INDUCED QUASI-ISOMETRIES OF HYPERBOLIC SPACES, MARKOV CHAINS, AND ACYLINDRICAL HYPERBOLICITY

ANTOINE GOLDSBOROUGH, MARK HAGEN, HARRY PETYT, AND ALESSANDRO SISTO

With an appendix by Jacob Russell

ABSTRACT. We show that quasi-isometries of (well-behaved) hierarchically hyperbolic groups descend to quasi-isometries of their maximal hyperbolic space. This has two applications, one relating to quasi-isometry invariance of acylindrical hyperbolicity, and the other a linear progress result for Markov chains. The appendix, by Jacob Russell, contains a partial converse under the (necessary) condition that the maximal hyperbolic space is one-ended.

CONTENTS

1. INTRODUCTION

In this paper we consider groups G acting on hyperbolic spaces X such that every quasi-isometry of G induces a quasi-isometry of X. We have two main motivations for considering this property, one relating to quasi-isometry invariance of acylindrical hyperbolicity, and the other to Markov chains. We discuss these separately below.

Examples of group actions with this induced quasi-isometry property include relatively hyperbolic groups whose peripheral subgroups are not relatively hyperbolic: the space X is the coned-off Cayley graph. This can be deduced from [BDM09, Thm 4.1], which implies that quasi-isometries map peripheral subgroups into uniform neighbourhoods of peripheral subgroups. Results of this type originated in [DS05].

Our first result, essentially a special case of Theorem 6.10, shows that many *hierarchically hyperbolic groups* (HHGs) also fit into this framework. That is, their quasi-isometries descend to their *maximal hyperbolic space*. We refer to the HHGs in question as *well behaved* in this introduction—all naturally occurring HHGs are well behaved with the right choice of structure. See Definition 6.9.

Theorem 1. *Let $(G, \mathfrak{S})$ be a well-behaved HHG, with maximal hyperbolic space $\mathcal{C}S$, and let $\pi_S : G \to \mathcal{C}S$ be the associated projection. Every quasi-isometry $f : G \to G$ induces a quasi-isometry $\bar{f} : \mathcal{C}S \to \mathcal{C}S$ such that $\pi_S f$ and $\bar{f}\pi_S$ coarsely agree.*

In Appendix A, Jacob Russell proves a partial converse to Theorem 1. This is Theorem A.1, which says that if $\mathcal{C}S$ is one-ended then all quasi-isometries of $\mathcal{C}S$ come from quasi-isometries of G. This generalises a result of Rafi–Schleimer [RS11]; see the appendix for more discussion.

Theorem 1 is a direct consequence of Lemma 7.11 and Corollary 6.11, and provides many examples of pairs (G, X) where quasi-isometries of G induce quasi-isometries of X. Indeed, the class of HHGs includes many: extensions and quotients of mapping class groups [BHS17a, BHMS20, DDLS21, Rus21]; cubical groups, such as special groups [HS20, Che20]; 3–manifold groups [HRSS22]; Artin groups [HMS21]; and combinations thereof [BR20a, BR20b] (in each case, either the HHG structure in the literature is already well-behaved, or it is easily modified to be so). In the case of mapping class groups, the fact that quasi-isometries descend to the curve graph is a consequence of QI rigidity results from [BKMM12, Bow18], but Theorem 1 is a significantly more direct route to the conclusion in this case. For right-angled Artin groups, the fact that quasi-isometries descend to the contact graph of the universal cover of the Salvetti complex is in many cases a consequence of results of Huang [Hua17, Hua16] showing that quasi-isometries descend to the *extension graph*, along with an observation of Kim-Koberda relating this to the contact graph [KK14], but it appears to be new for right-angled Artin groups where the outer automorphism group is sufficiently complicated.

Groups quasi-isometric to acylindrically hyperbolic groups. It is arguably the main open question on acylindrically hyperbolic groups whether being acylindrically hyperbolic is a quasi-isometry–invariant property of groups, as asked, for example, in [DGO17, Problem 9.1] and [Osi18, Question 2.20]. In fact, it is not even known whether acylindrical hyperbolicity is a commensurability invariant, and to highlight how little is known about this question, even the following are unanswered:

- Let G be acylindrically hyperbolic and let H be of the form $G \rtimes \mathbb{Z}/2$. Is H necessarily acylindrically hyperbolic?
- Can a torsion group be quasi-isometric to an acylindrically hyperbolic group?
- Let G be a group quasi-isometric to $Out(F_n)$, with $n \geqslant 3$. Is G necessarily acylindrically hyperbolic?

It is quite possible that acylindrical hyperbolicity is not a quasi-isometry, or even commensurability, invariant, making it quite interesting to obtain partial results in this direction. We isolate properties leading to such partial results in our set-up, where we consider pairs (G, X) where quasi-isometries of the group G induce quasi-isometries of the hyperbolic space X being acted on. The existence of induced quasi-isometries in itself is not sufficient, as the hyperbolic space might be a point. However, it can be combined with some additional restrictions into a short list of assumptions from which proving the next theorem is a fairly routine matter. We state the assumptions on (G, X) informally here, referring the reader to Definition 7.1 for the precise version.

(1) The action of G on X is nonelementary and acylindrical.
(2) Quasi-isometries of G induce quasi-isometries of X.
(3) (Morse detectability) A geodesic in G is Morse if and only if it maps to a parametrised quasi-geodesic in X.

We note that Morse detectability was abstracted in [RST22] (inspired by [ABD21, KL08] among others) as a sufficient condition to show the Morse local-to-global property.

The following combines Corollary 7.5 and Theorem 7.12 (see Remark 7.13).

Theorem 2. *Let G be a group satisfying Definition 7.1 (for example, let G be a well-behaved HHG which is non-elementary and has unbounded maximal hyperbolic space). Then any group quasi-isometric to G is acylindrically hyperbolic.*

Application to Markov chains. The study of Markov chains on groups as a "quasi-isometry invariant" generalisation of random walks was initiated in [GS21]. Given quasi-isometric groups, G and H, and a simple random walk on H, there is no meaningful notion of a random walk on G induced by the one on H. Markov chains resolve this issue for non-amenable groups since the push-forward (see Section 2 for definitions) of a tame Markov chain by a bijective quasi-isometry is again a tame Markov chain, and quasi-isometries between non-amenable groups are bounded distance from bijective ones [Why99].

For a group G acting on a hyperbolic space X with basepoint x_0, we say that a Markov chain (or random walk) $(w_n^o)_n$ in G makes *linear progress with exponential decay in X* if there is a constant $C > 0$ such that for all n and $o \in G$ we have $\mathbb{P}[d_X(ox_0, w_n^o x_0) \geqslant n/C] \geqslant Ce^{-n/C}$. In the case of random walks, establishing this property was done in [MT18] for weakly hyperbolic groups. This property feeds into the proof of several results for random walks; such as a Central Limit Theorem for the random walk [MS20], genericity of loxodromic elements [MT18] or that random subgroups of weakly hyperbolic groups are free [TT15], among many others.

In a similar spirit, the main result of [GS21] was establishing that tame Markov chains in G make *linear progress* in the hyperbolic space X for many groups G acting on a hyperbolic space X. Examples of groups admitting such an action on a hyperbolic space include (non-elementary) relatively hyperbolic groups and acylindrically hyperbolic 3-manifold groups. This enabled showing a Central Limit Theorem for random walks on groups quasi-isometric to such groups, and this property is used in [GS23] to study random divergence (see below), More generally, linear progress is intended as a crucial starting point for further study of Markov chains.

One of the applications of Theorem 1 is in establishing linear progress in the case where G is a hierarchically hyperbolic group and X is the maximal hyperbolic space in the HHG structure of G. We work with a more restrictive class of Markov chains than in [GS21], as we require a property that we call *quasi-homogeneity*, see Section 2. Importantly, this property is satisfied by simple random walks and their push-forwards by bijective quasi-isometries. The following is Theorem 5.1 (see Remark 7.13 for the verification of Assumption 2.1 for HHGs, part of which is Theorem 6.10).

Theorem 3. *Let G be a group acting on a hyperbolic space X and satisfying Assumption 2.1 (for example, let G be a well-behaved HHG which is non-elementary and has unbounded maximal hyperbolic space $\mathcal{CS}$). Then any tame, quasi-homogeneous Markov chain on G makes linear progress with exponential decay in X.*

One consequence is that, for a group G as in the theorem, the random divergence defined in [GS23] (and chosen according to tame, quasi-homogeneous Markov chains) is the same as the divergence of G. This means that generic points, chosen according to these Markov chains, realise the worst-case scenario for divergence, see [GS23, Thm 1.1]. Note that if G satisfies the stronger Definition 7.1, then using only Theorem 3 and [GS21, Thm 7.7] one can deduce a Central Limit Theorem on groups quasi-isometric to G. This also follows from Theorem 2 and [MS20].

Questions. Several questions arise. First, it would be interesting to find more classes of group actions on hyperbolic spaces such that quasi-isometries of the group descend to quasi-isometries of the space, in the sense of Theorem 1. For short, we will say in this case that *quasi-isometries descend*.

Problem 1.1. Find more classes of groups admitting non-elementary actions on hyperbolic spaces with the property that quasi-isometries descend.

We highlight two specific instances of Problem 1.1.

Question 1.2. Do CAT(0) groups with rank-one elements admit actions on hyperbolic spaces such that quasi-isometries descend? Do small-cancellation groups (of various flavours)?

A candidate hyperbolic space for CAT(0) groups could be the hyperbolic model from [PSZ22], while for -cancellation groups (say for $C'(1/6)$ groups) it could be the space constructed in [GS18]. For the CAT(0) question, any use of the model from [PSZ22] would have to use the geometric group action on the CAT(0) space in an essential way in view of the example in [Ves23].

More strongly, it would be useful to identify more classes of groups such that Theorems 2 and 3 apply. But we would like to emphasise that variations of said theorems should be possible. For instance, Theorem 2 does not apply to $C'(1/6)$ groups for the candidate space given above as the action is not acylindrical in general, but a more general theorem might. Also, it is not known whether the action of a CAT(0) group on the hyperbolic model is acylindrical. A positive answer to either of the following questions might come from a more general version of Theorem 2.

Question 1.3. Let G be a group quasi-isometric to an acylindrically hyperbolic CAT(0) group. Is G necessarily acylindrically hyperbolic?

Let G be a group quasi-isometric to an infinitely presented $C'(1/6)$ group. Is G necessarily acylindrically hyperbolic?

Of course, one can also ask about analogues of Theorem 3, whose conclusion we refer to as the linear progress for Markov chains property.

Question 1.4. Do CAT(0) groups with rank-one elements admit non-elementary actions on hyperbolic spaces with the linear progress for Markov chains property? Do small-cancellation groups (of various flavours)?

Finally, motivated by the appendix, it is natural to ask the following question.

Question 1.5. Which HHGs have an HHG structure with unbounded products and one-ended $\sqsubseteq-$ maximal hyperbolic space?

Specifically, we believe that the answer is not known for extra-large type Artin groups and extensions of lattice Veech groups, for instance. In fact, there is even no known classification of right-angled Coxeter groups admitting an HHG structure as in the question. However, for right-angled Artin groups, one can use [ABD21] together with the HHG structure with unbounded products and maximal hyperbolic space a quasi-tree [BHS17b], to see that no HHG structure as in the question has one-ended maximal hyperbolic space.

Outline of paper and proofs. Section 2 contains some general geometric group theory preliminaries, as well as preliminaries on Markov chains. We also state the relevant assumptions on group actions and Markov chains that we will need later.

The hardest theorem in this paper is Theorem 3, whose proof is contained in Sections 3–5. In particular, in Sections 3 and 4, which contain the core of the geometric arguments involved, we will show roughly that with positive probability the Markov chain has a bounded projection to the axis of a fixed WPD element. This will be then used in Section 5 to check a criterion for linear progress from [GS21].

The rough idea to show the bounded projection property is the following, and the reader might want to look at Figure 1 for reference. If the property fails, then with overwhelming probability the Markov chain creates a very large projection onto the axis. However, there is a positive probability that the Markov chains starts off in a Morse direction different from the axis. If it does, it needs to undo this second projection first, before creating the projection on the axis. We can also repeat this argument with another Morse direction which has a very different Morse gauge. Therefore, the Markov chain (from any basepoint due to quasi-homogeneity) has a large probability of creating a large projection in two directions with very different Morse gauges. This is not yet a contradiction, because we need to know that the two directions have different Morse gauges "close to the basepoint" rather than, say, starting out in the same way and then diverging later. This is what the notion of incompatible Morse rays from Section 3 is supposed to capture, and in that section we study it and prove the required preliminary results.

In Section 6 we show that quasi-isometries of well-behaved HHS descend to their maximal hyperbolic space, and related results. There are two main ideas here. The first one is to use a result from [BHS21] which says, roughly, that quasi-isometries between HHSs descend to quasi-isometries of "simpler" HHSs (certain so-called factored spaces) obtained coning-off certain product regions. One can repeat the procedure until the "simpler" HHSs are actually just hyperbolic spaces, but those will not in general be the maximal hyperbolic spaces of the HHSs as there might be further (quasiconvex) subspaces to cone-off to get there; this happens even for mapping class groups. The additional idea allows us to recognise these subspaces, and roughly we show in Proposition 6.8 that two points are in the same subspace to be coned-off if and only if their coarse fibres in the original HHS are "parallel".

In Section 7 we show our results related to quasi-isometry invariance of acylindrical hyperbolicity. The idea here is the following. If one has a group G acting on a hyperbolic space X with the property that quasi-isometries descend, and if H is quasi-isometric to G, then H quasi-acts on G and therefore on X. The quasi-action on X can be promoted to an action on a space Y quasi-isometric to X, which is also hyperbolic. This action admits a loxodromic element due to the classification of actions on hyperbolic spaces, and what is left to show is that any loxodromic is WPD. This comes from acylindricity of the original action, which can be translated into a geometric property about preimages of balls being geometrically separated as in [Sis16].

Finally, the appendix by Jacob Russell contains the result about quasi-isometries of the maximal hyperbolic space of an HHG coming from quasi-isometries of the HHG, under suitable conditions.

2. Background and assumptions

To set notation, we recall some standard notions from geometric group theory. All hyperbolic spaces considered in this paper will be geodesic.

2.1. **Geometric group theory notions and definitions.** A map $f : Y \to X$ between metric spaces is called a (λ, ϵ)-*quasi-isometric embedding*, with $\lambda \geqslant 1, \epsilon \geqslant 0$, if for all $x, y \in Y$ we have: $\lambda^{-1} d_Y(x, y) - \epsilon \leqslant d_X(f(x), f(y)) \leqslant \lambda d_Y(x, y) + \epsilon$. We say that f is a (λ, ϵ)–quasi-isometry if, in addition, for all $x \in X$, there exists an element $y \in Y$ such that $d_X(f(y), x) \leqslant \epsilon$. If Y is a segment of $\mathbb{R}$, we call the image of f in X a (λ, ϵ)-*quasi-geodesic*.

We will call a (λ, λ)-quasi-geodesic a λ-*quasi-geodesic*, and similarly for quasi-isometric embeddings and quasi-isometries. A subset Y of a geodesic metric space X is *quasi-convex* if there is a constant $C \geqslant 0$ such that all geodesics with endpoints in Y stay within the C-neighbourhood of Y. Further, we say that a map $f : X \to Y$ between metric spaces is R-*coarsely Lipschitz* if $d_Y(f(x), f(y)) \leqslant R d_X(x, y) + R$ for all $x, y \in X$.

Let M be a function $[1, \infty) \times [0, \infty) \to [0, \infty)$. We say a (quasi)-geodesic γ is M-*Morse* if any (λ, ϵ)-quasi-geodesic with endpoints on γ stays within the $M(\lambda, \epsilon)$-neighbourhood of γ. We call M the *Morse gauge of γ*. We can always assume that M takes values in $\mathbb{N}$. Note that a Morse quasi-geodesic is quasi-convex.

If a group G acts on a hyperbolic metric space X (with basepoint x_0), we say an element g is *loxodromic* if the map $\mathbb{Z} \to X$ given by $n \to g^n x_0$ is a quasi-isometric embedding. In this case, $\langle g \rangle x_0$ is quasi-convex in X. We say that g satisfies the *weak proper discontinuity condition*, or that g is *WPD*, if for all $\kappa > 0$ and $x_0 \in X$ there exists $N \in \mathbb{N}$ such that

$$\#\{h \in G \mid \quad d_X(x_0, hx_0) < \kappa, \quad d_X(g^N x_0, h g^N x_0) < \kappa\} < \infty.$$

Each loxodromic WPD element g is contained in a unique maximal elementary subgroup of G, denoted $E(g)$ and called the *elementary closure* of g, see [DGO17, Lemma 6.5].

Finally, all groups we consider are finitely generated, and whenever we consider a group G we automatically fix a word metric d_G on G coming from a finite generating set.

2.2. **Assumptions on the group action.** Now that we recalled the relevant notions, we can state our assumptions on the group action on a hyperbolic space that we will use for our result on linear progress. The slightly different assumptions needed for the results on acylindrical hyperbolicity are postponed to Section 7.

Let G be a group acting on a hyperbolic space X and fix $x_0 \in X$. We write $\rho : G \to X$ for the corresponding orbit map, though sometimes we shall suppress this.

Assumption 2.1. We assume the following.

(1) Some element of G acts on X as a loxodromic WPD.
(2) *Quasi-isometries of G descend to X.* This means that, for each ν there exists λ such that if $\phi : G \to G$ is a ν–quasi-isometry, then there is some λ–quasi-isometry $\bar{\phi} : X \to X$ such that $d_X(\phi(g)x, \bar{\phi}(gx)) \leqslant \lambda$ for all $g \in G$ and $x \in X$.
(3) *Partial (Morse) detectability.* This means: for every Morse gauge M there exists λ such that if $\gamma \subset G$ is an M–Morse geodesic, then $\rho\gamma \subset X$ is a λ–quasi-geodesic.

The notion of *Morse detectability* of G is [RST22, Def. 4.17], which is stated as an equivalence of two properties, and the partial detectability hypothesis above is one of the two implications.

2.3. **Projections.** Given a group G acting on a hyperbolic space X, we will make use of X to define 'projection maps' to subsets of G. The following definition makes this precise.

Definition 2.2. ([GS21, Definition 3.2]) Let G act on a hyperbolic space X, and fix $x_0 \in X$. For $A \subseteq G$, an *X-projection* is a retraction $\pi_A : G \to A$ such that for all $g \in G$ the point $\pi_A(g)x_0$ is a closest point in Ax_0 to gx_0.

Given a subset B of a metric space X, we call a map $\pi : X \to B$ a *closest-point projection* if $d_X(x, \pi(x)) = d_X(x, B)$ for all $x \in X$. The following lemma is a well-known exercise in hyperbolic geometry, see e.g. [GS21, Lemma 2.1].

Lemma 2.3. *Let X be a δ-hyperbolic space. Let Q be a quasi-convex set and $\pi_Q : X \to Q$ a closest-point projection. There exists a constant $R > 0$ depending only on δ and the quasi-convexity constant such that the following hold.*

(1) *π_Q is R-coarsely Lipschitz.*
(2) *For all $x, y \in X$ with $d_X(\pi_Q(x), \pi_Q(y)) \geqslant R$ and for any geodesic $[x, y]$ from x to y, there are points $m_1, m_2 \in [x, y]$ such that $d_X(m_1, \pi_Q(x)) \leqslant R$ and $d_X(m_2, \pi_Q(y)) \leqslant R$. Furthermore, the subgeodesic of $[x, y]$ from m_1 to m_2 lies in the R-neighbourhood of Q.*
(3) *If $\pi'_Q : X \to Q$ is another closest-point projection, then $d_X(\pi_Q(x), \pi'_Q(x)) \leqslant R$ for all $x \in X$.*

A particular consequence of the final part of Lemma 2.3 is that if a group G acts on a hyperbolic space X, then whenever $A \subset G$ has quasiconvex orbit and π_A is an X–projection, we have a uniform bound $d_X(\pi_A(p)x_0, \pi_{Ax_0}(px_0)) \leqslant R$.

The following lemma is well known. Inequalities of this type are often referred to as "Behrstock inequalities".

Lemma 2.4. *Let X be a δ-hyperbolic space and let $Q_1 \neq Q_2$ be two quasi-convex sets and $\pi_{Q_i} : X \to Q_i$ be closest point projections. There is a constant B only depending on δ and the quasi-convexity constants of Q_1 and Q_2 such that for all $x \in X$ we have*

$$d_X\left(\pi_{Q_1}(x), \pi_{Q_1}(Q_2)\right) > B \implies d_X\left(\pi_{Q_2}(x), \pi_{Q_2}(Q_1)\right) \leqslant B.$$

The following lemma is an exercise in hyperbolic geometry and it states that in a hyperbolic space quasi-isometries and closest-point projections are coarsely compatible.

Lemma 2.5. *Let X be a δ-hyperbolic space and $f : X \to X$ a λ-quasi-isometry. Let $Q \subseteq X$ be a quasi-convex subspace and let $\pi_Q : X \to Q$ be a closest point projection. Then there exists a constant N, depending only on λ, δ, and the quasiconvexity constant, such that for all $x \in X$,*

$$d_X\left(\pi_{f(Q)}(f(x)), f(\pi_Q(x))\right) \leqslant N.$$

Notation 2.6. For a group G acting on a hyperbolic space with fixed basepoint x_0 and corresponding orbit map ρ, given a subset $\alpha \subseteq G$ (usually a quasi-geodesic), we will always implicitly fix an X-projection $\pi_\alpha : G \to \alpha$. We also abbreviate the diameter of the union of two projections as measured in X by

$$d_\alpha(x,y) \ = \ \operatorname{diam}_X \left(\rho\big(\pi_\alpha(x) \cup \pi_\alpha(y) \big) \right),$$

where x and y can be either points or subsets of G.

2.4. **Markov chains.** We refer the reader to [GS21] for more background information on Markov chains, while here we describe the notion informally. A Markov chain on a group arises when transition probabilities $p(g,h)$ are assigned for all $g, h \in G$. These encode the probability that the Markov chain "jumps" from g to h in one step. The probability of going from g to h in n steps is a sum over all possible trajectories, that is, sequences of jumps, of length n of executing that exact sequence of jumps, and this is the product of the relevant transition probabilities.

We usually denote a Markov chain on a group by (w_n^*), where w_n^o denotes the position of the Markov chain starting at o after n steps, and we are usually interested in quantities such as $\mathbb{P}[w_n^o = g]$, the probability of getting from o to g in n steps.

The following notion of tameness was defined in [GS21].

Definition 2.7 (Tame). A Markov chain on G is *tame* if it satisfies the following:

   (1) *Bounded jumps:* There exists a finite set $S \subseteq G$ such that $\mathbb{P}[w_1^g = h] = 0$ if $h \notin gS$.
   (2) *Non-amenability:* There exist $A > 0$ and $\rho < 1$ such that for all $g, h \in G$ and $n \geqslant 0$ we have $\mathbb{P}[w_n^g = h] \leqslant A\rho^n$.
   (3) *Irreducibility:* For each $s \in G$ there exist constants $\epsilon_s, K_s > 0$ such that for all $g \in G$ we have $\mathbb{P}[w_k^g = gs] \geqslant \epsilon_s$ for some $k \leqslant K_s$.

For a bijection $\phi : G \to H$ and a Markov chain (w_n^*) on G, there is a natural push-forward Markov chain on H, which we denote by $\phi_\#(w_n^*)$. This is the Markov chain such that $\mathbb{P}\big[\phi_\#(w_n^o) = h\big] = \mathbb{P}\big[w_n^o = \phi^{-1}(h)\big]$ for all $h, o \in H$ and $n \geqslant 0$. To clarify, $\phi_\#(w_n^o)$ is the instance of the Markov chain starting at $\phi(o)$.

Definition 2.8 (Quasi-homogeneous). A Markov chain (w_n^*) is *quasi-homogeneous* if it has the following property for some ν. For every $p, q \in G$ there is a bijective ν–quasi-isometry $\phi : G \to G$ with $\phi(p) = q$ and $\phi_\#(w_n^o) = (w_n^{\phi(o)})$ for all $o \in G$.

Remark 2.9. A random walk driven by a measure whose support is bounded and generates the group as a semi-group is a tame Markov chain, and moreover any push-forward of such a random walk by a bijective quasi-isometry is a tame Markov chain by [GS21, Lemma 2.8]. Since any random walk is group-invariant, it is readily seen that such a push-forward is in fact also quasi-homogeneous.

3. INCOMPATIBLE MORSE RAYS

In this section we consider a group G acting on a hyperbolic space X such that Assumption 2.1.(3) holds, that is to say, Morse geodesics in G map to parametrised quasi-geodesics in X.

In particular, we are interested in a criterion to guarantee that two Morse rays in G travel in genuinely distinct directions in X. The key definition to achieve this is the following, which roughly describes a ray which, while being Morse, has an initial subgeodesic of controlled length which is "not very Morse".

Definition 3.1. Let Z be a metric space, let $M : [1, \infty) \times [0, \infty) \to \mathbb{N}$ be a Morse gauge, and let $\kappa, L \geqslant 0$. We say a Morse ray $\beta : [0, \infty) \to Z$ is (M, κ, L)-*incompatible* if there exists a (k,c)-quasi-geodesic μ with endpoints on $\beta|_{[0,L]}$, such that there is a point $p \in \mu$ with $d(p, \beta) > M(k, c + 2\kappa) + 2\kappa$.

The main result in this section is Lemma 3.5, guaranteeing the existence of incompatible rays. We need two preliminary results. The first one says that, given a Morse ray γ and a quasi-geodesic ray α, either α is contained in a controlled neighbourhood of γ, or α at some point starts diverging from γ at a linear rate.

Lemma 3.2 ([CM19], Corollary 4.3)**.** *Let α be a M-Morse quasigeodesic ray and let β be an (a, b)-quasi-geodesic ray, both in some proper geodesic metric space Z. There exist $\kappa_1 = \kappa_1(M, a, b)$ and $\kappa_2 = \kappa_2(\kappa_1, a, b)$ such that the following holds. If $d(\alpha(0), \beta) \leqslant \kappa_1$ then there are two possibilities.*

- *The ray β is contained in the κ_2-neighbourhood of α.*
- *There exists a last time T_0 such that for all $t \in [0, \infty)$ we have*

$$d(\beta(t), \alpha) \geqslant \frac{1}{2a}(t - T_0) - 2(b + \kappa_1).$$

We note that in [CM19] α is a geodesic ray, but since our α is Morse it lies within finite Hausdorff distance of a geodesic ray.

The following result says that a geodesic triangle where two sides are Morse is thin.

Lemma 3.3. [Cor17, Lemma 2.2] *For all Morse gauges M_1, M_2 there is a constant $\Delta = \Delta(M_1, M_2)$ such that every geodesic triangle where two of the sides are respectively M_1- and M_2-Morse is Δ-thin.*

The following lemma states that an M-Morse ray and an M-incompatible Morse ray have bounded projection onto each other. The constant dependencies are rather involved, but they are crucial for the proof of the key Proposition 4.1. Recall the notation that the orbit map $\rho : G \to X$ given by $g \mapsto gx_0$ is K–Lipschitz.

Lemma 3.4. *Suppose that G acts on a δ–hyperbolic space X and that Assumption 2.1.3 holds. For every Morse gauge M and constant ν there exists $\kappa = \kappa(M, \nu)$ such that for every Morse gauge M' and constant L there exists $D = D(M, \nu, M', L)$ such that the following holds.*

Let $\alpha \subset G$ be an M-Morse ν-quasi-geodesic ray issuing from $1 \in G$. If $\beta \subset G$ is an M'-Morse ray issuing from $1 \in G$ that is (M, κ, L)-incompatible, then

$$\mathrm{diam}_G\left(\pi_\alpha(\beta)\right) \leqslant D \quad and \quad \mathrm{diam}_G\left(\pi_\beta(\alpha)\right) \leqslant D.$$

Proof. Let κ_1 and κ_2 be the constants given by applying Lemma 3.2 to the M–Morse quasigeodesic α and the geodesic β. According to Assumption 2.1.(3), there are $\lambda_\alpha = \lambda_\alpha(M)$ and $\lambda_\beta = \lambda_\beta(M')$ such that $\rho\alpha \subset X$ is a λ_α–quasigeodesic and $\rho\beta \subset X$ is a λ_β–quasigeodesic. Let $R = R(\lambda_\alpha)$ be given by applying Lemma 2.3 to $\rho\alpha \subset X$. and let $\Delta = \Delta(M, M')$ be the constant from Lemma 3.3.

We first bound $\mathrm{diam}_G(\pi_\alpha(\beta))$. For this, we shall consider sufficiently large constants

$$\epsilon = \epsilon(\lambda_\alpha, \lambda_\beta, \delta) \quad and \quad U = U(M, M', R, \epsilon, \Delta).$$

From these, we define

$$D_1 = (U + L + R + \epsilon + 2\Delta + 2M(1, 0) + 6\kappa_1)K\lambda_\alpha + \lambda_\alpha^2.$$

Suppose that, contrary to the desired result, there is some $v \in \beta$ with the property that $d_G(1, \pi_\alpha(v)) > D_1$. Since $\rho\alpha$ is a λ_α–quasigeodesic, we have $d_X(x_0, \pi_\alpha(v)x_0) > K(U + L + R + \epsilon)$. As noted after Lemma 2.3, $\pi_\alpha(v)x_0$ differs from $\pi_{\rho\alpha}(vx_0)$ by at most R.

Since $\rho\beta$ is a λ_β–quasigeodesic, the Morse lemma implies that every geodesic from x_0 to vx_0 stays uniformly close to $\rho\beta$, so Lemma 2.3 provides a point $u \in \beta$ such that $d(ux_0, \pi_{\rho\alpha}(vx_0)) \leqslant \epsilon$, where $\epsilon = \epsilon(\lambda_\alpha, \lambda_\beta, \delta)$ is a uniform constant. Let us write $u = \beta(t)$. By the construction of u, we have

$$\begin{aligned}
t - U = d_G(1, u) - U &\geqslant \frac{1}{K}d_X(x_0, ux_0) - U \\
&\geqslant \frac{1}{K}\big(d_X(x_0, \pi_\alpha(v)x_0) - d_X(\pi_\alpha(v)x_0, \pi_{\rho\alpha}(vx_0)) - d_X(\pi_{\rho\alpha}(vx_0), ux_0)\big) - U \\
&> \frac{1}{K}(K(U + L + R + \epsilon + 2\Delta + 2M(1, 0) + 6\kappa_1) - R - \epsilon) - U \\
&= L + 2\Delta + 2M(1, 0) + 6\kappa_1 > L.
\end{aligned}$$

Claim 1. *We have $\beta[0, t - U] \subset \mathcal{N}_{\Delta + M(1, 0)}^G(\alpha)$.*

Proof. Consider a geodesic $[1, \pi_\alpha(v)]$ in G. Because α is M–Morse, $[1, \pi_\alpha(v)]$ lies in the $M(1, 0)$–neighbourhood of α, and is M_+–Morse, where $M_+ = M + M(1, 0)$. Now consider a geodesic $[u, \pi_\alpha(v)]$ in G. The geodesic triangle formed by $[1, \pi_\alpha(v)]$, $[\pi_\alpha(v), u]$, and $\beta[0, t]$ has two sides that are M'–

and M_+–Morse, so, according to [Cor17, Lem. 2.3], the geodesic $[u, \pi_\alpha(v)]$ is uniformly Morse. Assumption 2.1.(3) then tells us that $\rho[u, \pi_\alpha(v)]$ is a λ'–quasigeodesic, where $\lambda' = \lambda'(M', M_+)$. Since u was constructed so that $d_X(ux_0, \pi_\alpha(v)x_0) \leqslant \epsilon + R$, this implies that $d_G(u, \pi_\alpha(v)) \leqslant \lambda'(R + \epsilon) + \lambda'$. Ensure that U is larger than this bound by at least Δ.

From Lemma 3.3, we know that the above geodesic triangle is Δ–thin. Thus, for $s < t - U$, the point $\beta(s)$ must be Δ–close to $[1, \pi_\alpha(v)]$. In turn, this means that it is $(\Delta + M(1,0))$–close to α. $\diamond$

We write $t' = t - U - 2\Delta - 2M(1,0) - 6\kappa_1$. Note that by the above computation, $t' > L$.

Claim 2. *We have $\beta[0, t'] \subseteq \mathcal{N}_{\kappa_2}^G(\alpha)$.*

Proof. If not, then by Lemma 3.2, there exists $T_0 \leqslant t'$ such that for all $s \geqslant 0$ we have $d_G(\beta(s), \alpha) \geqslant \frac{1}{2}(s - T_0) - 2\kappa_1$. In particular, for $s = t - U$, Claim 1 leads to

$$\Delta + M(1,0) \geqslant d_G(\beta(t-U), \alpha)) \geqslant \frac{1}{2}(t - U - T_0) - 2\kappa_1$$

$$\geqslant \frac{1}{2}(t - U - t') - 2\kappa_1 \geqslant \Delta + M(1,0) + \kappa_1,$$

a contradiction. $\diamond$

Since α is a ν–quasigeodesic and β is a geodesic, it follows from Claim 2 that there is an initial subsegment α' of α that stays κ–close to $\beta[0, t']$, where κ is a uniform constant depending only on κ_2 and ν. That is, $\kappa = \kappa(M, \nu)$. In other words, the Hausdorff-distance between α' and $\beta[0, t']$ is at most κ.

Since $t' > L$, the fact that β is (M, κ, L)–incompatible means that there is some (k, c)–quasigeodesic μ with endpoints μ^- and μ^+ on $\beta[0, t']$ such that there is some $p \in \mu$ with $d_G(p, \beta) > M(k, c+2\kappa) + 2\kappa$. Let y^- and y^+ be closest points in α to μ^- and μ^+, respectively. We have $d(y^\pm, \mu^\pm) \leqslant \kappa$. It follows that the path

$$[y^-, \mu^-] \cup \mu \cup [\mu^+, y^+]$$

is a $(k, c+2\kappa)$–quasigeodesic with endpoints on the M–Morse quasigeodesic α'. But this path contains $p \in \mu$, which is at a distance of more than $M(k, c + 2\kappa) + 2\kappa$ from $\beta[0, t']$. Since $\beta[0, t']$ and α' are at Hausdorff-distance at most κ, this contradicts the fact that α is M–Morse. We conclude from this contradiction that $d_G(1, \pi_\alpha(v)) \leqslant D_1$ for all $v \in \beta$.

We now turn to the other inequality. Given a point $v \in \alpha$, consider $w = \pi_\beta(v)$ and $y = \pi_\alpha(w)$. We aim to bound $d_G(1, w)$. From the above, we know that $d_G(1, y) \leqslant D_1$.

Letting R_β be given by applying Lemma 2.3 to $\rho\beta$, we have $d_X(wx_0, \pi_{\rho\beta}(vx_0)) \leqslant R_\beta$ and $d_X(yx_0, \pi_{\rho\alpha}(wx_0)) \leqslant R$ by the comment after that lemma. Since $\rho\alpha$ and $\rho\beta$ are uniform quasigeodesics in the δ–hyperbolic space X issuing from the common point x_0, there is a uniform bound on $d_X(\pi_{\rho\beta}(vx_0), \pi_{\rho\alpha}\pi_{\rho\beta}(vx_0))$. Lemma 2.3 also states that $\pi_{\rho\alpha}$ is R–coarsely Lipschitz, so by combining these bounds we obtain a uniform bound on $d_X(wx_0, yx_0)$.

Because the restrictions of α and β between 1 and y and between 1 and w are, respectively, $M-$ and M'–Morse, [Cor17, Lem. 2.3] tells us that a geodesic $[w, y]$ in G is uniformly Morse. By Assumption 2.1.(3), it follows that $\rho[w, y]$ is a uniform quasigeodesic, and hence $d_G(w, y)$ is bounded by some uniform constant ϵ'. We can now compute

$$d_G(1, w) \leqslant d_G(1, y) + d_G(y, w) \leqslant D_1 + \epsilon'. \qquad \square$$

We now show the existence of an (M, κ, L)-incompatible ray, which combined with Lemma 3.4 will be enough to show that we can find two Morse rays that diverge after a short distance.

Lemma 3.5. *Let G be a non-hyperbolic group acting on a hyperbolic space X such that Assumption 2.1.(3) holds. For all Morse gauges M and constants κ, there exist M', L such that there is an M'-Morse ray which is (M, κ, L)-incompatible.*

Proof. Given M and κ, let $M_1(k, c) = M(k, c + 2\kappa) + 2\kappa$. Since G is not hyperbolic, its Morse boundary is not compact [CD19, Cor. 1.17], so there must be some Morse ray β that is not M_1-Morse. Let M' be the Morse gauge of β. As β is not M_1-Morse, there are some $k, c > 0$ such that there is a (k, c)-quasi-geodesic μ with endpoints on β that contains a point p with $d_G(p, \beta) > M_1(k, c)$. Let L be such that the endpoints of μ lie in $\beta[0, L]$. By definition, β is (M, κ, L)–incompatible. $\square$

Corollary 3.6. *Let G be a non-hyperbolic group acting on a hyperbolic space X such that Assumption 2.1.(3) holds. For every Morse gauge M there is a Morse gauge M', a constant $D > 0$, and an M'–Morse ray $\beta \subset G$, issuing from $1 \in G$, such that if α is an M-Morse ray issuing from $1 \in G$ then $\mathrm{diam}\,(\pi_\alpha(\beta)) \leqslant D$ and $\mathrm{diam}\,(\pi_\beta(\alpha)) \leqslant D$.*

Proof. Let $\kappa = \kappa(M)$ be given by Lemma 3.4. By Lemma 3.5, there are M' and L such that there is an M'–Morse ray β that is (M, κ, L)–incompatible. Now apply Lemma 3.4. $\square$

4. BOUNDED PROJECTIONS TO AXES

The goal of this section is to prove Proposition 4.1, which is an important part of the proof of Theorem 5.1 on linear progress in Section 5. It asserts that there is a definite probability that tame, quasi-homogeneous Markov chains have bounded projection to any given loxodromic WPD.

Proposition 4.1. *Let G be a non-hyperbolic group satisfying Assumption 2.1 and let (w_n^*) be a tame, quasi-homogeneous Markov chain. Fix a loxodromic WPD element $g \in G$. There exist $C, \epsilon > 0$ such that for all $p, h \in G$ and for all $n \in \mathbb{N}$, we have*

$$\mathbb{P}\Big[d_{hE(g)}(p, w_n^p) \leqslant C\Big] > \epsilon.$$

Recall from Section 2.1 that $E(g)$ is the elementary closure of g. See Notation 2.6 for the definition of $d_{hE(g)}$. For ease of notation, we write $\gamma = E(g)$ for the remainder of this section. Let us write ν for the quasi-homogeneity constant of (w_n^*). Recall that $\rho : G \to X$ is the orbit map with basepoint x_0.

Lemma 4.2. *For each Morse gauge M and $\nu \geqslant 1$, there is a constant $A > 0$ such that the following holds for all $C > 1$. Let $\xi \subset G$ be an M–Morse ray, and let $p' \in \xi$, $p \in G$. If $\phi : G \to G$ is a bijective ν–quasi-isometry with $\phi(p') = p$, then*

$$d_{\phi\xi}(p, \phi(h)) \; > \; \frac{1}{A}d_\xi(p', h) - A$$

for all $h \in G$. Moreover,

$$\mathbb{P}\Big[d_{\phi\xi}(p, \phi_\#(w_n^{p'})) > C\Big] \; \geqslant \; \mathbb{P}\Big[d_\xi(p', w_n^{p'}) > A(C + A)\Big].$$

Proof. The first statement is essentially a consequence of Assumption 2.1.(2), that ϕ descends to a quasi-isometry of X, and Lemma 2.5, which states that "the translate (by a quasiisometry) of the projection of a point to a quasigeodesic coarsely agrees with the projection of the translate (of the point) to the translate (of the quasigeodesic)". The nature of X–projections is such that a careful argument requires additional small errors. We now give details.

Let $\lambda \geqslant 1$, given by Assumption 2.1, be such that ξx_0 and $\phi(\xi)x_0$ are λ-quasi-geodesics. As noted after Lemma 2.3, if $Y \subset G$ has λ–quasiconvex orbit then there is a uniform constant R such that the maps $\rho\pi_Y$ and $\pi_{Y x_0}\rho$ differ by at most R. Thus

$$d_{\phi\xi}(p, \phi(h)) \; \geqslant \; d_X(\pi_{\phi(\xi)x_0}(px_0), \pi_{\phi(\xi)x_0}(\phi(h)x_0)) - 2R.$$

According to Assumption 2.1.(2), the maps $\rho\phi$ and $\bar{\phi}\rho$ differ by at most λ, so there is a uniform constant R' such that

$$d_{\phi\xi}(p, \phi(h)) \; \geqslant \; d_X\big(\pi_{\bar{\phi}(\xi x_0)}(px_0), \pi_{\bar{\phi}(\xi x_0)}(\bar{\phi}(hx_0))\big) - 2R'.$$

Using Lemma 2.5, there is now a uniform constant N such that

$$d_{\phi\xi}(p, \phi(h)) \; \geqslant \; d_X\big(\bar{\phi}\pi_{\xi x_0}(\bar{\phi}^{-1}(px_0)), \bar{\phi}\pi_{\xi x_0}(hx_0)\big) - 2R' - 2N.$$

As $\rho\phi$ and $\bar{\phi}\rho$ coarsely agree, $\bar{\phi}^{-1}(px_0)$ is uniformly close to $p'x_0$, because $\phi(p') = p$. Also note that, by Lemma 2.3, the map $\pi_{\xi x_0}$ is uniformly coarsely Lipschitz. Hence the fact that $\bar{\phi}$ is a λ–quasi-isometry means that there is a uniform constant R'' such that

$$d_{\phi\xi}(p, \phi(h)) \; \geqslant \; \frac{1}{\lambda}d_X(\pi_{\xi x_0}(p'x_0), \pi_{\xi x_0}(hx_0)) - R''.$$

To complete the proof of the first statement, observe that since $\rho\pi_\xi$ and $\pi_{\xi x_0}\rho$ differ by at most R, we have

$$d_X(\pi_{\xi x_0}(p'x_0), \pi_{\xi x_0}(hx_0)) \; \geqslant \; d_\xi(p', h) - 2R.$$

For the second statement, let $C > 1$. By the first statement, any $h \in G$ with $d_\xi(p', h) > A(C + A)$ also satisfies $d_{\phi\xi}(p, \phi(h)) > C$. We can therefore use quasi-homogeneity and the fact that ϕ is a bijection to compute

$$
\begin{aligned}
\mathbb{P}\big[d_\xi(p', w_n^{p'}) > A(C + A)\big] &= \sum_{d_\xi(p', h) > A(C+A)} \mathbb{P}\big[w_n^{p'} = h\big] \\
&\leqslant \sum_{d_{\phi\xi}(p, \phi(h)) > C} \mathbb{P}\big[w_n^{p'} = h\big] \\
&= \sum_{d_{\phi\xi}(p, h') > C} \mathbb{P}\big[w_n^{p'} = \phi^{-1}(h')\big] \\
&= \mathbb{P}\big[d_{\phi\xi}(p, \phi_\#(w_n^{p'})) > C\big]. \qquad \square
\end{aligned}
$$

The following lemma says that a tame Markov chain has a (small but) positive probability of reaching a given point at distance d within a number of steps linear in d.

Lemma 4.3 ([GS21, Lemma 2.9]). *If (w_n^*) is a tame Markov chain on a finitely generated group G, then there are constants $U, \epsilon_0 > 0$ such that the following holds. For each $p, q \in G$ with $d_G(p, q) = d$ there exists $t \leqslant dU$ such that*

$$
\mathbb{P}[w_t^q = p] \geqslant \epsilon_0^d.
$$

As a final preliminary lemma, the following will allow us to "pivot away" certain quasi-geodesics in X from a given translate $h\rho\gamma$ of $\rho\gamma = E(g)x_0$. The lemma is inspired by [MS20, Lemma 9.5] (but the proof is different and simpler).

Lemma 4.4. *For all θ there exist s, E such that the following holds. Given a θ-quasi-geodesic $\alpha \subset X$ from px_0 that passes through qx_0, and given $h \in G$, there exists $q' \in G$ with $d_G(p, q') \leqslant s$ such that*

$$
d_{h\rho\gamma}(q'q^{-1}\alpha, px_0) \leqslant E \quad and \quad d_{q'q^{-1}\alpha}(h\rho\gamma, q'x_0) \leqslant E.
$$

Proof. To simplify notation, let us write $\alpha' = pq^{-1}\alpha$. The desired element $q' \in G$ will be of the form $q' = pf$, where $d_G(1, f) \leqslant s$. We shall then have $q'q^{-1}\alpha = pfp^{-1}\alpha'$.

Fix any two loxodromic WPD elements $g_1, g_2 \in G$ such that $E(g_1)$, $E(g_2)$, and $E(g) = \gamma$ are pairwise distinct, the existence of which is given by the arguments in [BF02, Proposition 6] (see also [DGO17, Corollary 6.12]). Observe that all three of these contain $1 \in G$. We shall choose $f = g_i^j$ for some i and some bounded j.

At most one direction of $E(g_i)x_0$ can have large coarse intersection with the initial subsegment of $p^{-1}\alpha'$ from $q^{-1}px_0$ to x_0, so after inverting the g_i, we may assume that the positive direction has small coarse intersection with it. Write γ_i for this positive direction of $E(g_i)$. There is some constant C_0 such that for any $k, k' \in G$, both $\pi_{k\rho\gamma}(k'\rho\gamma_i)$ and $\pi_{k'\rho\gamma_i}(k\rho\gamma)$ have diameter at most C_0; see [AMS16, Thm 3.9], for instance (said reference gives geometric separation, which is equivalent to bounded projections). The conclusion of Lemma 2.4 (the Behrstock inequality) therefore holds, with some constant B, for all distinct pairs of translates of $\rho\gamma$, $\rho\gamma_1$, and $\rho\gamma_2$.

We now show that there exists C_1 such that a least one of the $\pi_{p\rho\gamma_i}(\alpha')$ has diameter at most C_1. If this were not the case, then there would exist $x_1, x_2 \in \alpha'$ such that $d_{p\rho\gamma_i}(px_0, x_i)$ were large. By Lemma 2.3, we then also have that $d_X(px_0, x_i)$ is large. Because the $p\rho\gamma_i$ have small coarse intersection with the initial subsegment of α' between $pq^{-1}px_0$ and px_0, we can relabel so that x_1 lies between px_0 and x_2. Because the conclusion of Lemma 2.4 holds for $p\rho\gamma_1$ and $p\rho\gamma_2$, the set $\gamma_p' = p\rho\gamma_1 \cup p\rho\gamma_2$ is quasiconvex. Lemma 2.3 therefore tells us that any geodesic from x_1 to x_2 must pass uniformly close to px_0. By the Morse lemma, this contradicts the fact that α' is a quasigeodesic.

After relabelling, we therefore have that $\mathrm{diam}_X(\pi_{p\rho\gamma_1}(\alpha')) \leqslant C_1$. Because pg_1p^{-1} acts on $pE(g_1)x_0$ with positive translation length independently of α, there is some uniformly bounded j such that the set $\pi_{p\rho\gamma_1}(pg_1^jp^{-1}\alpha') = pg_1^jp^{-1}\pi_{p\rho\gamma_1}(\alpha')$ lies at distance greater than B from $\pi_{p\rho\gamma_1}(h\rho\gamma)$. Setting $q' = pg_1^j$ (i.e. $f = g_1^j$), the desired inequalities follow from the Behrstock inequality, Lemma 2.4.

$\square$

We are now ready to prove the main proposition. Before doing so, we fix various constants and Morse gauges. The dependencies are rather delicate, so we have to be quite careful here. The reader

may prefer to skip this and refer back when checking that the various constants and Morse gauges have the claimed properties.

Constants and Morse gauges. We fix the following.

- M_0 is such that G contains an M_0–Morse ray (Assumption 2.1.(1)).
- ν is the quasi-homogeneity constant (Definition 2.8).
- M is the minimal Morse gauge such that whenever ϕ, ψ are bijective ν–quasi-isometries of G, any M_0–Morse ray gets mapped by $\phi^{-1}\psi$ to an M–Morse ray.
- M' and D are given by applying Corollary 3.6 to the gauge M.
- A is given by applying Lemma 4.2 to the gauge $\max\{M_0, M, M'\}$.
- U and ϵ_0 are as in Lemma 4.3.
- θ is such that the X–orbit of any $\max\{M_0, M, M'\}$–Morse geodesic in G is a θ–quasiconvex θ quasi-geodesic (Assumption 2.1.(3)).
- s and E are given by applying Lemma 4.4 to θ.
- R and B are given by applying Lemmas 2.3 and 2.4, respectively, with quasi-convexity constant θ.
- $C_1 = A(B + A + D + 1)$ and $C_2 = A(C_1 + A)$.
- $d = C_2 + \theta(B + E + 2R + \theta)$.
- $J > 0$ is a constant such that for any $h, p \in G$, consecutive points in $(\pi_{h\gamma} w_n^p)$ have distance at most J. This exists by tameness of the Markov chain and Lemma 2.3.
- $C = d + JU(d + s)$.

Proof of Proposition 4.1. For a contradiction, assume that Proposition 4.1 does not hold. Hence there exist $p, h \in G$ and $n \in \mathbb{N}$ such that:

$$(1) \qquad \mathbb{P}\left[d_{h\gamma}(p, w_n^p) \leqslant C\right] \leqslant \frac{1}{3}\epsilon_0^{d+s},$$

where C is as above. Let β be an M'-Morse ray issuing from p that satisfies the conclusion of Corollary 3.6, and let α be an M_0-Morse ray issuing from p.
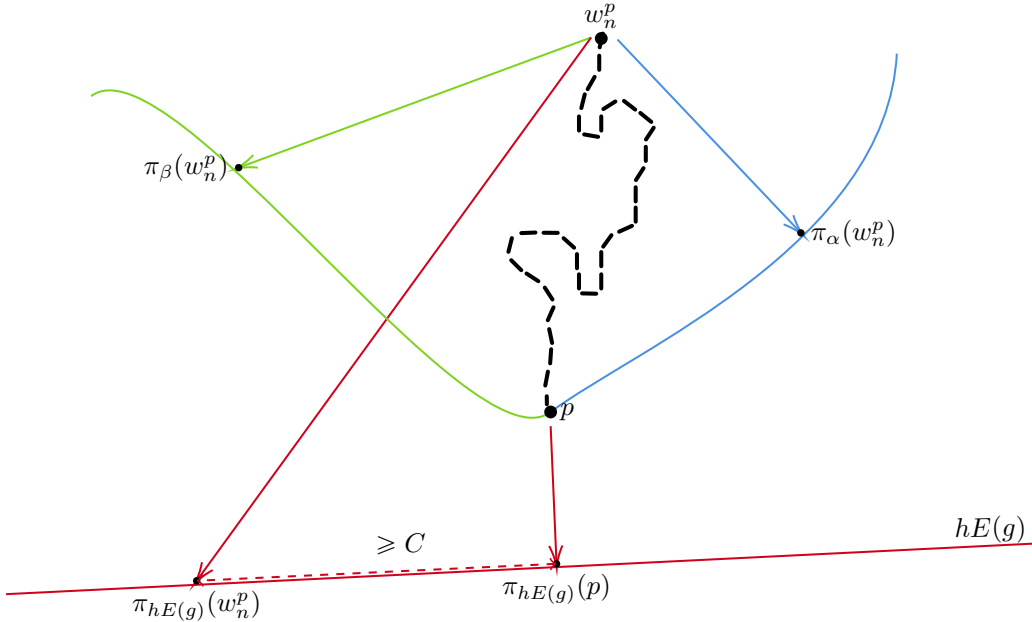

FIGURE 1. For a contradiction, we assume that with overwhelming probability the Markov chain creates a large projection on $hE(g)$. The dotted line is meant to represent the sample path of the Markov chain and α, β are the two Morse rays with very different Morse gauges.

The main ingredient in producing a contradiction is the following claim, which states that if (1) holds then in a definite number of steps, the Markov chain (w_n^p) simultaneously travels a long way along two Morse rays that are in very different directions.

Claim. *There exists a natural number m such that there are bijective ν–quasi-isometries $\phi_\alpha, \phi_\beta : G \to G$ with $\phi_\alpha(p) = \phi_\beta(p) = p$ such that*

$$\mathbb{P}\big[d_{\phi_\alpha\alpha}(p, w_m^p) > C_1\big] \geqslant \frac{2}{3} \quad and \quad \mathbb{P}\big[d_{\phi_\beta\beta}(p, w_m^p) > C_1\big] \geqslant \frac{2}{3}.$$

Proof. We first fix the value of m to be considered. Let $q \in \beta$ have $d_G(p, q) = d$, where d is as above. According to Lemma 4.4, there is some $q' \in G$, with $d_G(p, q') \leqslant s$, such that $\operatorname{diam}_X(\pi_{h\rho\gamma}(q'q^{-1}\rho\beta)) \leqslant E$. Let $p' = q'q^{-1}p$, so that $d_G(p', q') = d(p, q) = d$.

Because $d_G(p, p') \leqslant d + s$, Lemma 4.3 tells us that there is some $t_\beta \leqslant (d + s)U$ such that $\mathbb{P}[w_{t_\beta}^p = p'] \geqslant \epsilon_0^{d+s}$. An analogous construction yields a corresponding number $t_\alpha \leqslant (d + s)U$. Observe that, by (1) and the choice of J,

$$\mathbb{P}\Big[d_{h\gamma}(p, w_{n+t_\beta-t_\alpha}^p) \leqslant C - JU(d + s) = d\Big] \leqslant \frac{1}{3}\epsilon_0^{d+s}.$$

Set $m = n - t_\alpha$. Let us write $\hat{n} = n + t_\beta - t_\alpha$ and $t = t_\beta$. We have $m = \hat{n} - t$.

The above probability that $d_{h\gamma}(p, w_{\hat{n}}^p)$ is at most d is bounded below by the probability of the Markov chain going to p' in exactly t steps and then making little distance along $h\gamma$ for a further $\hat{n} - t$ steps. This gives us

$$\epsilon \geqslant \mathbb{P}\Big[w_t^p = p'\Big] \cdot \mathbb{P}\Big[d_{h\gamma}(p, w_{\hat{n}-t}^{p'}) \leqslant d\Big] \geqslant \epsilon_0^{d+s}\mathbb{P}\Big[d_{h\gamma}(p, w_{\hat{n}-t}^{p'}) \leqslant d\Big],$$

from which we deduce that $\mathbb{P}[d_{h\gamma}(p, w_{\hat{n}-t}^{p'}) \leqslant d] \leqslant \frac{1}{3}$. Since $m = \hat{n} - t$, we have established that

$$\mathbb{P}\Big[d_{h\gamma}(p, w_m^{p'}) > d\Big] \geqslant \frac{2}{3}.$$

Now let us consider an arbitrary point $k \in G$ satisfying $d_{h\gamma}(p, k) > d$. Since $h\rho\gamma$ is a θ–quasigeodesic, Lemma 2.3 tells us that $d_{h\rho\gamma}(px_0, kx_0) > \frac{d}{\theta} - \theta - R$. By the choice of q' via Lemma 4.4, we therefore have

$$d_{h\rho\gamma}(kx_0, q'q^{-1}\rho\beta) > \frac{d}{\theta} - \theta - R - E > B.$$

By the Behrstock inequality, Lemma 2.4, we have $d_{q'q^{-1}\rho\beta}(h\rho\gamma, kx_0) \leqslant B$. Now, using Lemmas 4.4 and 2.3, the fact that $q'q^{-1}\rho\beta$ is a θ–quasigeodesic means that this leads to

$$\begin{aligned}
d_{q'q^{-1}\beta}(p', k) &\geqslant d_{q'q^{-1}\beta}(p', q') - d_{q'q^{-1}\beta}(q', k) \\
&\geqslant d - \theta(d_{q'q^{-1}\rho\beta}(q'x_0, kx_0) + 2R + \theta) \\
&\geqslant d - \theta(B + E + 2R + \theta) = C_2.
\end{aligned}$$

In particular, our previous estimate on the Markov chain implies that

$$\mathbb{P}\Big[d_{q'q^{-1}\beta}(p', w_m^{p'}) > C_2\Big] \geqslant \frac{2}{3}.$$

By quasi-homogeneity, Definition 2.8, there is a bijective ν–quasiisometry $\phi : G \to G$ that respects the Markov chain and has $\phi(p') = p$. Applying Lemma 4.2 with $\xi = q'q^{-1}\beta$, we therefore have

$$\begin{aligned}
\mathbb{P}\Big[d_{\phi q'q^{-1}\beta}(p, w_m^p) > C_1\Big] &= \mathbb{P}\Big[d_{\phi q'q^{-1}\beta}(p, \phi_\#(w_m^{p'})) > C_1\Big] \\
&\geqslant \mathbb{P}\Big[d_{q'q^{-1}\beta}(p', w_m^{p'}) > A(C_1 + A)\Big] \\
&\geqslant \mathbb{P}\Big[d_{q'q^{-1}\beta}(p', w_m^{p'}) > C_2\Big] \geqslant \frac{2}{3}.
\end{aligned}$$

The statement follows by letting ϕ_β be the quasiisometry $\phi q'q^{-1}$ of X, which depends on β because q and q' do.

$\diamondsuit$

In view of the claim, there must be some $q \in G$ such that $d_{\phi\alpha}(p, q) > C_1$ and $d_{\phi\beta}(p, q) > C_1$. Let $\alpha' = \phi_\beta^{-1}\phi_\alpha\alpha$ and $q' = \phi_\beta^{-1}(q)$. Because $\phi_\beta(p) = p$, the first part of Lemma 4.2 tells us that

$$d_{\alpha'}(p, q') > \frac{1}{A}d_{\phi_\alpha\alpha}(p, q) - A > \frac{C_1}{A} - A.$$

The choice of β via Corollary 3.6 means that $\text{diam}(\pi_\alpha(\beta)) \leqslant D$, so since $p \in \beta$, we have $d_{\alpha'}(q', \beta) > B$ by definition of C_1. Similarly, $d_\beta(q', \alpha) > B$. This contradicts the Behrstock inequality, Lemma 2.4. $\square$

5. Linear progress

Our main theorem on Markov chains is the following.

Theorem 5.1. *Let G be a group acting on a hyperbolic space X with basepoint x_0 and satisfying Assumption 2.1. If $(w_n^*)_n$ is a tame, quasi-homogeneous Markov chain on G, then there exists a constant C such that for all $o \in G$ and $n \in \mathbb{N}$ we have:*

$$\mathbb{P}\Big[d_X(ox_0, w_n^o x_0) \geqslant n/C\Big] \geqslant 1 - Ce^{-n/C}.$$

Recall that a Markov chain satisfying the conclusion of the above theorem is said to make *linear progress with exponential decay* in the hyperbolic space X.

In this section, we prove Theorem 5.1. The main technical result that we establish in order to do this is Proposition 5.6 below, the conclusion of which is the same as that of [GS21, Proposition 5.1]. Theorem 5.1 is a consequence of this by the following.

Theorem 5.2 ([GS21, §6]). *Let G be a group acting on a hyperbolic space X. Every tame Markov chain on G satisfying the conclusion of [GS21, Prop. 5.1] makes linear progress in X with exponential decay.*

Thus, in order to prove Theorem 5.1 it suffices to establish Proposition 5.6, because Proposition 4.1 shows that, under the assumptions of Theorem 5.1, the hypotheses of Proposition 5.6 are met. Before we can state Proposition 5.6, we need to introduce some notation.

5.1. **Notation and preliminary lemmas.** Fix a group G acting on a hyperbolic space X, and a basepoint $x_0 \in X$. We extend the notation used in Notation 2.6 with the following lemma, which states that, in the case of an axis of a WPD element, one can perturb closest-point projections to make them satisfy a stronger version of the Behrstock inequality.

Lemma 5.3 ([BBFS20, Thm 4.1]). *Let $g \in G$ be a loxodromic WPD. Writing $\gamma = E(g)$, there is a constant B and a g-equivariant map $\pi_\gamma : G \to \mathcal{P}(\gamma)$ with the following property, where $\mathcal{P}(\gamma)$ is the set of all subsets of γ. For all $x \in G$ and distinct translates $h\gamma \neq h'\gamma$,*

$$\text{if } d_X(\pi_{h\gamma}(x), \pi_{h\gamma}(h'\gamma)) > B, \text{ then } \pi_{h'\gamma}(x) = \pi_{h'\gamma}(h\gamma),$$

where we define $\pi_{k\gamma}(z) = k\pi_\gamma(k^{-1}z)$. Moreover, for all $x \in G$ the Hausdorff distance between $\pi_\gamma(x)$ and any X-projection of x to $\langle g \rangle$ is bounded by B.

In view of this lemma and with an abuse of notation, for $h, x, y \in G$ we will denote

$$d_{h\gamma}(x, y) = \text{diam}(\pi_{h\gamma}(x) \cup \pi_{h\gamma}(y)),$$

where $\pi_{h\gamma}$ is the equivariant map of Lemma 5.3.

The following definition captures the set of cosets where two given elements have far away projections.

Definition 5.4 ([GS21, Def. 3.11]). Let g be a loxodromic WPD element of a group G acting on a hyperbolic space. Given $x, y \in G$ and $T \geqslant 0$, we write

$$\mathcal{H}_T(x, y) = \{hE(g) : d_{hE(g)}(x, y) \geqslant T\}.$$

With g fixed, let $o, p, x, y \in G$. We define the following "distance formula" expression:

$$\sum_{\mathcal{H}_T(o,p)} [x, y] := \sum_{\substack{hE(g) \in \mathcal{H}_T(o,p) \\ \pi_{hE(g)}(x) \neq \pi_{hE(g)}(y)}} d_{hE(g)}(x, y).$$

Remark 5.5 ([GS21, Rem. 3.12]). Since projection distances satisfy the triangle inequality, so do these distance-formula expressions. That is, for all $o, p, x, y, z, \in G$ we have

$$\sum_{\mathcal{H}_T(o,p)} [x, z] \leqslant \sum_{\mathcal{H}_T(o,p)} [x, y] + \sum_{\mathcal{H}_T(o,p)} [y, z].$$

We can now state the main result of this section, the conclusion of which is the same as that of [GS21, Prop. 5.1]. As discussed, Theorem 5.1 follows from it, Theorem 5.2, and Proposition 4.1.

Proposition 5.6. *Let G be a group acting on a hyperbolic space X and fix a basepoint $x_0 \in X$. Suppose that the conclusion of Proposition 4.1 holds for some loxodromic WPD $g \in G$. There exist T_0, C' such that the following holds for each $T \geqslant T_0$. For all $o, p \in G$, $n \in \mathbb{N}$, and $t > 0$,*

$$\mathbb{P}\Big[\exists r \leqslant n : \sum_{\mathcal{H}_T(o,p)} [p, w_r^p] \geqslant t\Big] \leqslant 2e^{-t/C'}.$$

Before proceeding to the proof, we collect two more preliminary lemmas that are needed in order to choose the constant T_0.

Lemma 5.7 ([BBFS20, Thm 4.1], [BBFS20, Thm 3.3(G)]). *Fix a loxodromic WPD element $g \in G$, and let $o, p \in G$. Write $\gamma = E(g)$. For any sufficiently large T, the set $\mathcal{H}_T(o, p) \cup \{o\gamma, p\gamma\}$ is totally ordered with least element $o\gamma$ and greatest element $p\gamma$. The order is given by $h\gamma \prec h'\gamma$ if any one of the following equivalent conditions holds, where B is as in Lemma 5.3.*

- $d_{h\gamma}(o, h'\gamma) > B$.
- $\pi_{h'\gamma}(o) = \pi_{h'\gamma}(h\gamma)$.
- $d_{h'\gamma}(p, h\gamma) > B$.
- $\pi_{h\gamma}(p) = \pi_{h\gamma}(h'\gamma)$.

Lemma 5.8 ([GS21, Lem. 3.14]). *Let $g \in G$ be a loxodromic WPD and let $x_0 \in X$. For all sufficiently large T, we have the following for all $a, b \in G$:*

$$d_X(ax_0, bx_0) \geqslant \frac{1}{2} \sum_{\mathcal{H}_T(a,b)} [a, b].$$

5.2. **Proof of Proposition 5.6.** Fix an element $g \in G$ satisfying the conclusion of Proposition 4.1. Throughout the proof, we write $\gamma = E(g)$.

Constants. We fix the following quantities.

- B is given by Lemma 5.3.
- $T' \geqslant 10B$ is sufficiently large that the conclusions of Lemmas 5.7 and 5.8 hold.
- $L = L(T')$ is given by [GS21, Lem. 3.16].
- J is such that for any $p, h \in G$, consecutive points of both $(w_n^p x_0)$ and $(\pi_{h\gamma} w_n^p)$ have distance at most J. This exists by tameness, Definition 2.7.
- C and ϵ are as in Proposition 4.1. We can and will assume that $C > B$.
- $T_0 = \max\{T', 2(C + B + J + 1)\}$.
- $D = T_0 + LJ + L + C$.

Fix $o, p, q \in G$ and $n \in \mathbb{N}$. For ease of notation, given $t > 0$ we define

$$g(t) = \mathbb{P}\Big[\exists r \leqslant n : \sum_{\mathcal{H}_T(o,p)} [p, w_r^p] \geqslant t\Big], \quad \text{and} \quad f(t) = \mathbb{P}\Big[\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \geqslant t\Big].$$

Thus our goal is to find a constant C_2, independent of t, such that $g(t) \leqslant C_2 e^{-t/C_2}$. Fix $T \geqslant T_0$.

The main technical step in our proof of Proposition 5.6 is the following, which relates the probability functions $f(t)$ and $g(t)$.

Lemma 5.9. *For all $t \geqslant D$, we have $\epsilon g(t) \leqslant f(t - D) - f(t + D)$.*

Proof. For ease of notation, set $D' = LJ + L + 2C + 2B$. For $t \geqslant D$, let $\mathcal{A}^t$ denote the set of all $x \in G$ such that $\sum_{\mathcal{H}_T(o,p)} [p, x] \geqslant t - D'$ and $d_{h\gamma}(p, x) \geqslant 2C$ and $d_{h\gamma}(o, x) \geqslant 2C$, where $h\gamma \in \mathcal{H}_T(o, p)$ is the minimal coset contributing to the sum $\sum_{\mathcal{H}_T(o,p)} [p, x]$, with respect to the linear order from Lemma 5.7 for $\mathcal{H}_T(o, p)$. For $k \leqslant n$ and $x \in G$, let $\mathcal{A}_{k,x}^t$ denote the event

"*We have $w_k^p = x$ and $x \in \mathcal{A}^t$ and $\forall i < k \; : w_i^p \notin \mathcal{A}^t$.*"

That is, $\mathcal{A}_{k,x}^t$ is the event that $w_k^p = x$ is the first time that $w_k^p \in \mathcal{A}^t$.

Claim 3. *For any $k \leqslant n$, $x \in G$, and $t \geqslant D$, we have*

$$\mathbb{P}\left[\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \in [t-D, t+D] \,\Big|\, \mathcal{A}_{k,x}^t \right] \geqslant \mathbb{P}[d_{h_1\gamma}(x, w_{n-k}^x) \leqslant C].$$

Proof of Claim. We fix $k \leqslant n$, $x \in G$ and $t \geqslant D$ and assume that $\mathcal{A}_{k,x}^t$ holds. Let $h\gamma$ be the minimal coset, with respect to the linear order from Lemma 5.7 for the set $\mathcal{H}_T(o,p)$ that contributes to the sum $\sum_{\mathcal{H}_T(o,p)}[p, w_k^p]$.

Therefore

(2) $$\pi_{h\gamma}(p) \neq \pi_{h\gamma}(w_k^p) \quad \text{and} \quad \pi_{h'\gamma}(p) = \pi_{h'\gamma}(w_k^p) \quad \forall h'\gamma < h\gamma.$$

By the Behrstock inequality (Lemma 5.3), any point $y \in G$ with $d_{h\gamma}(p, y) > B$ has $\pi_{h''\gamma}(y) = \pi_{h''\gamma}(h\gamma) = \pi_{h''\gamma}(o)$ for all $h''\gamma > h\gamma$. Similarly, any point $y \in G$ with $d_{h\gamma}(o, y) > B$ has $\pi_{h'\gamma}(y) = \pi_{h'\gamma}(h\gamma) = \pi_{h'\gamma}(p)$ for all $h'\gamma < h\gamma$ (and in $\mathcal{H}_T(o,p)$). Hence, if $d_{h\gamma}(w_k^p, w_n^p) \leqslant C$, then $\sum_{\mathcal{H}_T(o,p)}[w_k^p, w_n^p] = d_{h\gamma}(w_k^p, w_n^p) \leqslant C$ as $w_k^p \in \mathcal{A}^t$. Hence if $\mathcal{A}_{k,x}^t$ holds and if $d_{h\gamma}(w_k^p, w_n^p) \leqslant C$ then

$$\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \geqslant \sum_{\mathcal{H}_T(o,p)} [p, w_k^p] - \sum_{\mathcal{H}_T(o,p)} [w_k^p, w_n^p] \geqslant t - D' - C \geqslant t - D,$$

by the triangle inequality (Remark 5.5).

Now, by the triangular inequality (Remark 5.5) we also have that if $d_{h\gamma}(w_k^p, w_n^p) \leqslant C$, then

$$\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \leqslant \sum_{\mathcal{H}_T(o,p)} [p, w_k^p] + C.$$

Therefore, in order to complete the proof of the claim we need to bound $\sum_{\mathcal{H}_T(o,p)}[p, w_k^p]$.

Subclaim 1. *We have $\sum_{\mathcal{H}_T(o,p)}[p, w_k^p] \leqslant t - D' + 4C + 2J$.*

Proof of Subclaim 1 Let $h_1\gamma$ be the second minimal coset (with respect to the linear order from Lemma 5.7 for $\mathcal{H}_T(o,p)$) contributing to the sum $\sum_{\mathcal{H}_T(o,p)}[p, w_k^p]$. Let $r < k$ be the last time such that $d_{h_1\gamma}(o, w_r^p) > 2C$. Then we have that $d_{h_1\gamma}(p, w_r^p) \geqslant d_{h_1\gamma}(o, p) - d_{h_1\gamma}(o, w_{r+1}^p) - d_{h_1\gamma}(w_{r+1}^p, w_r^p) \geqslant T - 2C - J > 2C$. Therefore, $\sum_{\mathcal{H}_T(o,p)}[p, w_r^p] < t - D'$ otherwise $w_r^p \in \mathcal{A}^t$ contradicting the minimality of k. By the strong Behrstock inequality, we have that $\pi_{h'''\gamma}(o) = \pi_{h'''\gamma}(w_r^p) = \pi_{h'''\gamma}(w_k^p)$ for all $h'''\gamma > h_1\gamma$. We also have that $d_{h\gamma}(o, w_r^p) > B$ and hence $\pi_{h'\gamma}(p) = \pi_{h'\gamma}(w_r^p) = \pi_{h'\gamma}(w_k^p)$ for all $h'\gamma < h\gamma$. As $d_{h_1\gamma}(o, w_r^p) > B$ and $h\gamma < h_1\gamma$ we have that $\pi_{h\gamma}(p) = \pi_{h\gamma}(w_r^p)$.

Therefore,

$$\begin{aligned}
\sum_{\mathcal{H}_T(o,p)} [p, w_k^p] &\leqslant \sum_{\mathcal{H}_T(o,p)} [p, w_r^p] + \sum_{\mathcal{H}_T(o,p)} [w_r^p, w_r^p] \\
&\leqslant t - D' + d_{h\gamma}(w_r^p, w_k^p) + d_{h_1\gamma}(w_r^p, w_k^p) \\
&\leqslant t - D' + d_{h\gamma}(w_r^p, w_k^p) + d_{h_1\gamma}(w_{r+1}^p, w_r^p) + d_{h_1\gamma}(w_{r+1}^p, w_k^p) \\
&\leqslant t - D' + d_{h\gamma}(p, w_k^p) + J + 2C.
\end{aligned}$$

By definition of $\mathcal{A}_{k,x}^t$, there is a last time $s < k$ such that $d_{h\gamma}(p, w_s^p) \leqslant 2C$. If $s + 1 = k$ then $d_{h\gamma}(p, w_k^p) \leqslant d_{h\gamma}(p, w_s^p) + J \leqslant 2C + J$ and this proves the subclaim.

If $s < k - 1$ then let $u \in (s+1, k)$ be the last time such that $\sum_{\mathcal{H}_T(o,p)}[p, w_u^p] < t - D'$. If $u + 1 = k$ then

$$\sum_{\mathcal{H}_T(o,p)} [p, w_k^p] \leqslant \sum_{\mathcal{H}_T(o,p)} [p, w_u^p] + \sum_{\mathcal{H}_T(o,p)} [w_u^p, w_{u+1}^p] \leqslant t - D' + J$$

and we have proved the subclaim. If not, then we consider w_i^p, for $i = k - 1$ or $i = u + 1$ (these might be equal but that's fine), we have that $k > i > u$. By definition of w_k^p, we must have $w_i^p \notin A^t$, but $i > \max\{s, u\}$ hence $\sum_{\mathcal{H}_T(o,p)}[p, w_i^p] \geqslant t - D'$ and $d_{h\gamma}(p, w_i^p) \geqslant 2C$. Therefore $d_{h\gamma}(o, w_i^p) < 2C$ for

$i = k - 1$ or $i = u + 1$. Now, by the fact that the projections of w_k^p and w_u^p coincide on cosets which are not $h\gamma$ and by the triangular inequality (Remark 5.5), we get

$$\sum_{\mathcal{H}_T(o,p)} [p, w_k] \leqslant \sum_{\mathcal{H}_T(o,p)} [p, w_u^p] + \sum_{\mathcal{H}_T(o,p)} [w_u^p, w_k^p]$$
$$\leqslant t - D' + d_{h\gamma}(w_u^p, w_k^p)$$
$$\leqslant t - D' + d_{h\gamma}(w_u^p, w_{u+1}^p) + d_{h\gamma}(w_{u+1}^p, o) + d_{h\gamma}(o, w_{k-1}^p) + d_{h\gamma}(w_{k-1}^p, w_k^p)$$
$$< t - D' + J + 2C + 2C + J$$

and this proves the subclaim. $\blacklozenge$

Therefore, by Subclaim 1, if $A_{k,x}^t$ holds then the event "$d_{h\gamma}(w_k^p, w_{n-k}^p) \leqslant C$" implies that

$$\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \leqslant \sum_{\mathcal{H}_T(o,p)} [p, w_k^p] + \sum_{\mathcal{H}_T(o,p)} [w_k^p, w_n^p] \leqslant t - D' + 4C + 2J + C \leqslant t + D.$$

Combining this with the result obtained earlier (above the subclaim), we get that

$$\mathbb{P}\Big[\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \in [t - D, t + D] | A_{k,x}^t \Big] = \mathbb{P}\Big[d_{h\gamma}(w_k^p, w_n^p) \leqslant C | A_{k,x}^t \Big]$$
$$= \mathbb{P}\Big[d_{h\gamma}(x, w_{n-k}^x) \leqslant C \Big]$$

by the strong Markov property [GS21, Lemma 2.2]. $\diamond$

Since the conclusion of Proposition 4.1 holds for g, with $\gamma = E(g)$, we have $\mathbb{P}[d_{h_1\gamma}(x, w_{n-k}^x) \leqslant C] > \epsilon$. Claim 3 above and the law of total probability therefore gives us

$$f(t - D) - f(t + D) = \mathbb{P}\Big[\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \in [t - D, t + D] \Big]$$
$$\geqslant \sum_{k \leqslant n} \sum_{x \in G} \Big(\mathbb{P}\Big[\sum_{\mathcal{H}_T(o,p)} [p, w_n^p] \in [t - D, t + D] \Big| \mathcal{A}_{k,x}^t \Big] \cdot \mathbb{P}\big[\mathcal{A}_{k,x}^t \big] \Big)$$
$$> \epsilon \sum_{k \leqslant n} \sum_{x \in G} \mathbb{P}\big[\mathcal{A}_{k,x}^t \big]$$
$$\geqslant \epsilon \mathbb{P}\big[\exists k \leqslant n : w_k^p \in \mathcal{A}^t \big].$$

To complete the proof of the lemma, we show that $g(t) \leqslant \mathbb{P}\big[\exists k \leqslant n : w_k^p \in \mathcal{A}^t \big]$. For this, suppose that the defining event of $g(t)$ holds. That is, suppose that there exists $r \leqslant n$ such that $S = \sum_{\mathcal{H}_T(o,p)} [p, w_r^p] \geqslant t$. Let $h\gamma, h'\gamma \in \mathcal{H}_T(o,p)$ be the minimal and second-minimal elements contributing to the sum S, respectively. Note that $\pi_{h''\gamma}(o) = \pi_{h''\gamma}(w_r^p)$ for every $h''\gamma > h\gamma$.

If $d_{h\gamma}(p, w_r^p) \geqslant 2C$ and $d_{h\gamma}(o, w_r^p) \geqslant 2C$, then $w_r^p \in \mathcal{A}^t$. In particular, there exists $k \leqslant n$ such that $w_k^p \in \mathcal{A}^t$ and we are done.

Otherwise, $d_{h\gamma}(p, w_r^p) < 2C$ or $d_{h\gamma}(o, w_r^p) < 2C$. First, say that $d_{h\gamma}(p, w_r^p) < 2C$. Since $h\gamma$ contributes to the sum S, we have $\pi_{h\gamma}(p) \neq \pi_{h\gamma}(w_r^p)$, so by Lemma 5.3 we must have $d_{h'\gamma}(o, w_r^p) \leqslant B$, and in particular $d_{h'\gamma}(p, w_r^p) > B$. Because $T > 2(C + B + J)$, there is a maximal $r' \leqslant r$ such that $d_{h'\gamma}(o, w_{r'}^p) > 2C > B$.

By the choice of J, we have $d_{h'\gamma}(p, w_{r'}^p) > 2C > B$. By the assumption that $d_{h\gamma}(p, w_r^p) < 2C$, we have $d_{h\gamma}(o, w_r^p) > B$. By applying Lemma 2.4 using these various estimates, we can use the definition of the linear order on $\mathcal{H}_T(o,p)$ to obtain:

- $\pi_{h^+\gamma}(w_r^p) = \pi_{h^+\gamma}(o)$ for all $h^+\gamma > h'\gamma$;
- $\pi_{h^-\gamma}(w_{r'}^p) = \pi_{h^-\gamma}(p)$ for all $h^-\gamma < h'\gamma$;
- $\pi_{h^+\gamma}(w_{r'}^p) = \pi_{h^+\gamma}(o)$ for all $h^+\gamma > h'\gamma$;
- $\pi_{h''\gamma}(w_r^p) = \pi_{h''\gamma}(p)$ for all $h''\gamma < h\gamma$.

In particular, $h'\gamma$ is the minimal element of $\mathcal{H}_T(o,p)$ that contributes to $\sum_{\mathcal{H}_T(o,p)}[p,w^p_{r'}]$, and we have $d_{h'\gamma}(p,w^p_{r'}) > 2C$ and $d_{h'\gamma}(o,w^p_{r'}) > 2C$. Moreover, Remark 5.5 lets us compute

$$\sum_{\mathcal{H}_T(o,p)} [p,w^p_{r'}] \geq \sum_{\mathcal{H}_T(o,p)} [p,w^p_r] - \sum_{\mathcal{H}_T(o,p)} [w^p_r,w^p_{r'}]$$
$$\geq t - \left(d_{h\gamma}(w^p_r,w^p_{r'}) + d_{h'\gamma}(w^p_r,w^p_{r'})\right)$$
$$\geq t - d_{h\gamma}(w^p_r,p) - \left(d_{h'\gamma}(w^p_r,o) + d_{h'\gamma}(o,w^p_{r'})\right)$$
$$\geq t - 2C - B - (B + J) \geq t - D'.$$

We have shown that $w^p_{r'} \in \mathcal{A}^t$. Thus, whenever the defining event of $g(t)$ holds, there is some $k \leq n$ such that $w^p_k \in \mathcal{A}^t$.

If, instead, it is the case that $d_{h\gamma}(o,w^p_r) < 2C$ then let r' be the last time that $d_{h\gamma}(o,w^p_{r'}) \geq 2C$. Then we have that $d_{h\gamma}(o,w^p_{r'}) \leq 2C+J$ by the choice of J, hence $d_{h\gamma}(p,w^p_{r'}) \geq d_{h\gamma}(p,o)-d_{h\gamma}(o,w^p_{r'}) \geq 2C$ by the choice of T. Further, by Remark 5.5 we have

$$\sum_{\mathcal{H}_T(o,p)} [p,w^p_{r'}] \geq \sum_{\mathcal{H}_T(o,p)} [p,w^p_r] - d_{h\gamma}(w^p_r,w^p_{r'}) \geq t - (6C + J) \geq t - D'.$$

Thus, $w^p_{r'} \in \mathcal{A}^t$. Hence, we have shown that the defining event of $g(t)$ implies that there exists a $k \leq n$ such that $w^p_k \in \mathcal{A}^t$. This completes the proof of the lemma. $\qquad\square$

We have shown that the probability $g(t)$ is bounded above by a probability $f(t - D)$. In order to prove Proposition 5.6 it suffices to show that the function f decays exponentially.

Proof of Proposition 5.6. By definition, $g(t) \geq f(t) \geq f(t + D)$ for all t. Therefore for $t \geq D$, Lemma 5.9 gives us $f(t-D) - f(t+D) \geq \epsilon g(t) \geq \epsilon f(t+D)$. Hence $f(t+D) \leq \frac{f(t-D)}{1+\epsilon}$. For simplicity, set $\epsilon_2 = \frac{1}{1+\epsilon} < 1$. Rephrasing this, for every $s \geq 0$ we have $f(s+2D) \leq \epsilon_2 f(s)$. If we write $s = 2qD+r$ with $q \in \mathbb{N}$ and $r \in [0,2D)$, then iterating this estimate yields

$$f(s) \leq \epsilon_2^q f(r) \leq \epsilon_2^q = \epsilon_2^{\frac{s-r}{2D}} \leq \frac{1}{\epsilon_2}\epsilon_2^{\frac{s}{2D}} = (1+\epsilon)\epsilon_2^{\frac{s}{2D}}.$$

Changing the base of the exponential completes the proof, with $C' = \frac{2D}{\log(1+\epsilon)}$.

$\qquad\square$

6. Descent of quasi-isometries in HHSs

This section and the next one concern hierarchically hyperbolic spaces and groups. For more detailed background on hierarchical hyperbolicity the reader is directed to any of several expository treatments; see e.g. [Sis19] for a detailed conceptual explanation of the definition or [CHK22, Part 2] for a technical overview. For present purposes, we refer the reader to [BHS19, Def. 1.1] for the definition of a hierarchically hyperbolic space (HHS).

We will sometimes require the following properties of an HHS $(X,\mathfrak{S})$:

- $(X,\mathfrak{S})$ is *normalised* if all maps $\pi_U : X \to \mathcal{C}U$ to the various hyperbolic spaces $\mathcal{C}U$ are uniformly coarsely surjective; see [BHS19, Rem. 1.3] for why this can always be assumed.
- $(X,\mathfrak{S})$ has the *bounded domain dichotomy* if there is a constant B such that $\mathrm{diam}(\mathcal{C}U) \leq B$ for any $U \in \mathfrak{S}$ with $\mathcal{C}U$ bounded. For this and the next definition, see also [ABD21, Sec. 3].
- $(X,\mathfrak{S})$ has *unbounded products* if it has the bounded domain dichotomy and for all $U \in \mathfrak{S}\setminus\{S\}$ such that $\mathcal{C}U$ is unbounded, there exists $V \in \mathfrak{S}$ such that $U \perp V$ and $\mathcal{C}V$ is unbounded. (Recall that S denotes the unique $\sqsubseteq$–maximal element of $\mathfrak{S}$.)

Note that every HHG satisfies the bounded domain dichotomy, as there are finitely many isometry classes of domains. Other hierarchical hyperbolicity notions will be used incidentally and we will refer to the relevant literature as we go.

Two key notions for us are *standard product regions* and *factored spaces*. The standard product region P_U associated to a domain U is a subspace of the HHS under consideration which is naturally quasi-isometric to a product $F_U \times E_U$, where moving in the F_U (resp. E_U) factor only changes projections to domains nested into (resp. orthogonal to) U; see [BHS19, Sec. 5] and [CHK22, Sec.

15]. Similarly, there are product regions associated to collections of pairwise orthogonal domains (which is the coarse intersection of the various standard product regions).

Regarding factored spaces, the idea is that starting with an HHS, we can cone off standard product regions and obtain another HHS, where the associated index set is obtained from the original index set by removing all indices "relevant" for the product regions we coned-off (and the rest of the structure, in particular the hyperbolic spaces, remains the same). In particular, the new HHS is somewhat "simpler". We now make this more precise.

Let $(X, d, \mathfrak{S})$ be an HHS and let $\mathcal{V} \subset \mathfrak{S}$ be downwards-closed in the nesting poset $(\mathfrak{S}, \sqsubseteq)$. In [BHS17a], a metric $d^{\mathcal{V}} \leqslant d$ is constructed on X such that $(X, d^{\mathcal{V}}, \mathfrak{S} \smallsetminus \mathcal{V})$ is an HHS. We call $d^{\mathcal{V}}$ the *factored metric* and $(X, d^{\mathcal{V}})$ the *factored space* of $(d, \mathcal{V})$. HHS automorphisms descend to factored spaces [CHK22, Proposition 19.1].

Definition 6.1. Given an HHS $(X, \mathfrak{S})$, let $\Gamma_{\mathfrak{S}}$ be the graph with a vertex for each element of $\mathfrak{S}$ and an edge joining U to V whenever $U \perp V$ and both $\mathcal{C}U$ and $\mathcal{C}V$ are unbounded.

Remark 6.2 (Arranging unbounded products). If $(X, \mathfrak{S})$ has the bounded domain dichotomy, then $\Gamma_{\mathfrak{S}}$ has an edge whenever $\mathcal{C}U$ and $\mathcal{C}V$ have diameter more than B and $U \perp V$. We mainly use the bounded domain dichotomy indirectly, via the construction in [ABD21, Sec. 3]: if $(X, \mathfrak{S}_0)$ is an HHS with the bounded domain dichotomy, then there is an HHS structure $(X, \mathfrak{S})$ with *unbounded products.*

The following proposition says roughly that quasi-isometries of HHSs descend to quasi-isometries of the factored spaces where we cone-off the product regions with the maximal number of factors. A clique in a graph is said to be *largest* if its cardinality agrees with the clique number of the graph.

Proposition 6.3 ([BHS21, Cor. 6.3])**.** *Let $(X, d, \mathfrak{S})$ be an HHS with the bounded domain dichotomy, and let $\mathcal{V}$ be the downwards-closure of the union of all largest cliques of $\Gamma_{\mathfrak{S}}$. If there is some D such that for each $\sqsubseteq$–maximal $U \in \mathcal{V}$, every pair of points in F_U lies on some biinfinite D–quasigeodesic of F_U, then every quasiisometry of (X, d) is a quasiisometry of $(X, d^{\mathcal{V}})$.*

After applying the above proposition, we have an HHS $(X, d^{\mathcal{V}}, \mathfrak{S} \smallsetminus \mathcal{V})$. Since $\mathcal{V}$ contained all $U \in \mathfrak{S}$ belonging to a largest clique in $\Gamma_{\mathfrak{S}}$, the maximum size of a clique in $\Gamma_{\mathfrak{S} \smallsetminus \mathcal{V}}$ is strictly smaller than in $\Gamma_{\mathfrak{S}}$, while the bounded domain dichotomy persists since the cone-off construction does not modify the $\mathcal{C}U$.

Applying the proposition repeatedly, we therefore obtain HHSs where the maximal number of factors of a product regions decreases until there are no products with at least two factors left. The space obtained in this way is an HHS $\widehat{X}$ that is hyperbolic in view of [BHS21, Cor. 2.16]. Moreover, quasi-isometries of the original HHS descend to this HHS.

However, this hyperbolic space is not necessarily the maximal hyperbolic space associated to the original HHS. This is explained more concretely in the following remark.

Remark 6.4. Let $S \in \mathfrak{S}$ be the maximal element. By [BHS17a, Cor. 2.9], the maximal hyperbolic space $\mathcal{C}S$ is naturally quasi-isometric to the factored space $(X, d^{\mathfrak{S}-\{S\}})$. This is in general not naturally quasi-isometric to $\widehat{X}$, as the HHS structure of the latter contains more domains than just S.

To illustrate this, skipping all technical details but hopefully still conveying the picture, consider the group

$$G = (\mathbb{Z}^2 * \mathbb{Z}) \times \mathbb{Z}.$$

Notice that G has a splitting with two vertex groups, a copy of $\mathbb{Z}^3$ and a copy of $\mathbb{Z}$, and two edge groups with vertex group $\mathbb{Z}$, one edge being a loop while the other connects the two vertices. The first step of the procedure cones-off quasiflats of rank 3, resulting in a space quasi-isometric to the Bass-Serre tree T for the aforementioned splitting; in particular this is already a hyperbolic space. The action on T is not acylindrical, so T is not the top-level hyperbolic space for any HHG structure on G.

To illustrate this in a different way, we describe a factored space on the mapping class group that, whilst not giving exactly $\widehat{X}$, gives the idea of how it can differ from $\mathcal{C}S$, where S is a closed connected oriented surface of genus at least 3. Namely, let $\mathcal{U}$ be the set of subsurfaces that are neither S nor the complement of a curve. Then $\Gamma_{\mathfrak{S} \smallsetminus \mathcal{U}}$ has no edges, so $(X, d^{\mathcal{U}})$ is hyperbolic. But it is not naturally quasiisometric to $\mathcal{C}S$, and it has loxodromic isometries that are not pseudo-Anosovs of S, namely

partial pseudo-Anosovs supported on the complement of a single curve. The iterative construction of $\widehat{X}$ gives an even "bigger" space than $(X, d^{\mathcal{U}})$ as above, meaning that $\widehat{X}$ has an HHS structure containing more domains than S and complements of single curves. For instance, complements of four-holed spheres are also part of the HHS structure for $\widehat{X}$ (four-holed spheres get "removed" at the first iteration of the construction, so after then their complements are never involved in quasiflats of dimension at least 2).

More abstractly, Figure 2 shows how applying the iterative procedure to $\mathfrak{S}$ can result in $\widehat{X}$ having a nontrivial HHS structure.
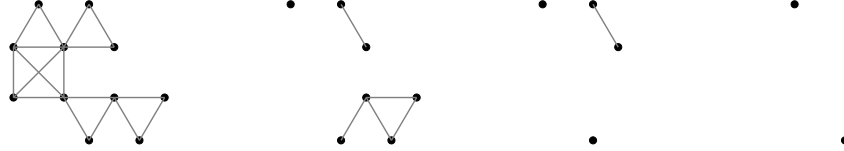

FIGURE 2. If the leftmost graph is part of $\Gamma_{\mathfrak{S}}$ (with non-adjacent vertices transverse), then it yields the second graph in $\Gamma_{\mathfrak{S} \smallsetminus \mathcal{V}}$, which is obtained by applying Proposition 6.3. Repeating yields the third graph, and we terminate with the fourth graph, which has vertices representing non-maximal unbounded domains. Hence $\widehat{X}$ is not quasiisometric to $\mathcal{C}S$ in this case.

We record the iterative construction described above here:

Lemma 6.5. *Let $(X, d, \mathfrak{S})$ be an HHS with the bounded domain dichotomy and the following property: if F_U is unbounded then every pair of points of F_U lies on a uniform biinfinite quasigeodesic of F_U. There is a downwards-closed set $\mathcal{U} \subset \mathfrak{S} - \{S\}$ such that $(X, d^{\mathcal{U}})$ is hyperbolic and every quasiisometry of (X, d) is a quasiisometry of $(X, d^{\mathcal{U}})$.*

Note that, in view of the HHS structure on $(X, d^{\mathcal{U}})$, the map $\pi_S : X \to \mathcal{C}S$ is still coarsely Lipschitz and coarsely surjective when X is endowed with the metric $d^{\mathcal{U}}$.

Proof. Let ω be the clique number of $\Gamma_{\mathfrak{S}}$. Let $\mathfrak{S}^0 = \mathfrak{S}$, $d^0 = d$. Given $\mathfrak{S}^i$ and d^i, let $\mathcal{U}^i$ be the downwards-closure of the union of all maximal cliques of $\Gamma_{\mathfrak{S}^i}$. Let d^{i+1} be the factored metric of $(d^i, \mathcal{U}^i)$, and let $\mathfrak{S}^{i+1} = \mathfrak{S}^i \smallsetminus \mathcal{U}^i$. Because $(X, d^i, \mathfrak{S}^i)$ is an HHS with the bounded domain dichotomy, so is $(X, d^{i+1}, \mathfrak{S}^{i+1})$. Moreover, the clique number of $\Gamma_{\mathfrak{S}^i}$ is at most $\omega - i$. Thus there is some minimal $n < \omega$ such that $\Gamma_{\mathfrak{S}^n}$ has no nontrivial cliques. Let $\mathcal{U} = \mathfrak{S} \smallsetminus \mathfrak{S}^n$. By Proposition 6.3, every quasiisometry of (X, d) is a quasiisometry of every (X, d^i), in particular of $(X, d^{\mathcal{U}}) = (X, d^n)$. $\square$

From now on, $\mathcal{U}$, $\mathcal{U}^i$, and $\mathfrak{S}^i$ denote the subsets of $\mathfrak{S}$ constructed in proving Lemma 6.5.

Remark 6.6. The proof of Lemma 6.5 did not need every $U \in \mathfrak{S}$ to have the property that each pair of points in F_U lies on a uniform quasigeodesic: it only needs this to hold for those U that are $\sqsubseteq$–maximal in some $\mathcal{U}^i$. Moreover, it does not need to hold for the full space F_U obtained by considering $U \in \mathfrak{S}$, only the F_U^i obtained by considering $U \in \mathfrak{S}^i$.

Lemma 6.7. *Let $(X, \mathfrak{S})$ be an HHS with unbounded products. If for all unbounded $U \in \mathfrak{S}$, every $x \in \mathcal{C}U$ lies on a uniform biinfinite quasigeodesic, then every quasiisometry of (X, d) is a quasiisometry of $(X, d^{\mathcal{U}})$.*

Proof. As noted in Remark 6.6, it suffices to show that if U is $\sqsubseteq$–maximal in some $\mathcal{U}^i$, then each pair of points of F_U^i lies on some uniform quasigeodesic of F_U^i, where F_U^i is the space of nested partial tuples in $(X, d^i, \mathfrak{S}^i)$.

Let $U \in \mathcal{U}^i$ be $\sqsubseteq$–maximal, and let $x_1, x_2 \in F_U^i$. Consider the set $\beta = H_{\theta_0}(x_1, x_2)$. By the definition of $\mathcal{U}^i$, no two unbounded elements of $\mathfrak{S}_U^i = \mathfrak{S}^i \cap \mathfrak{S}_U$ are orthogonal, so $\beta \subset F_U^i$ is a uniform quasigeodesic. Moreover, this implies that F_U^i is hyperbolic, for instance by [Bow13, Thm 2.1].

By [BHS19, Lem. 2.8], the set $\mathrm{Rel}_{100E}(x_1, x_2) \cap \mathfrak{S}_U^i$ is totally ordered. Let V_1 and V_2 be the minimal and maximal elements, respectively. By assumption, there are uniform biinfinite quasigeodesics $\alpha_j \subset$

CV_j through $\pi_{V_j}(x_j)$, respectively. Although the α_j might not be coarsely connected subsets of the hyperbolic space F_U^i, their convex hulls $\bar{\alpha}_j$ therein are uniform biinfinite quasigeodesics. By the distance formula [BHS19, Thm 4.5] and the choice of the V_j, the quasigeodesic $\bar{\alpha}_j$ passes uniformly close to $x_j \in F_U^i$. We can therefore take subrays $\bar{\alpha}_j'$ such that $\bar{\alpha}_1' \cup \beta \cup \bar{\alpha}_2'$ is a uniform biinfinite quasigeodesic in F_U^i through x_1 and x_2. $\qquad\square$

The hyperbolic space $(X, d^{\mathcal{U}})$ to which quasi-isometries descend is set-theoretically unchanged from X, so still admits a map $\pi_S : X \to \mathcal{C}S$, where $S \in \mathfrak{S}$ is the maximal element. Moreover, as explained in the factored space construction in [BHS17a], this map is coarsely lipschitz. However, it need not be a quasi-isometry, as explained in Remark 6.4. We will need to further cone off $(X, d^{\mathcal{U}})$ to get a quasi-isometry to $\mathcal{C}S$.

We now informally explain how to do this. We have to identify the (quasiconvex) subspaces that need to be coned off, which from the construction, come from product regions (with at least two factors) of the original HHS: they are isolated vertices as in Figure 2. That is, in the original HHS there is a subspace of the form $U \times V$, and only one of the factors, say V, gets coned off when constructing $(X, d^{\mathcal{U}})$. To recognise those subspaces, we look at preimages of points in $(X, d^{\mathcal{U}})$ and notice that if two points lie in the remnant of a product region, their preimages have unbounded coarse intersection. In the notation above, preimages are of the form $\{u\} \times V$, and any two such subsets lie at finite Hausdorff distance (we stated bounded coarse intersection because we will take preimages of balls, not points, and we emphasise that the Hausdorff distance is finite but depends on the subsets in question). In fact, this characterises pairs of points in remnants of product regions, once this is properly quantified. We make this precise after setting notation.

If $(X, \mathfrak{S})$ is an HHS with K–unbounded products, then we can and shall assume that $K \geqslant \max\{1, K_0, K_1, K_2, K_3\}$, where the K_i are the constants appearing in the statement of [BHS21, Lem. 1.20]. Recall the function θ_u from the uniqueness axiom of HHSs: if $(X, d, \mathfrak{S})$ is an HHS and $d(x, y) \geqslant \theta_u(r)$, then there is some $U \in \mathfrak{S}$ with $d_U(x, y) \geqslant r$. Also recall that every element $V \in \mathfrak{S} \smallsetminus \{S\}$ has an orthogonal container $V^\perp$. Finally, recall (see [BHS19, Lem. 6.2]) that for each HHS $(X, \mathfrak{S})$ there is a constant θ_0 such that for any $A \subset X$, the θ_0–hull $H_{\theta_0}A$, defined in [BHS19, Sec. 6], is hierarchically quasiconvex in the sense of [BHS19, Def. 5.1] (meaning roughly that its projections in all hyperbolic spaces are quasiconvex, and it is coarsely maximal with those projections). We will also use the notion of a *gate map* $\mathfrak{g}_H$ to a hierarchically quasiconvex set H [BHS19, Def. 5.4], defined by assembling all closest-point projections in the associated hyperbolic spaces.

Proposition 6.8. *Suppose that $(X, d, \mathfrak{S})$ is an HHS with K–unbounded products. There is a constant C such that the following holds. If $U \notin \mathcal{U} \cup \{S\}$ has F_U unbounded, then every $x, y \in P_U$ lie in unbounded subsets $I_x, I_y \subset (X, d)$, respectively, that are at finite Hausdorff distance and have diameter at most C in $(X, d^{\mathcal{U}})$.*

Conversely, for each C there exists D such that the following holds. Suppose that $x, y \in X$ lie in unbounded subsets $I_x, I_y \subset (X, d)$ that are at finite Hausdorff-distance and have diameter at most C in $(X, d^{\mathcal{U}})$. Either $d(x, y) \leqslant 2C + \theta_u(K)$, or there is some $V \in \mathcal{U}$ such that x and y are both D–close to P_V in $(X, d^{\mathcal{U}})$.

Proof. For the first statement, define $I_x = \{x\} \times E_U \subset P_U$ and $I_y = \{y\} \times E_U \subset P_U$. These are unbounded because $(X, \mathfrak{S})$ has unbounded products, and they are clearly at finite Hausdorff-distance. Moreover, the fact that $U \notin \mathcal{U}$ implies that every unbounded $W \perp U$ is in $\mathcal{U}$. There is some such W because $U \neq S$, so I_x and I_y have uniformly bounded diameter in $(X, d^{\mathcal{U}})$.

Now let us consider the converse statement, in which we start with subsets I_x and I_y containing x and y, respectively. Since I_x and I_y are at finite Hausdorff-distance, so are their projections to each $\mathcal{C}W$, and hence so are the convex hulls of these projections. This shows that the hierarchically quasiconvex hulls $H_{\theta_0}I_x$ and $H_{\theta_0}I_y$ are at finite Hausdorff-distance, and it follows that the gates $\mathfrak{g}_{H_{\theta_0}I_x}I_y$ and $\mathfrak{g}_{H_{\theta_0}I_y}I_x$ are unbounded. Let $I_x' = \mathfrak{g}_{H_{\theta_0}I_x}I_y$ and let $I_y' = \mathfrak{g}_{H_{\theta_0}I_y}I_x$.

Because I_x' is unbounded, the uniqueness axiom implies the existence of a pair of points $a, a' \in I_x'$ with $d(a, a') \leqslant 2\theta_u(K)$ and a domain $V \in \mathrm{Rel}_K(a, a')$ such that $\mathrm{Rel}_K(a, a') \subset \mathfrak{S}_V$. Consider $b = \mathfrak{g}_{H_{\theta_0}I_y}a$. If $d(a, b) \leqslant \theta_u(K)$, then

$$d^{\mathcal{U}}(x, y) \leqslant d^{\mathcal{U}}(x, a) + d^{\mathcal{U}}(a, b) + d^{\mathcal{U}}(b, y) \leqslant 2C + \theta_u(K).$$

Otherwise, the set $\mathrm{Rel}_K(a,b)$ is nonempty. According to [BHS21, Lem. 1.20(5)], we have $U \perp V$ for all $U \in \mathrm{Rel}_K(a,b)$. Now suppose that $W \in \mathfrak{S}$ has $V \subsetneq W$ or $V \pitchfork W$. Using the fact that $d_V(a,a') \geqslant K$, consistency and bounded geodesic image tell us that at least one of $\pi_W(a), \pi_W(a')$ is E–close to ρ_W^V. This shows that $d(a, P_V)$ is bounded by some uniform constant D' defined in terms of $d(a,a') \leqslant 2\theta_u(K)$ and the uniqueness function. Because of [BHS21, Lem 1.20], we can use a similar argument with slightly different constants to obtain a uniform bound D'' on $d(b, P_V)$ in terms of $d(b, \mathfrak{g}_{H_{\theta_0} I_y} a')$, which is itself bounded in terms of $d(a,a') \leqslant 2\theta_u(K)$. We can now compute

$$d^{\mathcal{U}}(x, P_V) \leqslant d^{\mathcal{U}}(x,a) + d^{\mathcal{U}}(a, P_V) \leqslant C + D',$$

and $d^{\mathcal{U}}(y, P_V)$ is bounded similarly. Since $V \in \mathrm{Rel}_K(a,a')$, it is unbounded, and the condition that $\mathrm{diam}_{(X,d^{\mathcal{U}})} I_x \leqslant C$ implies that $V \in \mathcal{U}$. $\qquad\square$

We can now prove the main theorem of this section, after collecting all required hypotheses in the following definition.

Definition 6.9. We call an HHS $(X, \mathfrak{S})$ *well behaved* if it is normalised, has unbounded products, and any one of the following holds for all $U \in \mathfrak{S}$:

- If F_U is unbounded then each pair of points of F_U lies on a uniform biinfinite quasigeodesic of F_U.
- If $\mathcal{C}U$ is unbounded, each $p \in \mathcal{C}U$ lies on a uniform biinfinite quasigeodesic of $\mathcal{C}U$.
- $\mathcal{C}U$ has uniformly cobounded isometry group.

Theorem 6.10. *Let $(X, \mathfrak{S})$ be a well-behaved HHS. Then every quasiisometry f of X induces a quasiisometry f' of $\mathcal{C}S$ such that $f'\pi_S$ and $\pi_S f$ coarsely agree.*

Proof. Let f be a quasiisometry of (X, d). Under the first assumption, Lemma 6.5 shows that f is a quasiisometry of $(X, d^{\mathcal{U}})$. Under the second assumption, this is given by Lemma 6.7, and the third assumption implies the second.

Let $\mathcal{V} = (\mathfrak{S} \smallsetminus \mathcal{U}) \smallsetminus \{S\}$. By the distance formula, the map π_S is a quasiisometry from the HHS $(X, (d^{\mathcal{U}})^{\mathcal{V}}, \{S\})$ to $\mathcal{C}S$ (here we use that the HHS is normalised, otherwise the map would only be a quasi-isometric embedding). It therefore suffices to show that f is a coarsely lipschitz map of $(X, (d^{\mathcal{U}})^{\mathcal{V}})$, for then the same will apply to a quasiinverse of f.

Let us write $d' = (d^{\mathcal{U}})^{\mathcal{V}}$. By the definition of d', we just have to check that if x and y both lie in some $F_U \subset (X, d^{\mathcal{U}}, \mathfrak{S} \smallsetminus \mathcal{U})$ with $U \in \mathcal{V}$ then $d'(f(x), f(y))$ is uniformly bounded. If $\mathrm{diam}_{(X,d^{\mathcal{U}})}(F_U) < \infty$, then $d^{\mathcal{U}}(x,y)$ is uniformly bounded by the bounded domain dichotomy, and hence so is $d'(f(x), f(y))$, because f is a quasiisometry of $(X, d^{\mathcal{U}})$ and $d' \leqslant d^{\mathcal{U}}$.

Otherwise, the forward direction of Proposition 6.8 provides a pair of unbounded subspaces $I_x, I_y \subset (X, d)$ that are at finite Hausdorff distance and lie in bounded $d^{\mathcal{U}}$–neighbourhoods of x and y, respectively. These properties are preserved by quasiisometries, so applying the reverse direction of Proposition 6.8 with fI_x and fI_y tells us that either $d(f(x), f(y))$ is uniformly bounded or there is a domain $V \in \mathfrak{S} \smallsetminus \{S\}$ such that $f(x)$ and $f(y)$ are uniformly close to P_V with respect to $d^{\mathcal{U}}$. In the latter case, the definition of d' gives a uniform bound on $d'(x,y)$. $\qquad\square$

In the case of hierarchically hyperbolic groups, we can remove the unbounded product assumption using results from [ABD21], at the expense of working in the slightly more general category of G–HHSes; see Definition 7.6.

Corollary 6.11. *Let $(G, \mathfrak{S})$ be an HHG and suppose that the action of $\mathrm{Stab}_G(U)$ on $\mathcal{C}U$ is cobounded for every $U \in \mathfrak{S}$, and that G is not wide. Then G admits a G–HHS structure $(G, \mathfrak{S}')$ such that the following holds, where $S \in \mathfrak{S}'$ is the unique $\sqsubseteq$–maximal element: $\mathcal{C}S$ is unbounded and every quasiisometry f of G induces a quasiisometry f' of $\mathcal{C}S$ such that $f'\pi_S$ and $\pi_S f$ coarsely agree.*

Proof. This is a consequence of Theorem 6.10 together with Lemma 7.11 below (which follows from [ABD21]), which can be used to remove the unbounded products hypothesis, and Lemma 7.9, which relates non-wideness to unboundedness of $\mathcal{C}S$. $\qquad\square$

The conclusion of Theorem 6.10 is not true without some extra hypothesis beyond unbounded products, as illustrated by the following example.

Example 6.12. Let X be the universal cover of the plane minus an open square, which is a CAT(0) cube complex whose hyperclosure (see [HS20]) $\mathfrak{S}$ is a factor system in the sense of [BHS17b], so that $(X, \mathfrak{S})$ is an HHS with unbounded products. However, there are self-quasiisometries of X, modelled on the logarithmic spiral map, that do not induce quasiisometries of the contact graph of X. See the discussion in [Cas16, §3].

7. Acylindrical hyperbolicity and quasi-isometries

Throughout this section, we still work with a fixed finitely generated group G and word metric d_G. We start with a stronger version of Assumption 2.1, which strengthens *partial detectability* to full Morse detectability and the WPD property to acylindricity.

Definition 7.1 (Geometrically faithful pair). Let G act on a δ–hyperbolic geodesic space (X, d_X). Fix a basepoint $x_0 \in X$ and let $\rho : G \to X$ be the corresponding orbit map. Assume that the action is cobounded, i.e. $X \subseteq N_\delta(\rho(G))$. Suppose that the following conditions all hold:

(1) (**Acylindricity.**) G acts on X acylindrically and non-elementarily.
(2) (**QIs descend.**) For each ν there exists some λ such that for each ν–quasi-isometry $\phi : G \to G$ there exists a λ–quasi-isometry $\bar\phi : X \to X$ such that $d_X(\phi(g)(x), \bar\phi(gx)) \leqslant \lambda$ for all $g \in G$ and $x \in X$.
(3) (**Morse detectable.**) For every Morse gauge M there exists λ such that if $\gamma \subset G$ is an M–Morse geodesic, then $\rho \circ \gamma$ is a λ–quasigeodesic. Conversely, for each λ there is a Morse gauge M such that if $\gamma \subset G$ is a geodesic such that $\rho \circ \gamma$ is a λ–quasigeodesic, then γ is M–Morse.

Then the pair (G, X) is *geometrically faithful*. We refer to the constant δ, the map $\nu \mapsto \lambda$ implicit in condition (2), and the maps $\lambda \mapsto M$ and $M \mapsto \lambda$ implicit in condition (3) as the *parameters* of the geometrically faithful pair (G, X).

A related coarse version (targeted at our applications) is as follows:

Definition 7.2 (Quasi-geometrically faithful pair). Let (X, d_X) be a δ–hyperbolic space and let $\rho : (G, d_G) \to (X, d_X)$ be a δ–coarsely surjective δ–coarsely lipschitz map. Suppose that the following all hold:

(1) (**Morse ray.**) G contains a Morse geodesic ray.
(2) (**QIs descend.**) For each ν there exists some λ such that for each ν–quasi-isometry $\phi : G \to G$ there exists a λ–quasi-isometry $\bar\phi : X \to X$ such that $d_X(\rho(\phi(g)), \bar\phi(\rho(g))) \leqslant \lambda$ for all $g \in G$ and $x \in X$.
(3) (**Morse detectable.**) For every Morse gauge M there exists λ such that if $\gamma \subset G$ is an M–Morse geodesic, then $\rho \circ \gamma$ is a λ–quasigeodesic. Conversely, for each λ there is a Morse gauge M such that if $\gamma \subset G$ is a geodesic such that $\rho \circ \gamma$ is a λ–quasigeodesic, then γ is M–Morse.
(4) (**Geometrically separated fibres.**) For all $r \geqslant 0$ there exists $R \geqslant 0$ such that for all $x, y \in X$ such that $d_X(x, y) \geqslant R$,
$$\mathrm{diam}\left(N_s^G(\rho^{-1}(N_r^X(x))) \cap \rho^{-1}(N_r^X(y))\right) < \infty$$
for all $s \geqslant 0$, where N^G and N^X respectively denote neighbourhoods in G and X.

Then the pair (G, X) is *quasi-geometrically faithful*. We refer to the constant δ, the map $\nu \mapsto \lambda$ implicit in condition (2), and the maps $\lambda \mapsto M$ and $M \mapsto \lambda$ implicit in condition (3) as the *parameters* of the quasi-geometrically faithful pair (G, X).

We relate the definitions, and show that Definition 7.2 is quasi-isometry invariant:

Lemma 7.3. *Let H be a finitely generated group such that there exists a quasi-isometry $\psi : H \to G$, and suppose that (G, X) is a quasi-geometrically faithful pair, with $\rho : G \to X$ the coarsely lipschitz map from Definition 7.2. Then (H, X) is a quasi-geometrically faithful pair, with map $\rho \circ \psi$ and parameters depending only on the parameters of (G, X) and the quasi-isometry constants of ψ.*

Also, if (G, X) is a geometrically faithful pair, then it is a quasi-geometrically faithful pair.

Proof. Let (G, X) be a quasi-geometrically faithful pair, with ρ as in the statement, and let $\psi : H \to G$ be a quasi-isometry. Hence $\rho\psi : H \to X$ is a coarsely surjective coarsely lipschitz map. Existence of a Morse ray is a quasi-isometry invariant property, so Definition 7.2.(1) holds for H.

Suppose that $\phi : H \to H$ is a quasi-isometry. Then by Definition 7.2.(2) we have a coarsely commutative diagram

$$
\begin{array}{ccccc}
H & \xrightarrow{\psi} & G & \xrightarrow{\rho} & X \\
{\scriptstyle\phi}\big\downarrow & & {\scriptstyle\phi'}\big\downarrow & & {\scriptstyle\bar\phi'}\big\downarrow \\
H & \xrightarrow{\psi} & G & \xrightarrow{\rho} & X
\end{array}
$$

(constants depending only on the parameters of (G, X) and the quasi-isometry constants for ϕ), where $\phi' = \psi\phi\hat\psi$ and $\hat\psi$ is a quasi-inverse of ψ. By construction, using $\rho\psi$ in place of ρ, we have verified Definition 7.2.(2) for (H, X).

Next, we check Definition 7.2.(3) for (H, X), using the map $\rho\psi$. This follows since the corresponding property holds for (G, X), and in fact the same property, with geodesics replaced by quasigeodesics, also holds in (G, X), using e.g. [RST22, Lem. 2.8], and we can move quasigeodesics between G and H using ψ.

Finally, we verify the geometric separation for fibers, i.e. that (H, X) satisfies Definition 7.2.(4). Let $s \geqslant 0$ and let $x, y \in X$. Let us fix points $a, b, a', b' \in H$ such that we have $d_H(a, a'), d_H(b, b') \leqslant s$ and $d_X(\rho\psi(a), x), d_X(\rho\psi(b), x) \leqslant r$ and $d_X(\rho\psi(a'), y), d_X(\rho\psi(b'), y) \leqslant r$. Then $d_G(\psi(a), \psi(a')), d_G(\psi(b), \psi(b'))$ are bounded in terms of s and ψ, and so Definition 7.2.(4), applied to (G, X), bounds $d_G(\psi(a'), \psi(b'))$ (the bound is allowed to depend on x, y, s), and we therefore get the required bound on $d_H(a', b')$.

Now we prove the second assertion, that geometrically faithful implies quasi-geometrically faithful. Suppose that (G, X) is a geometrically faithful pair, with orbit map ρ. Then Definition 7.2.(2),(3) are immediate from Definition 7.1.(2),(3). Let $g \in G$ be a loxodromic WPD element, so that $n \mapsto \rho(g^n)x_0 = \rho(g^n)$ is a quasi-isometric embedding $\mathbb{Z} \to X$, and hence $n \mapsto g^n$ is a Morse quasigeodesic $\mathbb{Z} \to G$, because of Definition 7.1.(3). Restricting to $\mathbb{N}$ gives a Morse ray in G, so Definition 7.2.(1) holds. Definition 7.2.(4) follows from Definition 7.1.(1) along with [Sis16, Lem. 3.3]. Indeed, the acylindricity assumption guarantees that G acts acylindrically along X in the sense of [Sis16, Def. 3.1] (viewing X as a subspace of itself), and the property of preimages of balls mentioned in Definition 7.2 coincides with geometric separation in the sense of [Sis16, Def. 2.1]. $\qquad\square$

The main technical result of this section is:

Proposition 7.4 (Acylindrical hyperbolicity from quasi-geometrically faithful pair)**.** *Let (G, X) be a quasi-geometrically faithful pair. Then there exists a hyperbolic geodesic space Y such that G acts on Y, there exists $g \in G$ acting on Y as a loxodromic element, and every loxodromic $g \in G$ is a WPD element for the G–action on Y. In particular, if G is nonelementary, then G is acylindrically hyperbolic.*

Before the proof, we note:

Corollary 7.5. *Let G be a nonelementary finitely generated group such that there is a hyperbolic space X making (G, X) a geometrically faithful pair, and let H be a finitely generated group quasi-isometric to G. Then H is acylindrically hyperbolic.*

Proof. By Lemma 7.3, (G, X), and hence (H, X), is a quasi-geometrically faithful pair, so applying Proposition 7.4 implies that H is acylindrically hyperbolic. $\qquad\square$

We now prove the proposition:

Proof of Proposition 7.4. Fix a quasi-geometrically faithful pair (G, X) with map ρ and parameters as in Definition 7.2.

Once we produce a hyperbolic space Y on which G acts with a loxodromic WPD element, then, under the additional assumption that G is nonelementary, acylindrical hyperbolicity follows from

[Osi16, Thm. 1.2]. In particular, the conclusion that every loxodromic $g \in G$ is WPD is stronger than necessary, but we included it in the statement since it introduces no extra complexity to the proof.

Quasi-action of G on X: Each $g \in G$ can be regarded as a left-multiplication isometry $g : G \to G$, so by Definition 7.2.(2), there is a constant λ such that for all g, we have a λ–quasi-isometry $A(g) : X \to X$ such that $d_X(A(g)(\rho(x)), \rho(gx)) \leqslant \lambda$ for all $x \in G$. Up to uniformly increasing λ, this implies that $d_X(A(gh)(y), A(g)(A(h)(y))) \leqslant \lambda$ for all $y \in X$ and $g, h \in G$. Here we have used coarse surjectivity of ρ.

Moreover, for any $y \in X$, we can choose $x \in G$ with $d_X(\rho(x), y) \leqslant \lambda$, and then $d_X(A(x^{-1})(\rho(x)), \rho(1)) \leqslant \lambda$, so since $A(x^{-1})$ is uniformly coarsely lipschitz, we have, up to uniformly enlarging λ, that $d_X(y, A(g)(\rho(1))) \leqslant \lambda$ for some $g \in G$. Thus far, we have produced a uniform λ such that $A : G \to X^X$ is a λ–cobounded λ–quasi-action of G on X by λ–quasi-isometries.

The space Y: Using any of the various "Milnor-Švarc for quasi-actions" statements in the literature, we now replace the quasi-action by an action. For instance, [Man05, Prop. 4.4] provides, up to uniform enlargement of δ and λ, a δ–hyperbolic graph Y, a λ–quasi-isometry $q : X \to Y$, and a homomorphism $C : G \to \mathrm{Isom}(Y)$ such that

- C is λ–cobounded, and
- $\sup_{x \in X} d_Y(C(g)(q(x)), q(A(g)(x))) \leqslant \lambda$ for all $g \in G$.

Since the action is cobounded, in particular the orbits are quasiconvex and from, e.g., [CCMT15, Prop. 3.2], there exists $g \in G$ acting on Y loxodromically, because Y is unbounded in view of Definition 7.2.(1),(3). Let $\tau : G \to Y$ be $\tau(g) = C(g)(q(\rho(1)))$, which λ–coarsely coincides with $q(A(g)(\rho(1)))$ and so with $q(\rho(g))$.

Checking WPD: We now verify that g is a WPD element. In our notation, this means we must show that for each $\epsilon > 0$, there exists $n \in \mathbb{Z}$ such that $|\mathcal{H}(n, \epsilon)| < \infty$, where

$$\mathcal{H}(n, \epsilon) = \{h \in G : d_Y(C(h)(\tau(1)), \tau(1)) < \epsilon, d_Y(C(h)(\tau(g^n)), \tau(g^n)) < \epsilon\}.$$

There exists r, depending only on ϵ and q, but not on n, such that $h \in \mathcal{H}(n, \epsilon)$ implies $d_X(\rho(1), A(h)(\rho(1))) \leqslant r$ and $d_X(\rho(g^n), A(h)(\rho(g^n))) \leqslant r$. So by the triangle inequality,

$$d_X(\rho(1), \rho(h)) \leqslant r + \lambda$$

and

$$d_X(\rho(g^n), \rho(hg^n)) \leqslant r + \lambda.$$

Definition 7.2.(4) provides an $R = R(r + \lambda)$ such that $d_X(\rho(1), \rho(g^n)) \geqslant R$ implies that the ρ–preimages of $N^X_{r+\lambda}(\rho(1))$ and $N^X_{r+\lambda}(\rho(g^n))$ are geometrically separated.

For $n \in \mathbb{N}$, let γ_n be a geodesic in G joining 1 to g^n. Since g is loxodromic, the composition $\tau\gamma_n$ is a quasi-isometric embedding with constants depending on λ, δ, and g but not on n. Hence the map $\rho\gamma_n$ is also a quasi-isometric embedding with constants independent of n, so for sufficiently large n we have $d_X(\rho(1), \rho(g^n)) > R$. Fix such an n.

If $h \in \mathcal{H}(n, \epsilon)$, then from the earlier discussion we have $h \in \rho^{-1}(N^X_{r+\lambda}(\rho(1)))$ and $hg^n \in \rho^{-1}(N^X_{r+\lambda}(\rho(g^n)))$. Thus $h \in N^G_s(\rho^{-1}(N^X_{r+\lambda}(\rho(g^n))))$, where $s = d_G(1, g^n)$ is independent of h. So Definition 7.2.(4) allows only finitely many possibilities for h, as required. $\square$

7.1. Application to hierarchically hyperbolic groups.
We now use Theorem 6.10 to apply Proposition 7.4 to hierarchically hyperbolic groups. We first recall the definition of an HHG (see e.g. [PS20] for the following modern formulation and [DHS20] for why it is equivalent to the original definition from [BHS19]):

Definition 7.6 (Hierarchically hyperbolic group, G–HHS)**.** Let G be a finitely generated group. Suppose that G, with the quasi-isometry class of word metrics associated to finite generating sets, admits an HHS structure $(G, \mathfrak{S})$. Suppose, moreover, that G acts cofinitely on $\mathfrak{S}$ preserving $\sqsubseteq, \perp, \pitchfork$, and that the following hold for all $U, V \in \mathfrak{S}$:

- for each $g \in G$, there is an isometry $g : \mathcal{C}U \to \mathcal{C}gU$ such that for all $g, h \in G$,
- the composition $\mathcal{C}U \xrightarrow{h} \mathcal{C}hU \xrightarrow{g} \mathcal{C}ghU$ agrees with the isometry gh, and
- $\pi_{gU}(gx) = g(\pi_U(x))$ for all $x \in G$ and
- $\rho^{gU}_{gV} = g(\rho^U_V)$ whenever $U \pitchfork V$ or $U \subsetneq V$.

Then $(G, \mathfrak{S})$ is a *hierarchically hyperbolic group*.

Following [ABR22, Def. 3.3], we say the pair $(G, \mathfrak{S})$ is a *G–HHS* if it has all of the above properties, except weakened in the following way: we *do not* require $\mathfrak{S}$ to contain only finitely many *G*–orbits, but we do require $(G, \mathfrak{S})$ to satisfy the bounded domain dichotomy.

Remark 7.7. The reason to work with *G*–HHSes is twofold: first, various results about HHGs hold in the slightly more general context of *G*-HHSes, and second, the *maximisation* procedure from [ABD21], used to produce an HHS with the bounded domain dichotomy into one with unbounded products, will in general convert an HHG structure on *G* into a *G*–HHS structure that does not necessarily have finitely many orbits. See [ABR22, Rem. 3.4].

Our main result is about acylindrically hyperbolic HHGs where the top-level hyperbolic space witnesses acylindrical hyperbolicity. We formulate this is follows:

Definition 7.8 (Irreducible). The *G*–HHS $(G, \mathfrak{S})$ is *irreducible* if *G* has unbounded orbits in $\mathcal{C}S$, where $S \in \mathfrak{S}$ is the unique $\sqsubseteq$–maximal element.

Irreducibility is related to acylindrical hyperbolicity, and various other properties by assembling various results in the literature:

Lemma 7.9 (HHG irreducibility criteria). *The following are equivalent, for a normalised G–HHS* $(G, \mathfrak{S})$ *with unbounded products and G nonelementary:*

(1) $(G, \mathfrak{S})$ *is irreducible.*
(2) *G is acylindrically hyperbolic.*
(3) *G has a Morse element.*
(4) *G is not* wide, *i.e. asymptotic cones of G have cut-points.*
(5) *G is not quasi-isometric to the product of two unbounded metric spaces.*
(6) *G has no nonempty finite orbit in* $\mathfrak{S} - \{S\}$.

Proof. Since the action of *G* on $\mathcal{C}S$ is always acylindrical [BHS17b, Thm. 14.3], and we are assuming *G* is nonelementary, (1) implies (2). Item (2) implies (3) and (4) by [Sis16], and both of the latter imply (5). If there exists $U \in \mathfrak{S} - \{S\}$ such that $G \cdot U$ is finite, then *G* virtually stabilises the standard product region P_U which, by the unbounded products assumption, is quasi-isometric to the product of two unbounded metric spaces (see [CHK22, Sec. 15] for the quasi-isometry computation). Hence (5) implies (6). Finally, using the normalisation assumption, [DHS17, Cor. 9.9] says that (6) implies (1). The latter implication also follows from [PS20, Thm. 5.1]. □

For what follows, in view of results from [ABD21] we can work with more general HHSs than well-behaved ones. Essentially, we can drop the unbounded products assumption:

Definition 7.10 (Reasonably behaved). The HHS $(X, \mathfrak{S})$ is *reasonably behaved* if it is normalised, it satisfies the *bounded domain dichotomy* from [ABD21, Def. 3.2] and for all $U \in \mathfrak{S}$, $\mathrm{Isom}(\mathcal{C}U)$ acts on $\mathcal{C}U$ coboundedly.

The next lemma is an application of results in [ABD21], used to avoid having to add the hypothesis of *unbounded products* to the final statements:

Lemma 7.11. *Let* $(G, \mathfrak{S})$ *be a reasonably behaved irreducible G–HHS. Then there is a G–HHS structure* $(G, \mathfrak{T})$ *that is well behaved and irreducible.*

Proof. Applying [ABD21, Cor. 3.8] yields a *G*–HHS structure $(G, \mathfrak{T})$ with unbounded products. Examining the proof of the aforementioned theorem, each element $U \in \mathfrak{T}$ is of one of three types.

First, *U* can be a *dummy domain* coming from the application of [ABD21, Thm. A.1], and in this case $\mathcal{C}U$ is a single point.

Second, *U* can correspond to an element of $\mathfrak{S}$, with stabiliser $\mathrm{Stab}_G(U)$ and associated hyperbolic space $\mathcal{C}U$ unchanged from the original HHG/HHS structure.

Third, we could have $U = S$. In this case, $\mathcal{C}S$ is replaced by a hyperbolic space $\mathcal{T}S$ which is an electrification of *G*, so there is a coarsely lipschitz map $G \to \mathcal{T}S$, and the construction makes the original map $\pi_S : G \to \mathcal{C}S$ factor as $G \to \mathcal{T}S \to \mathcal{C}S$, where the latter map is also coarsely lipschitz. So $\mathcal{T}S$ is unbounded if $\mathcal{C}S$ was. All of the maps involved are *G*–equivariant, and the action of *G* on

$\mathcal{T}S$ is cobounded since the latter is an equivariant electrification. We take $T = S$ and $\mathcal{C}T = \mathcal{T}S$ in the new HHG structure.

Thus $(G, \mathfrak{T})$ has unbounded products and is therefore a well-behaved HHG structure, with unbounded maximal hyperbolic space by Lemma 7.9. $\qquad\square$

Now we can state and prove the HHG part of Theorem 2 from the introduction, along with a slightly more general version:

Theorem 7.12 (AH from QI to G–HHS). *Let G be a nonelementary finitely generated group that is not quasi-isometric to the product of two unbounded spaces. Suppose that G admits a reasonably behaved G–HHS structure. Then any finitely-generated group H quasi-isometric to G is acylindrically hyperbolic.*

Proof. First suppose that $(G, \mathfrak{S})$ is a reasonably behaved G–HHS structure. Using Lemma 7.11 and Lemma 7.9, we can assume that $(G, \mathfrak{S})$ is well behaved and irreducible. Let $S \in \mathfrak{S}$ be the maximal element. By [BHS17b, Thm. 14.3], G acts on $\mathcal{C}S$ acylindrically. By Theorem 6.10, along with the well-behaved assumption, quasi-isometries of G descend through π_S to $\mathcal{C}S$ (since, by Definition 7.6, π_S is an orbit map), in the sense of Definition 7.1.(2). Finally, equivalence of Morseness of geodesics in G and quasi-geodesicity of their compositions with π_S is given by [ABD21, Cor. 6.2]. Hence $(G, \mathcal{C}S)$ is a geometrically faithful pair as in Definition 7.1. By Corollary 7.5, any H as in the statement is acylindrically hyperbolic. $\qquad\square$

Remark 7.13. The proof of Theorem 7.12 shows that any group G that admits a reasonably behaved irreducible HHG structure $(G, \mathfrak{S})$ admits an action on a hyperbolic space X such that (G, X) is a geometrically faithful pair, and in particular, (G, X) satisfies Assumption 2.1. If the original HHG structure had unbounded products, then we can take $X = \mathcal{C}S$. And, moreover, the preceding holds in the slightly more general setting of a G–HHS, not just an HHG.

A. Appendix: Quasi-isometries of HHGs with one-ended maximal hyperbolic spaces, by Jacob Russell

In this appendix, we prove a converse of Theorem 6.10 when $(X, \mathfrak{S})$ is a hierarchically hyperbolic group and the $\sqsubseteq$-maximal space $\mathcal{C}S$ is one-ended.

Theorem A.1. *Let $(G, \mathfrak{S})$ be a hierarchically hyperbolic group with unbounded products. If $\mathcal{C}S$ is one-ended, then every quasi-isometry f of $\mathcal{C}S$ induces a quasi-isometry F of G so that $f \circ \pi_S$ and $\pi_S \circ F$ coarsely agree. In particular, if $(G, \mathfrak{S})$ is well-behaved as described in Definition 6.9, then $f \to F$ induces a group isomorphism $\mathrm{QI}(\mathcal{C}S) \to \mathrm{QI}(G)$.*

Remark A.2. We prove Theorem A.1 under looser hypotheses than an HHG. All one needs is that the Cayley graph has an HHS structure on which the group acts by automorphisms. See Section A.6.

Our proof of Theorem A.1 uses the machinery of quasi-möbius maps on the Morse boundary developed by Charney, Cordes, and Murray [CCM19] and independently by Mousley and Russell in the case of HHGs [MR19]. The idea is that the quasi-isometry f of $\mathcal{C}S$ will induce a quasi-möbius map on the boundary $\partial\mathcal{C}S$. Using a result of Abbott, Behrstock, and Durham about unbounded products, this quasi-möbius map on $\partial\mathcal{C}S$ can be upgraded to a quasi-möbius map on the Morse boundary of G. The results of Charney, Cordes, and Murray (or Mousley and Russell in the HHG case), then say that this quasi-möbius map on the Morse boundary is induced by a quasi-isometry of the group.

One-endedness of $\mathcal{C}S$ comes into play in being able to upgrade the map on $\partial\mathcal{C}S$ to a map on the Morse boundary. Being one-ended allows us to adopt an idea of Rafi and Schleimer for the curve graph and the mapping class group [RS11].

Theorem A.1 can fail when $\mathcal{C}S$ is not one-ended. For example, both the the fundamental group of a non-geometric graph manifold and any non-relatively hyperbolic right angled Artin group have HHG structures with unbounded products and where the $\sqsubseteq$-maximal hyperbolic space is a quasi-tree. However, these examples are known to not all be quasi-isometric [BN12].

A.1. **The Morse Boundary.** We briefly recall the relevant properties of the Morse boundary that we shall need. We direct the reader to [Cor17] for a more detailed account.

Let X be a proper geodesics metric space and fix a basepoint $x_0 \in X$. As a set, the Morse boundary of X is the collection of all Morse geodesic rays based at x_0 up to asymptotic equivalence. We will topologize this boundary with the topology introduced by Cordes and denote this topological space by $\partial_* X_{x_0}$; see [Cor17] for details. The primary fact we need about this topology is that any quasi-isometry $f \colon X \to Y$ of proper geodesic metric spaces has a continuous extension to a homeomorphism $\partial f \colon \partial_* X_{x_0} \to \partial_* Y_{f(x_0)}$. This shows that the boundary is not affected by the choice of basepoint. Moreover, there is a well defined Morse boundary for a finitely generated group G by identifying it with the Morse boundary of any proper geodesic space on which it acts geometrically—usually its Cayley graph. We denote the Morse boundary of G by $\partial_* G$ and will always assume it is identified with the Morse boundary of some finitely generated Cayley graph. The Morse boundary is visual in the sense that for any two points $p, q \in \partial_* X_{x_0}$, there is some bi-infinite Morse geodesic between p and q (the specific Morse gauge will depend on p and q).

A.2. **Cross-ratio in hyperbolic spaces.** Let X be a (not necessarily proper) δ-hyperbolic space. Every pair of distinct points $a, b \in \partial X$ is joined by a bi-infinite $(1, 20\delta)$-quasi-geodesic. A point $x \in X$ is a K-*center* for the triple $(a, b, c) \in (\partial X)^3$ if x is within K of any $(1, 20\delta)$-quasi-geodesic between any two of a, b, c. There exists K_δ, so that for every triple $(a, b, c) \in (\partial X)^3$ the set of K_δ-centers is non-empty provided a, b, c are all distinct. Let $m(a, b, c)$ be the set of K_δ-centers for (a, b, c). There exists $D = D(\delta)$ so that $\operatorname{diam}(m(a, b, c)) \leqslant D$ whenever $m(a, b, c)$ is non-empty.

For distinct points $a, b, c, d \in \partial X$ define the *cross-ratio* $[a, b, c, d]$ to be

$$[a, b, c, d] := \operatorname{diam}(m(a, b, c) \cup m(a, d, c)).$$

This cross ratio is an additive error away from the absolute value of the cross ratio defined by Paulin; see [Pau96, Lemma 4.2] for the proper case and [MR19, Proposition 4] for the non-proper case.

The next lemma is proved by Paulin when X is proper [Pau96]. The same proof works in the non-proper case if you replace geodesics with $(1, 20\delta)$-quasi-geodesics.

Lemma A.3. *Let X and Y be δ-hyperbolic spaces where ∂X has at least 4 points. If $f \colon X \to Y$ is a (λ, ϵ)-quasi-isometric embedding, then there exist $\lambda' \geqslant 1$ and $\epsilon' \geqslant 0$ determined by λ, ϵ, and δ so that*

$$[\partial f(a), \partial f(b), \partial f(c), \partial f(d)] \leqslant \lambda'[a, b, c, d] + \epsilon'$$

for all $a, b, c, d \in \partial X$.

A.3. **Cross-ratio on the Morse boundary.** Let G be a finitely generated group and X a finitely generated Cayley graph. Let $\partial_* G$ denote the Morse boundary of a group identified with $\partial_* X$. For $k \in \mathbb{N}$, let $\partial_*^{(k,M)} G$ denote k-tuples of distinct elements of the Morse boundary so that every pair of points in the k-tuple are joined by a bi-infinite M-Morse geodesic in X (this set is G-invariant). The next set of definitions are from [CCM19] or [MR19].

Let H be a second finitely generated group. A map $h \colon \partial_* G \to \partial_* H$ is *2-stable* if for each Morse M there is a Morse gauge M' so that $h(\partial_*^{(2,M)} G) \subseteq \partial_*^{(2,M')} H$.

For each triple $(a, b, c) \in \partial_*^{(3,M)} G$ a point $x \in G$ is a K-*center* for (a, b, c) if x is within K of all three sides of any M-Morse ideal triangle with endpoints a, b, c. For each M, there is number K_M so that for any $(a, b, c) \in \partial_*^{(3,M)} G$ the set of K_M-centers is non-empty. Moreover, there exist $D = D(M)$ so that for each $(a, b, c) \in \partial_*^{(3,M)} G$, the set of K_M-centers, $m(a, b, c)$, has diameter at most D. For any tuple $(a, b, c, d) \in \partial_*^{(4,M)} G$, the M-*cross ratio* is

$$[a, b, c, d]_M := \operatorname{diam}(m(a, b, c) \cup m(a, d, c)).$$

A 2-stable map $h \colon \partial_* G \to \partial_* H$ is *quasi-möbius* if for every pair of Morse gauges M and M' with $h(\partial_*^{(2,M)} G) \subseteq \partial_*^{(2,M')} H$ there exists an increasing function ψ so that

$$[h(a), h(b), h(c), h(d)]_{M'} \leqslant \psi([a, b, c, d]_M).$$

Charney, Cordes, and Murray established that quasi-möbius maps on the Morse boundary characterize quasi-isometries of the group.

Theorem A.4 ([CCM19, Theorems 3.7 and 4.5])**.** *Let G and H be finitely generated groups so that $\partial_* G$ contains at least 3 points.*

(1) *If $F\colon G \to H$ is a quasi-isometry, then the induced map $\partial F\colon \partial_* G \to \partial_* H$ is a quasi-möbius homeomorphism with quasi-möbius inverse.*

(2) *If $h\colon \partial_* G \to \partial_* H$ is a quasi-möbius homeomorphism with quasi-möbius inverse, then there exist a quasi-isometry $F\colon G \to H$ so that $\partial F = h$.*

A.4. **The Morse boundary and unbounded products.** For the remainder of the appendix, we fix a finitely generated group G and a finitely generated Cayley graph, X, of G. We assume X has an HHS structure $\mathfrak{S}$ with unbounded products and let $S \in \mathfrak{S}$ be the $\sqsubseteq$-maximal domain. As described in [ABD21, Section 3], we can assume that $\mathcal{C}S$ is a graph obtained from X by adding additional edges between the vertices. In this case, $\pi_S\colon X \to \mathcal{C}S$ is taken to be the inclusion map.

Abbott, Behrstock, and Durham showed that in the presence of unbounded products, Morse geodesic are characterized by projecting to (parameterized) quasi-geodesics in $\mathcal{C}S$.

Theorem A.5 ([ABD21, Corollary 6.2])**.** *Let $(X, \mathfrak{S})$ be an HHS with unbounded product and $S \in \mathfrak{S}$ be the $\sqsubseteq$-maximal domain. Let γ be a geodesic in X.*

(1) *If $\pi_S \circ \gamma$ is a parameterized (λ, λ)-quasi-geodesic in $\mathcal{C}S$, then γ is M-Morse for some Morse gauge M determined by λ and $\mathfrak{S}$.*

(2) *If γ is M-Morse, then there exists $\lambda = \lambda(M, \mathfrak{S}) \geqslant 1$ so that $\pi_S \circ \gamma$ is a parameterized (λ, λ)-quasi-geodesic in $\mathcal{C}S$.*

Theorem A.5 implies there is a continuous injection $\partial \pi_S\colon \partial_* X \to \partial \mathcal{C}S$, which is a continuous extension of $\pi_S\colon X \to \mathcal{C}S$, see [Rus21, Lemma A.6].

A.5. **Downward relative projections and coboundedness.** For each $W \in \mathfrak{S} - \{S\}$ define $\rho_W^S\colon \mathcal{C}S^{(0)} \to \mathcal{C}W$ by $\rho_W^S(v) = \pi_W \circ \pi_S^{-1}(v)$. To extend this to points in $\partial \mathcal{C}S$ we need the next lemma, which is a basic consequence of the bounded geodesic image axiom of an HHS.

Lemma A.6. *Let $(X, \mathfrak{S})$ be an HHS with hierarchy constant E and S be the $\sqsubseteq$-maximal domain of $\mathfrak{S}$. For each $\kappa \geqslant 1$ there exist $\nu = \nu(\kappa, E)$ so that for all $W \in \mathfrak{S} - \{S\}$ we have:*

(1) *if γ is a (κ, κ)-quasi-geodesic in $\mathcal{C}W$ so that $\gamma \cap \mathcal{N}_\nu(\rho_S^W) = \varnothing$, then*

$$\operatorname{diam}(\rho_W^S(\gamma)) \leqslant E.$$

(2) *if $\gamma_1, \gamma_2 \subset \mathcal{C}S$ are (κ, κ)-quasi-geodesic rays that both represent $p \in \partial \mathcal{C}S$, and $\gamma_i \cap \mathcal{N}_\nu(\rho_S^W) = \varnothing$ for $i = 1$ and 2, then $d_{Haus}(\rho_W^S(\gamma_1), \rho_W^S(\gamma_2)) \leqslant \nu$.*

Now for each $p \in \partial \mathcal{C}S$, define $\partial \rho_W^S$ as follows: let Z be the set of $(1, 20E)$-quasi-geodesics from the basepoint of $\partial \mathcal{C}S$ to p. Let $\nu = \nu(E)$ be the constant from Lemma A.6 for $\kappa = 20E$. For $W \in \mathfrak{S} - \{S\}$, let Z_W be the subset of Z that is at least ν far from ρ_S^W. Define $\partial \rho_W^S(p)$ to be the ν-bounded diameter set $\rho_W^S(Z_W)$.

If p is a point in the Morse boundary $\partial_* X$, we define $\partial \pi_W(p)$ as $\partial \rho_W^S(\partial \pi_S(p))$.

If x, y are are points in any combination of X, $\partial_* X$, $\mathcal{C}S$, or $\partial \mathcal{C}S$, we say x and y are C-cobounded if the union of their projections to $\mathcal{C}W$ for each $W \in \mathfrak{S} - \{S\}$ has diameter at most C. Here the projection is under π_W, $\partial \pi_W$, ρ_W^S, or $\partial \rho_W^S$ depending on which space x and y are in.

In the language of coboundedness, Corollary 6.2 of [ABD21] becomes the following.

Theorem A.7 (Restatement of [ABD21, Corollary 6.2])**.** *Let $(X, \mathfrak{S})$ be a proper HHS with unbounded products and hierarchy constant E. Let $S \in \mathfrak{S}$ be the $\sqsubseteq$–maximal domains and $x, y \in X \cup \partial_* X$.*

(1) *For all Morse gauges M, there exist a constant $C \geqslant 0$, depending on M and E, so that if x and y are joined by an M-Morse geodesic, then x and y are C-bounded.*

(2) *For all $C \geqslant 0$, there exists a Morse gauge M, depending on C and E, so that if x and y are C-cobounded, then there exists an M-Morse geodesic from x to y.*

(3) *For all $C \geqslant 0$ and $p \in \partial \mathcal{C}S$, if $\pi_S(x)$ and p are C-cobounded, then there exists a Morse gauge M, depending on C and E, and $z \in \partial_* X$ so that $\partial \pi_S(z) = p$ and the geodesic from x to z is M-Morse.*

A.6. **Quasi-möbius maps induced by quasi-isometry of $\mathcal{C}S$.** For our final section, we require that G is a group of *automorphisms* of the HHS $(X, \mathfrak{S})$. This means that G acts on $\mathfrak{S}$ by $\sqsubseteq$-, $\perp$-, and $\pitchfork$-preserving bijections, and for each $W \in \mathfrak{S}$ and $g \in G$, there exists an isometry $g_W \colon \mathcal{C}W \to \mathcal{C}gW$ satisfying the following for all $V, W \in \mathfrak{S}$ and $g, h \in G$.

- The map $(gh)_W \colon \mathcal{C}W \to \mathcal{C}ghW$ is equal to the map $g_{hW} \circ h_W \colon \mathcal{C}W \to \mathcal{C}ghW$.
- For each $x \in X$, $g_W(\pi_W(x))$ uniformly coarsely equals $\pi_{gW}(g \cdot x)$.
- If $V \pitchfork W$ or $V \subsetneq W$, then $g_W(\rho_W^V)$ uniformly coarsely equals ρ_{gW}^{gV}.

This equivariance hypothesis is required to establish the next lemma, which is the important technical step in establishing that quasi-isometries of $\mathcal{C}S$ preserve being cobounded when $\mathcal{C}S$ is one-ended.

Lemma A.8. *Let G be a finitely generated group and X any Cayley graph with respect to a finite generating set. Let $\mathfrak{S}$ be an HHS structure for X with hierarchy constant E and suppose G is a group of automorphisms of $(X, \mathfrak{S})$. Let $S \in \mathfrak{S}$ be the $\sqsubseteq$–maximal domains and assume $\mathcal{C}S$ is one-ended. For all $C, R \in \mathbb{N}$, there exists $c, r \in \mathbb{N}$ so that the following holds: Suppose x, y are elements of G that are C-bounded, and let z be the midpoint of a M-Morse geodesic from x to y in X, where M is determined by C and E as in Theorem A.7. If $d_X(x, y) = 2r$, then*

(1) *$\pi_S(x)$ and $\pi_S(y)$ are at least $R + 1$ far from $\pi_S(z)$ and*
(2) *there is a path η in $\mathcal{C}S$ of length less than c that connects $\pi_S(x)$ to $\pi_S(y)$ in $\mathcal{C}S - B_R(\pi_S(z))$.*

Proof. Let x and y be elements of G that are C-bounded and z be the the midpoint of a M-Morse geodesic from x to y, where M is determined by C and E as in Theorem A.7. Thus, x, z and y, z are C'-cobounded for some C' depending on C. Hence the uniqueness axiom of an HHS implies that for each $R \geqslant 0$, there exists $r = r(R, E, C)$ so that $d_X(x, z) = d_X(y, z) \geqslant r$ implies $d_S(x, z), d_S(y, z) \geqslant R + 1$. Since $\mathcal{C}S$ is one-ended, we can connect $\pi_S(x)$ to $\pi_S(y)$ by a path η_{xy} that avoids $B_R(\pi_S(z))$.

Now, since balls in X contain finitely many elements of G, for each r, there only exists finitely many G-orbits of triples of x, y, z so that $d_X(x, y) = 2r$ and x, y are joined by an M-Morse geodesic with midpoint z. Since π_S is coarsely equivariant with respect to the the action of G, for each R and C, we can pick η_{xy} so that its length in $\mathcal{C}S$ is at most $c = c(R, C, E)$. $\square$

We now fix two finitely generated groups G and H, and let X and Y, respectively, be Cayley graphs with respect to finite generating sets. We assume there are HHS structures $\mathfrak{S}$ and $\mathfrak{T}$ for X and Y respectively so that $\mathfrak{S}$ and $\mathfrak{T}$ have unbounded products and that G and H are groups of automorphism of $(X, \mathfrak{S})$ and $(Y, \mathfrak{T})$ respectively. Let $S \in \mathfrak{S}$ and $T \in \mathfrak{T}$ be the respective $\sqsubseteq$-maximal domains, and let E be the hierarchy constant for both $\mathfrak{S}$ and $\mathfrak{T}$.

We now prove that when $\mathcal{C}S$ is one-ended, quasi-isometries $\mathcal{C}S \to \mathcal{C}T$ will preserve coboundedness and hence induce a quasi-möbius map on the Morse boundary. Our proof of Theorem A.9 is inspired by work of Rafi and Schleimer in the case of the mapping class group [RS11].

Theorem A.9. *Suppose $f \colon \mathcal{C}S \to \mathcal{C}T$ is a (λ, ϵ)-quasi-isometric embedding, and let $\partial f \colon \partial \mathcal{C}S \to \partial \mathcal{C}T$ be the topological embedding induced by f. If $\mathcal{C}S$ is one-ended, then:*

(1) *There exists a topological embedding $h \colon \partial_* G \to \partial_* H$ so that*

$$\partial \pi_T \circ h(p) = \partial f \circ \partial \pi_S(p)$$

for all $p \in \partial_* G$.

(2) *If $p, q \in \partial \mathcal{C}S$ are C-cobounded, then $\partial f(p)$ and $\partial f(q)$ are C'-cobounded for some $C' = C'(C, \lambda, \epsilon, E)$. In particular, the map h from (1) is 2-stable.*

(3) *The map h from (1) is quasi-möbius with increasing function ψ determined by λ, ϵ, and E.*

Proof. Without loss of generality, we can assume the image of f is contained in the vertices of $\mathcal{C}T$. Let e be the identity in G and b be the element of H so that $\pi_T(b) = f \circ \pi_S(e)$. Let e and b be the base points for the Morse boundaries of G and H respectively.

Proof of (1)**:** Let $p \in \partial_* G$ and $p_S = \partial \pi_S(p)$. We will first show that $\partial f(p_S) = \partial \pi_T(q_p)$ for some point $q_p \in \partial_* G$. By Theorem A.7, it suffices to show that $\partial f(p_S)$ and $\pi_T(b)$ are cobounded (although not necessarily uniformly).

Let $\gamma \colon [0, \infty) \to G$ be the M-Morse geodesic ray from e to p and let $W \in \mathfrak{T}$. There exist $C = C(M, E)$ so that any two points on γ are C-cobounded by Theorem A.7. By Theorem A.5, $\bar{\gamma} =$

$f \circ \pi_S \circ \gamma \colon [0, \infty) \to \mathcal{C}T$ is a (κ, κ)-quasi-geodesic for some $\kappa = \kappa(\lambda, \epsilon, M, E) \geqslant 1$. Without loss of generality, we can assume $\kappa \geqslant \max\{\lambda, \epsilon, 20E\}$. Let $\nu = \nu(E, \kappa) \geqslant 0$ be the constant from Lemma A.6 so that $\operatorname{diam}(\rho_W^T(\bar{\gamma})) \leqslant E$ whenever $\bar{\gamma}$ is not within ν of ρ_T^W. Let $R = \lambda(2\nu + \epsilon + 1)$ and let $r, c \in \mathbb{N}$ be as in Lemma A.8 for this choice of R and our given C. Hence, we can assume there exists $t \in [0, \infty) \cap \mathbb{Z}$ so that $d_T(\bar{\gamma}(t), \rho_T^W) \leqslant \nu$.

To start, assume $t > 2r + 1$. By Lemma A.8, there exists a path η from $\pi_S \circ \gamma(t-r)$ to $\pi_S \circ \gamma(t+r)$ in $\mathcal{C}S$ so that η has length at most c and η does not intersect the ball of radius R around $\pi_S \circ \gamma(t)$. Now consider a single edge η_0 of η. Since $d_S(\eta_0, \pi_S \circ \gamma(t)) > R = \lambda(2\nu + \epsilon + 1)$ and $\bar{\gamma} = f \circ \pi_S \circ \gamma$, we have

$$(\ast) \qquad d_T(f(\eta_0), \rho_T^W) \geqslant d_T\left(f(\eta_0), \bar{\gamma}(t)\right) - d_T\left(\bar{\gamma}(t), \rho_T^W\right) > \frac{1}{\lambda}R - \epsilon - \nu = \nu + 1.$$

Since $\kappa \geqslant \max\{\lambda, \epsilon\}$, $f(\eta_0)$ is a (κ, κ)-quasi-geodesic in $\mathcal{C}T$, thus $(\ast)$ implies that $\operatorname{diam}(\rho_W^S(f(\eta_0))) \leqslant E$. Since η has length c, this implies $\operatorname{diam}(\rho_W^T(f(\eta))) \leqslant cE$. From the proof of Lemma A.8, r is chosen so that for any $s \in [0, \infty)$ we have

$$|s - t| \geqslant r \implies d_S(\pi_S \circ \gamma(s), \pi_S \circ \gamma(t)) \geqslant R + 1.$$

Thus, our choice of R ensures that $\bar{\gamma}|_{[0, t-r]}$ and $\bar{\gamma}|_{[t+r, \infty)}$ do not intersect $\mathcal{N}_\nu(\rho_T^W)$. Thus, $\operatorname{diam}(\rho_W^T(\bar{\gamma}([0, t-r])))$ and $\operatorname{diam}(\rho_W^T(\bar{\gamma}([t+r, \infty))))$ are both at most E by Lemma A.6. Since $\operatorname{diam}(\rho_W^T(f(\eta))) \leqslant cE$, we have $\operatorname{diam}(\rho_W^T(\bar{\gamma}|_{[0,t-r]}) \cup \rho_W^T(\bar{\gamma}|_{[t+r,\infty)})) \leqslant 2E + cE$.

Now $\pi_W(b) \subseteq \rho_W^T(\bar{\gamma}|_{[0,t-r]})$ because $\pi_W(b) = \rho_W^T(f \circ \pi_S(e))$. Since $\bar{\gamma}|_{[t+r,\infty)}$ represents $\partial f(p_S)$ and does not intersect the ν-neighborhood of ρ_W^T, Lemma A.6 says $\partial \rho_W^S(\partial f(p_S))$ is ν-close to $\rho_W^T(\bar{\gamma}|_{[t+r,\infty)})$. Hence

$$\operatorname{diam}(\rho_W^T(\pi_T(b)) \cup \partial \rho_W^T(\partial f(p_S))) \leqslant 2E + cE + \nu.$$

Now assume $0 \leqslant t \leqslant 2r + 1$. As in the previous case, our choice of R and r ensures that $\bar{\gamma}|_{[2r+1,\infty)}$ does not intersect $\mathcal{N}_\nu(\rho_T^W)$, which implies $\operatorname{diam}(\rho_W^T(\bar{\gamma}|_{[2r+1,\infty)})) \leqslant E$ by Lemma A.6. Since $\rho_W^T \circ \pi_T = \pi_W$ for each $W \in \mathfrak{T} - \{T\}$, the distance formula for an HHS implies there is $A = A(E) \geqslant 0$ so that set of possible domains where $\operatorname{diam}(\rho_W^T(\bar{\gamma}([0, 2r + 1]))) \geqslant A$ is finite. Hence there is some bound $D \geqslant 0$ (depending on p and b) on $\operatorname{diam}(\rho_W^T(\bar{\gamma}))$. As in the previous case, this implies

$$\operatorname{diam}(\rho_W^T(\pi_T(b)) \cup \partial \rho_W^T(\partial f(p_S))) \leqslant D + \nu.$$

Taking $C' = \max\{D + \nu, 2E + cE + \nu\}$, the above shows that $\pi_T(b)$ and $\partial f(p_S)$ are C'-bounded. Hence, by Theorem A.7, there exists $q_p \in \partial_* H$ so that $\partial \pi_T(q_p) = \partial f(\partial_S(p))$.

We now show that the map $h \colon \partial_* G \to \partial_* H$ defined by $h(p) = q_p$ is a topological embedding. By construction $\partial \pi_T \circ h = \partial f \circ \partial \pi_S$. Since ∂f, $\partial \pi_S$, and $\partial \pi_T$ are all topological embeddings, h must also be a topological embedding.

Proof of (2): Let $p, q \in \partial \mathcal{C}S$ be C-cobounded. By Theorem A.7, there exist $p', q' \in \partial_* G$ so that $\partial \pi_S(p') = p$ and $\partial \pi_S(q') = q$. We will show $f(p)$ and $f(q)$ are C'-cobounded for some C' depending on λ, ϵ, C, and E.

By Theorem A.7, there exists a Morse gauge $M = M(C, E)$ so that p' and q' are connected by a bi-infinite M-Morse geodesic in G. Let $\gamma \colon (-\infty, \infty) \to G$ be such an M-Morse geodesic and let $\bar{\gamma} = f \circ \pi_S \circ \gamma$. Let $W \in \mathfrak{T} - \{T\}$. By Theorem A.5, $\pi_S \circ \gamma$ is a (κ, κ)-quasi-geodesic in $\mathcal{C}S$ for $\kappa = \kappa(M, E)$. Lemma A.6 says there is $\nu = \nu(\lambda, \epsilon, \kappa, E)$ so that if $\mathcal{N}_\nu(\rho_T^W) \cap \bar{\gamma} = \varnothing$, then $\operatorname{diam}(\rho_W^T(\bar{\gamma})) \leqslant E$. Hence, we can assume there is $t \in \mathbb{Z}$ so that $\bar{\gamma}(t) \in \mathcal{N}_\nu(\rho_T^W)$. Let $R = \lambda(2\nu + \epsilon + 1)$ and let $r, c \in \mathbb{N}$ be as in Lemma A.8 for this choice of R and our given C. Thus, there exists a path η in $\mathcal{C}S$ connecting $\pi_S \circ \gamma(t-r)$ and $\pi_S \circ \gamma(t+r)$ so that η avoids the R-ball around $\pi_S \circ \gamma(t)$ and η has length at most c.

By arguing exactly as we did in the $t > 2r + 1$ case of the proof of (1), we have $\operatorname{diam}(\rho_W^T(\eta)) \leqslant cE$ plus

$$\operatorname{diam}(\rho_W^T(\bar{\gamma}|_{(-\infty, r-t]})) \leqslant E \text{ and } \operatorname{diam}(\rho_W^T(\bar{\gamma}|_{[r+t, \infty)})) \leqslant E.$$

Continuing as in the proof of (1), this will imply $\partial f(p)$ and $\partial f(q)$ are $(2E + cE + 2\nu)$-cobounded.

The fact that the homeomorphism h from (1) is 2-stable now follows due to the correspondence in Theorem A.7 between being cobounded and being joined by a Morse geodesic.

Proof of (3): For distinct $a, b, c, d \in \partial \mathcal{C}S$, let $[a, b, c, d]_S$ denote the cross-ratio in $\partial \mathcal{C}S$. By Lemma A.3, f being a quasi-isometric embedding implies ∂f is quasi-möbius with respect to the cross ratio on

$\partial \mathcal{C}S$. Proposition 7 of [MR19] shows that for each Morse gauge M and each tuple $(a, b, c, d) \in \partial_*^{(4,M)}G$, we have

$$[a, b, c, d]_M \asymp [\partial \pi_S(a), \partial \pi_S(b), \partial \pi_S(c), \partial \pi_S(d)]_S$$

with constants determined by E and M. Hence the map h from (1) is quasi-möbius. $\qquad \square$

Our main result (Theorem A.1) is now a corollary of Theorem A.9.

Corollary A.10. *If $\mathcal{C}S$ is one-ended and there is a (λ, ϵ)-quasi-isometry $f : \mathcal{C}S \to \mathcal{C}T$, then there exists a quasi-isometry $F : G \to H$ so that for all $x \in G$, $\pi_T \circ F(x)$ is uniformly close to $f \circ \pi_S(x)$.*

Proof. Theorem 14.3 of [BHS17b] says G acts cobounded and acylindrically on $\mathcal{C}S$. Thus, $\mathcal{C}S$ being one-ended implies that ∂_*G contains at least 3 points.

Let $f^{-1} : \mathcal{C}T \to \mathcal{C}S$ be a quasi-inverse for f. Applying Theorem A.9 to both f and f^{-1}, we produce a quasi-möbius homeomorphism $h : \partial_*G \to \partial_*H$ with quasi-möbius inverse so that $\partial f \circ \partial \pi_S = \partial \pi_T \circ h$. By Theorem A.4 there is a quasi-isometry $F : G \to H$ so that $\partial F = h$.

Because G acts cocompactly on X, there is a Morse gauge M and a constant $K_M \geqslant 0$ so that for all $x \in G$ there is a triple $(p, q, z) \in \partial_*^{(3,M)}G$ so that x is a K_M-center for (p, q, z). Now $(\partial F(p), \partial F(q), \partial F(z)) \in \partial_*^{(3,M')}H$ for some M' determined by M, λ, ϵ, and E. Moreover, $F(x)$ is uniformly close to all three sides of any ideal M'-Morse triangle with vertices $\partial F(p), \partial F(q), \partial F(z)$. Since $\partial f \circ \partial \pi_S = \partial \pi_T \circ \partial F$, this implies $\pi_T \circ F(x)$ is uniformly close to $f \circ \pi_S(x)$. $\qquad \square$

Remark A.11. There is a proof of Theorem A.1 that does not rely on the Morse boundary and quasi-Möbius maps, but instead directly invokes the quasi-isometry on the $\sqsubseteq$-maximal hyperbolic space $\mathcal{C}S$. This argument uses the fact that quasi-isometries of $\mathcal{C}S$ preserve the set of cobounded pairs in $\partial \mathcal{C}S$ to show that quasi-isometries of $\mathcal{C}S$ produce "quasi-Möbius" maps on the set of cobounded tuples in $\partial \mathcal{C}S$. Thus, one can essentially repeat the arguments used in [CCM19] to build a quasi-isometry of G with coboundedness and the distance formula replacing the role of Morse geodesics. As this approach would result in a lengthier proof without any fundamentally new ideas, we have elected to give the shorter proof using the established results from the literature.

References

[ABD21]   Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham. Largest acylindrical actions and stability in hierarchically hyperbolic groups. *Trans. Amer. Math. Soc. Ser. B*, 8:66–104, 2021. With an appendix by Daniel Berlyne and Jacob Russell.

[ABR22]   Carolyn Abbott, Jason Behrstock, and Jacob Russell. Structure invariant properties of the hierarchically hyperbolic boundary. *arXiv preprint arXiv:2208.07930*, 2022.

[AMS16]   Yago Antolín, Ashot Minasyan, and Alessandro Sisto. Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups Geom. Dyn.*, 10(4):1149–1210, 2016.

[BBFS20]   Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. Acylindrical actions on projection complexes. *Enseign. Math.*, 65(1-2):1–32, 2020.

[BDM09]   Jason Behrstock, Cornelia Druţu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Math. Ann.*, 344(3):543–595, 2009.

[BF02]   Mladen Bestvina and Koji Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geom. Topol.*, 6:69–89, 2002.

[BHMS20]   Jason Behrstock, Mark Hagen, Alexandre Martin, and Alessandro Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. *arXiv preprint arXiv:2005.00567*, 2020.

[BHS17a]   Jason Behrstock, Mark Hagen, and Alessandro Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3)*, 114(5):890–926, 2017.

[BHS17b]   Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[BHS19]   Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299(2):257–338, 2019.

[BHS21]   Jason Behrstock, Mark Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):909–996, 2021.

[BKMM12]   Jason Behrstock, Bruce Kleiner, Yair N. Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.

[BN12]   Jason A. Behrstock and Walter D. Neumann. Quasi-isometric classification of non-geometric 3-manifold groups. *J. Reine Angew. Math.*, 669:101–120, 2012.

[Bow13]   Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013.

[Bow18]   Brian H. Bowditch. Large-scale rigidity properties of the mapping class groups. *Pacific J. Math.*, 293(1):1–73, 2018.

[BR20a]   Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic groups. *Groups Geom. Dyn.*, 14(4):1127–1203, 2020.

[BR20b]   Daniel Berlyne and Jacob Russell. Hierarchical hyperbolicity of graph products. *arXiv preprint arXiv:2006.03085*, 2020.

[Cas16]   Christopher H. Cashen. Quasi-isometries need not induce homeomorphisms of contracting boundaries with the Gromov product topology. *Anal. Geom. Metr. Spaces*, 4(1):278–281, 2016.

[CCM19]   Ruth Charney, Matthew Cordes, and Devin Murray. Quasi-Mobius homeomorphisms of Morse boundaries. *Bull. Lond. Math. Soc.*, 51(3):501–515, 2019.

[CCMT15]   Pierre-Emmanuel Caprace, Yves Cornulier, Nicolas Monod, and Romain Tessera. Amenable hyperbolic groups. *J. Eur. Math. Soc. (JEMS)*, 17(11):2903–2947, 2015.

[CD19]   Matthew Cordes and Matthew G. Durham. Boundary convex cocompactness and stability of subgroups of finitely generated groups. *Int. Math. Res. Not. IMRN*, 6:1699–1724, 2019.

[Che20]   Marissa Chesser. Stable subgroups of the genus two handlebody group. *arXiv preprint arXiv:2009.05067*, 2020.

[CHK22]   Montserrat Casals-Ruiz, Mark Hagen, and Ilya Kazachkov. Real cubings and asymptotic cones of hierarchically hyperbolic groups. *Preprint available at https://www.wescac.net/cones.html*, 2022.

[CM19]   Christopher H. Cashen and John M. Mackay. A metrizable topology on the contracting boundary of a group. *Trans. Amer. Math. Soc.*, 372(3):1555–1600, 2019.

[Cor17]   Matthew Cordes. Morse boundaries of proper geodesic metric spaces. *Groups Geom. Dyn.*, 11(4):1281–1306, 2017.

[DDLS21]   Spencer Dowdall, Matthew G. Durham, Christopher J. Leininger, and Alessandro Sisto. Extensions of Veech groups II: hierarchical hyperbolicity and quasi-isometric rigidity. *arXiv preprint arXiv:2111.00685*, 2021.

[DGO17]   F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.

[DHS17]   Matthew Gentry Durham, Mark Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.

[DHS20]   Matthew Gentry Durham, Mark Hagen, and Alessandro Sisto. Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 24(2):1051–1073, 2020.

[DS05]   Cornelia Druţu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis Osin and Mark Sapir.

[GS18]   Dominik Gruber and Alessandro Sisto. Infinitely presented graphical small cancellation groups are acylindrically hyperbolic. *Ann. Inst. Fourier (Grenoble)*, 68(6):2501–2552, 2018.

[GS21]   Antoine Goldsborough and Alessandro Sisto. Markov chains on hyperbolic-like groups and quasi-isometries. *arXiv:2111.09837. To appear in J. Reine Angew. Math.*, 2021.

[GS23]   Antoine Goldsborough and Alessandro Sisto. Random divergence of groups. *arXiv preprint arXiv:2303.09943*, 2023.

[HMS21]   Mark Hagen, Alexandre Martin, and Alessandro Sisto. Extra-large type Artin groups are hierarchically hyperbolic. *arXiv preprint arXiv:2109.04387*, 2021.

[HRSS22]   Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano. Equivariant hierarchically hyperbolic structures for 3–manifold groups via quasimorphisms. *arXiv preprint arXiv:2206.12244*, 2022.

[HS20]   Mark Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *Israel J. Math.*, 236(1):45–89, 2020.

[Hua16]   Jingyin Huang. Quasi-isometry classification of right-angled artin groups ii: several infinite out cases. *arXiv preprint arXiv:1603.02372*, 2016.

[Hua17]   Jingyin Huang. Quasi-isometric classification of right-angled artin groups, i: The finite out case. *Geometry & Topology*, 21(6):3467–3537, 2017.

[KK14]   Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled Artin group. *Internat. J. Algebra Comput.*, 24(2):121–169, 2014.

[KL08]   Autumn E. Kent and Christopher J. Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008.

[Man05]   Jason Fox Manning. Geometry of pseudocharacters. *Geom. Topol.*, 9:1147–1185, 2005.

[MR19]   Sarah C. Mousley and Jacob Russell. Hierarchically hyperbolic groups are determined by their Morse boundaries. *Geom. Dedicata*, 202:45–67, 2019.

[MS20]   P. Mathieu and A. Sisto. Deviation inequalities for random walks. *Duke Math. J.*, 169(5):961–1036, 2020.

[MT18]   Joseph Maher and Giulio Tiozzo. Random walks on weakly hyperbolic groups. *J. Reine Angew. Math.*, 742:187–239, 2018.

[Osi16]   D. Osin. Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.*, 368(2):851–888, 2016.

[Osi18]   Denis V. Osin. Groups acting acylindrically on hyperbolic spaces. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pages 919–939. World Sci. Publ., Hackensack, NJ, 2018.

[Pau96]   Frédéric Paulin. Un groupe hyperbolique est déterminé par son bord. *J. London Math. Soc. (2)*, 54(1):50–74, 1996.

[PS20]   Harry Petyt and Davide Spriano. Unbounded domains in hierarchically hyperbolic groups. *arXiv:2007.12535. To appear in Groups Geom. Dyn.*, 2020.

[PSZ22]   Harry Petyt, Davide Spriano, and Abdul Zalloum. Hyperbolic models for CAT(0) spaces. *In preparation*, 2022.

[RS11]   Kasra Rafi and Saul Schleimer. Curve complexes are rigid. *Duke Math. J.*, 158(2):225–246, 2011.

[RST22]   Jacob Russell, Davide Spriano, and Hung Cong Tran. The local-to-global property for Morse quasi-geodesics. *Math. Z.*, 300(2):1557–1602, 2022.

[Rus21]   Jacob Russell. Extensions of multicurve stabilizers are hierarchically hyperbolic. *arXiv preprint arXiv:2107.14116*, 2021.

[Sis16]   Alessandro Sisto. Quasi-convexity of hyperbolically embedded subgroups. *Math. Z.*, 283(3-4):649–658, 2016.

[Sis19]   Alessandro Sisto. What is a hierarchically hyperbolic space? In *Beyond hyperbolicity*, volume 454 of *London Math. Soc. Lecture Note Ser.*, pages 117–148. Cambridge Univ. Press, Cambridge, 2019.

[TT15]   Samuel Taylor and Giulio Tiozzo. Random extensions of free groups and surface groups are hyperbolic. *International Mathematics Research Notices*, 2016, 01 2015.

[Ves23]   Elliott Vest. The curtain model is not a quasi-isometry invariant of cat (0) spaces. *arXiv preprint arXiv:2312.16325*, 2023.

[Why99]   Kevin Whyte. Amenability, bilipschitz equivalence, and the von neumann conjecture. *Duke Mathematical Journal*, 99:93–112, 1999.

Maxwell Institute and Department of Mathematics, Heriot-Watt University, Edinburgh, UK
Email address: ag2017@hw.ac.uk

School of Mathematics, University of Bristol, Bristol, UK
Email address: markfhagen@posteo.net

Mathematical Institute, University of Oxford, UK
Email address: petyt@maths.ox.ac.uk

Maxwell Institute and Department of Mathematics, Heriot-Watt University, Edinburgh, UK
Email address: a.sisto@hw.ac.uk

Department of Mathematics & Statistics, Swarthmore College, Swarthmore, PA, USA
Email address: jrussel2@swarthmore.edu