ON WALL STRUCTURES AND COARSE EMBEDDINGS

A. KHUKHRO

ABSTRACT. Property A is a property of metric spaces that was designed to imply the existence of a coarse embedding into Hilbert space. However, there are metric spaces which do not have property A but nevertheless admit such a coarse embedding. The first example of such a metric space with bounded geometry was given by Arzhantseva, Guentner and Špakula, in the form of a metrized sequence of finite quotients of a finite rank free group. This example uses wall structures to construct the embedding. We present this, and other examples which rely instead on other easily-embeddable spaces to construct a coarse embedding into Hilbert space.

INTRODUCTION

A *coarse embedding* of one metric space into another generalizes the notion of a quasi-isometric embedding, by allowing the functions which control how the metric is distorted to be non-linear.

Definition. Let (X, d_X) and (Y, d_Y) be metric spaces. X *coarsely embeds* into Y if there is a map $F : X \longrightarrow Y$ such that there exist non-decreasing functions $\rho_\pm : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ with $\lim_{t\to\infty} \rho_\pm(t) = \infty$ and

$$\rho_-(d_X(x, x')) \leq d_Y(F(x), F(x')) \leq \rho_+(d_X(x, x'))$$

for all $x, x' \in X$. X and Y are *coarsely equivalent* if there exists a coarse embedding $F : X \longrightarrow Y$ and $C > 0$ such that for each $y \in Y$, $d_Y(y, F(X)) < C$.

For Cayley graphs of finitely generated groups (and, more generally, all quasi-geodesic spaces), if two such spaces are coarsely equivalent, then they are necessarily quasi-isometric.

We are interested in spaces which admit a coarse embedding into Hilbert space, since the existence of such an embedding implies the coarse Baum–Connes conjecture and, in the case that the space is a Cayley graph of a finitely generated group, the Novikov conjecture [Yu]. Along with these remarkable results, Yu defines in [Yu] a geometric property which implies coarse embeddability into Hilbert space.

Definition ([Yu])**.** A uniformly discrete metric space (X, d) is said to have property A if for all $R, \varepsilon > 0$ there exists a family of non-empty subsets $\{A_x\}_{x\in X}$ of $X \times \mathbb{N}$ such that

- for all x, y in X with $d(x, y) < R$ we have $\frac{|A_x \Delta A_y|}{|A_x \cap A_y|} < \varepsilon$,
- there exists S such that for all x in X and (y, n) in A_x we have $d(x, y) \leq S$.

Theorem ([Yu])**.** *If a discrete metric space has property A, then it admits a coarse embedding into Hilbert space.*

One of the few examples of spaces without property A is given by metrized sequences of bounded-degree graphs with girth (i.e. the length of the shortest loop) tending to infinity [Wil]. Perhaps the richest source of examples is provided by *box spaces*, which are spaces with bounded geometry constructed using quotients of residually finite groups.

Given a sequence of metric spaces $\{X_n, d_n\}$, their *coarse disjoint union* is the space $\sqcup_n X_n$ with metric d such that d is d_n when restricted to each component X_n, and the distance between two distinct components is chosen to be greater than the maximum of their diameters (note that any two such choices of metric result in coarsely equivalent spaces). The coarse disjoint union allows us to study properties that the sequence of spaces $\{X_n\}$ has *uniformly*.

Let G be a finitely generated residually finite group and let $\{N_i\}$ be a collection of finite index nested normal subgroups of G, for which the intersection $\cap_{i\in\mathbb{N}} N_i$ is trivial. Note that for a finitely generated residually finite group G, there always exists a such a collection of subgroups $\{N_i\}$.

Definition. The *box space* $\square_{\{N_i\}}G$ of G corresponding to $\{N_i\}$ is the coarse disjoint union $\sqcup_i G/N_i$ of finite quotient groups of G, where each quotient is endowed with the Cayley graph metric induced by the image of the generating set of G.

For a residually finite group G, and $\{N_i\}$ a sequence of subgroups as above, we have the following links between analytic properties of G and geometric properties of the box space:

$$\begin{aligned} G \text{ amenable } &\iff \square_{\{N_i\}}G \text{ property A},\\ G \text{ Haagerup } &\impliedby \square_{\{N_i\}}G \text{ coarsely embeddable into Hilbert space},\\ G \text{ property (T) } &\implies \square_{\{N_i\}}G \text{ expander.} \end{aligned}$$

The proofs of the first two statements can be found in [Roe], and the third in [Mar]. Note that the last two implications are not reversible.

The interested reader may want to bear in mind some further connections: when property (T) above is replaced by a weaker property called property (τ) [LZ], the implication becomes an equivalence; in [WY], Willett and Yu define a new geometric property called *geometric property (T)*, such that having a box space with this property is equivalent to the group having property (T); a recent result of Chen, Wang and Wang characterizes the Haagerup property for residually finite groups in terms of their box spaces admitting a *fibred coarse embedding* into Hilbert space [CWW].

The equivalence between amenability of a group and property A of its box space provides us with a source of examples of spaces without property A.

Box spaces are also the first examples of spaces which coarsely embed into Hilbert space but do not have property A. Such a discrete space with unbounded geometry was first exhibited by Nowak [Now], in the form of a disjoint union of cubes of increasing dimension.

The first bounded geometry example was given by Arzhantseva, Guentner and Spakula, and it is this example that we shall explore here in detail.

Definition. Given $m \in \mathbb{N}$ and a group G, the *derived m-series* of G is a sequence of subgroups defined inductively by $G_1 = G$, $G_{i+1} = [G_i, G_i]G_i^m$, where G_i^m is the subgroup of G generated by mth powers of elements of G_i.

When G is free, the intersection $\cap G_i$ of all the G_i is trivial by a theorem of Levi (see Proposition 3.3 in Chapter 1 of [LS]), since each G_i is a proper characteristic subgroup of the previous G_{i-1}. For free groups it thus makes sense to talk about the box space corresponding to the derived m-series, for $m \geq 2$.

Theorem ([AGS]). *Given a finitely generated free group, the box space corresponding to the derived* 2*-series coarsely embeds into Hilbert space.*

We will also show that the example of [AGS] is one of a family of examples, as follows.

Theorem ([Khu]). *Given a finitely generated free group and an integer $m \geq 2$, the box space corresponding to the derived m-series coarsely embeds into Hilbert space.*

1. Walls and embeddings

We begin by noting that the existence of coarse embeddings into ℓ^1 and into ℓ^2 are equivalent, and so we will be content with embedding into ℓ^1 whenever it is more natural to do so.

When property A was first defined in [Yu], it was unclear to what extent it captured the notion of being coarsely embeddable into Hilbert space. The first example that showed that property A is in fact a stronger property was given by Nowak in [Now].

Given a finite group F with a fixed generating set S, consider the coarse disjoint union $\sqcup_{n\in\mathbb{N}} \oplus^n F$, where $\oplus^n F$ is the direct sum of n copies of F, and the metric on each $\oplus^n F$ is taken to be the standard direct sum metric induced by S, namely the metric with respect to the generating set

$$S \times \{1\} \times \cdots \times \{1\} \cup \{1\} \times S \times \cdots \times \{1\} \cup \cdots \cup \{1\} \times \cdots \times S.$$

Theorem 1 ([Now]). *Given any finite group F, the (locally finite) metric space $\sqcup_{n\in\mathbb{N}} \oplus^n F$ admits a bi-Lipschitz embedding into ℓ^1. This space does not have property A.*

Proof. For the proof of the second part of the statement, we refer the reader to [Now]. We will only prove the (easier) first part.

We need only show that each of the spaces in the coarse disjoint union can be embedded into Hilbert space with the same distortion functions ρ_-, ρ_+. Since F is finite, there is a *bilipschitz* map $\phi : F \to \ell^1(\mathbb{N})$ such that $\forall g, h \in F$

$$\frac{1}{C} d_F(g,h) \leq \|\phi(g) - \phi(h)\|_1 \leq C d_F(g,h),$$

for some $C > 0$, where d_F denotes the Cayley graph metric on F with respect to the generating set S. Now for any n, taking the map $\phi^n = \phi \times \cdots \times \phi : \oplus^n F \to (\oplus_{i=1}^n \ell^1)$, where $(\oplus_{i=1}^n \ell^1)$ is the ℓ^1-sum, we still have

$$\frac{1}{C} d_{\oplus^n F}(g,h) \leq \|\phi(g) - \phi(h)\|_1 \leq C d_{\oplus^n F}(g,h)$$

for every $g, h \in \oplus^n F$. Since $(\oplus_{i=1}^n \ell^1)$ is isometrically isomorphic to $\ell^1(\mathbb{N})$, we are done. $\square$

When the finite group F is taken to be $\mathbb{Z}_2$, there is another way to construct an embedding into ℓ^1. The space $\sqcup_{n\in\mathbb{N}} \oplus^n \mathbb{Z}_2$ is now a coarse disjoint union of n-dimensional cubes, whose special structure allows us to easily construct an embedding. First, we need some definitions.

Definition. Given a graph Γ, a *wall* (sometimes also called a *cut*) in Γ is a subset of the edges of Γ whose removal yields exactly two remaining connected components. A *wall structure* $\mathcal{W}$ on Γ is a set of walls in Γ such that each edge in Γ is contained in exactly one wall in $\mathcal{W}$.

We will write $\mathcal{W}(x|y)$ for the set of walls in $\mathcal{W}$ that, when removed, separate x and y, i.e. x and y end up in different connected components. A wall structure gives rise to a *wall metric* $d_{\mathcal{W}}$ on the graph, defined by $d_{\mathcal{W}}(x,y) := |\mathcal{W}(x|y)|$.

Given a graph Γ equipped with a wall structure $\mathcal{W}$, one can easily embed the metric space $(\Gamma, d_{\mathcal{W}})$ into $\ell^1(\mathcal{W})$, via $\phi : (\Gamma, d_{\mathcal{W}}) \to \ell^1(\mathcal{W})$, $\phi(x) = 1_{\mathcal{W}(x|x_0)}$, for some fixed basepoint x_0. Moreover, this embedding is easily seen to be isometric. If the wall metric can be compared to the original graph metric via a coarse equivalence, this gives a method for coarsely embedding the graph into ℓ^1. For example, given a tree, the wall structure that has a wall for each edge of the tree gives rise to the same metric as the original graph metric.

In the case of $\oplus^n\mathbb{Z}_2$, consider the wall structure $\mathcal{W}$ with a wall for each of the n generators of $\oplus^n\mathbb{Z}_2$ consisting of the edges labelled by that generator. This is clearly a wall structure and, in addition, the associated wall metric is precisely the Cayley graph metric on $\oplus^n\mathbb{Z}_2$ with respect to the given generating set. Thus, taking the isometric embedding into ℓ^1 induced by the wall structure on each component of $\sqcup_{n\in\mathbb{N}} \oplus^n \mathbb{Z}_2$, we get the desired embedding of the whole space.

While the above examples of spaces which are embeddable but do not have property A are uniformly discrete and locally finite, they do not have bounded geometry.

Definition. A metric space has *bounded geometry* if for each $R > 0$ there is an M such that the cardinality of each ball of radius R is bounded above by M.

For example, finitely generated groups and their box spaces have bounded geometry. The question of whether property A and coarse embeddability into Hilbert space are equivalent for bounded geometry metric spaces was answered in [AGS], where the above example of a space without bounded geometry was encoded in the structure of a box space of the free group $\mathbb{F}_n$ ($n \geq 2$). This space automatically doesn't have property A, since $\mathbb{F}_n$ is non-amenable. We will now look at such bounded geometry examples in detail.

2. Covers

Let us first describe the general construction of the cover $\widetilde{X}$ of a finite graph X corresponding to a finite quotient K of $\pi_1(X)$. Throughout, we will assume that X is 2-connected, i.e. removing any edge leaves X connected. Let ρ be the surjective homomorphism $\rho : \pi_1(X) \twoheadrightarrow K$.

Denote the vertex set of X by $V(X)$ and the edge set by $E(X)$. Choose a maximal tree $T \subset X$. The set of edges $\{e_1, e_2, ..., e_r\}$ which are not in the maximal tree T correspond to free generators of $\pi_1(X)$, and so we can consider their image under the quotient map ρ. The cover of X corresponding to ρ is the finite graph $\widetilde{X}$ with vertex set given by

$$V(\widetilde{X}) = V(X) \times K$$

and edge set given by

$$E(\widetilde{X}) = E(X) \times K.$$

We now just need to specify the vertices which are connected by each edge in $E(\widetilde{X})$.

Given an edge $(e, k) \in E(\widetilde{X})$ (where $e \in E(X)$ and $k \in K$), let v and w be the vertices of X connected by e. There are two cases: $e \in T$ and $e \notin T$. If $e \in T$, let (e, k) connect the vertices (v, k) and (w, k). If $e \notin T$, let (e, k) connect (v, k) and $(w, \rho(e)k)$. The graph $\widetilde{X}$ defined in this way is the cover of X corresponding to $\rho : \pi_1(X) \twoheadrightarrow K$. Note that the cover we obtain does not depend on the choice of spanning tree or on the chosen orientation of edges, i.e. it is unique up to graph isomorphism commuting with the covering projections (see Proposition 2.2 of [AGS]).

The covering map $\pi : \widetilde{X} \to X$ is given by $(e, k) \mapsto e$ and $(v, k) \mapsto v$. We can consider the subgraphs $V(X) \times k$ as k ranges over the elements of K. Following [AGS], we will call these subgraphs *clouds*. Note that collapsing the clouds to points yields the Cayley graph of the group K with respect to the generating set consisting of the images of the free generating set of $\pi_1(X)$.

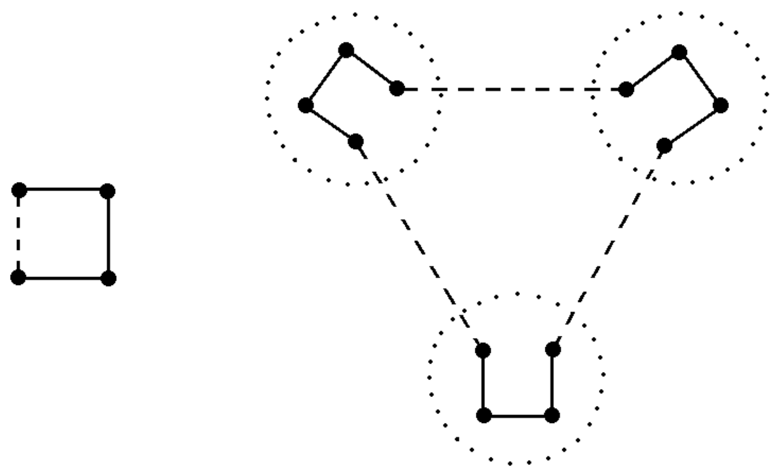

X $\widetilde{X}$

In the example above, $\widetilde{X}$ is the cover of X corresponding to the quotient $\rho : \pi_1(X) \cong \mathbb{Z} \twoheadrightarrow \mathbb{Z}_3$. We see copies of the (solid line) maximal tree of X inside the three clouds corresponding to elements of $\mathbb{Z}_3$, with edges which are lifts of the edge not in the maximal tree of X (represented by broken lines) connecting the clouds according to the quotient map ρ.

We will concentrate on the case where the cover $\widetilde{X}$ corresponds to the quotient

$$\pi_1(X) \longrightarrow \pi_1(X)/[\pi_1(X), \pi_1(X)]\pi_1(X)^m \cong \oplus^r \mathbb{Z}_m,$$

which we call the $\mathbb{Z}_m$-*homology cover* of X. Note that the Cayley graph of $\oplus^r \mathbb{Z}_m$ (where r is the free rank of $\pi_1(X)$) with respect to the image of the free generating set of $\pi_1(X)$ is the same as taking the natural generating set for $\oplus^r \mathbb{Z}_m$, namely, one generator $1 \in \mathbb{Z}_m$ for each copy of $\mathbb{Z}_m$. We will refer to the corresponding word metric as d_T.

For each element $x \in \widetilde{X}$, denote by C_x^T the cloud (with respect to the maximal spanning tree T) containing x. Since collapsing the clouds of $\widetilde{X}$ to points gives us the space $(\oplus^r \mathbb{Z}_m, d_T)$, the clouds are in one-to-one correspondence with elements of $\oplus^r \mathbb{Z}_m$. We will refer to clouds and points in $\oplus^r \mathbb{Z}_m$ interchangeably.

3. The case $m = 2$

We will begin by considering separately the case $m = 2$, which is the main result of [AGS]. This example is easier thanks to the aforementioned benefits of wall structures, and will be illuminating to us in the discussion of the general case $m \geq 2$.

Letting X and $\widetilde{X}$ be as in the previous section, we can define a wall structure on $\widetilde{X}$ as follows. For each edge e of X, consider the set of edges w_e of $\widetilde{X}$ given by $w_e := \pi^{-1}(e)$ (recalling that $\pi : \widetilde{X} \to X$ is the covering map). Defining $\mathcal{W} := \{w_e : e \in E(X)\}$, it is not difficult to see that this is a wall structure. In fact, following on from the discussion in the previous section, given an edge e of X, we can consider a maximal spanning tree T of X which does not contain e (this exists since X is assumed to be 2-connected). Considering $\widetilde{X}$ as the cover corresponding to this choice of maximal spanning tree, we can view it as clouds (corresponding to elements of $\oplus^r \mathbb{Z}_2$) which are connected to each other via edges exactly as the elements in the Cayley graph of $\oplus^r \mathbb{Z}_2$ are connected, with respect to the standard generating set.

We now see that the edges of $\widetilde{X}$ in w_e correspond exactly to edges between these clouds labelled by a particular generator of $\oplus^r \mathbb{Z}_2$ (namely, the generator $\rho(e)$). Thus, removing these edges yields exactly two connected components, just as the removal of edges labelled by a particular generator in the r-dimensional cube $\oplus^r \mathbb{Z}_2$ would leave two connected components. It is clear that each edge of $\widetilde{X}$ lies in precisely one wall of $\mathcal{W}$, and so $\mathcal{W}$ is a wall structure.

The corresponding wall metric $d_{\mathcal{W}}$ satisfies

$$d_{\mathcal{W}}(x, y) \leq d(x, y)$$

for all $x, y \in \widetilde{X}$, where d is the natural graph metric on $\widetilde{X}$. This is easy to see, since the walls are disjoint and given a d-geodesic from x to y (i.e. a path in $\widetilde{X}$ which realizes the distance $d(x, y)$), such a geodesic must traverse all the walls separating x and y at least once.

In [AGS], Arzhantseva, Guentner and Špakula go on to show that for every $x, y \in \widetilde{X}$,

$$d_{\mathcal{W}}(x, y) < \mathrm{girth}(X) \iff d(x, y) < \mathrm{girth}(X)$$

and if the above inequalities hold, then $d_{\mathcal{W}}(x, y) = d(x, y)$. It is this comparison between the metrics which eventually allows us to conclude that a particular box space of the free group coarsely embeds into ℓ^1 (see the section "Box spaces"). We will leave the proof of this relationship between the metrics until the next section, where we will prove a more general result which is based on a similar idea (see Remark 5 and Proposition 6).

4. Metrics on covers ($m \geq 2$)

Consider now $m \geq 2$. In this section, we will introduce a new metric on the cover $\widetilde{X}$, which generalises the situation for $m = 2$. This metric, just as the wall metric was designed to provide an embedding into ℓ^1, is designed to help us make use of the $\oplus^r \mathbb{Z}_m$ structure in the cover, in order to eventually achieve an embedding with the help of Nowak's result (Theorem 1).

We will deal with the following two metrics on $\widetilde{X}$. The first, d, is the natural graph metric on $\widetilde{X}$. The second, d_Q, is a metric which we will see comes from the $\oplus^r \mathbb{Z}_m$ structure present in $\widetilde{X}$. For each $x, y \in \widetilde{X}$, choose a geodesic geod$[x, y]$ between x and y with respect to the metric d, such that the geodesic geod$[y, x]$ is the geodesic geod$[x, y]$ travelled backwards.

For each edge $e \in E(X)$, choose an orientation. Define a function

$$\phi : E(X) \times \widetilde{X} \times \widetilde{X} \longrightarrow \mathbb{N}$$

by setting $\phi(e, x, y)$ to be the smallest non-negative residue modulo m of

$$\big| |\pi^{-1}(e) \cap \mathrm{geod}[x, y]| - |\pi^{-1}(e^{-1}) \cap \mathrm{geod}[x, y]| \big|,$$

where by $|\pi^{-1}(e) \cap \mathrm{geod}[x, y]|$, we mean the number of times that a lift of e occurs in the geodesic (with positive orientation), and by $|\pi^{-1}(e^{-1}) \cap \mathrm{geod}[x, y]|$, we mean the number of times a lift of the edge e occurs with reversed orientation. Let N_e be the number of maximal spanning trees which do not contain a given edge e.

Definition 2. *For $x, y \in \widetilde{X}$, define the metric d_Q by*

$$d_Q(x, y) := \sum_{e \in E(X)} \sum_{T : e \notin T} \frac{1}{N_e} \min\{\phi(e, x, y), m - \phi(e, x, y)\},$$

where the first sum ranges over all edges in X and the second sum ranges over all maximal spanning trees T of X which do not contain a given edge e.

Since X is always assumed to be 2-connected, the second sum is never empty and N_e is non-zero for all e. Note that the sum over $\{T : e \notin T\}$ is in fact a "dummy-sum" here, since the sum does not depend on T and we immediately go on to divide by N_e, the cardinality of this set. We include this sum so that we can later easily rearrange this definition to see how this metric helps us to embed in ℓ^1.

Proposition 3. *For all $x, y \in \widetilde{X}$, we have $d_Q(x, y) \leq d(x, y)$.*

Proof. We note that in Definition 2, the terms in the sum do not depend on T. Since N_e is the cardinality of the set $\{T : e \notin T\}$ by definition, we have

$$\begin{aligned} d_Q(x,y) &= \sum_{e\in E(X)} \sum_{T:e\notin T} \frac{1}{N_e} \min\{\phi(e,x,y), m-\phi(e,x,y)\} \\ &= \sum_{e\in E(X)} \min\{\phi(e,x,y), m-\phi(e,x,y)\} \\ &\leq \sum_{e\in E(X)} \phi(e,x,y) \\ &\leq \sum_{e\in E(X)} \left| |\pi^{-1}(e)\cap \mathrm{geod}[x,y]| - |\pi^{-1}(e^{-1})\cap \mathrm{geod}[x,y]| \right| \\ &\leq \sum_{e\in E(X)} (|\pi^{-1}(e)\cap \mathrm{geod}[x,y]| + |\pi^{-1}(e^{-1})\cap \mathrm{geod}[x,y]|) \\ &= |\,\mathrm{geod}[x,y]| = d(x,y). \end{aligned}$$

□

Suppose that X is such that $N_e = N$ is independent of the choice of edge e (which we will show to be the case in the situation we will be interested in). Then we have

$$d_Q(x,y) = \frac{1}{N}\sum_T \sum_{e\notin T} \min\{\phi(e,x,y), m-\phi(e,x,y)\}.$$

Now note that given a maximal spanning tree T and $x,y\in \widetilde{X}$, the sum

$$\sum_{e\notin T} \min\{\phi(e,x,y), m-\phi(e,x,y)\}$$

is exactly the distance between the clouds C_x^T and C_y^T containing x and y respectively in $\oplus^r \mathbb{Z}_m$ with the metric d_T. This is because given an element $(z_1, ..., z_r)$ in $\oplus^r\mathbb{Z}_m$, written additively, the geodesics from the identity to $(z_1, ..., z_r)$ are exactly those paths which for each i, contain z_i edges corresponding to the generator of the ith factor if $z_i \leq m/2$, or $m - z_i$ edges corresponding to the inverse of the generator of the ith factor if $z_i > m/2$, in any order. The minimum in the sum above therefore ensures that we get the d_T geodesic. Thus, we have proved the following.

Proposition 4. *Let X and $\widetilde{X}$ be as described above, with the additional requirement that X is such that $N_e = N$ is independent of the choice of edge e. Then*

$$d_Q(x,y) = \frac{1}{N}\sum_T d_T(C_x^T, C_y^T).$$

Note that this in particular proves that d_Q is a pseudometric, since the triangle inequality is obvious from the above. We will see that $d_Q(x,y) = 0$ if and only if $x = y$ later: it will follow from Proposition 6, when we compare d_Q with the metric d. Written like this, the metric d_Q will help us embed our space into ℓ^1, as it can be seen as a sum of metrics d_T on the embeddable space $\oplus^r\mathbb{Z}_m$ (see Proposition 10).

Remark 5. *In the case $m = 2$, the metric d_Q is exactly the wall metric discussed in the previous section. This follows from the fact that for $m = 2$, $\phi(e,x,y)$ is always either 0 or 1, depending on whether preimages of the edge e (with either* *orientation) occur an even or an odd number of times, respectively, in* geod$[x, y]$. *This means that in this case the definition of d_Q can be rewritten as*

$$\begin{aligned} d_Q(x,y) &:= \sum_{e\in E(X)} \text{number of times } e \text{ or } e^{-1} \text{ occurs in } \mathrm{geod}[x,y] \bmod 2 \\ &= \sum_{e\in \mathcal{W}(x|y)} 1 \;=\; d_{\mathcal{W}}(x,y). \end{aligned}$$

Thus, for $m = 2$, there is no need to rewrite the metric as we did in Proposition 4 for the general case, as it is clear from the above that $d_Q = d_{\mathcal{W}}$ is embeddable into ℓ^1.

For us, d_Q will be the metric which we will use to embed $\widetilde{X}$ into ℓ^1. However, we are interested in embedding $\widetilde{X}$ with its original metric, so we now need a way to compare the two metrics. The following proposition, which is inspired by Proposition 3.11 in [AGS], will do this for us.

Proposition 6. *If X and $\widetilde{X}$ are as described above, then for every $x, y \in \widetilde{X}$, we have*

$$d_Q(x,y) < \mathrm{girth}(X) \iff d(x,y) < \mathrm{girth}(X)$$

and if the above inequalities hold, then $d_Q(x,y) = d(x,y)$.

Proof. First assume that $d(x,y) < \mathrm{girth}(X)$. We have seen that $d_Q \leq d$ and so we have that $d_Q(x,y) < \mathrm{girth}(X)$. We now prove that when $d(x,y) < \mathrm{girth}(X)$, we have $d_Q(x,y) = d(x,y)$.

Projecting a geodesic geod$[x,y]$ for the metric d onto X, we see that there can be no repeated edges in this path since $d(x,y) < \mathrm{girth}(X)$. We now make the remark that in $\oplus^r \mathbb{Z}_m$, any path without multiple occurrences of edges labelled by the same generator must be a geodesic. Hence, for any choice of maximal tree T, the path between C_x^T and C_y^T induced in the corresponding $\oplus^r \mathbb{Z}_m$ by the d-geodesic must be a d_T-geodesic. We thus have (using the fact that a lift of either e or e^{-1} appears in the d-geodesic at most once):

$$\begin{aligned} d_Q(x,y) &= \frac{1}{N}\sum_T d_T(C_x^T, C_y^T) \\ &= \frac{1}{N}\sum_T \sum_{e\notin T} (|\pi^{-1}(e)\cap \mathrm{geod}[x,y]| + |\pi^{-1}(e^{-1})\cap \mathrm{geod}[x,y]|) \\ &= \sum_{e\in E(X)} \sum_{T: e\notin T} \frac{1}{N}(|\pi^{-1}(e)\cap \mathrm{geod}[x,y]| + |\pi^{-1}(e^{-1})\cap \mathrm{geod}[x,y]|) \\ &= \sum_{e\in E(X)} (|\pi^{-1}(e)\cap \mathrm{geod}[x,y]| + |\pi^{-1}(e^{-1})\cap \mathrm{geod}[x,y]|) \\ &= |\,\mathrm{geod}[x,y]| = d(x,y). \end{aligned}$$

Now it remains to show that $d_Q(x,y) < \mathrm{girth}(X)$ implies $d(x,y) < \mathrm{girth}(X)$. Assume that $d_Q(x,y) < \mathrm{girth}(X)$ and consider the projection $p[\pi(x), \pi(y)]$ of a d-geodesic onto X. Note that this is a shortest path in X which lifts to a d-geodesic in $\widetilde{X}$ (i.e. no shorter path in X can be lifted to a path in $\widetilde{X}$ between x and y), and that this path does not contain backtracks.

If there are no repeated edges in $p[\pi(x), \pi(y)]$, the path between C_x^T and C_y^T in $\oplus^r \mathbb{Z}_m$ induced by the d-geodesic is necessarily a geodesic. Hence if $d(x,y) \geq \text{girth}(X)$, we have

$$\begin{aligned}
d_Q(x,y) &= \frac{1}{N} \sum_T d_T(C_x^T, C_y^T) \\
&= \frac{1}{N} \sum_T \sum_{e \notin T} (|\pi^{-1}(e) \cap \text{geod}[x,y]| + |\pi^{-1}(e^{-1}) \cap \text{geod}[x,y]|) \\
&= \sum_{e \in E(X)} \sum_{T: e \notin T} \frac{1}{N} (|\pi^{-1}(e) \cap \text{geod}[x,y]| + |\pi^{-1}(e^{-1}) \cap \text{geod}[x,y]|) \\
&= \sum_{e \in E(X)} (|\pi^{-1}(e) \cap \text{geod}[x,y]| + |\pi^{-1}(e^{-1}) \cap \text{geod}[x,y]|) \\
&= |\,\text{geod}[x,y]| = d(x,y) \\
&\geq \text{girth}(X),
\end{aligned}$$

and so $d(x,y) \geq \text{girth}(X)$ would imply that $d_Q(x,y) \geq \text{girth}(X)$ which is a contradiction.

Now let us consider what happens when the path $p[\pi(x), \pi(y)]$ contains repeated edges. The arguments used are of a *modular graph theory* flavour. For simplicity, we will use the following definitions and prove a simple result before continuing with the proof of Proposition 6.

Definition 7. *Given a path $p[a,b]$ from $a \in V(X)$ to $b \in V(X)$ in a graph X, write $|e \cap p[a,b]|$ for the number of times $p[a,b]$ traverses e in the positive direction and $|e^{-1} \cap p[a,b]|$ for the number of times $p[a,b]$ traverses e in the opposite direction. We will call an edge e on this path m-*repeated *if*

$$0 \neq |e \cap p[a,b]| - |e^{-1} \cap p[a,b]| \equiv 0 \mod m.$$

*We say that two paths $p_1[a,b]$ and $p_2[a,b]$ from a to b are m-*congruent *if for all $e \in E(X)$,*

$$|e \cap p_1[a,b]| - |e^{-1} \cap p_1[a,b]| \equiv |e \cap p_2[a,b]| - |e^{-1} \cap p_2[a,b]| \mod m.$$

*A shorter path which is m-congruent to a path $p[a,b]$ will be called an m-*shortcut *for $p[a,b]$.*

Lemma 8. *Given two paths $p_1[a,b]$ and $p_2[a,b]$ from a to b in X which are m-congruent, take the lifts of these paths in the $\mathbb{Z}_m$-homology cover $\widetilde{X}$ of X, both starting at a point $x \in \pi^{-1}(a)$. Then both of these lifts end at the same point $y \in \pi^{-1}(b)$ of $\widetilde{X}$.*

Proof. Pick some maximal spanning tree T in X, and consider the clouds in $\widetilde{X}$ corresponding to T. Taking lifts of both of the paths starting at $x \in \pi^{-1}(a)$, the condition

$$|e \cap p_1[a,b]| - |e^{-1} \cap p_1[a,b]| \equiv |e \cap p_2[a,b]| - |e^{-1} \cap p_2[a,b]| \mod m$$

for the edges not in T implies that both of the lifts terminate in the same cloud, since the cloud depends only on the number of times the edges not in T are traversed modulo m. This now uniquely determines the point $y \in \pi^{-1}(b)$ and so we are done. □

We now continue with the proof of Proposition 6. Recall that we are assuming that $d_Q(x,y) < \mathrm{girth}(X)$, and aiming to show that this implies $d(x,y) < \mathrm{girth}(X)$. We have proved this in the case that the path $p[\pi(x),\pi(y)]$ has no repeated edges. Observe that repeated edges in $p[\pi(x),\pi(y)]$ necessarily imply that $d(x,y) \geq \mathrm{girth}(X)$ and so it remains to prove that the assumptions of repeated edges in $p[\pi(x),\pi(y)]$ and $d_Q(x,y) < \mathrm{girth}(X)$ lead to a contradiction.

We will write $\mathbb{Z}_m E(X)$ to mean the set of functions from $E(X)$ to $\mathbb{Z}_m$. Writing δ_e for the function which takes the value $1 \in \mathbb{Z}_m$ on the edge e and the value $0 \in \mathbb{Z}_m$ elsewhere, we can express any function in $\mathbb{Z}_m E(X)$ as a sum $\sum_{e\in E(X)} \alpha_e \delta_e$, where the $\alpha_e \in \mathbb{Z}_m$ are the values of the function on each edge e.

Consider the element

$$\sum_{e\in E(X)} \left(|e \cap p[\pi(x),\pi(y)]| - |e^{-1} \cap p[\pi(x),\pi(y)]|\right) \delta_e \in \mathbb{Z}_m E(X).$$

Note that if we remove all m-repeated edges from the path $p[\pi(x),\pi(y)]$, we still obtain the same element of $\mathbb{Z}_m E(X)$. For each edge e, let v_e denote the initial vertex and let w_e denote the terminal vertex. We can now apply the *boundary operator modulo m*, i.e. $\partial_m : \mathbb{Z}_m E(X) \longrightarrow \mathbb{Z}_m V(X)$ given by

$$\partial_m \left(\sum_{e\in E(X)} \alpha_e \delta_e \right) = \sum_{e\in E(X)} \alpha_e (\delta_{w_e} - \delta_{v_e}) = \sum_{v\in V(X)} \beta_v \delta_v.$$

Since we started with a path $p[\pi(x),\pi(y)]$, we have that

$$\partial_m \left(\sum_{e\in E(X)} \left(|e \cap p[\pi(x),\pi(y)]| - |e^{-1} \cap p[\pi(x),\pi(y)]|\right) \delta_e \right) = \delta_{\pi(y)} - \delta_{\pi(x)}.$$

Remove the m-repeated edges from $p[\pi(x),\pi(y)]$. This yields the same element of $\mathbb{Z}_m E(X)$, and so when we apply the boundary operator ∂_m, we still get $\delta_{\pi(y)} - \delta_{\pi(x)}$.

Assume first that $\pi(x) \neq \pi(y)$. Considerations of the boundary operator above imply that $p[\pi(x),\pi(y)]$ with the m-repeated edges removed is the union of a path $p'[\pi(x),\pi(y)]$ from $\pi(x)$ to $\pi(y)$ and possibly some loops which are disjoint from this path.

If the graph of the path $p[\pi(x),\pi(y)]$ with the m-repeated edges removed is connected, i.e. there are no loops disjoint from $p'[\pi(x),\pi(y)]$, then $p[\pi(x),\pi(y)]$ once the m-repeated edges have been removed is simply the path $p'[\pi(x),\pi(y)]$, which is m-congruent to the path $p[\pi(x),\pi(y)]$.

If $p'[\pi(x),\pi(y)]$ does not contain loops, it is also an m-shortcut for $p[\pi(x),\pi(y)]$, since we obtained it by removing m-repeated edges. Now Lemma 8 tells us that two m-congruent paths can be lifted to the same path in $\widetilde{X}$, and so the existence of an m-shortcut for the path $p[\pi(x),\pi(y)]$ is a contradiction to this path being the shortest path in X which can be lifted to a path in $\widetilde{X}$ between x and y.

If $p'[\pi(x),\pi(y)]$ contains loops, then each of the edges on such a loop must contribute at least 1 to

$$\sum_{T:a\notin T} \frac{1}{N} \min\{\phi(a,x,y), m - \phi(a,x,y)\}$$

in the total sum

$$d_Q(x,y) = \sum_{a \in E(X)} \sum_{T : a \notin T} \frac{1}{N} \min\{\phi(a,x,y), m - \phi(a,x,y)\},$$

whence $d_Q(x,y)$ is at least the length of this loop, which contradicts our assumption that $d_Q(x,y) < \operatorname{girth}(X)$.

If the graph of the path $p[\pi(x), \pi(y)]$ has become disconnected upon removal of the m-repeated edges, this means that it is the disjoint union of a path and a non-zero number of loops. Thus, in our original path $p[\pi(x), \pi(y)]$, we have found at least one loop of edges which are traversed non-zero modulo m times and so the same conclusion holds, as above.

Assume now that $\pi(x) = \pi(y)$. In this case, the graph of the path $p[\pi(x), \pi(y)]$ with the m-repeated edges removed must be trivial, or must be a union of a non-zero number of loops. If it is trivial, then the trivial path consisting of the vertex $\pi(x)$ is an m-shortcut for the path $p[\pi(x), \pi(y)]$ which is a contradiction, as above. If it is a union of a non-zero number of loops, we again deduce the contradiction $d_Q(x,y) \geq \operatorname{girth}(X)$.

This completes the proof of the implication

$$d_Q(x,y) < \operatorname{girth}(X) \Rightarrow d(x,y) < \operatorname{girth}(X)$$

and so Proposition 6 is proved. □

Corollary 9. *d_Q is a metric.*

Proof. We have already seen that d_Q is a pseudometric from the description of d_Q in Proposition 4. That $d_Q(x,y) = 0$ if and only $x = y$ is now obvious from Proposition 6. □

Proposition 10. *Let X be a finite 2-connected graph such that the number $N_e = N$ of maximal spanning trees not containing a given edge e is independent of the edge chosen. Its $\mathbb{Z}_m$-homology cover $\widetilde{X}$ admits a coarse embedding into ℓ^1 with respect to the metric d_Q such that the functions $\rho_\pm$ depend only on m, and not on X.*

Proof. Let d_m be the metric on $\sqcup_{n \in \mathbb{N}} \oplus^n \mathbb{Z}_m$ which on each component is the Cayley graph metric coming from taking one generator for each factor of $\oplus^n \mathbb{Z}_m$, and such that distances between components tend to infinity. Nowak's theorem (Theorem 1) tells us that there is a bi-Lipschitz embedding

$$\phi : \sqcup_{n \in \mathbb{N}} \oplus^n \mathbb{Z}_m \longrightarrow \ell^1,$$

i.e. there exists a $C > 0$ such that

$$\frac{1}{C} d_m(a,b) \leq \|\phi(a) - \phi(b)\|_1 \leq C d_m(a,b)$$

for all $a, b \in \sqcup_{n \in \mathbb{N}} \oplus^n \mathbb{Z}_m$. Let r denote the m-rank of $\pi_1(X)$. Recall that for a point x in the cover $\widetilde{X}$, C_x^T denotes the cloud (corresponding to some point in $\oplus^r \mathbb{Z}_m$) containing x with respect to the maximal tree T. Define $\psi : \widetilde{X} \longrightarrow \oplus_T \ell^1$ by

$$x \longmapsto \oplus_T \frac{1}{N} \phi(C_x^T).$$

Here $\oplus_T \ell^1$ is equipped with the ℓ^1-norm given by the sum of the norms on the factors, together with which $\oplus_T \ell^1$ is isometric to ℓ^1. This embedding satisfies

$$\|\psi(x)-\psi(y)\|_1 = \| \oplus_T \frac{1}{N}\phi(C_x^T) - \oplus_T \frac{1}{N}\phi(C_y^T)\|_1 = \frac{1}{N}\sum_T \|\phi(C_x^T)-\phi(C_y^T)\|_1$$

and thus we have

$$\begin{aligned}\frac{1}{C}d_Q(x,y) &= \frac{1}{C}\frac{1}{N}\sum_T d_T(C_x^T, C_y^T)\\ &\le \frac{1}{N}\sum_T \|\phi(C_x^T)-\phi(C_y^T)\|_1\\ &= \|\psi(x)-\psi(y)\|_1\\ &\le \frac{1}{N}C\sum_T d_T(C_x^T, C_y^T)\\ &= Cd_Q(x,y).\end{aligned}$$

Note that C only depends on m, and so the proof is complete. □

5. Box spaces

We can now add all of the ingredients of the previous sections to get the following general result for $m \geq 2$, which we will then apply to box spaces of free groups.

Theorem 11. *Let $\{X_n\}$ be a sequence of 2-connected finite graphs such that for each n, the number of maximal spanning trees in X_n not containing a given edge does not depend on the edge. Given $m \in \mathbb{N}$, let $\{\widetilde{X}_n\}$ be the sequence of $\mathbb{Z}_m$-homology covers of the X_n. If $\mathrm{girth}(X_n) \to \infty$ as $n \to \infty$, then the coarse disjoint union $\sqcup_n \widetilde{X}_n$ coarsely embeds into Hilbert space.*

Let N_n be the number of maximal spanning trees of X_n not containing a given edge. Using the method of the previous section, one can define a metric d_{Q_n} on each $\widetilde{X}_n$ by $d_{Q_n}(x,y) = \frac{1}{N_n}\sum_T d_T(C_x^T, C_y^T)$, where the sum is taken over all maximal spanning trees of X_n. We will write d_Q to mean the coarse disjoint union metric which is d_{Q_n} on each component $\widetilde{X}_n$. Let d denote the coarse disjoint union metric on $\sqcup_n \widetilde{X}_n$ which restricts to the natural graph metric on each component.

We will first need the following proposition, which is proved exactly as Proposition 4.5 of [AGS], using the comparison of metrics on the scale of the girth that we proved in the previous section.

Proposition 12. *The identity map between the metric spaces formed by taking $\sqcup_n \widetilde{X}_n$ with the two different metrics d and d_Q is a coarse equivalence, i.e. the identity map and its inverse are both coarse embeddings.*

Proof. It was proved in the previous section that d_{Q_n} is always smaller than the natural graph metric on each $\widetilde{X}_n$ and so we need only prove that for each $R > 0$ there is an $S > 0$ such that $d_Q(x,y) \leq R$ implies $d(x,y) \leq S$ for all $x, y \in \sqcup_n \widetilde{X}_n$.

Given $R > 0$, take $M > 0$ such that for all $n, m \geq M$, we have $\mathrm{girth}(X_n) > R$ and the distance d_Q between components $\widetilde{X}_n$ and $\widetilde{X}_m$ is greater than R. Let $S := \max\{R, d(x,y) : x, y \in \sqcup_{n<M} \widetilde{X}_n\}$.

Now if $x, y \in \sqcup_n \widetilde{X}_n$ with $d_Q(x,y) \leq R$, then either $x, y \in \sqcup_{n<M} \widetilde{X}_n$ whence $d(x,y) \leq S$ by the definition of S, or $x, y \in \widetilde{X}_n$ for $n \geq M$. For this $\widetilde{X}_n$, Proposition 6 tells us that the restriction of the identity map to balls of radius $R < \text{girth}(X_n)$ in $\widetilde{X}_n$ is an isometry. This means we have $d(x,y) = d_Q(x,y) \leq R \leq S$ and so the proof is complete. □

Proof of Theorem 11. Proposition 10 in the previous section tells us that each $\widetilde{X}_n$, being a $\mathbb{Z}_m$-homology cover, admits a coarse embedding into ℓ^1 with respect to the metric d_Q. Moreover, the embedding functions $\rho_\pm$ depend only on $m \in \mathbb{N}$ and not on X_n, and so the embedding is uniform over all n. Since ℓ^1 coarsely embeds into ℓ^2 (and, in fact, into any ℓ^p space for $1 \leq p \leq \infty$) by a result of Nowak [Now2], $\sqcup_n \widetilde{X}_n$ with the metric d_Q coarsely embeds into Hilbert space. Now the theorem is proved since by Proposition 12 above the metric d_Q is coarsely equivalent to d, the metric arising from the natural graph metrics on $\sqcup_n \widetilde{X}_n$. □

We now apply Theorem 11 to certain box spaces of free groups. Let G be the free group on n generators, and let $m \geq 2$. Let $\{G_i\}$ denote the derived m-series of G. Note that this is a nested sequence of finite index characteristic subgroups of G, so it makes sense to talk about the box space $\square_{\{G_i\}} G$. Each successive quotient G/G_{i+1} is the $\mathbb{Z}_m$-homology cover of G/G_i.

To apply the above, set X_n to be G/G_n. We then have that $\sqcup_n \widetilde{X}_n$ is equal to $\sqcup_n G/G_{n+1}$ and so to show that the box space $\square_{\{G_i\}} G$ is embeddable, we need to show that the assumptions of Theorem 11 are satisfied. It is clear that the graphs G/G_n are 2-connected, and since $\sqcup_n G/G_{n+1}$ is a box space of the free group, we have that $\text{girth}(G/G_n) \to \infty$ as $n \to \infty$. It remains to show that for each n, the number of maximal spanning trees in G/G_n not containing a given edge does not depend on the edge chosen. This is clear, since permuting the generators of G induces an isomorphism of G/G_n, and hence a graph isomorphism of the corresponding Cayley graph.

We have obtained the following.

Corollary 13. *Given a finitely generated free group $\mathbb{F}_n$ and $m \geq 2$, the box space corresponding to the derived m-series of $\mathbb{F}_n$ coarsely embeds into Hilbert space.*

In particular, taking such a box space of a free group $\mathbb{F}_n$ with $n \geq 2$ gives a family of examples of bounded geometry metric spaces which coarsely embed into Hilbert space but do not have property A, since non-abelian free groups are not amenable.

References

[AGS] G. Arzhantseva, E. Guentner and J. Špakula, *Coarse non-amenability and coarse embeddings*, Geom. Funct. Anal. 22, 2012

[CWW] X. Chen, Q. Wang and X. Wang, *Characterization of the Haagerup property by fibred coarse embedding into Hilbert space*, Bull. Lond. Math. Soc. 45, 2013

[Khu] A. Khukhro, *Embeddable box spaces of free groups*, Math. Ann. 360, 2014

[LZ] A. Lubotzky and A. Żuk, *On property (τ)*, preliminary version available online at `www.ma.huji.ac.il/~alexlub/`, 2003

[LS] R. Lyndon and P. Schupp, *Combinatorial group theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete 89, 1977

[Mar] G. Margulis, *Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators*, Problems Inform. Transmission 24, 1988

[Now2] P. Nowak, On coarse embeddability into ℓ_p-spaces and a conjecture of Dranishnikov, Fund. Math. 189, 2006

[Now] P. Nowak, *Coarsely embeddable metric spaces without property A*, J. Funct. Anal. 252, 2007

[Roe] J. Roe, *Lectures on coarse geometry*, AMS University Lecture Series, Volume 31, 2003

[Wil] R. Willett, *Property A and girth of graphs*, Journal of Topology and Analysis 3, 2011

[WY] R. Willett and G. Yu, *Higher index theory for certain expanders and Gromov monster groups II*, Adv. Math. 229, 2012

[Yu] G. Yu, *The coarse Baum–Connes conjecture for spaces which admit a uniform embedding into Hilbert space*, Inventiones Mathematicae 139, 2000

Université de Neuchâtel, anastasia.khukhro@unine.ch