Zeitschrift: L'Enseignement Mathématique

Herausgeber: Commission Internationale de l'Enseignement Mathématique

Band: 65 (2019)

Heft: 1-2

Artikel: Acylindrical actions on projection complexes

Autor: Bestviona, Mladen / Bromberg, Ken / Fujiwara, Koji

DOI: https://doi.org/10.5169/seals-869344

Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften auf E-Periodica. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. Das Veröffentlichen von Bildern in Print- und Online-Publikationen sowie auf Social Media-Kanälen oder Webseiten ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. Mehr erfahren

Conditions d'utilisation

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. La reproduction d'images dans des publications imprimées ou en ligne ainsi que sur des canaux de médias sociaux ou des sites web n'est autorisée qu'avec l'accord préalable des détenteurs des droits. En savoir plus

Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. Publishing images in print and online publications, as well as on social media channels or websites, is only permitted with the prior consent of the rights holders. Find out more

Download PDF: 03.11.2025

ETH-Bibliothek Zürich, E-Periodica, https://www.e-periodica.ch

Acylindrical actions on projection complexes

Mladen Bestvina, Ken Bromberg, Koji Fujiwara and Alessandro Sisto

Abstract. We simplify the construction of projection complexes from [BBF2]. To do so, we introduce a sharper version of the Behrstock inequality, and show that it can always be enforced. Furthermore, we use the new setup to prove acylindricity results for the action on the projection complexes.

We also treat quasi-trees of metric spaces associated to projection complexes, and prove an acylindricity criterion in that context as well.

Mathematics Subject Classification (2010). Primary: 20F65; Secondary: 20E08

Keywords. Quasi-trees, acylindricity, hyperbolically embedded subgroups, Mapping Class Groups.

1. Introduction

1.1. Quasi-trees. A *quasi-tree* is a geodesic metric space quasi-isometric to a tree; in particular quasi-trees are hyperbolic spaces. The study of actions on (simplicial) trees, or Bass-Serre theory, is a very well-developed and established theory. One might expect that actions on quasi-trees can be promoted to actions on trees, but it turns out that this is only possible under rather restrictive hypotheses [MSW1, MSW2]. In fact, in recent years it has emerged that, somewhat surprisingly, actions on quasi-trees are ubiquitous, much more so than actions on trees. The first evidence in this direction can be found in [Man], where Manning constructs actions on quasi-trees starting from quasi-morphisms. In [BBF2] an axiomatic setup was introduced for producing many actions of groups on quasi-trees (including groups, such as mapping class groups, that do not admit actions on trees without fixed points). The goal of this paper is to streamline the construction and arguments from that paper, which will in turn allow us to get more precise results about the actions one obtains. But first, we discuss [BBF2]; the expert reader may skip ahead to the next section.

A typical example where [BBF2] applies is when $G = \pi_1(\Sigma)$ where Σ is a closed, hyperbolic surface. One then takes a simple, closed geodesic on Σ and lets Y be the set of components of the pre-image of the geodesic in the universal cover \mathbb{H}^2 , each of which is a geodesic line. Given distinct geodesics $X, Y \in Y$ we let $\pi_Y(X)$ be the nearest point projection of X to Y and observe that the diameters of these sets will be uniformly bounded as X and Y vary through Y. More interestingly, if we have distinct elements $X, Y, Z \in Y$ and the projections $\pi_Y(X)$ and $\pi_Y(Z)$ are far apart on Y then the projections $\pi_Z(X)$ and $\pi_Z(Y)$ will be coarsely the same on Z (meaning that the diameter of the union is uniformly bounded). This is illustrated in Figures 1 and 2 below. Also, another elementary example, in a tree, satisfying the axiomatic setup is provided in Remark 2.1, right after (stronger versions of the) axioms are stated.

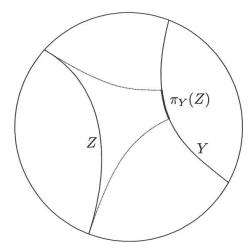


FIGURE 1
Negative curvature forces the lines to have bounded projection onto each other.

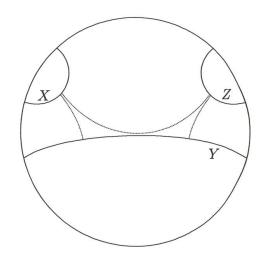


FIGURE 2 If X and Z project far on Y, then X and Y project close on Z.

This is, in a nutshell, the setup one needs to construct quasi-trees, according to [BBF2].

1.2. Acylindricity. One particular case where this construction turned out to be extremely useful is in the development of the theory of acylindrically hyperbolic groups, as defined in [Osi]. These are groups G admitting an "interesting enough" action on a hyperbolic space \mathcal{X} ; more precisely the requirements are that G is not virtually cyclic, has unbounded orbits in \mathcal{X} (the action is non-elementary), and that the action is acylindrical. The action of a group G on a metric space \mathcal{X} is acylindrical if for all D>0 there exist L>0 and B>0 such that if $x,y\in\mathcal{X}$ and $d_{\mathcal{X}}(x,y)\geq L$ then there are at most B elements $g\in G$ with $d_{\mathcal{X}}(x,gx)\leq D$ and $d_{\mathcal{X}}(y,gy)\leq D$. In [Osi], Osin gives a number of different

characterizations of acylindrical hyperbolicity which is very useful in generating examples. (See Theorem 5.1 below.) In particular, there are seemingly weaker conditions that imply that a group is acylindrically hyperbolic. In Section 5 we will derive Osin's results more directly and will further prove a stronger result due to Balusubramanya [Bal]: Every acylindrically hyperbolic group has an acylindrical action on a quasi-tree.

Acylindrical hyperbolicity has strong consequences, including: Every acylindrically hyperbolic group is SQ-universal (in particular it has uncountably many pairwise non-isomorphic quotients), it contains free normal subgroups [DGO], it contains Morse elements and hence all its asymptotic cones have cut-points [Sis], and its bounded cohomology is infinite dimensional in degrees 2 [HO] and 3 [FPS]. Moreover, if an acylindrically hyperbolic group does not contain finite normal subgroups, then its reduced C^* -algebra is simple [DGO] and every commensurating endomorphism is an inner automorphism [AMS].

1.3. Axioms. We now formally state the axioms from [BBF2] that give rise to quasi-trees.

Let $\mathbf{Y} = \{(Y, \rho_Y)\}$ be a collection of metric spaces. Let $\theta \ge 0$ be a fixed constant and assume that for all $X, Y, Z \in \mathbf{Y}$ with $Y \ne X, Z$ we are given numbers $d_Y^{\pi}(X, Z) \in [0, \infty]$ (referred to as "projection distance") satisfying:

(P1)
$$d_Y^{\pi}(X, Z) = d_Y^{\pi}(Z, X)$$

(P2)
$$d_Y^{\pi}(X, Z) + d_Y^{\pi}(Z, W) \ge d_Y^{\pi}(X, W)$$

(P3)
$$d_Y^{\pi}(X, X) \leq \theta$$
.

(P4) if
$$d_Y^{\pi}(X, Z) > \theta$$
 then $d_X^{\pi}(Y, Z) \le \theta$.

(P5)
$$\{W \neq X, Z : d_W^{\pi}(X, Z) > \theta\}$$
 is finite.

Families of metric spaces with projection distances satisfying (P1)–(P5) occur naturally in many contexts. See the introduction to [BBF2] for some examples. In most cases there are subsets $\pi_Y(Z) \subseteq Y$ for $Y \neq Z$ so that

$$d_Y^{\pi}(X,Z) := \operatorname{diam} \pi_Y(X) \cup \pi_Y(Z)$$

satisfy (P1)–(P5). Note that in this case (P1) and (P2) are automatic, and (P3) amounts to the requirement that the diameter of $\pi_Y(Z)$ is uniformly bounded. Frequently, the spaces $Y \in Y$ are subspaces of some ambient metric space and $\pi_Y(Z)$ is defined as the nearest point projection of Z to Y. The stronger axiom (SP4) (see next section) follows from the following additional requirement:

If
$$d_Y^{\pi}(X, Z) > \theta$$
 then $\pi_X(Y) = \pi_X(Z)$.

1.4. Projection complexes. We now discuss the projection complex construction (in a slightly more restrictive setting than necessary). One starts with a collection Y of metric spaces, and bounded sets $\pi_Y(Z) \subseteq Y$ (which one thinks of as projections) for all $Y \neq Z$. Then, one has to perform a preparatory step in which one perturbs the projections up to finite Hausdorff distance, which we explain below. At this point, one is ready to build the *projection complex* $\mathcal{P}_K(Y)$, whose vertex set is Y, by fixing a large constant K and adding an edge between distinct vertices X and X if the diameter of $\pi_Y(X) \cup \pi_Y(X)$ is at most X for all $X \in Y \setminus \{X, Z\}$. The central result of [BBF2] is that $\mathcal{P}_K(Y)$ is a quasi-tree. In certain situations, the perturbation of the projections is necessary.

One of the technical challenges of [BBF2] is that when the diameter of $\pi_Y(X) \cup \pi_Y(Z)$ is large the projections $\pi_Z(X)$ and $\pi_Z(Y)$ are coarsely equal, but are not exactly equal. This causes problems in induction arguments, because of constants that might get worse at every step. We will see here that by assuming equality (instead of just coarse equality) the proof that the projection complex is a quasi-tree (Theorem 3.5) vastly simplifies. Unfortunately, in most naturally occurring situations we do not have equality. In the second part of the paper we introduce the notion of a *forcing sequence* and use it to show that the projection maps can be modified (coarsely) so that we have the desired equality (Theorem 4.1). In this way one can replace the work of Section 3 of [BBF2] with the much simpler arguments in this paper.

1.5. Main Results. Besides simplifying the approach from [BBF2], the new setup allows us to obtain acylindricity results, which are unknown for the original version of the construction. In fact, under some simple conditions, it is straightforward to show that the G-action on the projection complex will be acylindrical, see Theorem 3.9. We summarize the results described so far in the following statement (the condition for acylindricity in Theorem 3.9 is less restrictive).

Theorem 1.1 (Theorems 4.1, 3.5, 3.9). Assume that \mathbf{Y} is a set, and that the functions $d_Y : \mathbf{Y}^2 \setminus \{(Y,Y)\} \to [0,\infty]$ satisfy axioms (P1)–(P5) above (or from Section 4). After perturbing the projections as in Theorem 4.1, consider the graph $\mathcal{P}_K(\mathbf{Y})$ for a fixed $K \geq 0$ with vertex set \mathbf{Y} and vertices X, Z connected by an edge if $d_Y(X,Z) \leq K$, for every $Y \neq X,Z$. The graph $\mathcal{P}_K(\mathbf{Y})$ is equipped with the path metric where every edge has length 1. Then:

- (1) If K is sufficiently large, then $\mathcal{P}_K(\mathbf{Y})$ is a quasi-tree.
- (2) If the group G acts on Y preserving d_Y (meaning that $d_{gY}(gX, gZ) = d_Y(X, Z)$ whenever the right-hand side is defined), then G acts by isometries on $\mathcal{P}_K(Y)$.

(3) If in addition there exist integers B and N so that the common stabilizer of any collection of N elements of Y has cardinality at most B, then the action is acylindrical.

As in [BBF2] we can also build a *quasi-tree of metric spaces*. In the final section we give necessary and sufficient conditions for it to be a quasi-tree or a hyperbolic space (Theorem 6.6 and Corollary 6.8-(2)) and, more generally, prove that it is in a natural way a tree-graded space. Furthermore, we show that (under some natural conditions) the group action on the quasi-tree of metric spaces is also acylindrical (Theorem 6.4).

While many of the arguments here follow a similar outline to what is in [BBF2] (a notable exception being the part about forcing sequences) this paper is completely self-contained, with the only exception of the use of Manning's bottleneck criterion and, in the last section, its generalization from [Hum], and does not require any of the results from [BBF2]. We also note that the section on hyperbolically embedded subgroups does not require the forcing sequence technology from the previous section as the projection maps defined there satisfy the equality condition without modification.

An abridged version [BBFS] of this paper, only containing the proof that projection complexes are quasi-trees and the perturbation of the projection distances, is available on the authors' websites. This shorter version already contains most of the ideas and techniques that we use in this paper.

2. Axioms

Let Y be a set and for each $Y \in Y$ assume that we have a function

$$d_Y: \mathbf{Y}^2 \setminus \{(Y,Y)\} \to [0,\infty]$$

such that the following *strong projection axioms* are satisfied for some $\theta \ge 0$ and all $X, Y, Z, W \in \mathbb{Y}$ when expressions are defined:

- **(SP1)** $d_Y(X, Z) = d_Y(Z, X);$
- (SP1) $d_Y(X,Z) + d_Y(Z,W) > d_Y(X,W)$;
- **(SP1)** $d_Y(X, X) \leq \theta$;
- (SP1) if $d_Y(X, Z) > \theta$ then $d_Z(X, W) = d_Z(Y, W)$ for all $W \in Y \setminus \{Z\}$;
- (SP1) $\#\{Y|d_Y(X,Z)>\theta\}$ is finite for all $X,Z\in\mathbf{Y}$.

The constant θ is the *projection constant*. Note that we allow $d_Y(X, Z) = \infty$. Thus d_Y is not a distance function in the usual sense, e.g., we also allow

 $d_Y(X, X) > 0$. Axioms (SP1)–(SP3) formalize this weak notion of distance and in practice are easy to verify. Axioms (SP4) and (SP5) are the key properties that make the construction possible and need to be checked in applications.

Example 2.1. Consider a tree T and a locally finite family of disjoint geodesic lines Y in it. Given distinct $X, Y \in Y$ there are unique points $x \in X$, $y \in Y$ so that the geodesic [x, y] intersects $X \cup Y$ at its endpoints only, and in fact x is the nearest-point projection of any $y' \in Y$ to X, and similarly for y. Hence, in this situation set $x = \pi_X(Y)$. Set $d_Y(X, Z) = d(\pi_Y(X), \pi_Y(Z))$. Then the axioms above are satisfied with $\theta = 0$. (SP1), (SP2), and (SP3) are clear.

Let us argue that (SP4) holds. For X, Y, Z as in (SP4), notice that if [x, y] (resp. [z, y']) intersects $X \cup Y$ (resp. $Z \cup Y$) only at its endpoints, then the concatenation $[x, y] \cup [y, y'] \cup [y', z]$ is so that consecutive geodesics intersect only at the common endpoint, and [y, y] is not a single point (this is what $d_Y(X, Z) > 0$ yields). Since T is a tree, this implies that the given concatenation is in fact the geodesic [x, z]. Hence, [x, z] only intersects $X \cup Z$ at its endpoints, showing that $\pi_Z(X) = \{z\} = \pi_Z(Y)$. Hence, for any $W \in Y \setminus \{Z\}$, we have $d_Z(X, W) = d_Z(\{z\}, \pi_Z(W)) = d_Z(Y, W)$, as required.

Notice that the argument we just gave shows that if $d_Y(X, Z) > 0$, then the geodesic from $\pi_X(Z)$ to $\pi_Z(X)$ intersects Y non-trivially. This fact and local finiteness of Y imply (SP5).

The most important axiom is arguably (SP4), which is a version of the Behrstock inequality [Beh] in the context of Masur-Minsky subsurface projections. As in [BBF2], we will use it to order certain subsets of Y, the idea being that if $d_Y(X, Z)$ is large, then Y is between X and Z. We note that (SP4) is in fact a more precise version of the Behrstock inequality because the conclusion is an actual equality, not an approximate one. This allows us to know the exact value of certain d_Y , and it is the key to our much simpler proofs, compared to [BBF2].

Lemma 2.2. (*SP3*) and (*SP4*) imply

$$\min \left\{ d_Y(X, Z), d_Z(X, Y) \right\} \le \theta$$

Proof. If $d_Y(X, Z) > \theta$ then letting W = Y in (SP4) we have $d_Z(X, Y) = d_Z(Y, Y) \le \theta$ by (SP3).

For a constant K > 0 define $Y_K(X, Z)$ to be the collection of $Y \in Y \setminus \{X, Z\}$ such that $d_Y(X, Z) > K$.

Lemma 2.3 and Proposition 2.4 below say that, for large enough K, $Y_K(X, Z)$ can be totally ordered using the idea, as mentioned above, that if $d_Y(X, Z)$ is large

then Y is between X and Z. The order has several equivalent characterizations, which is good for applications, and they are listed in Lemma 2.3:

Lemma 2.3. For $Y_0, Y_1 \in \mathbf{Y}_{2\theta}(X, Z)$ the following conditions are equivalent:

- (1) $d_{Y_0}(X, Y_1) > \theta$;
- (2) $d_{Y_1}(Y_0, W) = d_{Y_1}(X, W)$ for all $W \neq Y_1$;
- (3) $d_{Y_1}(X, Y_0) \leq \theta$;
- (4) $d_{Y_1}(Y_0, Z) > \theta$;
- (5) $d_{Y_0}(W, Y_1) = d_{Y_0}(W, Z)$ for all $W \neq Y_0$;
- (6) $d_{Y_0}(Y_1, Z) \leq \theta$.

Proof. By Lemma 2.2, (1) \Rightarrow (3) and (4) \Rightarrow (6). By (SP2), (3) \Rightarrow (4) and (6) \Rightarrow (1). By (SP4), (1) \Rightarrow (2) and (4) \Rightarrow (5). Since $Y_1 \in \mathbf{Y}_{2\theta}(X, Z)$ by letting W = Z we have (2) \Rightarrow (4) and similarly (5) \Rightarrow (1).

Given $Y_0, Y_1 \in \mathbf{Y}_{2\theta}(X, Z)$ we define $Y_0 < Y_1$ if any one of (1)–(6) hold.

Proposition 2.4. The relation < defines a total order on $\mathbf{Y}_{2\theta}(X, Z)$ that extends to a total order on $\mathbf{Y}_{2\theta}(X, Z) \cup \{X, Z\}$ with least element X and greatest element Z. Furthermore if $Y_0 < Y_1 < Y_2$ then $d_{Y_1}(Y_0, Y_2) = d_{Y_1}(X, Z)$.

Notice that with a coarse version of (SP4) there would be no hope to obtain the last conclusion as stated.

Proof. By swapping Y_0 and Y_1 we see that $Y_0 < Y_1$ if and only if $Y_1 \not< Y_0$. So any two elements of $\mathbf{Y}_{2\theta}(X, Z)$ can be compared.

Now we check transitivity of the order. If $Y_0 < Y_1$ and $Y_1 < Y_2$ we apply (2) for $Y_0 < Y_1$ with $W = Y_2$ and then again to $Y_1 < Y_2$ with W = Z to see that $d_{Y_1}(Y_0, Y_2) = d_{Y_1}(X, Y_2) = d_{Y_1}(X, Z) > 2\theta$. Applying (SP4) and then (5) we have $d_{Y_2}(Y_0, Z) = d_{Y_2}(Y_1, Z) = d_{Y_2}(X, Z) > 2\theta$. Therefore $Y_0 < Y_2$ and the total order is well defined on $\mathbf{Y}_{2\theta}(X, Z)$. We can extend it to a total order on $\mathbf{Y}_{2\theta}(X, Z) \cup \{X, Z\}$ by declaring X to be the least element and Z the greatest element.

Observe that we have also shown that
$$d_{Y_1}(Y_0, Y_2) = d_{Y_1}(X, Z)$$
 if $Y_0 < Y_1 < Y_2$.

The main use of the following lemma will be to construct free groups (for other purposes, simpler versions would suffice). It is a kind of local-to-global principle.

Lemma 2.5. For any $K \geq 2\theta$ the following holds. Let Q be a connected simplicial graph and $\phi: Q^{(0)} \to \mathbf{Y}$ a map such that adjacent vertices are mapped to distinct elements of \mathbf{Y} , and if x, y and z are distinct vertices with x and z adjacent to y then $d_{\phi(y)}(\phi(x), \phi(z)) > K$. Then for any immersed path $\{x_0, \ldots, x_k\}$ in Q we have $\{\phi(x_0) < \cdots < \phi(x_k)\} \subset \mathbf{Y}_K(\phi(x_0), \phi(x_k)) \cup \{\phi(x_0), \phi(x_k)\}$. In particular, ϕ is injective and Q is a tree.

Proof. For $k \leq 2$ the conclusion clearly holds. We proceed by induction on k. Let x_0, \ldots, x_k be an immersed path and let $X_i = \phi(x_i)$. We first show $X_i \in \mathbf{Y}_K(X_0, X_k)$ for 0 < i < k. If $1 \leq i < k-1$ then $d_{X_i}(X_0, X_{k-1}) > K$ and $d_{X_{k-1}}(X_i, X_k) > K$ by induction. Then by (SP4), $d_{X_i}(X_0, X_k) = d_{X_i}(X_0, X_{k-1}) > K$ so $X_i \in \mathbf{Y}_K(X_0, X_k)$.

If i = k - 1 we reverse the roles of X_0 and X_k , and of X_1 and X_{k-1} .

For the order we have that $d_{X_i}(X_0, X_j) > K$ if 0 < i < j < k since $X_i \in \mathbf{Y}_K(X_0, X_j)$ and therefore $X_i < X_j$.

Corollary 2.6. Let $K \ge 2\theta$. If $Y_0, Y_1 \in \mathbf{Y}_K(X, Z) \cup \{X, Z\}$ and $d_Y(Y_0, Y_1) > K$ then $Y \in \mathbf{Y}_K(X, Z)$.

Proof. We assume $Y_i \notin \{X, Z\}$, since in those cases the proof is similar and easier.

We can assume that $Y_0 < Y_1$. We then apply Lemma 2.5 where Q is a closed interval subdivided into 4 segments with vertices labeled X, Y_0, Y, Y_1, Z , in order. By assumption $d_Y(Y_0, Y_1) > K$ so we only need to check that $d_{Y_0}(X, Y)$ and $d_{Y_1}(Y, Z)$ are greater than K. By (SP4) $d_{Y_0}(X, Y) = d_{Y_0}(X, Y_1)$. By Proposition 2.4, $d_{Y_0}(X, Y_1) = d_{Y_0}(X, Z) > K$. Therefore $d_{Y_0}(X, Y) > K$. Similarly, $d_{Y_1}(Y, Z) > K$.

3. The projection complex

Fix $K \ge 2\theta$ and define the graph $\mathcal{P}_K(Y)$ with vertex set Y and an edge between any two vertices X and Z with $Y_K(X, Z) = \emptyset$. We denote the distance in $\mathcal{P}_K(Y)$ simply by d, even though it depends on K.

We first note that $\mathcal{P}_K(Y)$ is connected.

Lemma 3.1. If $K \ge 2\theta$ and $X, Z \in Y$ then $Y_K(X, Z) \cup \{X, Z\} = \{X < X_1 < \dots < X_k < Z\}$ is a path in $\mathcal{P}_K(Y)$.

Proof. By Corollary 2.6, if Y is the immediate predecessor of Y' in the total order on $\mathbf{Y}_K(X,Z) \cup \{X,Z\}$ then $\mathbf{Y}_K(Y,Y') = \emptyset$ and therefore d(Y,Y') = 1. \square

The path $Y_K(X, Z) \cup \{X, Z\} = \{X < X_1 < \cdots < X_k < Z\}$ is the *standard* path from X to Z.

The following lemma says that, when moving outside the ball of radius 2 around a vertex Z of $\mathcal{P}_K(Y)$, the projection to Z varies slowly, where slowly is independent of K.

Lemma 3.2. If $K \ge 2\theta$ then the following holds. Let $X_0, X_1, Z \in \mathbf{Y}$ with $d(X_0, X_1) = 1$ and $d(X_0, Z) \ge 2$. Then $d_Z(X_0, X_1) \le \theta$.

Proof. Since $d(X_0, Z) \ge 2$ there exists a $Y \in \mathbf{Y}_K(X_0, Z)$ and therefore by (SP4) $d_Z(X_0, X_1) = d_Z(Y, X_1)$. If $d_Z(Y, X_1) = d_Z(X_0, X_1) > \theta$ then by (SP4) $d_Y(X_0, X_1) = d_Y(X_0, Z) > K$, a contradiction with $\mathbf{Y}_K(X_0, X_1) = \emptyset$.

The following lemma and its corollary are the key to proving that $\mathcal{P}_K(Y)$ is a quasi-tree. They say that, when moving outside the ball of radius 3 around a vertex Z of $\mathcal{P}_K(Y)$, the projection to Z basically does not change.

Lemma 3.3. If $K \ge 3\theta$ then the following holds. Let X_0, \ldots, X_k be a path in $\mathcal{P}_K(\mathbf{Y})$ and $Z \in \mathbf{Y}$ with $d(X_i, Z) \ge 3$ for all i. Then the greatest elements of $\mathbf{Y}_{3\theta}(X_i, Z)$ all agree.

Proof. We can assume k=1. Let Y_0 and Y_1 be the corresponding greatest elements and assume they are distinct. By Corollary 2.6, $\mathbf{Y}_{3\theta}(Y_i,Z)=\varnothing$ so $d(Y_i,Z)=1$ and $d(X_i,Y_i)\geq 2$. Applying Lemma 3.2 we see that $d_{Y_i}(X_0,X_1)\leq \theta$ and therefore by (SP2), $d_{Y_i}(X_{1-i},Z)>2\theta$. In particular, both Y_0 and Y_1 are in $\mathbf{Y}_{2\theta}(X_i,Z)$ for i=0,1. We can assume that $Y_0<Y_1$ in $\mathbf{Y}_{2\theta}(X_0,Z)$. By Lemma 2.3(6) this means that $d_{Y_0}(Y_1,Z)\leq \theta$ and so we also have $Y_0<Y_1$ in $\mathbf{Y}_{2\theta}(X_1,Z)$. In particular, $d_{Y_1}(X_0,Z)=d_{Y_1}(Y_0,Z)=d_{Y_1}(X_1,Z)>3\theta$, contradicting the assumption that Y_0 is the greatest element of $\mathbf{Y}_{3\theta}(X_0,Z)$. \square

Corollary 3.4. If $K \ge 3\theta$ then the following holds. Let X_0, \ldots, X_k be a path in $\mathcal{P}_K(\mathbf{Y})$ and $Z \in \mathbf{Y}$ with $d(X_i, Z) \ge 3$. Then $d_Z(X_i, X_j) \le \theta$ for all i, j.

Proof. By Lemma 3.3, there exists a $Y \in \mathbb{Y}$ that is the greatest element of all of the $\mathbb{Y}_{3\theta}(X_i, Z)$. We now have $d_Z(X_i, X_j) = d_Z(X_i, Y) \leq \theta$ by (SP4) and Lemma 2.2.

We can now use Manning's bottleneck condition [Man] to show that $\mathcal{P}_K(\mathbf{Y})$ is a quasi-tree. We will use a variant of Manning's condition that is described in [BBF2]: Let X be a connected simplicial graph with its usual combinatorial metric and $D \geq 0$. Assume that for all vertices $v_0, v_1 \in X^{(0)}$ there is a path p such that the D-Hausdorff neighborhood of any path from v_0 to v_1 contains p. Then X is a quasi-tree.

Theorem 3.5. For $K \geq 3\theta$, $\mathcal{P}_K(\mathbf{Y})$ is a quasi-tree.

Proof. By Lemma 3.1 $Y_K(X,Z) \cup \{X,Z\} = \{X < X_1 < \cdots < X_k < Z\}$ is a path in $\mathcal{P}_K(Y)$. Let $X = Y_0, Y_1, \ldots, Y_n = Z$ be an arbitrary path from X to Z. Since $d_{X_i}(X,Z) > K \geq 3\theta$ by Corollary 3.4 there must be a Y_j such that $d(Y_j,X_i) \leq 2$. Therefore $\mathcal{P}_K(Y)$ satisfies the bottleneck condition and is a quasi-tree.

The following lemma is a variant of [HO, Lemma 3.9] proved by Hull and Osin in the context of hyperbolically embedded subgroups.

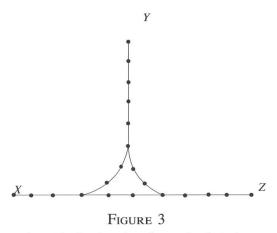
Lemma 3.6. If $K > 2\theta$ then the following holds. Given $X, Y, Z \in Y$ the union $\mathbf{Y}_K(X,Y) \cup \mathbf{Y}_K(Y,Z)$ contains all but at most two elements of $\mathbf{Y}_K(X,Z)$, and if there are two such elements they are consecutive.

In particular, $\mathbf{Y}_K(X, Z)$ is written as the disjoint union of three consecutive segments (some possibly empty) so that the initial segment (if not empty) is also initial in $\mathbf{Y}_K(X,Y)$, the second contains at most two elements, and the terminal segment (if not empty) is also terminal in $\mathbf{Y}_K(Y,Z)$.

Proof. Set $\mathbf{Y}_K(X,Z) = \{X_1 < \cdots < X_k\}$. Then for any X_i at least one of $d_{X_i}(X,Y)$ or $d_{X_i}(Y,Z)$ is $> \theta$. If it is the former then $X_j \in \mathbf{Y}_K(X,Y)$ for all j < i (proof: use (SP4) twice to see that $d_{X_j}(X,X_i) > K$ implies $d_{X_i}(X_j,Y) = d_{X_i}(X,Y) > K \ge \theta$ so $d_{X_j}(X,Y) = d_{X_j}(X,X_i) > K$) while if it is the latter then $X_j \in \mathbf{Y}_K(Y,Z)$ for all j > i.

Now assume that X_i is the smallest element not in the union. By the previous paragraph $d_{X_{i+1}}(X,Y) \leq \theta$, and in turn $d_{X_{i+1}}(Y,Z) > \theta$ by triangle inequality. But this implies that $X_j \in \mathbf{Y}_K(Y,Z)$ if j > i+1.

The rest is clear. See Figure 3.



A typical triangle of standard paths

Corollary 3.7. If $K \ge 3\theta$ then the following holds. Let $X \ne Z$ and $n = |\mathbf{Y}_K(X,Z)| + 1$. Then $\lfloor \frac{n}{2} \rfloor + 1 \le d(X,Z) \le n$.

Proof. The second inequality follows from the fact that $Y_K(X, Z) \cup \{X, Z\}$ is path from X to Z. The proof of the other inequality is by induction on d(X, Z), the case d(X, Z) = 1 being clear.

Suppose $d(X, Z) \ge 2$. Pick $Y \ne X, Z$ on a geodesic from X to Z. Setting $n_1 = |\mathbf{Y}_K(X, Y)| + 1, n_2 = |\mathbf{Y}_K(Y, Z)| + 1$, Lemma 3.6 gives $n_1 + n_2 \ge n - 1$. Using the induction hypothesis, we get $d(X, Z) = d(X, Y) + d(Y, Z) \ge \lfloor \frac{n_1}{2} \rfloor + \lfloor \frac{n_2}{2} \rfloor + 2 \ge \lfloor \frac{n_2}{2} \rfloor + 1$.

Corollary 3.8. Standard paths are quasi-geodesics.

Assume that a group G acts on Y and the functions d_Y are G-invariant. Then G acts isometrically on $\mathcal{P}_K(Y)$.

The following theorem gives a simple criterion for the action on $\mathcal{P}_K(Y)$ to be acylindrical, in terms of finiteness of the size of the stabilizer of several elements of Y. Roughly speaking, we look at far away X, Z and an element g that moves both a small distance, and deduce that a large middle interval in $Y_K(X, Z)$ is also contained in $Y_K(gX, gZ)$. With too many g, we would get too many elements stabilizing several elements of Y.

Theorem 3.9. If $K \geq 3\theta$ then the following holds. Assume that for some fixed N and B, for any N distinct elements of any $\mathbf{Y}_K(X,Z)$ the common stabilizer is a finite subgroup of size at most B. Then the action of G on $\mathcal{P}_K(\mathbf{Y})$ is acylindrical.

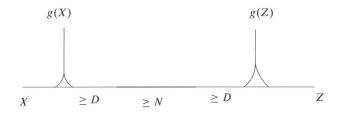


FIGURE 4
Proof of acylindricity

Proof. Fix D > 0 and assume that $X, Z \in Y$ with $d(X, Z) \ge N + 4D + 6$. Let $g \in G$ be such that $d(X, gX) \le D$ and $d(Z, gZ) \le D$. Consider the quadrilateral of standard paths spanned by X, Z, gX, gZ. Let the segment I (and J) in $Y_K(X, Z)$ be obtained by removing initial and terminal segments of length 2D + 2 (and D + 2, respectively). Then I has length $|I| \ge N$ and J

has length |I| + 2D, and by Lemma 3.6 $J \subset Y_K(gX, gZ)$. Thus $g(I) \subset J$ and there are $\leq 2D + 1$ possible restrictions of g to I (translation by a number in [-D, D]). By the pigeon-hole principle and the assumption about stabilizers it follows that there are at most B(2D + 1) such elements g.

In the following theorem we construct free groups acting on $\mathcal{P}_K(Y)$. We use Lemma 2.5 to certify that certain elements generate a free group.

Theorem 3.10. Assume that $K \ge 3\theta$ and fix $Y \in Y$. If there exists $g_0, g_1, g_2 \in G$ such $d_Y(g_iY, g_jY) > K$ and $d_Y(g_i^{-1}Y, g_j^{-1}Y) > K$ when $i \ne j$ then the group generated by $g_1g_0^{-1}$ and $g_2g_1^{-1}$ is free and acts faithfully on $\mathcal{P}_K(Y)$ with orbit map a QI-embedding.

Proof. Let $F \subset G$ be the subgroup generated by $g_1g_0^{-1}$ and $g_2g_1^{-1}$. Consider the theta graph Θ with two vertices v_0, v_1 and three oriented edges labeled g_0, g_1, g_2 connecting v_0 to v_1 . Then $\pi_1(\Theta)$ can naturally be identified with the free group on $g_1g_0^{-1}$ and $g_2g_1^{-1}$ and there is a canonical epimorphism $\psi:\pi_1(\Theta)\to F$. Let Γ be the covering space of Θ corresponding to the kernel of ψ . Thus F acts on Γ as the deck group. The edges of Γ have an induced orientation and labeling by the g_i 's.

Define the F-equivariant map $\phi: \Gamma^{(0)} \to \mathbb{Y}$ by sending a base vertex w_0 to Y and if $u_1u_2\cdots u_k$ is a path from w_0 to w with $u_i\in\{g_0^{\pm 1},g_1^{\pm 1},g_2^{\pm 1}\}$ (with exponents necessarily alternating) then $\phi(w)=u_1u_2\cdots u_k(Y)$. Our assumption implies that the hypotheses of Lemma 2.5 are satisfied and so Γ is a tree and ψ is an isomorphism. The last sentence follows from Corollary 3.7.

In practice it is easy to verify the conditions of Theorem 3.10. In applications of the projection complex the set Y is a collection of infinite diameter metric spaces and the subgroup that fixes each metric space acts coarsely transitively. It is then relatively easy to find the necessary elements of G.

4. Forcing sequences

We start by recalling the axioms (P1)–(P5). Let $\mathbf{Y} = \{(Y, \rho_Y)\}$ be a collection of metric spaces. Let $\theta \ge 0$ be a fixed constant and assume that for all $X, Y, Z \in \mathbf{Y}$ with $Y \ne X, Z$ we are given numbers $d_Y^{\pi}(X, Z) \in [0, \infty]$ satisfying:

(P1)
$$d_Y^{\pi}(X, Z) = d_Y^{\pi}(Z, X)$$

(P2)
$$d_Y^{\pi}(X, Z) + d_Y^{\pi}(Z, W) \ge d_Y^{\pi}(X, W)$$

$$(\mathbf{P3}) \ d_Y^{\pi}(X,X) \leq \theta.$$

- **(P4)** if $d_Y^{\pi}(X, Z) > \theta$ then $d_X^{\pi}(Y, Z) \leq \theta$.
- **(P5)** $\{W \neq X, Z : d_W^{\pi}(X, Z) > \theta\}$ is finite.

The goal of this section is to prove the following theorem.

Theorem 4.1. Assume that $\mathbf{Y} = \{(Y, \rho_Y)\}$ is a collection of metric spaces with $\{d_Y^{\pi}\}$ satisfying (P0)–(P4) with constant θ . Then there are $\{d_Y\}$ satisfying (SP1)–(SP5) for the constant 11θ such that

$$d_Y^{\pi} - 2\theta \le d_Y \le d_Y^{\pi} + 2\theta.$$

Moreover, assume that d_Y^{π} are obtained from projections π_Y . Then there are projections $\pi_Z'(X)$ for $X \neq Z$ satisfying the following:

- (1) $\pi'_{\mathbf{Z}}(X) \subseteq N_{\theta}(\pi_{\mathbf{Z}}(X)),$
- (2) $d_Y(X, Z)$ is equal to the diameter of $\pi'_Y(X) \cup \pi'_Y(Z)$.
- (3) If a group G acts on Y preserving the metrics and the projections $\pi_Z(X)$, then G also preserves the new projections $\pi'_Z(X)$.

Thus G acts on a quasi-tree as in Theorem 3.5 and the action is frequently acylindrical as in Theorem 3.9.

Mimicking the earlier section we let $Y_K^{\pi}(X, Z)$ be the collection of $Y \in Y \setminus \{X, Z\}$ such that $d_Y^{\pi}(X, Z) > K$.

4.1. Modifying the distance d^{π} . We assume d^{π}_{Y} satisfy (P1)–(P5).

The first step is to modify d^{π} to achieve monotonicity (see Lemma 4.4). Recall from [BBF2] that for $X \neq Z$ we define $\mathcal{H}(X, Z)$ as the set of pairs $(X', Z') \in \mathbb{Y} \times \mathbb{Y}$ such that one of the following holds.

- $d_X^{\pi}(X', Z'), d_Z^{\pi}(X', Z') > 2\theta,$
- X = X' and $d_Z^{\pi}(X, Z') > 2\theta$,
- Z = Z' and $d_X^{\pi}(X', Z) > 2\theta$,
- X = X' and Z = Z'.

Lemma 4.2. If $d_Y^{\pi}(X,Z) > 2\theta$ and $(X',Z') \in \mathcal{H}(X,Z)$ then, after possibly permuting X' and Z', $d_Y^{\pi}(X,X'), d_Y^{\pi}(Z,Z') \leq \theta$. In particular, $|d_Y^{\pi}(X,Z) - d_Y^{\pi}(X',Z')| \leq 2\theta$.

Proof. By the triangle inequality (P2)

$$d_X^{\pi}(X',Y) + d_X^{\pi}(Y,Z') \ge d_X(X',Z') > 2\theta$$

and therefore $\max\{d_X^\pi(X',Y),d_X^\pi(Y,Z')\}>\theta$. After possibly permuting X' and Z' we can assume $d_X^\pi(X',Y)>\theta$ and therefore by (P4) we have $d_Y^\pi(X,X')\leq\theta$. By another application of the triangle inequality

$$d_Y^{\pi}(X, X') + d_Y^{\pi}(X', Z) \ge d_Y^{\pi}(X, Z) > 2\theta$$

and since $d_Y^\pi(X,X') \leq \theta$ this implies that $d_Y^\pi(X',Z) > \theta$. Therefore (P4) implies that $d_Z^\pi(X',Y) \leq \theta$. Now, replacing X with Z in the above application of the triangle inequality (P2) we have $\max\{d_Z^\pi(X',Y),d_Z^\pi(Y,Z')\} > \theta$ and therefore $d_Z^\pi(Y,Z') > \theta$ since we have just seen that the other term is $\leq \theta$. Another application of (P4) gives that $d_Y^\pi(Z,Z') \leq \theta$.

The final inequality follows from the triangle inequality (P2). \Box

Corollary 4.3. If $d_Y^{\pi}(X, Z) > 4\theta$ then $\mathcal{H}(X, Z) \subseteq \mathcal{H}(X, Y)$.

Proof. Suppose $(X', Z') \in \mathcal{H}(X, Z)$. We will again assume the first bullet holds and leave the other cases to the reader. To show $(X', Z') \in \mathcal{H}(X, Y)$ it suffices to argue that $d_Y^{\pi}(X', Z') > 2\theta$, and this follows from $d_Y^{\pi}(X, Z) > 4\theta$ and the lemma.

We now define the modified distance

$$\tilde{d}_Y(X,Z) = \sup_{(X',Z') \in \mathcal{H}(X,Z)} d_Y^{\pi}(X',Z')$$

if $d_Y^{\pi}(X,Z) > 2\theta$, and $\tilde{d}_Y(X,Z) = 2\theta$ otherwise. Thus

$$d_Y^{\pi}(X,Z) \le \tilde{d}_Y(X,Z) \le d_Y^{\pi}(X,Z) + 2\theta.$$

The triangle inequality for d holds only up to an error of 2θ . What we gain with this modification is the following monotonicity property.

Lemma 4.4. If $\tilde{d}_Y(X,Z) > 5\theta$ and $\tilde{d}_W(Y,Z) > 7\theta$ then $\tilde{d}_Y(X,W) \geq \tilde{d}_Y(X,Z)$.

Proof. We have $d_W^\pi(Y,Z) > 5\theta$ so $d_Y^\pi(W,Z) \leq \theta$. Likewise, $d_Y^\pi(X,Z) > 3\theta$ so $d_Y^\pi(X,W) \geq d_Y^\pi(X,Z) - d_Y^\pi(W,Z) > 2\theta$ and so $d_W^\pi(X,Y) \leq \theta$. Thus $d_W^\pi(X,Z) \geq d_W^\pi(Y,Z) - d_W^\pi(X,Y) > 4\theta$. Corollary 4.3 gives $\mathcal{H}(X,W) \supseteq \mathcal{H}(X,Z)$. We saw above that $d_Y^\pi(X,W) > 2\theta$ and the statement follows.

4.2. The second (and final) modification of d^{π} . To prove the theorem we need to modify the d_Y^{π} so that they satisfy the projection axioms (SP1)–(SP5). The key notion to do so is that of a forcing sequence, which uses the first modification \tilde{d} . In slightly different contexts, an idea similar to forcing sequences has appeared in [Bar] and [BB, Section 2.B].

Also, if d_Y^{π} are obtained from projections π_Y , this modification is realized by modifications of π_Y to π_Y' .

Definition 4.5. A *K-forcing sequence* is a sequence $Y = Y_0, ..., Y_n = Z$ of distinct elements of **Y** so that $\tilde{d}_{Y_i}(Y_{i-1}, Y_{i+1}) > K$ (for all i = 1, ..., n-1).

Notice that if $X \neq Z$ then $X = Y_0, Y_1 = Z$ is a (usually non-maximal) forcing sequence.

Lemma 4.6. Let Y_0, Y_1, \dots, Y_n be a 4θ -forcing sequence. Then

- (i) $d_{Y_n}^{\pi}(Y_0, Y_{n-1}) \leq \theta,$
- (ii) $|d_{Y_i}^{\pi}(Y_i, Y_k) d_{Y_i}^{\pi}(Y_{j-1}, Y_{j+1})| \le 2\theta$ for all i < j < k.

Proof. We prove (i) by induction on n, starting with the obvious case n=1. Suppose that it is true for a given n and let us prove it for n+1. Observe that $d_{Y_i}^{\pi}(Y_{i-1},Y_{i+1})>2\theta$.

Since $d_{Y_n}^{\pi}(Y_0, Y_{n-1}) \leq \theta$, by the triangle inequality we have $d_{Y_n}^{\pi}(Y_0, Y_{n+1}) > \theta$, so that $d_{Y_{n+1}}^{\pi}(Y_0, Y_n) \leq \theta$, as required.

To prove (ii) note that both $Y_i, Y_{i+1}, \ldots, Y_j$ and $Y_k, Y_{k-1}, \ldots, Y_j$ are 4θ -forcing sequences. We apply (i) to each of them and (ii) then follows from the triangle inequality.

The lemma below tells us when we can insert elements in forcing sequences, and it will be used to show that if $\tilde{d}_W(X,Z)$ is large, then any maximal forcing sequence from X to Z goes through W. Its proof uses the monotonicity of \tilde{d} .

Lemma 4.7. Let $Y_0, ..., Y_n$ be a K-forcing sequence with $K \ge 7\theta$ and $W \in Y$ such that $\tilde{d}_W(Y_i, Y_{i+1}) > K$. Then $Y_0, ..., Y_i, W, Y_{i+1}, ..., Y_n$ is a K-forcing sequence.

Proof. We need to argue that $\tilde{d}_{Y_i}(Y_{i-1}, W), \tilde{d}_{Y_{i+1}}(W, Y_{i+2}) > K$. Both follow from Lemma 4.4, e.g., $\tilde{d}_{Y_i}(Y_{i-1}, W) \geq \tilde{d}_{Y_i}(Y_{i-1}, Y_{i+1}) > K$.

Lemma 4.8. For $K \ge 7\theta$, any K-forcing sequence from X to Z can be refined into a maximal one.

Proof. The obvious process of refinement, using Lemma 4.7, must terminate by Lemma 4.6(ii) and (P5). \Box

Lemma 4.9. Let $Y_0, ..., Y_n$ be a maximal K-forcing sequence, $K \ge 7\theta$, and let $W \in \mathbf{Y}$ with $d_W^{\pi}(Y_0, Y_n) > K + 2\theta$. Then $W = Y_i$ for some $i \in \{1, ..., n-1\}$.

Proof. We assume that W is distinct from all the Y_i and derive a contradiction.

By Lemma 4.6, $d_{Y_i}^\pi(Y_0,Y_n)>2\theta$. We first observe that if $d_W^\pi(Y_i,Y_n)>\theta$ then $d_{Y_i}^\pi(Y_0,W)\geq d_{Y_i}^\pi(Y_0,Y_n)-d_{Y_i}^\pi(W,Y_n)>\theta$ since by (P4) $d_{Y_i}^\pi(W,Y_n)\leq\theta$. Again applying (P4) we have $d_W^\pi(Y_0,Y_i)\leq\theta$.

We have proved that for every i, either $d_W^\pi(Y_0,Y_i) \leq \theta$ or $d_W^\pi(Y_n,Y_i) \leq \theta$. There is some i such that $d_W^\pi(Y_0,Y_i) \leq \theta$ and $d_W^\pi(Y_n,Y_{i+1}) \leq \theta$. From the triangle inequality and our assumption, we have $d_W^\pi(Y_i,Y_{i+1}) > K$ and therefore $Y_0,\ldots,Y_{i-1},W,Y_i,\ldots,Y_n$ is a K-forcing sequence by Lemma 4.7, contradicting our maximality assumption.

Lemma 4.10. Let $Y_0, \ldots, Y_{n-1}, Y_n$ be a maximal K-forcing sequence with $K \geq 7\theta$ and suppose that $X \in \mathbf{Y}$ satisfies $d_{Y_0}^{\pi}(X, Y_n) > K + \theta$. Then there exists a maximal K-forcing sequence from X to Y_n with penultimate element Y_{n-1} .

Proof. By Lemma 4.6, $d_{Y_0}^{\pi}(Y_1, Y_n) \leq \theta$ so

$$\tilde{d}_{Y_0}(X,Y_1) \ge d_{Y_0}^{\pi}(X,Y_1) \ge d_{Y_0}^{\pi}(X,Y_n) - d_{Y_0}^{\pi}(Y_1,Y_n) > K.$$

Therefore X, Y_0, \ldots, Y_n is a K-forcing sequence. Any maximal refinement will have the required property for if in the refinement an element appeared between Y_{n-1} and Y_n then the original sequence would not be maximal.

Definition 4.11 (Penultimate elements). For distinct elements $X, Z \in Y$ define a subset $P_Z(X) = \{W\} \subset Y$, where W are all penultimate elements of maximal 7θ -forcing sequences from X to Z. Note that $P_Z(X)$ is not empty.

When $Y \neq X, Z$, we define

$$d_Y(X,Z) = \sup d_Y^{\pi}(W_1,W_2),$$

where $W_1 \in P_Y(X), W_2 \in P_Y(Z)$.

If projection maps π_Z are defined, set

$$\pi'_Z(X) = \bigcup \pi_Z(W),$$

where $W \in P_Z(X)$.

Note that $d_Y(X, Z)$ is equal to the ρ_Y -diameter of $\pi'_Y(X) \cup \pi'_Y(Z)$. In other words, $d_Y = d_Y^{\pi'}$.

Also, if a group G acts on Y preserving the metrics and the projections $\pi_Z(X)$, then G preserves $\pi_Z'(X)$ by the construction.

Lemma 4.12. We have

$$d_Y^{\pi} - 2\theta \le d_Y \le d_Y^{\pi} + 2\theta.$$

Also,

$$\pi'_Z(X) \subseteq N_\theta(\pi_Z(X)),$$

where N_{θ} denotes the Hausdorff θ -neighborhood.

Proof. By Lemma 4.6 if W is the penultimate element of a 7θ -forcing sequence from X to Z then $d_Z^{\pi}(X,W) \leq \theta$. The inequalities follow from triangle inequality of d_Y^{π} . The second claim is clear.

Lemma 4.13. If
$$d_Y(X, Z) > 11\theta$$
 then $P_Z(X) = P_Z(Y)$, hence $\pi'_Z(X) = \pi'_Z(Y)$.

Proof. By Lemma 4.12 if $d_Y(X,Z) > 11\theta$ then $d_Y^{\pi}(X,Z) > 9\theta$. By Lemma 4.9 if $X = Y_0, \ldots, Y_n = Z$ is a maximal 7θ -forcing sequence then $Y = Y_i$ for some $i \in \{1, \ldots, n-1\}$. Then Y_i, \ldots, Y_n is a maximal 7θ -forcing sequence from Y to Z and it follows that $P_Z(Y) \supseteq P_Z(X)$.

By Lemma 4.10 any maximal 7θ -forcing sequence from Y to Z can be extended to a maximal 7θ -forcing sequence from X to Z with the same penultimate element so $P_Z(X) \supseteq P_Z(Y)$.

We showed
$$P_Z(X) = P_Z(Y)$$
. Then $\pi'_Z(X) = \pi'_Z(Y)$ follows. \square

Proposition 4.14. If d_Y^{π} satisfy (PI)–(P5), then d_Y satisfy the projection axioms (SPI)-(SP5) with projection constant 11θ .

Proof. (SP1) and (SP2) are trivial. (SP4) is exactly Lemma 4.13 and (SP5) follows from (P5) and Lemma 4.12. The other axioms are clear. \Box

Lemma 4.12 and Proposition 4.14 complete the proof of Theorem 4.1.

5. Acylindrical examples

In this section we apply Theorem 3.9 to prove in concrete examples that the action on one of the projection complexes we constructed is acylindrical.

5.1. Acylindrically hyperbolic groups. We have the following equivalent definitions of an acylindrically hyperbolic group:

Theorem 5.1 (Osin [Osi], Balasubramanya [Bal], Dahmani–Guirardel–Osin [DGO]). *The following conditions are equivalent:*

- (1) G has a non-elementary acylindrical action on a hyperbolic space.
- (2) G has a non-elementary acylindrical action on a quasi-tree.
- (3) G acts on a hyperbolic space with a WPD element.
- (4) G contains an infinite hyperbolically embedded subgroup.

Clearly (2) is a strengthening of (1) so $(2) \Rightarrow (1)$. As every hyperbolic element in an acylindrical action is a WPD element we also have $(1) \Rightarrow (3)$. The implication $(3) \Rightarrow (4)$ is proven in [DGO] and uses the quasi-tree construction of [BBF2]. In particular, using the methods of this paper we can directly prove $(3) \Rightarrow (2)$ and in practice this seems to be the most useful statement. The implication $(4) \Rightarrow (1)$ is one of the main results of [Osi] and $(1) \Rightarrow (2)$ is the main result of [Bal]. Again, the methods here can be used to directly prove $(4) \Rightarrow (2)$ and in fact we'll see that in this setting the strong projection axioms hold without applying the technology of forcing sequences.

5.2. WPD elements and B **-contracting geodesics.** We now prove that $(3) \Rightarrow (2)$. In fact we'll prove a stronger statement but we'll first need to make a few definitions so that we can set up the spaces and projections.

Let $\mathcal X$ be a geodesic metric space. Assume that a group G acts on $\mathcal X$ by isometries. Let $f \in G$ be a hyperbolic element (i.e., the translation length is positive). Then f is a WPD element if for all D>0 and $x \in \mathcal X$ there exists an integer N>0 such that the set

$$\left\{g \in G | d_{\mathcal{X}}(x, f(x)) \le D, d_{\mathcal{X}}(g^{N}(x), f(g^{N}(x))) \le D\right\}$$

is finite. For convenience, we will also assume that f acts as a translation on a geodesic line $\gamma \subset \mathcal{X}$ and that γ is *strongly contracting*, i.e., the image under the nearest point projection $p: \mathcal{X} \to \gamma$ (which is in general a multivalued map) of any metric ball disjoint from γ has uniformly bounded diameter. Note that if \mathcal{X} is hyperbolic then every geodesic strongly contracting where the diameter bound only depends on the hyperbolicity constant. Having a strong contracting axis implies that there is a subgroup EC(f), which is virtually cyclic, such that if $g \in EC(f)$ then $g(\gamma)$ and γ have finite Hausdorff distance, and if $g \notin EC(f)$ then $p(g(\gamma))$ has uniformly bounded diameter. For convenience we will assume that EC(f) leaves γ invariant. Both assumptions made for convenience can be removed, at the expense of making the definitions below more complicated, or else replacing \mathcal{X} by a quasi-isometric space where these assumptions hold.

Let **Y** be the set of G-translates of γ . For $A, B, C \in \mathbf{Y}$ define $d_B^{\pi}(A, C) = \operatorname{diam} p_B(A) \cup p_B(C)$ where $p_B : \mathcal{X} \to B$ is the nearest point projection to B.

Theorem 5.2. The set **Y** and the functions d_B^{π} , $B \in \mathbf{Y}$ satisfy (P1)–(P5) for some $\theta > 0$.

When the ambient space \mathcal{X} is hyperbolic then Theorem 5.2 is [DGO, Lemma 4.7]. The general case can be found in [BBF1, Section 4]. Note that there are many examples where \mathcal{X} is not hyperbolic but there are elements with strongly contracting axes. Two prominent cases being axes of pseudo-Anosovs in Teichmüller space and axes of fully irreducible elements in Outer Space.

Corollary 5.3. Assume that G acts on a geodesic metric space with a WPD element that has a strongly contracting axis. Then G has a non-elementary, acylindrical action on a quasi-tree.

- **5.3. Hyperbolically embedded subgroups.** Let G be a group and H a subgroup. Fix a (possibly infinite) set $S \subset G$ such that $S \cup H$ generates G. Let $\Gamma(G, S \sqcup H)$ be the Cayley graph for this generating set; more precisely, we introduce double edges corresponding to elements in $S \cap H$ and regard every edge as labelled by the corresponding copy of a generator. We define a function $\hat{d}: H \times H \to [0, \infty]$ as follows. If $x, y \in H$ are connected by a path in $\Gamma(G, S \sqcup H)$ that does not contain any edges from H we let $\hat{d}(x, y)$ be the length of the shortest such path. If there is no such path we let $\hat{d}(x, y) = \infty$. Then H is hyperbolically embedded in G (with respect to the generating set S) if
- (1) $\Gamma(G, S \sqcup H)$ is δ -hyperbolic;
- (2) \hat{d} is proper.

Each coset aH in $\Gamma(G, S \sqcup H)$ consists of the vertices of a complete graph and when we refer to aH as a subset of $\Gamma(G, S \sqcup H)$ we refer to this complete graph whose edges are labeled by the elements of H. A path p in $\Gamma(G, S \sqcup H)$ penetrates the coset aH if the intersection of p with aH contains a segment. Note that every coset has diameter 1 so a geodesic that penetrates a coset will intersect it in exactly one segment.

The following is a consequence of [DGO, Proposition 4.13], but we provide a proof in the interest of self-containment.

Lemma 5.4. There exists a C > 0 such that the following holds. Suppose that we have a geodesic quadrilateral in $\Gamma(G, S \sqcup H)$ with sides s, p_0, p_1, p_2 so that s is an edge in the coset aH with endpoints s_0, s_1 , and no p_i penetrates aH. Then $\hat{d}(s_0, s_1) \leq C$.

Proof. We label the sides so that p_2 is opposite to s, and $s_i \in p_i$. Let δ be an integer so that $\Gamma(G, S \sqcup H)$ is δ -hyperbolic. Recall that for a geodesic quadrilateral, the 2δ -neighborhood of any three sides contains the fourth.

We consider the case when the lengths of p_0 and p_1 are $> 2\delta + 2$, leaving the other (easier) cases to the reader. Let x_i be the vertex of p_i at distance $2\delta + 2$ from the endpoint of p_i that belongs to aH, and let y_i be a vertex at distance $\leq 2\delta$ from x_i on a side of the quadrilateral distinct from p_i . Note that necessarily $y_i \not\in s$ and further a geodesic $[x_i, y_i]$ does not intersect aH.

If both y_i belong to p_2 we have the path $s_0, x_0, y_0, y_1, x_1, s_1$ made of segments in the quadrilateral and geodesics $[x_i, y_i]$. It has length bounded by a function of δ and it is disjoint from aH except at the endpoints.

If say $y_0 \in p_1$, we have the path s_0, x_0, y_0, s_1 with the same conclusion. \square

Given subsets of vertices X and Y define $\pi_X(Y)$ to be the set of all $x \in X$ so that x is an endpoint of a geodesic that minimizes the distance between X and Y.

Lemma 5.5. Let aH, bH and cH be distinct cosets. If every geodesic that minimizes the distance between aH and cH penetrates bH then $\pi_{cH}(aH) = \pi_{cH}(bH)$.

Proof. Let p be a geodesic that minimizes the distance from aH to cH and that penetrates bH. In particular p contains a single segment in bH. Decompose p into three segments $p = p_0 p_1 p_2$ where p_1 is the segment in bH. Then p_0 is a geodesic that minimizes the distance from aH to bH and p_2 minimizes the distance from bH to cH. In particular the terminal endpoint of p will lie in $\pi_{cH}(bH)$ so $\pi_{cH}(aH) \subseteq \pi_{cH}(bH)$.

Given any other point $z \in \pi_{cH}(bH)$ we can find a minimizing geodesic p_2' from bH to cH that has terminal endpoint z. Let p_1' be a segment in bH that has the same initial endpoint as p_1 and whose terminal endpoint is the initial endpoint of p_2' . Then $p_0p_1'p_2'$ is a path from aH to cH that has the same length as p (because p penetrates bH) and is therefore a minimizing geodesic. Therefore $z \in \pi_{cH}(aH)$ and $\pi_{cH}(bH) \subseteq \pi_{cH}(aH)$.

Lemma 5.6. If $aH \neq bH$ and $x, x' \in \pi_{aH}(bH)$ then $\hat{d}(x, x') \leq C$, where C is the constant from Lemma 5.4.

Proof. Let p and p' be geodesics that minimize the distance from aH to bH and have initial endpoints x and x', respectively. Connect the terminal endpoints of p and p' with segments to form a 4-gon. Then Lemma 5.4 implies that $\hat{d}(x,x') \leq C$.

The above lemma shows that there is a coarsely well-defined projection $\Gamma \to aH$. Using Lemma 5.4 it is easy to see that geodesics that do not penetrate aH have uniformly bounded image in aH with respect to the \hat{d} -metric, a version of the Bounded Geodesic Image Theorem.

Let

$$d_{bH}(aH, cH) = \sup_{x \in \pi_{bH}(aH), z \in \pi_{bH}(cH)} \hat{d}(x, z).$$

Proposition 5.7. The collection of cosets $\{aH\}$ and the functions d_{aH} satisfy the projection axioms (SP1)–(SP5) with $\theta = 3C + 1$.

Proof. Both (SP1) and (SP2) are clear and (SP3) follows from Lemma 5.6.

For (SP4), assume $d_{bH}(aH,cH) \geq 3C+1$ and let p be a geodesic minimizing the distance between aH and cH with initial endpoint $x \in aH$ and terminal endpoint $z \in cH$. We will show that p penetrates bH and then (SP4) follows from Lemma 5.5. Assume not and let q_0 be a geodesic that minimizes the distance from x to bH and let x' be the terminal endpoint of q_0 . We can assume that $x' \in \pi_{bH}(aH)$ for if not there will be a geodesic from $\pi_{aH}(bH)$ to $\pi_{bH}(aH)$ that is shorter than q_0 and we can connect this geodesic to x with a segment in aH to form a path from x to bH that is at most as long as q_0 . As q_0 is minimizing this new path must be a geodesic. Similarly we can find a geodesic q_2 from z to a point $z' \in \pi_{bH}(cH)$ that minimizes the distance between z and bH. Note that neither q_0 nor q_2 have segments in bH. We then let q_1 be the segment in bH between x' and z' and $p^{-1}q_0q_1q_2$ is a 4-gon that satisfies the conditions of Lemma 5.4 so $\hat{d}(x',z') \leq C$. But since $d_{bH}(aH,cH) \geq 3C+1$ we have $\hat{d}(x',z') \geq C+1$ by Lemma 5.6, a contradiction.

Given cosets aH, bH and cH, by the previous paragraph if $d_{bH}(aH, cH) \ge 3C + 1$ then every geodesic that minimizes the distance between aH and cH penetrates bH. Since any geodesic can only penetrate a finite number of cosets this proves (SP5).

The group G acts on the set of cosets $\{aH\}$ by g(aH) = (ga)H. The coset aH is fixed by the subgroup $H^a = aHa^{-1}$. We will need the following result of Dahmani–Guirardel–Osin, which we prove for the convenience of the reader.

Proposition 5.8 ([DGO]). There is a uniform bound on $|H^a \cap H^b|$ over all distinct cosets aH, bH.

Proof. We can assume b=1. Also, up to multiplying a on the left by an element of H, we can assume that a geodesic γ in $\Gamma(G,S\sqcup H)$ from 1 to a does not intersect H except for 1. If $h\in H\cap H^a$ then there exists a quadrilateral where two opposite sides are H-translates of γ , one side consists of the edge in H from 1 to h, and a side is contained in aH. By Lemma 5.4, $\hat{d}(1,h)\leq C$, and we are done by local finiteness of \hat{d} .

This implies that the stabilizer of two distinct cosets will have uniformly bounded size.

The following theorem then follows from Theorems 3.5 and 3.9 and implies $(4) \Rightarrow (2)$ in Theorem 5.1.

Theorem 5.9 (Balasubramanya [Bal]). Suppose G contains an infinite hyperbolically embedded subgroup H of infinite index. Then G has a non-elementary acylindrical action on the quasi-tree $\mathcal{P}_K(\{aH\})$, where H is a hyperbolically embedded infinite subgroup of G of infinite index, K is large enough, and the projections are defined above.

Proof. Since we have Proposition 5.7 and Proposition 5.8, it follows from Theorems 3.5 and 3.9 that G has an acylindrical action on the quasi-tree $\mathcal{P}_K(\{aH\})$.

To see that the action is non-elementary, we observe that H acts transitively on itself and since H is infinite and \hat{d} is proper for any $K \geq 0$ we can find $h_0 = 1, h_1, h_2 \in H$ such that $\hat{d}(h_i, h_j) > K$, $\hat{d}(h_i^{-1}, h_j^{-1}) > K$.

Now choose $s \in S \setminus H$, where S is the set that is used for $\Gamma(G, S \sqcup H)$. Set $g_i = h_i s h_i$. Then $h_i^{\pm} \in \pi_H(g_i^{\pm}H)$. And we can apply Theorem 3.10 to H and g_0, g_1, g_2 .

5.4. Mapping class groups. In this section we assume that the reader is familiar with the theory of curve complexes and subsurface projections, as developed in [MM1, MM2]. Let Σ be closed connected oriented surface with finitely many punctures, and supporting a finite-area hyperbolic metric. We will consider the collection $\mathbf{Y} = \{\mathcal{C}(Y)\}$ of all curve complexes of isotopy classes of subsurfaces Y of Σ obtained cutting Σ along a non-separating simple closed curve. Two such subsurfaces Y, Z overlap if and only if they are not isotopic, so that, just as in [BBF2, Page 6], in view of of results in [MM2] and [Beh] we have that \mathbf{Y} with subsurface projections satisfies axioms (P1)–(P5).

Theorem 5.10. For \mathbf{Y} as above, $MCG(\Sigma)$ acts acylindrically and non-elementarily on $\mathcal{P}_K(\mathbf{Y})$ for all sufficiently large K.

Proof. By Theorem 3.9, it suffices to show that, for some sufficiently large K, if $\mathbf{Y}_K(X,Z)$ contains 4 distinct elements $\mathcal{C}(X_0) < \mathcal{C}(Y_0) < \mathcal{C}(Y_1) < \mathcal{C}(X_1)$ then the common stabilizer of X_0, Y_0, Y_1, X_1 is finite (finite subgroups of $MCG(\Sigma)$ have bounded cardinality). Since the stabilizer of the (isotopy class of the) subsurfaces we are considering coincides with the stabilizer of (the isotopy class of) either of its boundary components, it suffices to show that $\partial X_0, \partial Y_0, \partial Y_1, \partial X_1$ fill the subsurface (if K is large enough). Consider any essential simple closed curve c, and let us show that it intersects the boundary of one of the subsurfaces. Up to switching X and Z and re-indexing, we can assume that c is not parallel to ∂Y_0 , so that c has a well-defined subsurface projection to $\mathcal{C}(Y_0)$. Since ∂X_0 and ∂Y_1 have far away subsurface projection to $\mathcal{C}(Y_0)$, c must intersect one of them. The action is nonelementary by Theorem 3.10; the details are left to the reader.

Remark 5.11. We note that in this example the metric spaces in **Y** are not subspaces of a single ambient space as in our other examples. We also note that we are not using the hyperbolicity of the curve graph or the fact that the mapping class groups acts on it acylindrically. In fact, the current proofs of the projection axioms are fairly simple so this gives a reasonably quick proof that the mapping class group is acylindrically hyperbolic.

Remark 5.12. In general, the action of $MCG(\Sigma)$ on standard projection complexes is not acylindrical. For example, say Σ has genus 5 and consider the action of $MCG(\Sigma)$ on $\mathcal{P}_K(Y)$ where Y is the collection of genus 3 subsurfaces with 1 boundary component. Choose a nonseparating curve a on Σ . Then the Dehn twist in a and its powers fix all subsurfaces in the complement of a, so it suffices to show that the set Y_a of elements of Y disjoint from a form an unbounded set. Choose $Y \in Y_a$ and $f \in MCG(\Sigma)$ that fixes a and is pseudo-Anosov on $\Sigma \setminus a$ and so that the distance in Y between the stable and unstable laminations of f is large compared to K. Then the set $\{f^N(Y) \mid N \in \mathbb{Z}\}$ is unbounded in $\mathcal{P}_K(Y)$ (f acts as a loxodromic isometry).

6. Quasi-trees of metric spaces

We now return to the setup of Section 4 with a collection of metric spaces $\mathbf{Y} = \{(Y, \rho_Y)\}$, projections π_Y and metrics d_Y , obtained from π_Y , satisfying axioms (P1)–(P5).

Here, for simplicity, we will make the extra assumption that *metric spaces* are graphs with each edge having length one. We also assume that projections $\pi_Z(X)$ are subgraphs.

6.1. Construction. Using Theorem 4.1 and Lemma 4.13 we can modify the projections (and suitably increase θ) to replace (P4) with

(P4') if
$$d_Y(X, Z) > \theta$$
 then $\pi_X(Y) = \pi_X(Z)$.

The functions d_Y then satisfy the projection axioms (SP1)–(SP5).

As in [BBF2] we build the quasi-tree of metric spaces $\mathcal{C}_K(Y)$ by taking the union of the metric spaces in Y with an edge of length K > 0 connecting every pair of vertices in $\pi_Y(X)$ and $\pi_X(Y)$ if $d_{\mathcal{P}_K(Y)}(X,Y) = 1$. For convenience, we will assume that θ and K are integers.

We define a metric ρ (that is possibly infinite) on the disjoint union of elements of **Y** by setting $\rho(x_0, x_1) = \rho_X(x_0, x_1)$ if $x_0, x_1 \in X$, for some $X \in \mathbf{Y}$, and $\rho(x_0, x_1) = \infty$ if x_0 and x_1 are in different spaces in **Y**. Assume that the group G acts isometrically on **Y** with this metric and that the projections π_X are G-invariant, i.e., $\pi_{gX}(gY) = g(\pi_X(Y))$. Then G acts isometrically on $\mathcal{C}_K(\mathbf{Y})$. We will give conditions for this action to be acylindrical.

In what follows we will adopt the convention that lower case letters will refer to vertices in $C_K(Y)$ with the corresponding upper case letter denoting the metric space in Y that contains the vertex.

It will be convenient to extend the definition of the projections π_Y to vertices in $\mathcal{C}_K(Y)$. For a vertex $x \in \mathcal{C}_K(Y)$ we set $\pi_Y(x) = \pi_Y(X)$ if $X \neq Y$. If X = Y then $\pi_Y(x) = \{x\}$. We then set $d_Y(x, z) = \dim \pi_Y(x) \cup \pi_Y(z)$ for $x, z \in \mathcal{C}_K(Y)$. We also define

$$\mathbf{Y}_L(x,z) = \left\{ Y \in \mathbf{Y} | d_Y(x,z) > L \right\}$$

and observe that it is possible for X or Z to be in $Y_L(x, z)$.

To save notation, when x and z are vertices of $C_K(Y)$, we denote by $d_C(x, z)$ their distance in $C_K(Y)$.

First of all, we prove a coarsely Lipschitz property of projections:

Lemma 6.1. Assume that $K > \theta$. Let x, z be vertices of $C_K(Y)$ and let $Y \in Y$. If $d_C(x, z) \ge \theta$ then $d_Y(x, z) \le d_C(x, z)$. If $d_C(x, z) \le \theta$ then $d_Y(x, z) \le \theta$.

Proof. If $d_{\mathcal{C}}(x,z) \leq \theta$ then X = Z so $d_Y(x,z) = d_{\mathcal{C}}(x,z) \leq \theta$ if Y = X and $d_Y(x,z) \leq \theta$ by (SP3) otherwise.

In general, we induct on the the distance because our spaces are graphs. If $d_{\mathcal{C}}(x,z) \geq \theta + 1 > \theta$ let x_0 be a vertex adjacent to x such that $d_{\mathcal{C}}(x,z) = d_{\mathcal{C}}(x,x_0) + d_{\mathcal{C}}(x_0,z)$. If $X = X_0$ then $d_{\mathcal{C}}(x,x_0) = 1$ so $d_{\mathcal{C}}(x_0,z) \geq d_{\mathcal{C}}(x,z) - 1 \geq \theta$ since $d_{\mathcal{C}}(x,z) \geq \theta + 1$. Furthermore, since $X = X_0$ either $d_Y(x,z) = d_Y(x_0,z)$ if $X \neq Y$ or $d_Y(x,z) \leq 1 + d_Y(x_0,z)$ if X = Y. In both cases $d_{\mathcal{C}}(x,z) = d_{\mathcal{C}}(x_0,z) + 1 \geq d_Y(x_0,z) + 1 \geq d_Y(x,z)$.

If $X \neq X_0$ then $d_{\mathcal{C}}(x,x_0) = K$ and it may be that $d_{\mathcal{C}}(x_0,z) < \theta$. However, in this case the vertex adjacent to z will be in Z and we can apply the previous case, unless $x_0 = z$. But if $x_0 = z$, then $d_{\mathcal{C}}(x,z) = K$ and $d_Y(x,z) \leq \theta$, therefore this is fine too.

So we can assume that $d_{\mathcal{C}}(x_0, z) \geq \theta$ and $d_Y(x_0, z) \leq d_{\mathcal{C}}(x_0, z)$. Since x and x_0 are adjacent we have that $d_Y(x, x_0) \leq K$ and therefore $d_{\mathcal{C}}(x, z) = K + d_{\mathcal{C}}(x_0, z) \geq d_Y(x, x_0) + d_Y(x_0, z) \geq d_Y(x, z)$.

Lemma 6.2. Assume $K > 2\theta$. Given $x, y, z \in C_K(Y)$ with $Y_K(x, z) \cup \{X, Z\} = \{X = X_0 < X_1 < \dots < X_n = Z\}$ there exists $k \in \{0, \dots, n\}$ such that if i < k then $\pi_{X_i}(y) = \pi_{X_i}(z)$ and if i > k + 1 then $\pi_{X_i}(y) = \pi_{X_i}(x)$.

Proof. Let k be the smallest value such that $\pi_{X_k}(y) \neq \pi_{X_k}(z)$. If there is no such k, or $k \in \{n, n-1\}$, then we are done by setting k = n-1, so we now assume that k exists and k < n-1.

First we observe that $d_{X_{k+1}}(X_k, y) \leq \theta$ for otherwise $\pi_{X_k}(y) = \pi_{X_k}(X_{k+1}) = \pi_{X_k}(z)$ by (P4'). But then, by the triangle inequality, we have $d_{X_{k+1}}(y, X_i) \geq d_{X_{k+1}}(X_i, X_k) - \theta \geq \theta$ for all i > k+1. Therefore, by (P4'), $\pi_{X_i}(y) = \pi_{X_i}(X_{k+1}) = \pi_{X_i}(x)$.

Theorem 6.3. If $K \ge 4\theta$ then for all $x, z \in C_K(Y)$ we have

$$\frac{1}{4} \sum_{Y \in \mathbf{Y}_K(x,z)} d_Y(x,z) \le d_{\mathcal{C}}(x,z) \le 2 \sum_{Y \in \mathbf{Y}_K(x,z)} d_Y(x,z) + 3K.$$

Proof. Let $\mathbf{Y}_K(x,z) \cup \{X,Z\} = \{X = X_0 < X_1 < \cdots < X_n = Z\}$. We first prove the upper bound by finding a path from x to z. Fix points $x_i \in \pi_{X_i}(x) = \pi_{X_i}(X_{i-1})$ and $z_i \in \pi_{X_i}(z) = \pi_{X_i}(X_{i+1})$. Note that $x = x_0$ and $z = z_n$. We then have $d_{\mathcal{C}}(x_i, z_i) \leq d_{X_i}(x_i, z_i) \leq d_{X_i}(x, z)$. Since $d_{\mathcal{P}_K(Y)}(X_i, X_{i+1}) = 1$ we also have that $d_{\mathcal{C}}(z_i, x_{i+1}) = K$. Therefore

$$d_{\mathcal{C}}(x,z) \le \sum_{i} d_{X_{i}}(x,z) + nK.$$

If X_0 (or X_n) are not in $\mathbf{Y}_K(x,z)$ then $d_{X_0}(x,z) < K$ (or $d_{X_n}(x,z) \le K$). On the other hand for $X_i \in \mathbf{Y}_K(x,z)$ we have $d_{X_i}(x,z) + K \le 2d_{X_i}(x,z)$. The upper bound follows.

The lower bound is more involved. Let $x = y_0, y_1, \ldots, y_k = z$ be a geodesic. Fix i_1, \cdots, i_{n-1} such that $\pi_{X_{j-1}}(y_{i_j}) = \pi_{X_{j-1}}(z)$ but $\pi_{X_{j-1}}(y_i) \neq \pi_{X_{j-1}}(z)$ if $i < i_j$. Note that if $\pi_{X_j}(y_i) = \pi_{X_j}(z)$ then $d_{X_j}(X_{j-1}, y_i) \geq \theta$ so $\pi_{X_{j-1}}(y_i) = \pi_{X_{j-1}}(X_j) = \pi_{X_{j-1}}(z)$. Therefore $i_j \leq i_{j+1}$. (However, it is possible that $i_j = i_{j+1}$.)

Next we show that $\pi_{X_{j+3}}(y_{i_j}) = \pi_{X_{j+3}}(x)$. For this we observe that $\pi_{X_{j-1}}(y_{i_j-1}) \neq \pi_{X_{j-1}}(z)$ so by Lemma 6.2 we have that $\pi_{X_{j+1}}(y_{i_j-1}) = \pi_{X_{j+1}}(x)$. This implies that $\pi_{X_{j+1}}(y_{i_j}) \neq \pi_{X_{j+1}}(z)$ for if it did we would have $d_{X_{j+1}}(y_{i_j-1}, y_{i_j}) = d_{X_{j+1}}(x, z) \geq K$ contradicting that y_{i_j-1} and y_{i_j} are consecutive vertices in a geodesic. Another application of Lemma 6.2 implies that $\pi_{X_{j+3}}(y_{i_j}) = \pi_{X_{j+3}}(x)$.

By Lemma 6.1

$$d_{\mathcal{C}}(y_{i_{j-3}}, y_{i_{j+1}}) \ge d_{X_j}(y_{i_{j-3}}, y_{i_{j+1}}) = d_{X_j}(x, z)$$

and therefore

$$d_{\mathcal{C}}(x,z) \ge \sum_{i} d_{X_{j+4i}}(x,z)$$

where j = 0, 1, 2 or 3. Summing over j gives the lower bound.

6.2. Acylindricity.

Theorem 6.4. Let $K \ge 4\theta$. Assume that for each $Y \in \mathbf{Y}$ the stabilizer of Y acts acylindrically on Y with uniform constants independent of Y. Furthermore assume that for some fixed N and B, for any N distinct elements of \mathbf{Y} the common stabilizer is a finite subgroup of size at most B. Then the action of G on $\mathcal{C}_K(\mathbf{Y})$ is acylindrical.

Proof. Fix $D \ge \theta > 0$. By assuming that $D \ge \theta$, by Lemma 6.1 we have that if $d_{\mathcal{C}}(x,x') \le D$ then $d_{Y}(x,x') \le D$ for all $Y \in \mathbb{Y}$. We will use this repeatedly throughout the proof.

By Theorem 3.9, G acts on $\mathcal{P}_K(Y)$ acylindrically so there exists $L_{\mathcal{P}} > 0$ and $B_{\mathcal{P}} > 0$ such that if $d_{\mathcal{P}}(X,Z) > L_{\mathcal{P}}$ then there at most $B_{\mathcal{P}}$ elements $g \in G$ such that $d_{\mathcal{P}}(X,gX), d_{\mathcal{P}}(Z,gZ) < D$. By our assumption, there exist L_Y and B_Y such that for every $Y \in Y$ and $x,z \in Y$ with $d_Y(x,z) \geq L_Y$ there are at most B_Y elements $g \in G$ in the stabilizer of Y such that $d_Y(x,gx), d_Y(z,gz) \leq 2D$. It will be convenient to assume that $L_Y \geq K$.

Fix $X, Z \in Y$ and $x \in X$ and $z \in Z$. Note that it is possible that X = Z. Let $\mathcal{A} = \{g \in G | d_{\mathcal{C}}(x, gx) \leq D \text{ and } d_{\mathcal{C}}(z, gz) \leq D\}$. Using the distance formulas, Corollary 3.7 and Theorem 6.3, we have that there exists an $L_{\mathcal{C}}$ such that if $d_{\mathcal{P}}(x, z) \geq L_{\mathcal{C}}$ then either:

- (1) $d_{\mathcal{P}}(X,Z) \geq L_{\mathcal{P}}$ or
- (2) there exists a $Y \in \mathbf{Y}_{L_{\mathbf{Y}}+2D+4\theta}(x,z)$ and $|\mathbf{Y}_K(x,z)| \leq 2L_{\mathcal{P}}$.

Since the natural projection $C_K(Y) \to \mathcal{P}_K(Y)$ is 1-Lipschitz if $g \in \mathcal{A}$ then $d_{\mathcal{P}}(X, gX) \leq D$ and $d_{\mathcal{P}}(Z, gZ) \leq D$. Therefore if (1) holds there at most $B_{\mathcal{P}}$ elements in \mathcal{A} .

Now assume (2) holds. For any $g \in G$ we have that $d_Y(x,y) \leq d_Y(x,gx) + d_Y(gx,gz) + d_Y(z,gz)$. Therefore if $g \in \mathcal{A}$ we have $d_Y(gx,gz) \geq L_Y + 4\theta$. In particular, $Y \in \mathbf{Y}_{L_Y+4\theta}(gx,gz) \subset \mathbf{Y}_K(gx,gz)$. Now let \mathcal{A}_i be the set of all $g \in \mathcal{A}$ such that gY is the ith element of $\mathbf{Y}_K(gx,gz)$. Since $\mathbf{Y}_K(x,y)$ contains at most $2L_{\mathcal{P}}$ elements, if $i > 2L_{\mathcal{P}}$ then \mathcal{A}_i is empty. We will see that each \mathcal{A}_i has at most $2L_{\mathcal{P}}B_Y$ elements.

Fix $g \in A_i$ and pick an $x' \in \pi_Y(gx)$ and $z' \in \pi_Y(gz)$. Let

$$\mathcal{B} = \{ h \in G | h(Y) = Y, d_Y(x', hx') \le 2D \text{ and } d_Y(z', hz') \le 2D \}.$$

Then by our assumption, there is a constant B_Y that does not depend on Y such that $|\mathcal{B}| \leq B_Y$.

If $g' \in \mathcal{A}_i \subset \mathcal{A}$ then $d_Y(gx, g'x) \leq d_Y(x, gx) + d_Y(x, g'x) \leq 2D$. We also have $d_Y(x', g'g^{-1}x') \leq d_Y(gx, g'x) \leq 2D$ since $x' \in \pi_Y(gx)$ and $g'g^{-1}x' \in g'g^{-1}(\pi_Y(gx)) = \pi_{g'g^{-1}Y}(g'x) = \pi_Y(g'Y)$. Similarly $d_Y(z', g'g^{-1}z') \leq 2D$ and therefore $g'g^{-1} \in \mathcal{B}$. This gives the desired bound on the size of \mathcal{A}_i .

6.3. Tree-gradedness. In this section we study the geometry of $C_K(Y)$. In particular, we prove that it is a quasi-tree (resp. hyperbolic) when the elements of Y are uniform quasi-trees (resp. uniformly hyperbolic), provided that K is large.

Lemma 6.5. Assume $K \ge 4\theta$.

- (1) Let $x, z \in C_K(Y)$ be vertices connected by an edge, and let $Y \in Y$ be so that $d_C(x, \pi_Y(x)) > 3K$. Then $\pi_Y(z) = \pi_Y(x)$.
- (2) Let $x, z \in C_K(\mathbf{Y})$ be vertices with $d_Y(x, z) > \theta$ for $Y \in \mathbf{Y}$, then any path from x to z intersects the (closed) neighborhood of radius 3K around $\pi_Y(x)$ (and also $\pi_Y(z)$).
- *Proof.* (1) By Theorem 6.3, there exists $W \in Y_K(x, \pi_Y(x))$ (and necessarily $W \neq Y$). By Lemma 6.1 we have $W \in Y_{3\theta}(z, \pi_Y(x))$. Hence we have $\pi_Y(x) = \pi_Y(W) = \pi_Y(z)$, where each equality holds either because x and/or z lie in W or because of (P4').
- (2) By (P4'), we have $\pi_Y(x) \neq \pi_Y(z)$ therefore any path from x to z intersects the (closed) neighborhood of radius 3K by (1).

Let $x, z \in \mathcal{C}_K(Y)$. Let $\{X = X_0 < X_1 < \cdots < X_n = Z\}$ be the standard path in $\mathcal{P}_K(Y)$ from Lemma 3.1, as in the proof of Theorem 6.3. We also use the notation x_i, z_i from there. A *standard path* from x to z is a path joining $x = x_0, z_0, x_1, z_0, \cdots, x_n, z_n = z$ in this order such that between x_i, z_i it is a geodesic in $\mathcal{C}(Y_i)$ and that between z_i, x_{i+1} it is an edge.

The following theorem also follows from Theorem 6.7 below, but we give a separate proof since Theorem 6.7 uses results from the literature. Recall that we discussed (a variation of) the bottleneck property right before Theorem 3.5.

Theorem 6.6. Let $K \geq 4\theta$. Suppose every Y satisfies the bottleneck property with uniform constant D. Then $C_K(\mathbf{Y})$ is a quasi-tree.

Proof. First, we note that one may assume that each Y satisfies the (variant of the) bottleneck property for a geodesic between any given two points with respect to a uniform constant which is maybe larger than D, but we keep using D.

Let $x, z \in \mathcal{C}_K(Y)$ be any points. We will check the variant of the bottleneck property for a standard path between x, z. Let $[x_i, z_i]$ be the part (a geodesic) of the standard path in X_i . Let γ be any path from x to z. We want to show $[x_i, z_i]$ is contained in a Hausdorff neighborhood of γ .

For simplicity assume $X_i \neq X_0, X_n$. By Lemma 6.5 (2) applied to x, z, X_i , the path γ intersects each 4K-ball around x_i, z_i (we added a K to account for the diameter of $\pi_{X_i}(x), \pi_{X_i}(z)$). Let w_1, \dots, w_m be the vertices in a subpath of γ starting in the 4K-ball at x_i and ending in the 4K-ball at z_i . Set $y_j = \pi_{X_i}(w_j)$ then y_1, \dots, y_m is a (coarse) path in X_i , starting in the 4K-ball around x_i and ending in the 4K-ball around z_i by Lemma 6.1. (Here a coarse path means that $|y_j - y_{j+1}| \leq \theta$.) By the bottleneck property of X_i , $[x_i, z_i]$ is contained in the L-Hausdorff neighborhood of the path $\{y_i\}$, where L depends only on K, D.

But we now show that the path $\{y_j\}$ is contained in the 4K-Hausdorff neighborhood of the path $\{w_i\}$. To see that fix a vertex y_j . If $d_{X_i}(w_1,w_j)>\theta$, apply Lemma 6.5 (2) to w_1,w_j,X_i , then y_j is contained in the 4K-neighborhood of the path between w_1,w_j . Otherwise $d_{X_i}(w_j,w_m)>\theta$, then we apply the lemma to w_j,w_m and we are done too.

In conclusion, $[x_i, z_i]$ is contained in the (4K + L)-neighborhood of γ . We are left with the case that $X_i = X_0$ (or X_n). But if $|x_0 - z_0| \le 2\theta$ then $[x_0, z_0]$ is contained in the 2θ -neighborhood of γ , or the argument is same as above. \square

We now observe that $C_K(Y)$ is a *tree-graded space*. This notion was introduced in [DS] where tree-graded spaces arise as asymptotic cones of relatively hyperbolic groups, but a simpler example (which is more relevant for us) are Cayley graphs of free products A * B with respect to a generating set contained in $A \cup B$, which are tree-graded with respect to the copies of the Cayley graphs of A and B that they contain.

A geodesic metric space X is said to be tree-graded with respect to the collection of geodesic subspaces \mathcal{P} , called *pieces*, if distinct elements of \mathcal{P} intersect in at most one point, and every simple loop in X is contained in some $P \in \mathcal{P}$.

Theorem 6.7. Let $K \ge 4\theta$. Then there exists C so that $C_K(\mathbf{Y})$ is (C, C)-quasi-isometric to a tree-graded space each of whose pieces is (C, C)-quasi-isometric to some $Y \in \mathbf{Y}$.

Proof. Using Lemma 6.5, one can prove the relative bottleneck property from [Hum] just as in [Hum, Proposition 2.8], so that the conclusion follows from [Hum, Theorem 1].

We now collect some immediate consequences of the theorem and elementary properties of tree-graded spaces.

Corollary 6.8. *Let* $K \ge 4\theta$. *Then:*

- (1) If, for some C, each $Y \in Y$ is (C, C)-quasi-isometric to a tree, then $C_K(Y)$ is a quasi-tree.
- (2) If, for some δ , each $Y \in \mathbb{Y}$ is δ -hyperbolic, then $C_K(\mathbb{Y})$ is hyperbolic.
- (3) $C_K(\mathbf{Y})$ is hyperbolic relative to \mathbf{Y} (more precisely, the family of copies of the $Y \in \mathbf{Y}$ that it contains).

Proof. All properties follow from the analogous statements about tree-graded spaces, as outlined below.

It is readily checked that if the pieces of a tree-graded space X satisfy the bottleneck property uniformly, then the same holds for the whole space. In fact, consider a geodesic γ in X and a path α with the same endpoints. Up to replacing α with a path whose image is contained in α , we can assume that α is injective. The conclusion that γ is contained in a uniform neighborhood of α now easily follows from the fact that any simple loop consisting of a subpath of γ and a subpath of α is contained in a piece.

The proof that if the pieces of a tree-graded space X are hyperbolic then X is hyperbolic follows from a similar argument, where α is now a concatenation of two geodesics.

Finally, tree-graded spaces are hyperbolic relative to their pieces by [DS, Theorem 3.30].

6.4. Examples. Using Theorem 6.4 we briefly discuss the acylindricity of the action on quasi-trees of metric spaces for examples in 5.

First, we use the same setting as in the Section 5.4. Notice that action of $MCG(\Sigma)$ on $C_K(Y)$ is not acylindrical because the action of the stabilizer of any $C(Y) \in Y$ has infinite kernel (generated by a Dehn twist around a boundary component of Y).

¹ In the terminology of [DS], $\mathcal{C}_K(Y)$ is asymptotically tree-graded with respect to Y.

Next, in the same setup as in Theorem 5.2, again without assuming that an axis exists for f that is preserved by EC(f), we also obtain $C_K(Y)$, on which G acts. By Corollary 6.8, $C_K(Y)$ is a quasi-tree.

This action is also acylindrical. Although f fixes a point γ in $\mathcal{P}_K(Y)$, f is a hyperbolic isometry on $\mathcal{C}_K(Y)$ with an axis γ (γ is a subset in $\mathcal{C}_K(Y)$ that is invariant by f. Moreover γ is a geodesic in $\mathcal{C}_K(Y)$, which easily follows from Lemma 6.1).

We record it as a theorem. A similar statement appears as [BBF2, Theorem H], which is weaker since it is only stated that f is WPD in $\mathcal{C}_K(Y)$, but not the acylindricity of the action.

Theorem 6.9. Make the same assumptions as in Theorem 5.2. Then, for a sufficiently large K, $C_K(\mathbf{Y})$ is a quasi-tree on which G acts acylindrically such that the given hyperbolic, WPD element f with an axis γ , is hyperbolic with an axis γ in $C_K(\mathbf{Y})$.

Acknowledgements. The first two authors gratefully acknowledge the support by the National Science Foundation. The third author is supported in part by Grantin-Aid for Scientific Research (No. 15H05739). All four authors were supported by the National Science Foundation under Grant No. DMS-1440140 while the authors were in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester.

References

- [AMS] Y. Antolin, A. Minasyan and A. Sisto, Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups Geom. Dyn.* **10** (2016), 1149–1210. Doi 10.4171/GGD/379. Zbl 1395.20027 MR 3605030
- [Bal] S. H. Balasubramanya, Acylindrical group actions on quasi-trees. *Algebr. Geom. Topol.* **17** (2017), 2145–2176. Doi 10.2140/agt.2017.17.2145. Zbl 06762688 MR 3685605
- [Bar] A. Bartels, Coarse flow spaces for relatively hyperbolic groups. *Compositio Mathematica* **153** (2017), 745–779. Zbl 1366.18012 MR 3631229
- [BB] A. Bartels and M. Bestvina, The Farrell–Jones conjecture for mapping class groups. *Invent. Math.* **215** (2019), 651–712. Doi 10.1007/s00222-018-0834-9. Zbl 07023770 MR 3910072
- [BBF1] M. Bestvina, K. Bromberg and K. Fujiwara, Bounded cohomology via quasi-trees. arXiv:1306.1542v1.
- [BBF2] Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.* **122** (2015), 1–64. Doi 10.1007/s10240-014-0067-4. Zbl 1372.20029 MR 3415065

- [BBFS] M. Bestvina, K. Bromberg, K. Fujiwara and A. Sisto, Projection complexes revisited: A sample. www.math.utah.edu/~bestvina/eprints/acylind-super-short.pdf.
- [Beh] J. A. Венктоск, Asymptotic geometry of the mapping class group and Teichmüller space. *Geom. Topol.* **10** (2006), 1523–1578. Doi 10.2140/gt.2006.10.1523. Zbl 1145.57016 MR 2255505
- [DGO] F. Dahmani, V. Guirardel and D. Osin, Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.* 245 (2017). Doi:10.1090/memo/l156. Zbl:1396.20041 MR 3589159
- [DS] C. Druţu and M. Sapir, Tree-graded spaces and asymptotic cones of groups. *Topology* **44** (2005), 959–1058. With an appendix by D. Osin and Sapir. Zbl 1101.20025 MR 2153979
- [FPS] R. FRIGERIO, M.B. POZZETTI and A. SISTO, Extending higher-dimensional quasi-cocycles. *J. Topol.* **8** (2015), 1123–1155. Doi 10.1112/jtopol/jtv017. Zbl 1360.20034 MR 3431671
- [HO] M. HULL and D. Osin, Induced quasicocycles on groups with hyperbolically embedded subgroups. *Algebr. Geom. Topol.* **13** (2013), 2635–2665. Doi 10.2140/agt.2013.13.2635. Zbl 1297.20045 MR 3116299
- [Hum] D. Hume, Embedding mapping class groups into a finite product of trees. *Groups Geom. Dyn.* 11 (2017), 613–647. Doi 10.4171/GGD/410. Zbl 06759115 MR 3668054
- [Man] J. F. Manning, Geometry of pseudocharacters. *Geom. Topol.* **9** (2005), 1147–1185. Doi 10.2140/gt.2005.9.1147 Zbl 1083.20038 MR 2174263
- [MM1] H. A. MASUR and Y. N. MINSKY, Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.* 138 (1999), 103–149. Doi 10.1007/s002220050343. Zbl 0941.32012 MR1714338
- [MM2] Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.* **10** (2000), 902–974. Doi 10.1007/PL00001643. Zbl 0972.32011 MR 1791145
- [MSW1] L. Mosher, M. Sageev and K. Whyte, Quasi-actions on trees.

 I. Bounded valence. *Ann. of Math.* (2) **158** (2003), 115–164.

 Doi 10.4007/annals.2003.158.115. Zbl 1038.20016 MR 1998479
- [MSW2] Quasi-actions on trees II: Finite depth Bass—Serre trees. *Mem. Amer. Math. Soc.* **214** (2011). Doi 10.1090/S0065-9266-2011-00585-X. Zbl 1234.20034 MR 2867450
- [Osi] D. Osin, Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.* **368** (2016), 851–888. Doi 10.1090/tran/6343. Zbl 1380.20048 MR 3430352
- [Sis] A. Sisto, Quasi-convexity of hyperbolically embedded subgroups. *Math. Z.* **283** (2016), 649–658. Doi 10.1007/s00209-016-1615-z. Zbl 1380.20044 MR 3519976

Mladen Bestvina, Department of Mathematics, University of Utah, 155 South 1400 East, Salt Lake City, UT 84112, USA

e-mail: bestvina@math.utah.edu

Ken Bromberg, Department of Mathematics, University of Utah, 155 South 1400 East, Salt Lake City, UT 84112, USA *e-mail:* bromberg@math.utah.edu

Koji Fujiwara, Department of Mathematics, Kyoto University, Kitashirakawa Oiwake-cho, Sakyo-ku, Kyoto 606-8502, Japan *e-mail:* kfujiwara@math.kyoto-u.ac.jp

Alessandro Sisto, ETH Zurich, Department of Mathematics, Raemistrasse 101, 8092 Zurich, Switzerland

e-mail: sisto@math.ethz.ch