \Lambda \in \mathcal{G},\ \Lambda' \in \mathcal{G}',\ s \in \Gamma\}.$$

If Γ is bi-exact relative to $\mathcal{G}$ and to $\mathcal{G}'$, then Γ is bi-exact relative to $\mathcal{G} \wedge \mathcal{G}'$.

Proof. For notational simplicity, set $\tilde{\Gamma} = \Gamma \times \Gamma$, $A = \ell^\infty(\Gamma)$ and $I = c_0(\Gamma; \mathcal{G})$, $I' = c_0(\Gamma; \mathcal{G}')$. Then, the natural short exact sequence

$$0 \longrightarrow (I/(I\cap I')) \rtimes \tilde{\Gamma} \longrightarrow (A/(I\cap I')) \rtimes \tilde{\Gamma} \longrightarrow (A/I) \rtimes \tilde{\Gamma} \longrightarrow 0$$

is exact. By assumption, $(I/(I \cap I')) \rtimes \tilde{\Gamma} \cong ((I + I')/I') \rtimes \tilde{\Gamma} \lhd (A/I') \rtimes \tilde{\Gamma}$ and $(A/I) \rtimes \tilde{\Gamma}$ are nuclear. Hence the middle algebra $(A/(I \cap I')) \rtimes \tilde{\Gamma}$ is also nuclear. Therefore, it suffices to show that $I \cap I' = c_0(\Gamma; \mathcal{G} \wedge \mathcal{G}')$. We may assume that $\mathcal{G}$ is *saturated* in the sense that $s\Lambda s^{-1} \in \mathcal{G}$ for any $\Lambda \in \mathcal{G}$ and $s \in \Gamma$, and likewise for $\mathcal{G}'$. It is not hard to see that $I \cap I'$ is generated by a function whose support is contained in $\Lambda t \cap \Lambda' t'$. Pick any $x \in \Lambda t \cap \Lambda' t'$ (unless it is empty) and observe that $\Lambda t \cap \Lambda' t' = (\Lambda \cap \Lambda')x$. This completes the proof. □

Exercises

Exercise 17.2. Let X be a compact Γ-space and $\mathrm{Prob}(X)$ be the state space of $C(X)$ equipped with the natural Γ-action. Prove that X is amenable if and only if $\mathrm{Prob}(X)$ is amenable.

Exercise 17.3. Let X be a compact Γ-space and assume there is a Γ-equivariant u.c.p. map from $\ell^\infty(\Gamma)$ into $C(X)$. Prove that X is amenable provided that Γ is exact.

Examples

Direct product of hyperbolic groups.

Lemma 17.19. *Let $\Gamma_1, \dots, \Gamma_n$ be groups and $\Gamma = \prod_{i=1}^n \Gamma_i$ be the direct product. Let $\mathcal{G}_i$ be a family of subgroups of Γ_i and define a family $\mathcal{G}$ of subgroups of Γ by*

$$\mathcal{G} = \bigcup_i \{\Lambda \times \prod_{j \neq i} \Gamma_j : \Lambda \in \mathcal{G}_i\}.$$

If each of Γ_i is bi-exact relative to $\mathcal{G}_i$, then Γ is bi-exact relative to $\mathcal{G}$.

We leave the proof as an exercise.

Theorem 17.20. *Let $\Gamma_1, \dots, \Gamma_n$ be hyperbolic groups and $N_1, \dots, N_m$ be noninjective II_1-factors. If there exists an embedding*

$$N_1 \bar{\otimes} \cdots \bar{\otimes} N_m \hookrightarrow pL(\Gamma_1 \times \cdots \times \Gamma_n)p,$$

for some projection $p \in L(\Gamma_1 \times \cdots \times \Gamma_n)$, then $m \leq n$.

Proof. By Theorem 17.10 and Lemma 17.19, after permuting indices, one has

$$e_1 N_1 e_1 \bar{\otimes} \cdots \bar{\otimes} N_{m-1} \hookrightarrow p_0 L(\Gamma_1 \times \cdots \times \Gamma_{n-1}) p_0$$

for some nonzero projections $e_1 \in N_1$ and $p_0 \in L(\Gamma_1 \times \cdots \times \Gamma_{n-1})$. By induction, we are done. □

Semidirect products and wreath products.

Lemma 17.21. *Let $\Gamma = \Upsilon \rtimes \Lambda$ be a semidirect product of discrete groups. Let $\mathcal{G}_\Lambda$ be a family of subgroups of Λ and set $\mathcal{G} = \{\Upsilon \rtimes \Lambda_0 : \Lambda_0 \in \mathcal{G}_\Lambda\}$. If Υ is amenable and Λ is bi-exact relative to $\mathcal{G}_\Lambda$, then Γ is bi-exact relative to $\mathcal{G}$.*

Proof. Let $\mu\colon \Lambda \to \mathrm{Prob}(\Lambda)$ be a map as in Definition 17.7. It is not hard to see that the composition of μ with the quotient $\Gamma \to \Lambda = \Gamma/\Upsilon$ satisfies the conditions of Lemma 17.17. □

This is not so interesting unless $\mathcal{G}_\Lambda$ is very small (e.g., if Λ is hyperbolic). So, we consider another example.

Let us recall the definition of the wreath product $\Upsilon \wr \Lambda$ of a group Υ by another group Λ. To ease notation, denote by Υ_Λ the algebraic direct product group $\bigoplus_\Lambda \Upsilon$ and view an element $x \in \Upsilon_\Lambda$ as a finitely supported function $x\colon \Lambda \to \Upsilon$, where the support of x is $\mathrm{supp}(x) = \{p \in \Lambda : x(p) \neq e\}$. We note that $(xy)(p) = x(p)y(p) \in \Upsilon$ for $x, y \in \Upsilon_\Lambda$ and $p \in \Lambda$. Then, Λ acts on Υ_Λ by left translation: $\alpha_s(x)(p) = x(s^{-1}p)$. The *wreath product* $\Upsilon \wr \Lambda$ is defined to be the semidirect product $\Upsilon_\Lambda \rtimes_\alpha \Lambda$.

In what follows, we denote $\Upsilon \wr \Lambda$ by Γ and agree that p, s and t represent elements of Λ, while x, y and z represent elements of Υ_Λ (the group $\bigoplus_\Lambda \Upsilon$ of finitely supported functions from Λ into Υ). Hence a typical element of Γ will be denoted by xs or yt. In particular, $sx = \alpha_s(x)s$, where α is the left translation action of Λ on Υ_Λ.

Proposition 17.22. *Let $\Gamma = \Upsilon \wr \Lambda$ be the wreath product and let $\mathcal{G} = \{\Lambda\}$. If Υ is amenable and Λ is exact, then Γ is bi-exact relative to $\mathcal{G}$.*

The proof of this proposition requires several steps. We fix a proper length function $|\cdot|_\Lambda$ on Λ:

(1) $|s|_\Lambda = |s^{-1}|_\Lambda \in \mathbb{R}_{\geq 0}$ for $s \in \Lambda$ and $|s|_\Lambda = 0$ if and only if $s = e$;
(2) $|st|_\Lambda \leq |s|_\Lambda + |t|_\Lambda$ for every $s, t \in \Lambda$;

(3) the subset $B_\Lambda(R) = \{s \in \Lambda : |s|_\Lambda \le R\}$ is finite for every $R > 0$.

(Such a function exists – see [2, Proposition 5.5.2].) Likewise, fix a length function on Υ. For $yt \in \Gamma$, we define $\zeta(yt) \in \ell^1(\Lambda)$ by

$$\zeta(yt)(p) = \begin{cases} \min\{|p|_\Lambda, |t^{-1}p|_\Lambda\} + |y(p)|_\Upsilon & \text{if } p \in \operatorname{supp}(y), \\ 0 & \text{if } p \notin \operatorname{supp}(y). \end{cases}$$

Lemma 17.23. *For $\mathcal{G} = \{\Lambda\}$, one has*

$$\lim_{yt\to\infty/\mathcal{G}} \frac{|\operatorname{supp}(y)|}{\|\zeta(yt)\|} = 0.$$

Proof. We first claim that $\lim_{yt\to\infty/\mathcal{G}} \|\zeta(yt)\| = \infty$. Let $R > 0$ be given and suppose $yt \in \Gamma$ is such that $\|\zeta(yt)\| \le R$. Then $\operatorname{supp}(y) \subset B_R(\Lambda) \cup tB_R(\Lambda)$ and $y(p) \in B_R(\Upsilon)$ for every $p \in \Lambda$. Define $y' \in \Upsilon_\Lambda$ by $y'(p) = y(p)$ for $p \in B_R(\Lambda)$ and $y'(p) = e$ for $p \notin B_R(\Lambda)$. Then, $y = y'y''$ with $\operatorname{supp}(y') \subset B_R(\Lambda)$ and $\operatorname{supp}(y'') \subset tB_R(\Lambda)$. Hence, for the finite subset

$$E = \{z \in \Upsilon_\Lambda : \operatorname{supp}(z) \subset B_R(\Lambda) \text{ and } z(p) \in B_R(\Upsilon) \text{ for every } p \in \Lambda\}$$

of Υ_Λ, we have

$$yt = y't\alpha_{t^{-1}}(y'') \in \bigcup_{z',z''\in E} z'\Lambda z''.$$

This means that the subset $\Omega_R = \{yt \in \Gamma : \|\zeta(yt)\| \le R\}$ is small relative to $\mathcal{G}$ and the claim follows.

Let $C > 0$ be given and suppose $yt \in \Gamma$ is such that $\|\zeta(yt)\| \le C|\operatorname{supp}(y)|$. Since $\zeta(yt)(p) \ge 2C$ for $p \in \operatorname{supp}(y) \setminus \big(B_{2C}(\Lambda) \cup tB_{2C}(\Lambda)\big)$, we have

$$|\operatorname{supp}(y) \setminus \big(B_{2C}(\Lambda) \cup tB_{2C}(\Lambda)\big)| \le |\operatorname{supp}(y)|/2.$$

This implies $|\operatorname{supp}(y)| \le 4|B_{2C}(\Lambda)|$ and $yt \in \Omega_R$ for $R = 4C|B_{2C}(\Lambda)|$. By the first part of the proof, $\{yt \in \Gamma : |\operatorname{supp}(y)|/\|\zeta(yt)\| \ge C^{-1}\}$ is small relative to $\mathcal{G}$. □

Lemma 17.24. *The following hold:*

(1) $\|\zeta(xyt) - \zeta(yt)\| \le \|\zeta(x)\|$ *for every* $x, y \in \Upsilon_\Lambda$ *and* $t \in \Lambda$*;*
(2) $\|\zeta(syt) - s.\zeta(yt)\| \le |s|_\Lambda |\operatorname{supp}(y)|$ *for every* $y \in \Upsilon_\Lambda$ *and* $s, t \in \Lambda$*;*
(3) $\|\zeta(ytx) - \zeta(yt)\| \le \|\zeta(x)\|$ *for every* $x, y \in \Upsilon_\Lambda$ *and* $t \in \Lambda$*;*
(4) $\|\zeta(yts) - \zeta(yt)\| \le |s|_\Lambda |\operatorname{supp}(y)|$ *for every* $y \in \Upsilon_\Lambda$ *and* $s, t \in \Lambda$.

Proof. Note that $\zeta(xyt)(p) - \zeta(yt)(p)$ is nonzero only if $p \in \operatorname{supp}(x)$. Also,

$$\begin{aligned} |\zeta(xyt)(p) &- \zeta(yt)(p)| \\ &= \begin{cases} \big| |x(p)y(p)|_\Upsilon - |y(p)|_\Upsilon \big| & \text{if } p \in \operatorname{supp}(y) \cap \operatorname{supp}(xy), \\ \zeta(xt)(p) & \text{otherwise} \end{cases} \\ &\le \zeta(x)(p) \end{aligned}$$

for $p \in \operatorname{supp}(x)$. This yields the first assertion. For the second, observe that $\zeta(syt)(p)$ and $(s.\zeta(yt))(p)$ are nonzero only if $p \in s\operatorname{supp}(y)$ and that

$$\begin{aligned} |\zeta(syt)(p) &- (s.\zeta(yt))(p)| \\ &= |\min\{|p|_\Lambda, |(st)^{-1}p|_\Lambda\} - \min\{|s^{-1}p|_\Lambda, |(st)^{-1}p|_\Lambda\}| \\ &\le |s|. \end{aligned}$$

This yields the second assertion. For the third assertion, we observe that $\zeta(ytx)(p)-\zeta(yt)(p)$ is nonzero only if $p \in t\,\mathrm{supp}(x)$ and that for $q \in \mathrm{supp}(x)$, one has

$$\zeta(ytx)(tq) = \begin{cases} \min\{|tq|_\Lambda, |q|_\Lambda\} + |y(tq)x(q)|_\Upsilon & \text{if } y(tq) \neq x(q)^{-1}, \\ 0 & \text{if } y(tq) = x(q)^{-1}, \end{cases}$$

$$\zeta(yt)(tq) = \begin{cases} \min\{|tq|_\Lambda, |q|_\Lambda\} + |y(tq)|_\Upsilon & \text{if } tq \in \mathrm{supp}(y), \\ 0 & \text{if } tq \notin \mathrm{supp}(y). \end{cases}$$

Hence for $q \in \mathrm{supp}(x)$, one has

$$|\zeta(ytx)(tq) - \zeta(yt)(tq)| \leq |q|_\Lambda + |x(q)|_\Upsilon = \zeta(x)(q)$$

and the third assertion follows. Finally, since $\zeta(yts)(p)$ and $\zeta(yt)(p)$ are nonzero only if $p \in \mathrm{supp}(y)$ and

$$|\zeta(yts)(p) - \zeta(yt)(p)| = |\min\{|p|_\Lambda, |s^{-1}t^{-1}p|_\Lambda\} - \min\{|p|_\Lambda, |t^{-1}p|_\Lambda\}| \leq |s|,$$

for $p \in \mathrm{supp}(y)$, the fourth assertion follows. □

Proof of Proposition 17.22. With Lemmas 17.23 and 17.24 in hand, it is easy to verify the condition of Lemma 17.17. Indeed, one just has to check the condition separately for $x \in \Upsilon_\Lambda$ and $s \in \Lambda$, acting from the left or the right. □

Corollary 17.25. *Let $\Gamma = \Upsilon \wr \Lambda$ be the wreath product. Suppose that Υ is amenable and Λ is bi-exact relative to $\{\mathbf{1}\}$ (e.g., if Λ is hyperbolic). Then, Γ is bi-exact relative to $\{\mathbf{1}\}$.*

Proof. Combine Lemma 17.21, Proposition 17.22 and Proposition 17.18. □

Theorem 17.26. *Let $\Gamma = \Upsilon \wr \Lambda$ be the wreath product of an amenable group Υ by an exact group Λ. If $N \subset pL(\Gamma)p$ is a von Neumann subalgebra with a noninjective relative commutant, then N embeds in $L(\Lambda)$ inside $L(\Gamma)$.*

Proof. Combine Theorem 17.10 and Proposition 17.22. □

Corollary 17.27. *Let $\Gamma = \Upsilon \wr \Lambda$ be the wreath product of an amenable group Υ by an exact group Λ. If $N \subset L(\Gamma)$ is a noninjective nonprime factor whose relative commutant $N' \cap L(\Gamma)$ is a factor, then there exists a unitary element $u \in L(\Gamma)$ such that $uNu^* \subset L(\Lambda)$.*

Proof. Write N as a tensor product $N = N_1 \mathbin{\bar{\otimes}} N_2$ of type II_1-factors N_1 and N_2. Since N is noninjective, we may assume that N_2 is noninjective. By Theorem 17.26, N_1 embeds in $L(\Lambda)$ inside $L(\Gamma)$. By Lemma 17.4 and Theorem 17.5, we can find a unitary element $u \in L(\Gamma)$ such that $uN_1u^* \subset L(\Lambda)$. This implies $uN_2u^* \subset (uN_1u^*)' \cap L(\Gamma) \subset L(\Lambda)$, by Theorem 17.5. Therefore, $uNu^* \subset L(\Lambda)$. □

Amalgamated free products.

Proposition 17.28. *Let $\Gamma = \Gamma_1 *_\Lambda \Gamma_2$ be an amalgamated free product and let $\mathcal{G} = \{\Gamma_1, \Gamma_2\}$. If both Γ_i are exact and Λ is amenable, then $\bar{\Gamma}^{\mathcal{G}}$ is amenable and, in particular, Γ is bi-exact relative to $\mathcal{G}$.*

Before giving the proof, we point out that the amenability assumption on Λ is essential. Indeed, if $\Gamma_i = \Gamma_i' \times \Lambda$ and Λ is nonamenable, then $\Gamma = (\Gamma_1' * \Gamma_2') \times \Lambda$ and thus $L(\Gamma_1' * \Gamma_2')$ has a noninjective commutant in $L(\Gamma)$. More generally, if $s_i \in \Gamma_i \setminus \Lambda$ normalize Λ and $s_1 a s_1^{-1} = s_2 a s_2^{-1}$ for all $a \in \Lambda$, then $s = s_1 s_2^{-1} \in \Gamma$ has infinite order and commutes with Λ.

Proof. We first prove that the Γ-space $\bar{\Gamma}^{\mathcal{G}}$ is amenable as a Γ_i-space. We prove this for $i = 1$. Let $A \subset \ell^\infty(\Gamma_1)$ be the C*-subalgebra of those functions f such that $f = f^t$ for all $t \in \Lambda$. Averaging over the right Λ-action, we obtain a (left) Γ_1-equivariant conditional expectation from $\ell^\infty(\Gamma_1)$ onto A. By Exercise 17.3, it suffices to find a Γ_1-equivariant $*$-homomorphism π from A into $C(\bar{\Gamma}^{\mathcal{G}})$. Fix a system $\{e\} \sqcup S_i^0 \subset \Gamma_i$ of representatives of $\Lambda\backslash\Gamma_i$, and set

$$\mathfrak{X} = \{e\} \sqcup S_2^0 \sqcup S_2^0 S_1^0 \sqcup S_2^0 S_1^0 S_2^0 \sqcup \cdots \subset \Gamma.$$

Then, every $s \in \Gamma$ can be uniquely written in the form $s = s_1 x$, where $s_1 \in \Gamma_1$ and $x \in \mathfrak{X}$ (cf. [2, Appendix E]). We define $\pi\colon A \to \ell^\infty(\Gamma)$ by $\pi(f)(s_1 x) = f(s_1)$ for $s_1 \in \Gamma_1$ and $x \in \mathfrak{X}$. Our task is to show $\pi(f) - \pi(f)^t \in c_0(\Gamma; \mathcal{G})$ for every $t \in \Gamma_1 \cup \Gamma_2$. Suppose first that $t \in \Gamma_1$. Then, for every $s_1 x \in \Gamma$, one has either $s_1 x t^{-1} = s_1 t^{-1}$ (if $x = e$) or $s_1 x t^{-1} = s_1 a y$ for some $a \in \Lambda$ and $y \in \mathfrak{X}$ (if $x \neq e$). Since f is right Λ-invariant, $\pi(f) - \pi(f)^t$ has support in Γ_1. It follows that $\pi(f) - \pi(f)^t \in c_0(\Gamma; \mathcal{G})$. Suppose next that $t \in \Gamma_2$. Then, one has $\pi(f) - \pi(f)^t = 0$ by similar reasoning. Altogether, this implies $\bar{\Gamma}^{\mathcal{G}}$ is amenable as a Γ_i-space.

Let $\mathbf{T} = \Gamma/\Gamma_1 \sqcup \Gamma/\Gamma_2$ be the Bass-Serre tree on which $\Gamma = \Gamma_1 *_\Lambda \Gamma_2$ acts and let $\bar{\mathbf{T}}$ be its compactification,[35] as defined in Section 14. We will find a Γ-equivariant continuous map from $\bar{\Gamma}^{\mathcal{G}}$ into $\bar{\mathbf{T}}$, which suffices to show the amenability of $\bar{\Gamma}^{\mathcal{G}}$ by Proposition 14.1 and Lemma 14.6. Choose a base point $o \in \mathbf{T}$ and define a Γ-equivariant $*$-homomorphism $\sigma\colon C(\bar{\mathbf{T}}) \to \ell^\infty(\Gamma)$ by $\sigma(f)(s) = f(so)$. We will show $\sigma(f) - \sigma(f)^t \in c_0(\Gamma; \{\Lambda\})$ for every $f \in C(\bar{\mathbf{T}})$ and $t \in \Gamma$. Suppose by contradiction that this is not the case. Then, there exists $\varepsilon > 0$ such that the set

$$\Omega = \{s \in \Gamma : |f(so) - f(st^{-1}o)| \geq \varepsilon\} \subset \Gamma$$

is not small relative to $\{\Lambda\}$. Hence, there exists a net (s_n) in Ω such that $s_n \to \infty/\{\Lambda\}$. We may assume that $s_n o \to z$ for some $z \in \bar{\mathbf{T}}$. Since every edge stabilizer of the Γ-action on the Bass-Serre tree is an inner conjugate of Λ, we can apply Lemma 14.8 and deduce that $s_n t^{-1} o \to z$. Hence we obtain the contradiction

$$\varepsilon \leq \lim_n |f(s_n o) - f(s_n t^{-1} o)| = |f(z) - f(z)| = 0.$$

Therefore, $\sigma(f) - \sigma(f)^t \in c_0(\Gamma; \{\Lambda\}) \subset c_0(\Gamma; \mathcal{G})$ and we are done. □

Theorem 17.29. *Let $\Gamma = \Gamma_1 *_\Lambda \Gamma_2$ be an amalgamated free product such that both Γ_i are exact and Λ is amenable. If $N \subset L(\Gamma)$ is a von Neumann subalgebra with a noninjective relative commutant, then there exists i such that N embeds in $L(\Gamma_i)$ inside $L(\Gamma)$.*

Proof. Combine Theorem 17.10 and Proposition 17.28. □

Recall that a group Γ is said to have *infinite conjugacy classes (ICC)* if the sets $\{sts^{-1} : s \in \Gamma\}$ are infinite for every nonneutral element $t \in \Gamma$.

Corollary 17.30. *Let $\Gamma = \Gamma_1 * \Gamma_2$ be a free product of ICC exact groups. If $N \subset L(\Gamma)$ is a noninjective nonprime factor whose relative commutant $N' \cap L(\Gamma)$ is a factor, then there exist $i \in \{1, 2\}$ and a unitary element $u \in L(\Gamma)$ such that $uNu^* \subset L(\Gamma_i)$.*

We omit the proof of this corollary as it is very similar to the proof of Corollary 17.27.

We say Γ is a *product* group if it is isomorphic to a direct product of nontrivial groups. We note that if $\Gamma = \Gamma' \times \Gamma''$ is an ICC product group, then Γ' and Γ'' are also ICC and, in particular, infinite.

[35] Although we use the term "compactification", $\mathbf{T}$ is not open in $\bar{\mathbf{T}}$.

Corollary 17.31. *Let* $\Gamma_1, \ldots, \Gamma_n$ *and* $\Lambda_1, \ldots, \Lambda_m$ *be ICC nonamenable exact product groups. If*

$$M = L(\mathbb{F}_\infty * \Gamma_1 * \cdots * \Gamma_n) \cong L(\mathbb{F}_\infty * \Lambda_1 * \cdots * \Lambda_m),$$

then $n = m$ *and, modulo permutation of indices,* $L(\Gamma_i)$ *is unitarily conjugated to* $L(\Lambda_i)$ *inside* M *for every* $1 \leq i \leq n$.

Proof. It follows from Corollary 17.30 that there are maps $\imath, \jmath$ and unitary elements $u_1, \ldots, u_m$ and $v_1, \ldots v_n$ such that $u_j L(\Lambda_j) u_j^* \subset L(\Gamma_{\imath(j)})$ and $v_i L(\Gamma_i) v_i^* \subset L(\Lambda_{\jmath(i)})$. It follows that

$$v_{\imath(j)} u_j L(\Lambda_j) u_j^* v_{\imath(j)}^* \subset L(\Lambda_{\jmath(\imath(j))})$$

for every j. By Theorem 17.5 and Exercise 17.4, this implies $\jmath(\imath(j)) = j$ and $v_{\imath(j)} u_j \in L(\Lambda_j)$. In particular, the above inclusions are tight and $u_j L(\Lambda_j) u_j^* = L(\Gamma_{\imath(j)})$. Likewise, one has $\imath(\jmath(i)) = i$ for every $1 \leq i \leq n$. □

This corollary is an analogue of Kurosh's isomorphism theorem for groups and, like Kurosh's Theorem, it says almost nothing about the positions of the copies of $L(\mathbb{F}_\infty)$.

Exercise

Exercise 17.4. Let $\Lambda_1, \Lambda_2 \subset \Gamma$ be groups and suppose that for every $s \in \Gamma$ one has $s\Lambda_1 s^{-1} \cap \Lambda_2 = \{e\}$ (e.g., $\Gamma = \Lambda_1 * \Lambda_2$). Let $A_0 \subset L(\Lambda_1)$ be a diffuse von Neumann subalgebra. Prove that there is no unitary element $u \in L(\Gamma)$ such that $uA_0u^* \subset L(\Lambda_2)$.

References

1. C.A. Akemann and P.A. Ostrand, *On a tensor product C*-algebra associated with the free group on two generators*, J. Math. Soc. Japan **27** (1975), 589–599.
2. Nathanial P. Brown and Narutaka Ozawa, *C^*-algebras and finite-dimensional approximations*, Graduate Studies in Mathematics, **88**. American Mathematical Society, Providence, RI, 2008. xvi+509 pp.
3. A. Connes, *Classification of injective factors: cases II_1, II_∞, III_λ, $\lambda \neq 1$*, Ann. Math. **104** (1976), 73–115.
4. L. Ge, *Applications of free entropy to finite von Neumann algebras. II.* Ann. of Math. **147** (1998), 143–157.
5. M. Gromov, *Hyperbolic groups*, Essays in group theory, 75–263, Math. Sci. Res. Inst. Publ. 8, Springer, New York, 1987.

Department of Mathematics, Penn State University, State College, PA 16802
E-mail address: nbrown@math.psu.edu