

COARSE NON-AMENABILITY AND COARSE EMBEDDINGS

GOULNARA ARZHANTSEVA, ERIK GUENTNER, AND JÁN ŠPAKULA

ABSTRACT. We construct the first example of a coarsely non-amenable (= without Guoliang Yu's property A) metric space with bounded geometry which coarsely embeds into a Hilbert space.

1. INTRODUCTION

The purpose of this paper is to prove the following theorem:

Theorem 1.1. *There exists a uniformly discrete metric space with bounded geometry, which coarsely embeds into a Hilbert space, but does not have property A.*

The concept of coarse embedding was introduced by Gromov [3, p. 218] in relation to the Novikov conjecture (1965) on the homotopy invariance of higher signatures for closed manifolds.

Definition 1.2. A metric space X is said to be *coarsely embeddable* into a Hilbert space \mathcal{H} if there exists a map $f: X \rightarrow \mathcal{H}$ such that for any $x_n, y_n \in X$, $n \in \mathbb{N}$,

$$\text{dist}(x_n, y_n) \rightarrow \infty \iff \|f(x_n) - f(y_n)\|_{\mathcal{H}} \rightarrow \infty.$$

Yu established the coarse Baum-Connes conjecture (1995) in topology for every coarsely embeddable discrete space with bounded geometry [14, Theorem 1.1]. This implies the Novikov conjecture for all closed manifolds whose fundamental group, viewed with the word length metric, coarsely embeds into a Hilbert space. The result confirmed Gromov's intuition and sparked an intense study of groups and metric spaces which are coarsely embeddable into Hilbert space.

The coarse embeddability, which is a geometric property by nature, shares many features with Property A, a weak form of amenability, introduced by Yu by the following Følner-type condition.

Definition 1.3. A discrete metric space X has *Property A* if for every $\varepsilon > 0$ and every $R > 0$ there is a family $\{A_x\}_{x \in X}$ of finite subsets of $X \times \mathbb{N}$ and a number $S > 0$ such that

Date: January 9, 2011.

2000 Mathematics Subject Classification. Primary 20F69; Secondary 43A07, 05C25, 20F65, 46T99.

Key words and phrases. Amenability, coarse embeddings, graph coverings.

The research of the first author was partially supported by the Swiss NSF, Sinergia Grant CRSI22_130435 and CNRS, UMR 6632. The second author was partially supported by NSF Grant DMS-0349367. The third author was supported by the Deutsche Forschungsgemeinschaft (SFB 878).

- $\frac{|A_x \triangle A_y|}{|A_x \cap A_y|} < \varepsilon$ whenever $d(x, y) \leq R$,
- $A_x \subseteq B(x, S) \times \mathbb{N}$ for every $x \in X$.

Just as amenability, Property A has a large number of significant applications, see a survey [13]. For a countable discrete group Γ , it is equivalent to the existence of a topological amenable action of Γ on a compact Hausdorff space and to the C^* -exactness of the reduced C^* -algebra $C_r^*(\Gamma)$ [9, 7].

Property A implies coarse embeddability [14, Th.2.2]. Moreover, discrete spaces with property A provide the largest known class of spaces admitting such an embedding. Conversely, whether or not the existence of a coarse embedding of a bounded geometry discrete metric space guarantees property A is a crucial open problem that has attracted much research in the area.

Our main result, Theorem 1.1, yields a negative answer to this problem. Here is the construction of the counterexample. For a discrete group Γ , we denote by $\Gamma^{(2)}$ its normal (in fact characteristic) subgroup generated by all the squares of elements of Γ . Let \mathbb{F}_2 be the free group of rank two. We define inductively a sequence of normal subgroups of \mathbb{F}_2 by letting $\Gamma_0 = \mathbb{F}_2$ and $\Gamma_n = \Gamma_{n-1}^{(2)}$, $n \geq 1$. We denote by X_n the Cayley graph of \mathbb{F}_2/Γ_n with respect to the canonical image of free generators of \mathbb{F}_2 .

Theorem 1.4. *The space $X = \bigsqcup_{n=0}^{\infty} X_n$, (the box space of \mathbb{F}_2 associated with the family $(\Gamma_n)_{n \geq 1}$) is coarsely embeddable into a Hilbert space, but does not have property A.*

In outline the proof goes as follows. We describe the Cayley graphs X_n in a graph-theoretical way as a tower of successive $\mathbb{Z}/2$ -homology covers, starting with the “figure eight” graph. Next, we define a wall structure on each X_n , which gives rise to a wall metric on the graphs X_n . We show that these graphs, endowed with the wall metric, coarsely embed into a Hilbert space (uniformly), but their natural graph metric does not coincide with the wall metric. Nevertheless, we prove that on the small scale it does, which is then enough to conclude that the two metrics are coarsely equivalent. Consequently, X is also coarsely embeddable. On the other hand, it is a fact that $\bigcap_{n \geq 1} \Gamma_n = \{1\}$, so if X would have property A, the free group \mathbb{F}_2 would be amenable, which is certainly not the case.

The structure of the paper is as follows: in Section 2 we review the construction of graph coverings, in particular $\mathbb{Z}/2$ -homology coverings, and list their properties. In Section 3 we define a wall structure on the $\mathbb{Z}/2$ -homology cover of a graph and investigate its relation to the graph metric. In Section 4 we prove Theorem 1.4. In the last Section 5 we discuss applications to C^* -algebras, namely to variants of the uniform Roe C^* -algebras.

2. CONSTRUCTING GRAPH COVERINGS

2.1. Graph terminology. For a graph G , that is, an 1-dimensional simplicial complex, we denote by $V(G)$ the set of vertices and by $E(G)$ the set of edges of G . Our graphs will

be unoriented, with a couple of exceptions which are specified in the text later. We shall talk about an edge e being *between* two vertices x and y . We allow edges with endpoints being the same (sometimes in the literature called “loops”, but we shall not call them that as not to confuse them with loops defined below) and multiple edges between pairs of vertices.

A *path* p in G will mean a sequence $(v_0, e_1, v_1, e_2, v_2, \dots, v_n)$, such that $v_i \in V(G)$, $e_i \in E(G)$ and each e_i is an edge between v_{i-1} and v_i . We denote $E(p) = (e_1, e_2, \dots, e_n)$ the sequence of edges of p . We also write $\ell(p) = n$ for the *length* of the path p .

A path $l = (v_0, e_1, v_1, e_2, v_2, \dots, v_n)$ is called a *loop*, if $v_0 = v_n$ and $n > 0$ (i.e. we don't consider one vertex as a loop). The loop l is called *simple*, if there are no repeated vertices or edges, i.e. the vertices v_0, \dots, v_{n-1} are all different and the same holds for the edges e_1, \dots, e_n . The *girth* of a graph G is the length of a shortest simple loop.

When we talk about topological properties of G (e.g. connectedness), we consider G as a realization of a 1-dimensional simplicial complex determined by the structure of G .

2.2. Graph coverings: general case. The notion of graph covering is essentially a restriction to the case of graphs of the general topological notion of covering. For the sake of completeness, we recall this classical construction (see, for example, [11, Ch. 2] or [5, Ch. 1]) and subsequently specialize it to the $\mathbb{Z}/2$ -homology covering situation.

We start with the following data: a finite connected graph G and a surjective homomorphism $\rho : \pi_1(G) \twoheadrightarrow K$ from the fundamental group of G onto a finite group K .

We construct the corresponding *covering graph* \tilde{G} of G as follows. Choose a spanning tree T of G , that is, a contractible subgraph of G with $V(T) = V(G)$. Note that then necessarily $|E(T)| = |V(T)| - 1$. The remaining edges of G , denoted by $S = E(G) \setminus E(T)$, identify $\pi_1(G)$ with the free group $\mathbb{F}(S)$ on S : contracting T to a point gives a homotopy equivalence of G with a bouquet of $|S|$ circles. With this identification, we consider the domain of ρ to be $\mathbb{F}(S)$, that is, $\rho : \mathbb{F}(S) \twoheadrightarrow K$. Finally, we choose an orientation on each edge in S , we shall use the self-explanatory notation $x \xrightarrow{e} y$. Then the covering graph \tilde{G} is defined by:

- $V(\tilde{G}) = V(G) \times K$,
- $E(\tilde{G}) = E(G) \times K$, where the incidence of an edge $(e, k) \in E(\tilde{G})$ is as follows:
 - if $e \in E(T)$ is an edge between x and y , then (e, k) connects (x, k) with (y, k) ,
 - if $e \in S$ and $x \xrightarrow{e} y$, then (e, k) connects (x, k) with $(y, \rho(e)k)$.

The *covering projection* $\pi : \tilde{G} \rightarrow G$ is the obvious one $V(\tilde{G}) \rightarrow V(G)$, $E(\tilde{G}) \rightarrow E(G)$. We will refer to the full subgraphs of \tilde{G} spanned by the vertices $V(G) \times \{k\}$ as *clouds*.

Example 2.1. Let H be the graph given by Figure 1. We choose $S = \{f, g\}$, the dashed and the dotted edge; the remaining edges form the spanning tree T . Let us consider the surjective homomorphism $\rho : \mathbb{F}(S) \twoheadrightarrow \mathbb{Z}/2 \times \mathbb{Z}/2$ defined by assigning $\rho(f) = (1, 0)$ and $\rho(g) = (0, 1)$. Observe that this is in fact an example of a $\mathbb{Z}/2$ -homology covering as described below.




FIGURE 1. The graph H from Example 2.1 (left); the “figure eight” graph (right).

The covering graph \tilde{H} , see Figure 2, has four clouds $\boxed{(i, j)} = V(G) \times \{(i, j)\}$, where $i, j \in \{0, 1\}$. Within these clouds, one sees copies of the spanning tree T . The edges from S connect different clouds according to the action of ρ . Note that collapsing clouds to points yields the Cayley graph of $K = \mathbb{Z}/2 \times \mathbb{Z}/2$ with respect to the generating set $\rho(S)$.




FIGURE 2. The graph \tilde{H} from Example 2.1.

Let us summarize basic properties of this construction, see, for example, [11, Ch. 2].

Proposition 2.2. *Let G be a finite connected graph, $\rho : \pi_1(G) \twoheadrightarrow K$ be a surjective homomorphism, and let \tilde{G} be the covering graph constructed above. Then the following hold.*

- *The covering projection $\pi : \tilde{G} \rightarrow G$ enjoys the unique path lifting property and the homotopy lifting property.*
- *\tilde{G} does not depend on the choices made, i.e. on T and on the orientation of edges in S . It is unique, up to graph isomorphism commuting with the projections.*
- *There is a short exact sequence $1 \rightarrow \pi_1(\tilde{G}) \xrightarrow{\pi_*} \pi_1(G) \xrightarrow{\rho} K \rightarrow 1$ and the group of deck transformations of the covering is K .*

Observe that if G is the “figure eight”, see Figure 1, and we identify $\pi_1(G)$ with the free group \mathbb{F}_2 generated by a and b , then \tilde{G} is the Cayley graph of K with respect to the generators $\rho(a)$ and $\rho(b)$.

The following classical result follows from the fundamental theorem of Galois theory for graph coverings, see, for example, [5, Section 1.3].

Proposition 2.3 (Transitivity of covers). *Let G be a finite connected graph and $L = \pi_1(G)$ denote its fundamental group. Let M and N be two normal subgroups of L satisfying $M \subset N$. Let us denote by*

- *$\pi^M : G^M \rightarrow G$ and $\pi^N : G^N \rightarrow G$ the coverings of G corresponding to the surjective homomorphisms $\rho^M : L \twoheadrightarrow L/M$ and $\rho^N : L \twoheadrightarrow L/N$, respectively;*
- *$\pi_N^M : G_N^M \rightarrow G^N$ the covering of G^N corresponding to $\rho_N^M : N \twoheadrightarrow N/M$. (Observe that $\pi_1(G^N) \cong N$ by the basic properties of the construction.)*

Then, $\pi^M : G^M \rightarrow G$ is isomorphic to $\pi_N^M \circ \pi^N : G_N^M \rightarrow G$ as graph coverings of G .

$$\begin{array}{ccc}
 G^M & \xrightarrow{\cong} & G_N^M \\
 \downarrow \pi^M & & \downarrow \pi_N^M \\
 & & G^N \\
 & \swarrow \pi_N & \\
 & G &
 \end{array}$$

2.3. Graph coverings: $\mathbb{Z}/2$ -homology case. Let us now specialize to the $\mathbb{Z}/2$ -homology coverings. Given a finite graph G , there is always a surjective homomorphism ρ from the fundamental group $\pi_1(G)$ to the first $\mathbb{Z}/2$ -homology group of G . Indeed, it is the quotient map $\rho : \pi_1(G) \twoheadrightarrow \pi_1(G)/\pi_1(G)^{(2)}$. Recall that $\Gamma^{(2)}$ denotes the normal subgroup generated by all the squares of elements in a group Γ . The resulting quotient is always a finite direct sum of $\mathbb{Z}/2$'s.

It is convenient to label the elements of the group $K \cong \mathbb{Z}/2^{|S|}$ by the elements of the power set $\mathcal{P}(S)$ of S ; thinking of factors of $\mathbb{Z}/2^{|S|}$ indexed by elements of S , a subset $\tau \subset S$ denotes the element of $\mathbb{Z}/2^{|S|}$ having 1 at places indexed by elements from τ and 0

at the remaining places. In this description, an edge from S acts on them by “symmetric difference”: for $\tau \subset S$ and $e \in S$, we have $\rho(e)\tau = \tau \Delta \{e\}$.

Summarizing, given a finite connected graph G , choose a spanning tree T of G and orient the edges in $S = E(G) \setminus E(T)$ arbitrarily. The vertices of \tilde{G} will be indexed by $V(G) \times \mathcal{P}(S)$, the edges by $E(G) \times \mathcal{P}(S)$. An edge (e, τ) , where $e \in E(G)$ between $x, y \in V(G)$, $\tau \subset S$, connects

- (x, τ) and (y, τ) if $e \notin S$,
- (x, τ) and $(y, \tau \Delta \{e\})$ if $e \in S$ and $x \xrightarrow{e} y$.

Observe that if we collapse clouds of \tilde{G} to points, we obtain the Cayley graph of $\mathbb{Z}/2^{|S|}$ with respect to its natural set of generators; that is, an $|S|$ -dimensional cube.

Example 2.4. The graph \tilde{H} from Example 2.1 is an instance of a $\mathbb{Z}/2$ -homology cover; According to our specific labelling, the elements of $\mathbb{Z}/2 \times \mathbb{Z}/2$ are encoded by the subsets of $\{f, g\}$. Here this encoding is as follows: $\boxed{(0, 0)} \rightarrow \emptyset$, $\boxed{(1, 0)} \rightarrow \{f\}$, $\boxed{(0, 1)} \rightarrow \{g\}$, $\boxed{(1, 1)} \rightarrow \{f, g\}$.

3. WALLS

In this section, we first briefly recall the notions of walls and wall structures on a graph. Next, we construct a wall structure on $\mathbb{Z}/2$ -homology covers from the previous section, and establish a relationship between the wall metric and the graph metric in this context. Finally, in Proposition 3.11, we show that the two metrics agree on small scale (depending on the girth of the graphs involved).

Definition 3.1. Let G be a finite graph. A *wall* (or a *cut*) on G is a set of edges of G , such that removing them from G separates G into exactly two connected components. We shall refer to the components as to *half-spaces* associated to the wall.

A *wall structure* W on a graph G is a set of walls on G , such that each edge is contained in exactly one wall from W .

We say that a wall w *separates* two vertices of G , if they reside in different half-spaces associated to w .

Given a wall structure W on G , we denote by $d_W(x, y)$ the number of walls in W separating the vertices x and y of G . Then d_W is a pseudo-metric¹ on G , which we call the *wall metric* associated to the wall structure W .

Let us now specialize to the situation of the previous section.

Definition 3.2. A pair (\tilde{G}, G) is said to be a $\mathbb{Z}/2$ -pair, if G is a finite 2-connected² graph and \tilde{G} is its $\mathbb{Z}/2$ -homology cover.

¹It is in fact a metric, as it turns out later in this section. We shall call d_W a metric right away.

²A graph is 2-connected if every edge lies on a loop, or equivalently, if removing any single edge does not disconnect the graph.

Recall that T denotes a chosen spanning tree in a finite graph G and $S = E(G) \setminus E(T)$. For an edge $e \in E(G)$ and the covering map $\pi : \tilde{G} \rightarrow G$, we denote $w_e = \pi^{-1}(e) \subset E(\tilde{G})$ and $W = \{w_e \mid e \in E(G)\}$.

Lemma 3.3. *If (\tilde{G}, G) is a $\mathbb{Z}/2$ -pair, each w_e separates \tilde{G} into exactly two connected components. Thus, W is a wall structure on \tilde{G} .*

Proof. If $e \in S$, denote $A = \{\tau \subset S \mid e \notin \tau\}$ and $B = \{\tau \subset S \mid e \in \tau\}$. We claim that removing $w_e = \pi^{-1}(e)$ from \tilde{G} , we obtain two components, which happen to be the full subgraphs \tilde{A} and \tilde{B} of \tilde{G} spanned by the sets of vertices $V(G) \times A$ and $V(G) \times B$, respectively³. Indeed, if an edge in w_e connects two vertices (\cdot, τ) and (\cdot, τ') in $V(\tilde{G})$, then $\tau \Delta \tau' = \{e\}$, thus one of the vertices is in $V(\tilde{A})$ and the other in $V(\tilde{B})$. Conversely, in a similar way, any edge between \tilde{A} and \tilde{B} has to be in w_e . Observe that \tilde{A} is connected: we can move within the clouds along the copies of the spanning tree T , and from cloud to cloud within \tilde{A} because if $\tau, \tau' \subset A$, then $\tau \Delta \tau'$ never contains e . The same argument works for \tilde{B} . Thus, the set w_e is a wall whenever $e \in S$.

Since G is 2-connected, given any edge $e \in E(G)$, we may choose a spanning tree T of G which does not contain e (removing e from G leaves it connected, so choose a spanning tree there). Such a choice makes e belong to the corresponding S , so the argument above applies. Here we use the fact that the covering $\pi : \tilde{G} \rightarrow G$ is unique, the choices just give a different labelling.

In fact, one can avoid the preceding re-labelling trick and give a direct argument for edges in $E(T)$. It turns out that even if G is not 2-connected, w_e for $e \in E(T)$ separates \tilde{G} into at least two components; using 2-connectedness one gets that there are at most 2 components. We leave this as an exercise for the reader. \square

We emphasize that \tilde{G} is now endowed with two metrics: $d_{\tilde{G}}$, the usual graph metric and d_W , the wall metric with respect to the wall structure W constructed above. The crucial aspect is to compare these two metrics. This task will occupy the rest of this section. We begin with an easy observation.

Proposition 3.4. *If (\tilde{G}, G) is a $\mathbb{Z}/2$ -pair, we have $d_W(x, y) \leq d_{\tilde{G}}(x, y)$ for all $x, y \in V(\tilde{G})$.*

Proof. If a wall separates x and y , any path between x and y necessarily uses at least one of the edges from the wall. Since different walls are disjoint, any shortest path between x and y has to traverse at least $d_W(x, y)$ edges. \square

The following notion is used to characterize $d_{\tilde{G}}$ and d_W in terms of paths in G .

Definition 3.5. Let (\tilde{G}, G) be a $\mathbb{Z}/2$ -pair. Given $x = (x_0, \tau_x)$ and $y = (y_0, \tau_y)$ from $V(\tilde{G}) = V(G) \times \mathcal{P}(S)$, we say that a path p in G is (x, y) -admissible, if

³Thinking of \tilde{G} roughly as a cube (vertices being clouds), this splitting corresponds to choosing a coordinate direction and taking “the front” and “the back” of the cube. The wall consists of edges connecting going from front to back and vice versa.

- (i) p begins at x_0 and ends at y_0 ,
- (ii) for every $e \in S$, the parity of the number of times that e appears in $E(p)$ is equal

$$\text{to } \begin{cases} 0 & \text{if } e \notin \tau_x \triangle \tau_y \\ 1 & \text{if } e \in \tau_x \triangle \tau_y \end{cases}.$$

Lemma 3.6. *The (x, y) -admissible paths in G are in bijection with paths from x to y in \tilde{G} . Moreover, this correspondence preserves the path length.*

Proof. Take $x, y \in V(\tilde{G})$. We can represent $x = (x_0, \tau_x)$, $y = (y_0, \tau_y)$, where τ_x and τ_y are subsets of S . Let \tilde{p} be a path in \tilde{G} between x and y . The projection $\pi(\tilde{p})$ obviously satisfies (i).

Now consider the sequence of the second coordinates of the vertices appearing in \tilde{p} , say $(\tau_x, \tau_1, \tau_2, \dots, \tau_n, \tau_y)$. Then each two consecutive ones are either the same (if the edge between them is in $\pi^{-1}(E(T))$, i.e. we remain in the same cloud), or they differ by exactly one edge $\{e\}$ (if the edge between them is in $\pi^{-1}(e)$, $e \in S$, i.e. we move between the clouds). Since we begin with τ_x and end with τ_y , then the parity of the number of times that some edge from $\pi^{-1}(e)$ is used in \tilde{p} depends only on whether e belongs to $\tau_x \triangle \tau_y$ or not. Thus (ii) follows.

Conversely, if $p = (x_0, e_1, v_1, \dots, e_k, y_0)$ is an (x, y) -admissible path, it has a unique lift \tilde{p} in \tilde{G} starting at $x \in V(\tilde{G})$. By construction of \tilde{G} , the sequence $E(p) = (e_1, e_2, \dots, e_k)$ determines the sequence $(\tau_x = \tau_0, \tau_1, \dots, \tau_k)$ of the second coordinates of the vertices in \tilde{p} by the following inductive rule: if $e_i \in E(T)$, then $\tau_i = \tau_{i-1}$; if $e_i \in S$ then $\tau_i = \tau_{i-1} \triangle \{e_i\}$. Now the condition (ii) plus this rule implies that the last one, τ_k , must be equal to τ_y . Hence the endpoint of \tilde{p} is $(y_0, \tau_y) = y$ and we are done.

The claim about preserving the lengths is obvious. \square

Example 3.7. Referring again to Example 2.1 and Figure 1, the path

$$p_0 = (x, e_1, y, e_3, v, e_5, w, g, z, e_2, y, e_1, x)$$

is $((x, \emptyset), (x, \{g\}))$ -admissible. It is the projection under π of the path in \tilde{G} given by the sequence of vertices $(x_{00}, y_{00}, v_{00}, w_{00}, z_{01}, y_{01}, x_{01})$. In fact, p_0 is a shortest $((x, \emptyset), (x, \{g\}))$ -admissible path. So, using Propositions below, $d_{\tilde{G}}((x, \emptyset), (x, \{g\})) = 6$, but $d_W((x, \emptyset), (x, \{g\})) = 4$. The point is that the wall w_{e_1} does not separate (x, \emptyset) and $(x, \{g\})$, it is used on p_0 twice.

The next proposition follows readily from Lemma 3.6.

Proposition 3.8. *If (\tilde{G}, G) is a $\mathbb{Z}/2$ -pair, for any $x, y \in V(\tilde{G})$ we have*

$$d_{\tilde{G}}(x, y) = \min\{\ell(p) \mid p \text{ is an } (x, y)\text{-admissible path in } G\}.$$

Remark 3.9. Note that no shortest (x, y) -admissible path has a backtrack (i.e. a subpath $(v_i, e_{i+1}, v_{i+1}, e_{i+2}, v_{i+2})$, which satisfies $v_i = v_{i+2}$, $e_{i+1} = e_{i+2}$). The reason is that removing such a backtrack from the path does not change admissibility of the path, thus providing a shorter admissible path.

Proposition 3.10. *If (\tilde{G}, G) is a $\mathbb{Z}/2$ -pair, for any $x, y \in V(\tilde{G})$, $d_W(x, y)$ is equal to the number of edges appearing odd number of times in any (x, y) -admissible path in G .*

Proof. Pick any wall w in \tilde{G} . Label the half-spaces corresponding to w by H_-^w and H_+^w . Now if a path in \tilde{G} uses one of the edges of w , it is clear that it passes from H_-^w to H_+^w or vice versa. Thus, if a wall w separates x and y , then they do not belong to the same half-space of w , and so every path between them has to “cross” w an odd number of times. Likewise, if a wall w does not separate x and y , both vertices belong to the same half-space of w , and any path between them has to “cross” w an even number of times. By “crossing” w , we mean using some edge in w . To summarize, $d_W(x, y)$ is equal to the number of walls that are crossed odd number of times by any path from x to y in \tilde{G} .

The proof is now finished by referring to Lemma 3.6 and remembering that walls in \tilde{G} are projected via π exactly to edges in G . □

We now use the characterizations of d_W and $d_{\tilde{G}}$ to show that the two metrics agree on the scale of $\text{girth}(G)$.

Proposition 3.11. *If (\tilde{G}, G) is a $\mathbb{Z}/2$ -pair, then for every $x, y \in V(\tilde{G})$,*

$$d_W(x, y) < \text{girth}(G) \quad \text{if and only if} \quad d_{\tilde{G}}(x, y) < \text{girth}(G).$$

Furthermore, if the above inequalities hold, then $d_W(x, y) = d_{\tilde{G}}(x, y)$.

Proof. The implication “ \Leftarrow ” is trivial by Proposition 3.4. It remains to prove the implication “ \Rightarrow ” and the “Furthermore” part. We rely on the following lemma.

Lemma 3.12. *Given $x, y \in V(\tilde{G})$, every shortest (x, y) -admissible path in G either does not contain a loop, or else every edge on any loop it contains is traversed exactly once.*




FIGURE 3. A typical shortest (x, y) -admissible path; $\pi(x) = x_0 = \pi(y)$.

Taking the lemma for granted, we finish the proof of the Proposition. Assume that $d_W(x, y) < \text{girth}(G)$. Then any shortest (x, y) -admissible path q in G cannot contain a loop — otherwise the edges on it are traversed exactly once by the above Lemma, and thus $d_W(x, y)$ is at least the length of that loop by Proposition 3.10. But the length of a

loop in G is at least $\text{girth}(G)$, which is a contradiction. Thus we've shown that q does not contain a loop. It doesn't contain backtracks, either, see Remark 3.9. Thus, each edge of q is traversed exactly once, so $d_W(x, y) = \ell(q)$ by Proposition 3.10. Now the lift \tilde{q} of q to \tilde{G} is a path between x and y satisfying $\ell(\tilde{q}) = \ell(q)$, whence

$$d_{\tilde{G}}(x, y) \leq \ell(\tilde{q}) = \ell(q) = d_W(x, y) < \text{girth}(G).$$

Note that the reverse inequality between $d_{\tilde{G}}$ and d_W is asserted by Proposition 3.4, so we in fact have $d_W(x, y) = d_{\tilde{G}}(x, y)$. This proves also the ‘‘Furthermore’’ part. \square




FIGURE 4. Shortening admissible paths

Proof of Lemma 3.12. We argue by contradiction. Assume that for some $x, y \in V(\tilde{G})$ there is a shortest (x, y) -admissible path q in G , which contains a loop on which there is an edge which is traversed more than once. Let l be the shortest such loop, and denote by v the vertex through which the loop is entered and left. See Figure 4, the dashed loop. There are two possibilities: either there is an edge on it traversed *before* q enters into l (situations B1 and B2 in the figure), or *after* it has left l (situations A1 and A2). In each of the possibilities, select the first edge e in the respective direction which happens to be on l (if there are such in both directions, choose the closer one). Now in each case the edge e is traversed either in the same direction as l (situations A2 and B2) or in the opposite direction (A1 and B1). Let us denote the various parts of the path q as on the figures; we shall check each of the four possibilities and show that we can find a shorter

(x, y) -admissible path, which will be a contradiction. Let us remark that the parts of q outside l can cross l , or each other quite arbitrarily, it does not affect the argument.

In each of the cases, we show that there is a way to traverse all of q (some parts maybe in the opposite direction) except that we pass the edge e exactly two times less. Also, we preserve the beginning and end vertices x_0 and x_1 . It is clear from the definition of admissibility that any new path like this will also be (x, y) -admissible. Let us now finally describe the traversing; for brevity, we write “opp ...” for “... in the opposite direction”.

A1: We pass p , then opp l_2 , opp f_1 , l_1 and f_2 .

A2: We pass p , l_1 , opp f_1 , opp l_2 and finally f_2 .

B1: We pass p_2 , l_2 , l_1 , p_1 and f .

B2: We pass p_2 , then opp l_1 , opp p_1 , l_2 and finally f .

This finishes the proof. \square

4. RESULTS

In this section we prove three Propositions which yield the proof of Theorem 1.1.

Recall from the introduction that $\Gamma_0 = \mathbb{F}_2$ is the free group of rank 2, $\Gamma_n = \Gamma_{n-1}^{(2)}$ are defined inductively for $n \geq 1$, and X_n denotes the Cayley graph of Γ_0/Γ_n with respect to the canonical image of free generators of \mathbb{F}_2 .

The *box space* of \mathbb{F}_2 associated with the family $(\Gamma_n)_{n \geq 1}$ is, by definition, the *coarse union* $(X, d_X) = \sqcup(X_n, d_{X_n})$ of X_n 's. That is, $X = \sqcup X_n$ is the disjoint union with each X_n endowed with its natural graph metric d_{X_n} , and $d(X_n, X_m) \rightarrow \infty$ as $n + m \rightarrow \infty$. For the sake of concreteness, let us declare that $d_X(X_n, X_m) = \text{diam}(X_n) + \text{diam}(X_m) + n + m$ whenever $n \neq m$.

Lemma 4.1. *For the sequence $(\Gamma_n)_{n \geq 1}$ described above, we have $\bigcap_{n \geq 1} \Gamma_n = \{1\}$.*

Proof. The iterated squares are proper characteristic subgroups of the free group, hence, by Levi's theorem [8, Ch.I, Prop. 3.3], they have trivial intersection. \square

Proposition 4.2. *The metric space (X, d_X) does not have property A.*

Proof. Just use the previous Lemma, together with [10, Proposition 11.39]. \square

Note that X_0 is just the “figure eight” graph (Figure 1). It follows from the discussion in Subsection 2.3 and Proposition 2.3 that each X_n is the $\mathbb{Z}/2$ -homology cover of X_{n-1} , $n \geq 1$. Hence we may endow each X_n with the metric d_{W_n} , the wall metric that comes from X_n being the $\mathbb{Z}/2$ -homology cover of X_{n-1} . Let us denote by $(X, d_W) = \sqcup(X_n, d_{W_n})$ the coarse union of X_n 's with the wall metrics. As above, $d_W(X_n, X_m) = \text{diam}_W(X_n) + \text{diam}_W(X_m) + n + m$ whenever $n \neq m$ with diam_W being the corresponding diameter with respect to the wall metric.

Proposition 4.3. *The metric space (X, d_W) is coarsely embeddable into a Hilbert space.*

Proof. It is not hard to check that the wall metric d_{W_n} is a symmetric, normalized negative type kernel on each X_n , see [10, Chapter 11], for the terminology. This fact, together with the definition of $d_W(X_n, X_m)$ above for $n \neq m$, implies that, d_W is an effective, symmetric, normalized negative type kernel on X . The existence of such a kernel is equivalent to the coarse embeddability by [10, Theorem 11.16]. \square

In fact, Hilbert space embeds coarsely into any ℓ_p for $1 \leq p \leq \infty$ [13, Corollary 3.18]. So, does the metric space (X, d_W) .

Lemma 4.4. *Let $\Lambda = \mathbb{F}(S)$ be the free group on a finite set of generators S . Let $(\Lambda_n)_{n \geq 1}$ be a sequence of finite index normal subgroups of Λ satisfying $\Lambda_{n+1} \subset \Lambda_n$ and $\bigcap_{n \geq 1} \Lambda_n = \{1\}$. If we denote by Y_n the Cayley graph of Λ/Λ_n with respect to the image of the generating set S , then $\text{girth}(Y_n) \rightarrow \infty$ as $n \rightarrow \infty$.*

Proof. The generating set S determines a word metric on $\Lambda = \mathbb{F}(S)$, which is also the graph metric on the Cayley graph of $\mathbb{F}(S)$ with respect to the generating set S ; this graph is a tree. Take $r > 0$ and consider the ball $B_\Lambda(1, r)$ of radius r around 1 in $\mathbb{F}(S)$. Since it is finite, the conditions that $\bigcap_{n \geq 1} \Lambda_n = \{1\}$ and $\Lambda_{n+1} \subset \Lambda_n$ imply that there is some n_r , such that $B_\Lambda(1, 2r) \cap \Lambda_n = \{1\}$ for all $n > n_r$. That is, under the quotient map $\Lambda \rightarrow \Lambda/\Lambda_n$, the ball $B_\Lambda(1, r)$ bijectively corresponds to the ball $B_{Y_n}(1, r)$ in the Cayley graph Y_n , for $n > n_r$. Since Y_n is homogeneous, any ball of radius r in Y_n is actually a tree, as is $B_\Lambda(1, r)$. It follows that $\text{girth}(Y_n) > r$ for $n > n_r$ and we are done. \square

Proposition 4.5. *The identity map $\text{id} : (X, d_X) \rightarrow (X, d_W)$ is a coarse equivalence (that is, both id and id^{-1} are coarse embeddings).*

Proof. Since $\text{id} : (X, d_X) \rightarrow (X, d_W)$ is 1-Lipschitz (by Proposition 3.4), it is sufficient to show the following statement: For any $R \geq 0$ there exists $S \geq 0$, such that $d_W(x, y) \leq R$ implies $d_X(x, y) \leq S$, for all $x, y \in X$.

So take $R \geq 0$ and let $N \geq 0$ be such that $\text{girth}(X_n) > R$ for $n \geq N$; this is possible by Lemma 4.4. By possibly enlarging N , we may further assume that $d_W(X_n, X_m) > R$ for any $m, n \geq N$ with $m \neq n$, or for any $m < N$ and $n \geq N$. Now define $S = \max\{R\} \cup \{d_X(x, y) \mid x, y \in \cup_{n < N} X_n\}$ (note that the latter is a finite set).

Let us check the required implication. Take $x, y \in X$ with $d_W(x, y) \leq R$. By our setup, there are two possibilities: either both x and y belong to $\cup_{n < N} X_n$ and then $d_X(x, y) \leq S$ by the choice of S ; or $x, y \in X_n$ for some $n \geq N$. But for such n the restriction of the identity map $\text{id} : (X_n, d_{X_n}) \rightarrow (X_n, d_{W_n})$ to any ball $B_X(z, R)$, $z \in X_n$, of radius $R < \text{girth}(X_n)$, is an isometry onto the ball $B_W(z, R)$ by Proposition 3.11. Thus necessarily $d_X(x, y) = d_W(x, y) \leq R \leq S$ and we are done. \square

Observe that (X, d_X) consists of Cayley graphs of groups with a fixed number of generators, it follows that it has bounded geometry. Then the previous Proposition implies that (X, d_W) has also bounded geometry. This finishes the proof of our main result, Theorem 1.1.

5. C*-ALGEBRAIC CONTEXT AND QUESTIONS

In this section, we discuss our example in the C*-algebraic context. To every metric space, there are associated various Roe C*-algebras, coarse–geometric analogues of group C*-algebras.

Definition 5.1 (See [10]). Given a uniformly discrete metric space X with bounded geometry, its *algebraic uniform Roe algebra* (sometimes also called *Gromov’s translation algebra*) C_{alg}^*X is the algebra of X -by- X complex matrices $T = (t_{xy})_{x,y \in X}$, which have uniformly bounded entries (i.e. $\sup_{x,y} |t_{xy}| < \infty$) and finite propagation (i.e. $\sup\{d(x,y) \mid x,y \in X \text{ and } t_{xy} \neq 0\} < \infty$). This algebra is naturally represented (by matrix multiplication) on $\ell^2 X$; its closure in $\mathcal{B}(\ell^2 X)$ is called the *uniform Roe algebra* of X , denoted C_u^*X . The *maximal uniform Roe algebra* of X , denoted $C_{u,\text{max}}^*X$, is the closure of C_{alg}^*X in the biggest C*-norm that it admits.

5.1. Relation between the maximal uniform Roe algebra and its maximal version.

There is always a natural surjective map $\lambda : C_{u,\text{max}}^*X \rightarrow C_u^*X$. If X has Property A, this map is an isomorphism [2, 12]; the converse is an open question, a particular case of an open question about groupoid C*-algebras. In the paper [12] the authors prove that if X coarsely embeds into a Hilbert space, λ induces an isomorphism $\lambda_* : K_*(C_{u,\text{max}}^*X) \rightarrow K_*(C_u^*X)$. Thus our example from Theorem 1.4 is covered by this result. This leads to an interesting question: Is the map λ an isomorphism in this case?

5.2. Nuclearity and exactness. For a discrete group Γ endowed with a proper metric, a result of Ozawa [9] (see also Guentner–Kaminker [4]) says that the following are equivalent:

- Γ has Property A;
- Γ is C*-exact;
- $C_u^*|\Gamma|$ is nuclear.

For general bounded geometry metric spaces, the first and the third statements are also equivalent [2].

It follows from this result that for $C_u^*|\Gamma|$, exactness implies nuclearity (the converse is a general C*-algebraic fact). It is an open question whether this is true for general metric spaces with bounded geometry. A progress in the positive direction has been made by Brodzki, Niblo and Wright in [1]. If one were to look for a counterexample, one should take a space without Property A, so that its uniform Roe algebra is not nuclear, and prove that it is exact. Naturally, one can consider expanders, but it is known for a large class that their uniform Roe algebras are not exact [6].

From this perspective, one can ask whether the uniform Roe algebra of our example from Theorem 1.4 is exact or not. The criterion from [1] does not readily apply. However, an argument due to Willett (private communication) shows that in fact uniform Roe algebras of box spaces in general are either nuclear or not exact.

REFERENCES

- [1] Jacek Brodzki, Graham Niblo, and Nick Wright. Property A, partial translation structures, and uniform embeddings in groups. *J. Lond. Math. Soc. (2)*, 76(2):479–497, 2007.
- [2] Nathaniel P. Brown and Narutaka Ozawa. *C*-algebras and finite-dimensional approximations*, volume 88 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2008.
- [3] Mikhael Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.
- [4] Erik Guentner and Jerome Kaminker. Exactness and the Novikov conjecture and Addendum. *Topology*, 41(2):411–418 and 419–420, 2002.
- [5] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [6] Nigel Higson and Erik Guentner. *K*-theory of group C^* -algebras. In S. Doplicher and R. Longo, editors, *Noncommutative Geometry*, Lecture Notes in Mathematics, pages 253–262. Springer Verlag, Berlin, 2003.
- [7] Nigel Higson and John Roe. Amenable group actions and the Novikov conjecture. *J. Reine Angew. Math.*, 519:143–153, 2000.
- [8] Roger C. Lyndon and Paul E. Schupp. *Combinatorial Group Theory*, volume 89 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, New York, 1977.
- [9] Narutaka Ozawa. Amenable actions and exactness for discrete groups. *C. R. Acad. Sci. Paris Sér. I Math.*, 330(8):691–695, 2000.
- [10] John Roe. *Lectures on coarse geometry*, volume 31 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2003.
- [11] John Stillwell. *Classical topology and combinatorial group theory*, volume 72 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1993.
- [12] Ján Špakula and Rufus Willett. Maximal and reduced roe algebras of coarsely embeddable spaces. Preprint, 2010.
- [13] Rufus Willett. Some notes on property A. In *Limits of graphs in group theory and computer science*, pages 191–281. EPFL Press, Lausanne, 2009.
- [14] Guoliang Yu. The coarse Baum-Connes conjecture for spaces which admit a uniform embedding into Hilbert space. *Invent. Math.*, 139(1):201–240, 2000.

UNIVERSITY OF VIENNA, FACULTY OF MATHEMATICS, NORDBERGSTRASSE 15, 1090 WIEN, AUSTRIA

E-mail address: goulmara.arzhantseva@univie.ac.at

UNIVERSITY OF HAWAII AT MĀNOA, DEPARTMENT OF MATHEMATICS, 2565 MCCARTHY MALL, HONOLULU, HI 96822, USA

E-mail address: erik@math.hawaii.edu

MATHEMATISCHES INSTITUT, UNIVERSITÄT MÜNSTER, EINSTEINSTR. 62, 48149 MÜNSTER, GERMANY

E-mail address: jan.spakula@uni-muenster.de