

CONSTRUCTING METRIC SPACES FROM SYSTEMS OF WALLS

HARRY PETYT AND ABDUL ZALLOUM

With an appendix with Davide Spriano

ABSTRACT. We give a general procedure for constructing metric spaces from systems of partitions. This generalises and provides analogues of Sageev’s construction of dual $\text{CAT}(0)$ cube complexes for the settings of hyperbolic and injective metric spaces.

As applications, we produce a “universal” hyperbolic action for groups with strongly contracting elements, and show that many groups with “coarsely cubical” features admit geometric actions on injective metric spaces. In an appendix with Davide Spriano, we show that a large class of groups have an infinite-dimensional space of quasimorphisms.

CONTENTS

1. Introduction	1
2. Background on medians and $\text{CAT}(0)$ cube complexes	8
Part 1. Generalising Sageev’s construction	9
3. Ultrafilters and dualisable systems	9
4. Systems of chains	14
5. Producing hyperbolic spaces	18
Part 2. Applications	25
6. A universal hyperbolic space for contracting geodesics	25
7. Strong coarse median spaces	31
8. Other directions	38
Appendix. Quasimorphisms (with Davide Spriano)	40
References	42

1. INTRODUCTION

A guiding principle of geometric group theory is the idea that one can learn about a group from its actions on nice metric spaces. There are two aspects to implementing this: producing the actions, and studying the properties of groups that possess them. In this article we are interested in the former.

Approaches to producing spaces for groups to act on take many forms, such as local-to-global results [AB90, Lea13, Hae22, Bow21] and combination theorems [BF92, Dah03, MR08, HW15b]. However, amongst such techniques, the stand-out examples are those with inter-categorical features, such as: hyperbolisation procedures [Gro87, DJ91, CD95, Ont20], which convert simplicial complexes into aspherical manifolds of nonpositive curvature; Sageev’s construction [Sag95], which outputs a $\text{CAT}(0)$ cube complex from a system of partitions; and

the projection complex machinery [BBF15, BBFS20], which can (amongst many other things) produce actions on quasitrees from simple geometric conditions.

We shall work in the same vein as these latter two, and provide a general framework that one can use for producing group actions on different types of metric space. The framework takes the same form of input as Sageev’s construction, but we allow additional flexibility that allows for a much wider variety of metrics to be produced, and under weaker conditions. That is, we start with a set S together with a collection P of bipartitions. This pair has a naturally-defined dual median algebra associated with it, and by considering varying subsets of 2^P we obtain a range of metrics on this dual. Sageev’s construction is recovered as a degenerate case. The principal idea of this article can therefore be phrased as follows. See Section 3 (and also Definition 1.2 later in the introduction) for more precise formulations of some initial statements.

Construction 1.1. Let S be a set with a collection P of bipartitions. By choosing an appropriate subset $\mathcal{C} \subset 2^P$, one obtains a metric space X and a natural map $S \rightarrow X$.

To obtain an action on X from an action on (S, P) , one merely has to insist that $\mathcal{C}$ is preserved. The three main types of metric spaces that we systematically produce in this way are Helly graphs, coarsely injective spaces, and hyperbolic spaces. However, the set-up is quite general and can likely be used to provide actions on other interesting metric spaces.

Let us mention that Sageev’s construction has also been generalised by Chatterji–Druţu–Haglund, with improvements by Fioravanti [CDH10, Fio20]. The two generalisations are of rather different flavours: the one in [CDH10, Fio20] can be thought of as a “continuous version” of Sageev’s construction, and it firmly belongs in the *median* category; whereas our generalisation is still fundamentally discrete, but is not restricted to ℓ^1 geometry.

1.1. HYPERBOLIC DUALS

Group actions on hyperbolic spaces are a major theme of research in the field, especially in the wake of the result of Masur–Minsky that the curve graph of a surface is hyperbolic [MM99]. Indeed, this spawned the theory of acylindrical hyperbolicity [Sel97, Bow08, DGO17, Osi16]. In this context, the aforementioned projection complex machinery is a potent tool, because it allows one to produce actions on quasitrees such that a given *strongly contracting* (i.e. “negatively-curved”) group element has a so-called *WPD* action [BF02], which is sufficient to establish acylindrical hyperbolicity [DGO17]. Strongly contracting elements have been studied by many authors, including [ACT15, GY22, Alg11, Pap22, Cou23].

One unsatisfactory aspect of the actions on quasitrees coming from projection complexes is that they are rather unnatural, in that they have little to do with the original space. The procedure also has the limitation of only being able to handle a finite number of strongly contracting elements at once. By building an appropriate collection of partitions, we are able to apply Construction 1.1 to produce, for any geodesic metric space S , a hyperbolic space that is “universal”, in the sense that all strongly contracting geodesics are witnessed there. (See Theorem 6.8 for the precise statement.)

Theorem A. *Let S be a geodesic metric space. There is a hyperbolic space X and an Isom S –equivariant coarsely Lipschitz map $\pi : S \rightarrow X$ such that a geodesic $\gamma \subset S$ is strongly contracting if and only if $\pi\gamma$ is a quasigeodesic. Moreover, the correspondence is quantitative.*

This can be compared to (and, via [SZ22], be viewed as greatly extending) the situation for pseudo-Anosov mapping classes acting on the curve graph [MM99]; outer space and the

free factor complex [BF14, DT19]; and rank-one isometries acting on the curtain model of a CAT(0) space [PSZ22].

The partitions that are used to construct the above hyperbolic space are obtained by closest-point projecting to strongly contracting geodesics of S , similarly to the construction of the curtain model in [PSZ22], but the choice of $\mathcal{C} \subset 2^P$ is a little different to the notion of L -separation used there; see Definition 6.4.

The statement of Theorem A is purely about the metric space S and its isometries. Applying it to groups acting properly on S , we can obtain a simultaneous WPD action of *all* strongly contracting elements; see Corollary 6.11 and Proposition 6.13.

Theorem B. *Let S be a geodesic metric space, and let X be the hyperbolic space of Theorem A. For any group G acting properly on S , elements of G are strongly contracting exactly when they act loxodromically on X , and every such element is WPD.*

Furthermore, under the additional assumption that G acts properly coboundedly on a metric space where all Morse geodesics are strongly contracting, we show that the hyperbolic space from Theorem A is a *universal recognising space* for stable subgroups of G , in the sense of [BCK⁺23]. That is, every stable subgroup of G has quasiisometrically embedded orbits in X . The following is Theorem 6.15 in the text.

Theorem C. *Suppose that a group G acts properly coboundedly on a metric space where all Morse geodesics are (quantitatively) strongly contracting. There is a hyperbolic space X , on which G acts, such that a finitely generated subgroup of G is stable if and only if its orbits in X are quasiisometric embeddings.*

(Here “quantitatively” means that the contracting constant depends only on the Morse gauge, not the particular Morse geodesic.) The list of spaces covered by Theorem C includes CAT(0) spaces [CS15], injective spaces [SZ22], Garside groups [CW21b], certain small-cancellation groups [Zbi23], and weakly modular graphs with convex balls [SZ]. It therefore generalises known results for mapping class groups [KL08, DT15], hierarchically hyperbolic groups more generally [ABD21, HHP23, SZ22], and CAT(0) groups [PSZ22].

Another application of this universal hyperbolic space X , and of the particulars of its construction, is investigated in the appendix, together with Davide Spriano. A *quasimorphism* of a group G is a map $f : G \rightarrow \mathbf{R}$ that fails to be a homomorphism by a bounded amount, in the sense that there is some r such that $|f(gh) - f(g) - f(h)| \leq r$ for all $g, h \in G$. There are two “trivial” classes of quasimorphisms, namely the bounded functions and the homomorphisms. The quotient of the space of quasimorphisms by these trivial classes is denoted $\widetilde{\text{QM}}(G)$; it is intimately related with the second (real) cohomology of G .

For the most part, calculations of $\widetilde{\text{QM}}(G)$ (when it is nontrivial) rely on G acting with some form of properness on some nice metric space (a notable exception where properness is not needed is [CF10]). The main result of the appendix is the following, which has no properness assumption. See Theorem A.2.

Theorem D. *Let S be a non-hyperbolic geodesic space where all Morse geodesics are (quantitatively) strongly contracting. If S contains a biinfinite Morse geodesic, then $\widetilde{\text{QM}}(G)$ is infinite-dimensional for every group G acting coboundedly on S .*

Note that the conclusion of Theorem D can fail for automorphism groups of trees [IPS21].

On Zbinden’s work. Results similar to Theorems A and B have also been proved in upcoming work of Stefanie Zbinden [Zbi]. More precisely, Zbinden has a construction, rather

different to the one employed in this article and more tailored to the setting of strongly contracting geodesics, that can produce, for any geodesic space S , a hyperbolic space Z with the following features.

- A geodesic in S is strongly contracting if and only if it projects to a quasigeodesic of Z .
- Any group G acting properly on S acts *non-uniformly acylindrically* on Z , and strongly contracting elements are loxodromic.
- If S has the property that all Morse geodesics are strongly contracting, and if G acts properly coboundedly on S , then all generalised loxodromics of G are loxodromic and WPD on Z .

1.2. INJECTIVE DUALS

Construction 1.1 can also systematically produce Helly graphs and coarsely injective metric spaces. These are recent additions to the arsenal of geometric group theory that have proved to be rather powerful [HO21, Hae21, HHP23, Hae23, Zal23]. A geodesic metric space is (coarsely) injective if its metric balls satisfy a Helly property: any balls that pairwise intersect have nonempty total (coarse) intersection.

The theory appears to have much in common with that of CAT(0) cube complexes and CAT(0) spaces, but aside from the fact that every CAT(0) cube complex is injective when equipped with the ℓ^∞ metric [MT91, Mie14], this is largely heuristic at this stage. In the cubical setting, it is Sageev’s construction that is really the foundation on which many (often highly nontrivial) cubulation results are built, such as in [BW12, HW15b, HW15a].

This motivates the following, a precise version of which is stated as Theorem 4.9; also see Theorem K below for more information.

Theorem E. *Under simple conditions on $\mathcal{C}$, the metric space X produced by Construction 1.1 is coarsely injective.*

It should be emphasised that the assumptions in Theorem E are weaker than those of Sageev’s construction.

As a first implementation, we consider a class of groups that naturally generalises hyperbolic spaces and CAT(0) cube complexes. In δ -hyperbolic spaces, Gromov’s tree approximation lemma states that the quasiconvex hull of every finite set A is quasiisometric to a tree, where the quasiisometry constant depends only on δ and the cardinality of A [Gro87]. In view of results for mapping class groups and Teichmüller space [BM11, EMR17, Dur16], the fact that CAT(0) cube complexes can be considered as higher-dimensional versions of trees led Bowditch to introduce the class of *coarse median* spaces [Bow13a], which are defined by a *cubical* approximation property.

Such an approximation property could be formulated in multiple ways, and although Bowditch uses perhaps the most general possibility, which contains precious little metric information, he is able to obtain some remarkably strong results; see especially [Bow19]. To our knowledge, though, all coarse median spaces of prior interest satisfy a seemingly much stronger property that more faithfully generalises Gromov’s lemma [BHS21, Bow18a]. Moreover, this stronger property is sometimes very useful [HHP23, DMS23, DZ22].

We therefore make a small tweak to the definition of coarse median spaces, and introduce the class of *strong coarse median* spaces (these have also been called *locally quasicubical* spaces in [DZ22]). Essentially, a strong coarse median space is a geodesic space where every finite set A has an appropriate “coarsely convex” hull that is *median-preservingly* quasiisometric to a

CAT(0) cube complex, with the constants depending only on the cardinality of A . A *strong-coarse-median group* is then a finitely generated group possessing an equivariant strong coarse median structure. Aside from hyperbolic and cubical groups, the class includes toral relatively hyperbolic groups, mapping class groups, extra-large-type Artin groups, and hierarchically hyperbolic groups more generally [BHS19, BHS21, HMS21]. Note that some cubical groups do not *naturally* admit hierarchically hyperbolic structures [She22].

Our main result about strong-coarse median groups is the construction of good sets P and $\mathcal{C}$, leading to the following, which generalises and gives a hierarchy-free proof of the main result of [HHP23]; it appears as Corollary 7.14.

Theorem F. *For any strong coarse median space S , Theorem E produces a coarsely injective space $X_{\mathcal{C}}(S)$ such that the natural map $S \rightarrow X_{\mathcal{C}}(S)$ is a quasiisometry. In particular, if S is a strong-coarse-median group, then S acts properly coboundedly on an injective space.*

Since $\mathcal{C}$ is constructed from metric and median data, every median-preserving isometry of S induces an isometry of $X_{\mathcal{C}}(S)$. As a special case of this, known results about groups acting properly coboundedly on injective spaces lead to the following, which notably does not need the quasiisometry to be equivariant [Lan13, KMV22, AB95] (see also [Hae23]).

Corollary G. *Suppose that G is quasiisometric to a CAT(0) cube complex Q . If the induced quasiaction of G on Q is coarsely median-preserving, then:*

- G has finitely many conjugacy classes of finite subgroups;
- Q is not a subgroup of G ;
- polycyclic subgroups of G are virtually abelian and undistorted;
- G is semihyperbolic.

Already this corollary applies to mapping class groups, and many hierarchically hyperbolic groups more generally [Pet21], overlapping with [HHP23]. In particular, the statement about polycyclic subgroups provides a large generalisation of [FLM01, Thm 1.2]. In view of the fact that not all cubical groups are naturally hierarchically hyperbolic, it seems likely that there are non-cubical groups that are covered by Corollary G but not [HHP23].

In the same setting as Theorem E, we build a natural family of paths in the coarsely injective space X , which we call *normal wall paths*. These paths simultaneously behave well with respect to the metric on X and with respect to its structure as a *median algebra* (Propositions 4.6, 4.7). Degenerating to the setting of CAT(0) cube complexes exactly gives the Niblo–Reeves *normal cube paths* [NR98].

One simple use for normal wall paths is that they provide a nice melding of results from [DMS23] and [HHP23]. The main result of each of those articles is semihyperbolicity of mapping class groups, but this is achieved in different ways. Semihyperbolicity asks for an equivariant set of paths with good fellow-travelling properties. In [HHP23], these come from Lang’s bicombing of injective spaces [Lan13], whereas [DMS23] contains a direct construction that yields paths compatible with the Masur–Minsky hierarchy structure [MM00]. It is not clear that these systems of paths are related to each other: the former may not be hierarchical, and the latter may be unnatural to the injective space. The following shows that normal wall paths satisfy both. (See Propositions 4.6 and 4.7, as well as Lemma 7.10.)

Theorem H. *Let S be either the mapping class group $\mathrm{MCG}\Sigma$ of a finite-type surface Σ , or the corresponding Teichmüller space with the Teichmüller metric. Normal wall paths in the dual space $X_{\mathcal{C}}(S)$ are median paths, rough geodesics, and yield hierarchy paths of S : their images under subsurface projections are unparametrised quasigeodesics. Furthermore, they form a $\mathrm{MCG}\Sigma$ -invariant bicombing.*

The statement of Theorem H applies to all strong-coarse-median groups, whereas [DMS23] and [HHP23] cover (colourable) hierarchically hyperbolic groups. In that generality, normal wall paths witness semihyperbolicity. For the Teichmüller metric, a system of quasigeodesics with similar properties to those of Theorem H has been announced by Kapovich–Rafi.

The interplay between the median and the metric on $X_{\mathcal{C}}(S)$ has consequences for the geometry of its balls. For instance, following Cannon [Can87], we say that a metric space M is *almost convex* if for every k there exists $N(k)$ such that for all r , every pair of points in an r -sphere $S_r(m) \subset M$ can be joined by a path in the ball $B_r(m)$ of length at most $N(k)$. The following shows that this holds for $X_{\mathcal{C}}(S)$ in a strong way. (See Lemma 4.2 and Proposition 4.6.)

Theorem I. *If S is a strong coarse median space, then normal wall paths make $X_{\mathcal{C}}(S)$ almost convex, with $N(k) = k + 6$.*

In [Far06, Qns 3.4, 3.5], Farb asks whether there exist almost convex Cayley graphs of mapping class groups, and whether the Teichmüller metric is almost convex. Although the above theorem does not directly answer Farb’s questions, it shows that almost convexity holds in the space $X_{\mathcal{C}}(S)$, which is equivariantly quasiisometric to the desired spaces, and is witnessed by hierarchy paths.

Let us mention one more application of Construction 1.1, which both makes use of normal wall paths and brings us back to the setting of hyperbolic duals. By using the same set of walls as is used to establish Theorem F but a set $\mathcal{C}$ more along the lines of Theorem A, one can, for any strong coarse median space, produce a hyperbolic space $Y(S)$ associated with it. In the hierarchically hyperbolic setting, a proof involving normal wall paths shows that $Y(S)$ coarsely recovers the largest hyperbolic space of [ABD21]. The following summarises Section 7.3.

Theorem J. *For every strong coarse median space S , all median-preserving isometries of S are isometries of the hyperbolic space $Y(S)$. If S is the mapping class group of a finite-type surface Σ , or its Teichmüller space with the Teichmüller metric, then $Y(S)$ is MCG Σ -equivariantly quasiisometric to the curve graph of Σ .*

This provides a systematic construction of the curve graph from only median and metric information. It also lets the curve graph be viewed as a kind of quotient of the coarsely injective space $X_{\mathcal{C}}(S)$ of Theorem H, because the underlying walls are the same.

1.3. FURTHER DISCUSSION OF THE CONSTRUCTION

In this subsection, we discuss Construction 1.1 in more detail and provide several applications. Given a set S and a collection P of bipartitions, an *ultrafilter* is a “consistent orientation” of the elements of P ; see Definition 3.1. Each $s \in S$ determines a *principal* ultrafilter ϕ_s by orienting each $h \in P$ towards s . Let $\hat{X}$ be the set of all ultrafilters on P . We say that $x, y \in \hat{X}$ are *separated* by $c \in P$ if they orient every element of c differently. Given $\mathcal{C} \subset 2^P$, let $d_{\mathcal{C}}$ be the function on $\hat{X}$ given by $d_{\mathcal{C}}(x, y) = \sup\{|c| : c \in \mathcal{C} \text{ separates } x \text{ from } y\}$, which takes values in $\mathbf{N} \cup \{\infty\}$.

Definition 1.2 (Dualisable system). Let (S, P) be a set with walls, and let $\mathcal{C} \subset 2^P$ be closed under subsets, with $P \subset \mathcal{C}$. We say that $\mathcal{C}$ is a *dualisable system* on P if $d_{\mathcal{C}}(\phi_s, \phi_t) < \infty$ for all $s, t \in S$. The $\mathcal{C}$ -dual of S is $X_{\mathcal{C}} = (X, d_{\mathcal{C}})$, where $X = \{x \in \hat{X} : d_{\mathcal{C}}(x, \phi_s) < \infty\}$.

The following theorem summarises the correspondences between combinatorics of $\mathcal{C}$ and metric properties of $X_{\mathcal{C}}$ obtained in this paper; see Theorem 4.9 and Corollary 5.7.

Theorem K. *For any dualisable system $\mathcal{C}$ on (S, P) , the space $X_{\mathcal{C}}$ is a metric space, and*

- *if $\mathcal{C} = 2^P$, then $X_{\mathcal{C}}$ is Sageev’s dual $\text{CAT}(0)$ cube complex;*
- *if $\mathcal{C}$ consists of all chains in P , then $X_{\mathcal{C}}$ is the vertex set of a Helly graph.*

Moreover, assuming that $\mathcal{C}$ is m -gluable:

- *if each $c \in \mathcal{C}$ is a chain, then $X_{\mathcal{C}}$ is coarsely injective;*
- *if $\mathcal{C}$ is L -separated, then $X_{\mathcal{C}}$ is hyperbolic.*

The set $\mathcal{C}$ is said to be m -gluable if it is closed under taking certain unions (see Definition 4.5). It is L -separated if each $c \in \mathcal{C}$ is a chain and for any $\{h_1, h_2\} \in \mathcal{C}$, every $c \in \mathcal{C}$ whose elements all cross both h_1 and h_2 has $|c| \leq L$.

The final bullet of Theorem K bears similarity to the construction of [Bow98]. Bowditch starts with a group G acting properly discontinuously and cocompactly on the space of triples of a perfect metric compactum. He then constructs “annulus systems”, which are certain collections of bipartitions that play the role of $\mathcal{C}$, and shows that they are L -separated, using the geometry of the action.

Whilst Construction 1.1 is based on Sageev’s construction, one fundamental difference is that it does not require local finiteness of the collection of walls. Of course, there is a trade-off between how general the input can be and how strong the conclusions are. Consider hyperbolic groups, for instance. Adding the assumption of a cocompact cubulation to hyperbolicity has some very strong consequences [Ago13, Wis21], but there are many hyperbolic groups with property (T) [Ž03, KK13b], and these have no unbounded actions on $\text{CAT}(0)$ cube complexes [NR97]. Sageev’s construction itself (and even the continuous version of [CDH10, Fio20]) cannot, therefore, convey any of the geometry of such groups. Nevertheless, Theorem A outputs a metric space X that is quasiisometric to G and still has several fine properties in common with $\text{CAT}(0)$ cube complexes, as we now discuss.

For one, in the full generality of Construction 1.1, every finite group acting on S fixes a point in a “first subdivision” of X (Proposition 3.16); and so a standard argument shows that if G acts properly coboundedly on X then it has finitely many conjugacy classes of finite subgroups. Even in δ -hyperbolic spaces, one can only guarantee an orbit of diameter roughly δ . In this regard, X can be thought of as playing a similar role to the *injective hull* [Lan13]. That said, neither the median structure on X nor constructions such as normal wall paths have analogues in the injective hull, as they rely on the walls. Also, X better reflects the coarse geometry of S . For instance, if S is the standard cubulation of $\mathbf{R}^3$, then X is quasiisometric to S , whereas the injective hull is $\mathbf{R}^4$.

Furthermore, the dual space $X_{\mathcal{C}}(S)$ of a strong coarse median space S has a median algebra structure that is uniformly close to the coarse median structure on S (Lemma 7.10). Each *coarsely convex* set $Y \subset S$ (such as quasiconvex subsets of hyperbolic spaces) is at a uniform Hausdorff-distance from a wall-theoretically convex set Z in $X_{\mathcal{C}}(S)$: an intersection of half spaces (Definition 3.8). The median algebra structure on $X_{\mathcal{C}}(S)$ then allows us to produce a *gate* map $g : X_{\mathcal{C}}(S) \rightarrow Z$. This gate map is a closest-point projection in $X_{\mathcal{C}}(S)$ and is 1-Lipschitz. To summarise, one can replace a strong coarse median space S by a space $X_{\mathcal{C}}(S)$ equivariantly quasiisometric to S and with upgraded fine properties.

This highlights a philosophical difference between this article and [HHP23] compared to [DMS23, Dur23]. In the former articles, a given space is extended to a larger one with good properties, whereas in the latter ones, subspaces of a given group are accurately approximated by cube complexes.

As a final instance of the improved fine properties of $X_{\mathcal{C}}(S)$, we mention the following rank-rigidity result for hierarchically hyperbolic groups, which improves upon the coarse product structure of [DHS17, PS23].

Corollary L. *Let G be a hierarchically hyperbolic group. There is a natural choice of $\mathcal{C}$ such that $X_{\mathcal{C}}(G)$ is G -equivariantly quasiisometric to G and either:*

- G contains a Morse element; or
- $X_{\mathcal{C}}(G) = X_{\mathcal{C}_1}(S_1) \times X_{\mathcal{C}_2}(S_2)$, where S_1 and S_2 are unbounded hierarchically hyperbolic spaces and $\mathcal{C} = \mathcal{C}_1 \sqcup \mathcal{C}_2$.

The article is divided into two parts. The first part focuses on general properties of Construction 1.1, and natural conditions one can consider on $\mathcal{C}$. The intention is to be self-contained, with a view to being applicable more widely. The outcome is a number of recipes that, given collections of walls satisfying certain combinatorial conditions, will produce associated spaces with good properties.

In the second part, we apply these generalities in two main settings, in order to deduce the results discussed above. In a short final section we point to some potential further directions that could be pursued. The appendix with Davide Spriano concerns quasimorphisms.

Acknowledgements. We are happy to thank Stefanie Zbinden for friendly discussions about her work and ours. We thank Kasra Rafi for conversations about Teichmüller space, and Matt Durham, Anthony Genevois, and Mark Hagen for comments on an earlier version of this article. We are grateful to the organisers of the thematic program Geometric Group Theory in Montreal in 2023, where this work was started.

2. BACKGROUND ON MEDIANS AND CAT(0) CUBE COMPLEXES

An authoritative general reference for the material discussed here is [Bow22]. A *median algebra* is a set X together with a ternary operation μ such that for all $a, b, c, d, e \in X$ we have $\mu(a, b, c) = \mu(b, a, c) = \mu(c, a, b)$, $\mu(a, a, b) = a$, $\mu(a, b, \mu(c, d, e)) = \mu(\mu(a, b, c), \mu(a, b, d), e)$.

The latter equality is called the *five-point condition*. (See also [Rol98, BH83].) One family of median algebras is provided by *median graphs*. A graph X is median if for each three vertices v_1, v_2, v_3 there is a unique point lying on some geodesic from v_i to v_j for all i, j . By a result of Chepoi [Che00, Thm 6.1], a graph is median if and only if it is the 1-skeleton of some CAT(0) cube complex (see also [Gen23]).

CAT(0) cube complexes can equivalently be characterised in terms of their *hyperplanes*. (See [Wis21] for thorough information on cubical hyperplanes.) Indeed, *Sageev's construction* [Sag95], as clarified in [Nic04, CN05], shows how to reconstruct the one-skeleton from the combinatorics of the hyperplanes. Each hyperplane h has two corresponding *halfspaces*, h^- and h^+ , which are the components of its complement.

The median μ on a CAT(0) cube complex Q can be described in terms of hyperplanes as follows: given vertices $x_1, x_2, x_3 \in Q$, the point $\mu(x_1, x_2, x_3)$ is the unique vertex that, for every hyperplane h , lies on the same side of h as the majority of the x_i .

A subcomplex C of a CAT(0) cube complex Q is *convex* if any of the following equivalent conditions hold:

- $\mu(c_1, c_2, x) \in C$ for all $c_1, c_2 \in C$, $x \in Q$ (this is called *median-convexity*);
- C is the largest subcomplex contained in the intersection of some collection of halfspaces;
- C is (geodesically) convex with respect to either the CAT(0) or the ℓ^1 metric on Q .

The *convex hull*, $\text{Hull}_Q A$, of $A \subset Q$ is the smallest convex subcomplex of Q containing A .

Given a convex subcomplex C of a CAT(0) cube complex Q , there is a natural projection map $\mathbf{g}_C : Q \rightarrow C$, called the *gate* map. Here are three equivalent descriptions of the gate of a vertex $x \in Q$:

- $\mathbf{g}_C(x)$ is the unique closest point in C to x when Q is equipped with the combinatorial metric;
- $\mathbf{g}_C(x)$ is the unique vertex of Q such that a hyperplane h separates x from $\mathbf{g}_C(x)$ if and only if h separates x from C ;
- $\mathbf{g}_C(x)$ is the unique point of C with the property that $\mu(c, \mathbf{g}_C(x), x) = \mathbf{g}_C(x)$ for all $c \in C$.

Combining the first and third characterisations explains the terminology: every point $c \in C$ has a geodesic to x that passes through $\mathbf{g}_C(x)$.

The following technical lemma will not be needed until Section 7.

Lemma 2.1. *Let Q be a CAT(0) cube complex, and let $a_1, \dots, a_n \in Q$, for $n \geq 2$. The gate map $Q \rightarrow \text{Hull}_Q(a_1, \dots, a_n)$ can be expressed as $x \mapsto \mu(a_n, x, \mu(a_{n-1}, x, \mu(\dots, \mu(a_2, x, a_1) \dots)))$.*

Proof. Write $\mathbf{g}'(x)$ for the map in the statement of the lemma. Starting at the “inner level” of the expression for $\mathbf{g}'(x)$, observe that $\mu(a_2, x, a_1) \in \text{Hull}_Q(a_1, a_2) \subset \text{Hull}_Q(a_1, \dots, a_n)$. Proceeding outwards one level at a time, we see inductively that $\mathbf{g}'(x) \in \text{Hull}_Q(a_1, \dots, a_n)$.

Suppose that a hyperplane h of Q has $x \in h^-$, $\mathbf{g}'(x) \in h^+$. Consider the “outer level” of the expression for $\mathbf{g}'(x)$. The majority of the arguments of that median must lie on the same side of h as $\mathbf{g}'(x)$, so we must have $a_n \in h^+$, and also the nested expression must determine a point of h^+ . The same argument then applies at the next level of the expression to give $a_{n-1} \in h^+$, and successively we find that $a_i \in h^+$ for all i .

We have shown that the only hyperplanes separating x from the point $\mathbf{g}'(x) \in \text{Hull}_Q(a_1, \dots, a_n)$ are those that separate x from $\text{Hull}_Q(a_1, \dots, a_n)$ itself. By the second characterisation above, $\mathbf{g}'(x)$ is the gate of x to $\text{Hull}_Q(a_1, \dots, a_n)$. $\square$

PART 1. GENERALISING SAGEEV’S CONSTRUCTION

3. ULTRAFILTERS AND DUALISABLE SYSTEMS

In this section we introduce the core framework within which we operate.

A *set with walls* is a pair (S, P) , where S is a set and P is a set of bipartitions $h = \{h^+, h^-\}$ of S . That is, $h^-, h^+ \subset S$ have $S = h^- \cup h^+$ and $h^- \cap h^+ = \emptyset$. We refer to h as a *wall*, and to $h^\pm$ as the *halfspaces* of h .

Definition 3.1 (Ultrafilter). A *filter* ϕ on P consists of a subset $Q \subset P$ and a choice of halfspace $\phi(h) \in \{h^+, h^-\}$ for each $h \in Q$ such that:

if $h_1, h_2 \in Q$ have $h_1^+ \subset h_2^+ \subset S$, then $\phi(h_1) = h_1^+$ implies that $\phi(h_2) = h_2^+$.

We say that ϕ is *supported* on Q . An *ultrafilter* is a filter whose support is P .

The terminology comes from the fact that one can see these filters as being filters on the Boolean subalgebra of 2^S generated by the halfspaces of S . In a Boolean algebra every filter extends to an ultrafilter, giving the following.

Lemma 3.2 (Ultrafilter lemma). *Every filter on P extends to an ultrafilter.*

Space of ultrafilters. Write $\hat{X}$ for the set of ultrafilters on S defined by P . Every point $s \in S$ defines an ultrafilter by setting $\phi_s(h)$ to be the halfspace of h that contains s , for every $h \in P$. If each pair of points in S is separated by some element of P , then $s \mapsto \phi_s$ is injective.

Even when this is not the case we tend to abuse notation and fail to distinguish between S and its image in $\hat{X}$.

Another way of casting Lemma 3.2 is that it says that any collection of halfspaces that intersect pairwise in S has nonempty total intersection in $\hat{X}$: it is a kind of Helly property.

We say that $h_1, h_2 \in P$ *cross* if all four orientations of h_1 and h_2 are filters. Equivalently, all four quarterspaces $h_1^\pm \cap h_2^\pm$ are nonempty, in either S or $\hat{X}$. Equivalently, there is no pair of orientations, say h_1^+ and h_2^+ , such that $h_1^+ \subset h_2^+$ as subsets of S .

3.1. DUALISABLE SYSTEMS

Let $h \in P$. If s and t are points of S that lie in different halfspaces of h , then we say that h *separates* s from t . More generally, given $x, y \in \hat{X}$, we say that h *separates* x from y if $x(h) \neq y(h)$.

So far, the set $\hat{X}$ is both extremely large (in general) and lacking any metric structure. The following definition lets us deal with both of these issues simultaneously.

Definition 3.3. A *dualisable system* for (S, P) is a subset $\mathcal{C} \subset 2^P$ such that the following hold.

- $\mathcal{C}$ contains all singletons and is closed under taking subsets.
- For each pair $s, t \in S$, there is a number M_{st} such that $|c| \leq M_{st}$ for all $c \in \mathcal{C}$ with the property that every $h \in c$ separates s from t .

The dual space. Given a dualisable system $\mathcal{C}$ for (S, P) , consider the function on $\hat{X} \times \hat{X}$ given by

$$\mathrm{d}_{\mathcal{C}}(x, y) = \sup\{|c| : c \in \mathcal{C} \text{ and every } h \in c \text{ separates } x \text{ from } y\}.$$

The $\mathcal{C}$ -dual of S is the space $X_{\mathcal{C}} = \{x \in \hat{X} : \mathrm{d}_{\mathcal{C}}(x, s) < \infty \text{ for all } s \in S\}$, equipped with the function $\mathrm{d}_{\mathcal{C}}$.

Lemma 3.4. $\mathrm{d}_{\mathcal{C}}$ is an extended metric. Its restrictions to $X_{\mathcal{C}}$ and S (strictly speaking its image in $\hat{X}$) are metrics.

Proof. Because $\mathcal{C}$ contains all singletons, the function $\mathrm{d}_{\mathcal{C}}$ separates points of $\hat{X}$. It is evidently symmetric. The fact that the restrictions take only finite values follows from the second bullet of Definition 3.3. It remains to show that $\mathrm{d}_{\mathcal{C}}$ satisfies the triangle inequality.

Let $x, y, z \in \hat{X}$. Every $c \in \mathcal{C}$ separating x from y can be partitioned as $c = c_x \sqcup c_y$, where every element of c_x separates x from z and every element of c_y separates y from z . By definition, $\mathrm{d}_{\mathcal{C}}(x, z) \geq |c_x|$ and $\mathrm{d}_{\mathcal{C}}(z, y) \geq |c_y|$. Since this holds for every c separating x from y , we have $\mathrm{d}_{\mathcal{C}}(x, y) \leq \mathrm{d}_{\mathcal{C}}(x, z) + \mathrm{d}_{\mathcal{C}}(z, y)$. $\square$

By a *chain* in P , we mean a sequence $c = (h_i)_{i \in I}$, where I is some (finite or infinite) interval in $\mathbf{Z}$ and $h_i \in P$, such that $h_{i-1}^- \subset h_i^- \subset h_{i+1}^-$ for all i .

Example 3.5. The conditions under which Sageev's construction applies are exactly equivalent to the statement that 2^P is a dualisable system. If we then take $\mathcal{C} = 2^P$, then $X_{\mathcal{C}}$ is exactly the dual CAT(0) cube complex of Sageev. If $X_{\mathcal{C}}$ contains no infinite cubes, then taking $\mathcal{C}'$ to be the set of all chains in P we get that $X_{\mathcal{C}'}$ is exactly the Helly graph obtained from $X_{\mathcal{C}}$ by *thickening* each cube (i.e. replacing it by a complete graph) [BV91]. Otherwise $X_{\mathcal{C}'}$ will be bigger. Indeed, if $S = (\bigoplus_{\mathbf{N}} \{0, 1\}, \ell^1)$ is an infinite cube and P is the set of cubical walls, then $X_{\mathcal{C}} = S$, whereas $X_{\mathcal{C}'}$ is metrically an uncountable clique.

The set of all chains can be a dualisable system even if 2^P fails to be. For instance, in the above example, $\mathcal{C}'$ is dualisable for $(X_{\mathcal{C}'}, P)$, whereas $\mathcal{C}$ is not. The following simple example

illustrates how configurations of this type can arise naturally in continuous settings. Let $S = \mathbf{C}$ and let P be the set of walls induced by taking all lines through 0. The points 1 and -1 are separated by uncountably many walls, but no two walls form a chain.

Recall that a subset Y of a median algebra $\hat{X}$ is said to be *median-convex* if for any $x, y \in Y$ and $z \in \hat{X}$, we have $\mu(x, y, z) \in Y$.

Lemma 3.6 (Median algebra). *The set P gives $\hat{X}$ the structure of a median algebra. If $\mathcal{C}$ is a dualisable system, then $X_{\mathcal{C}}$ is median-convex. In particular, $X_{\mathcal{C}}$ is itself a median algebra.*

Proof. Given three ultrafilters $x_1, x_2, x_3 \in \hat{X}$, define an orientation ϕ of each $h \in P$ by setting $\phi(h)$ to be the halfspace that contains the majority of the x_i . This is clearly an ultrafilter on P , and it is straightforward to check that this assignment $\mu : (x_1, x_2, x_3) \mapsto \phi$ satisfies the five-point condition.

If $x_1, x_2 \in X_{\mathcal{C}}$ and $z \in \hat{X}$, then x_2 is at finite $d_{\mathcal{C}}$ -distance from x_1 . If $c \in \mathcal{C}$ separates x_1 from $\mu(x_1, x_2, x_3)$ then it separates x_1 from $\{x_2, x_3\}$. In particular $d_{\mathcal{C}}(x_1, \mu(x_1, x_2, x_3)) \leq d_{\mathcal{C}}(x_1, x_2)$ is finite. $\square$

The median defined in the proof of Lemma 3.6, which is also known as the “majority vote” median, will be denoted $\mu : \hat{X}^3 \rightarrow \hat{X}$ throughout.

Although it is both a median algebra and a metric space, in general $X_{\mathcal{C}}$ will not be a *median metric space*, as can be seen from Example 3.5. Note the following simple observation.

Lemma 3.7. *If a group G acts on S and preserves the dualisable system $\mathcal{C}$, then G acts by isometries on $X_{\mathcal{C}}$.*

3.2. CONVEXITY AND GATES

Given a dualisable system $\mathcal{C}$ on (S, P) , aside from the notion of median-convexity on $X_{\mathcal{C}}$ that is available thanks to Lemma 3.6, there is another version that is perhaps more natural.

Definition 3.8 (P -convex). Let $\mathcal{C}$ be a dualisable system for (S, P) . A subset $\hat{C} \subset \hat{X}$ is P -convex if there is a subset $Q \subset P$ and a choice $\hat{C}(h) \in \{h^+, h^-\}$ for all $h \in Q$ such that $\hat{C} = \{v \in \hat{X} : v(h) = \hat{C}(h) \text{ for all } h \in Q\}$. In other words, $\hat{C}$ is an intersection of P -halfspaces in $\hat{X}$.

Note that $\hat{C}$ is nonempty if and only if the choices $\hat{C}(h)$ define a filter with support Q .

Remark 3.9 (P -convexity vs median-convexity). Whilst it is true that every P -convex set in $\hat{X}$ is also median-convex with respect to μ , the converse can fail in general. For instance, the set $\bigoplus_{\mathbf{N}}\{0, 1\} \subset \prod_{\mathbf{N}}\{0, 1\}$ is median-convex but is not an intersection of halfspaces. In this example, $\bigoplus_{\mathbf{N}}\{0, 1\}$ is exactly the set of points in $\prod_{\mathbf{N}}\{0, 1\}$ that are at finite ℓ^1 distance from the zero sequence. Another, perhaps simpler, example appears in Remark 3.12.

We now define a projection map to P -convex sets inside $\hat{X}$.

Proposition 3.10 (Gates). *Given a nonempty P -convex set $\hat{C} \subset \hat{X}$ and an ultrafilter $z \in \hat{X}$, let w be the orientation of all elements of P obtained from z by switching the orientation of every element of P that separates z from $\hat{C}$. The orientation w is an ultrafilter, and $w \in \hat{C}$.*

Proof. Let $Q \subset P$ be the subset witnessing the P -convexity of $\hat{C}$. Note that $z \in \hat{C}$ if and only if no $h \in Q$ separates z from $\hat{C}$, in which case $w = z$. Otherwise, it is clear from the construction that if w is an ultrafilter then it lies in $\hat{C}$. It therefore suffices to show that w

is a filter, because its support is P . That is, supposing that $h_1, h_2 \in P$ satisfy $h_1^+ \subset h_2^+$, we must show that if $w(h_1) = h_1^+$, then $w(h_2) = h_2^+$.

First suppose that h_1 does not separate z from $\hat{C}$. By Lemma 3.2 there must be some $c \in \hat{C}$ such that $c(h_1) = z(h_1)$. From the construction of w , we have $c(h_1) = z(h_1) = w(h_1) = h_1^+$. Since c and z are ultrafilters, $c(h_2) = h_2^+ = z(h_2)$, so h_2 also does not separate z from $\hat{C}$. We therefore have $w(h_2) = z(h_2) = h_2^+$, as desired.

Alternatively, h_1 separates z from $\hat{C}$. Because $\hat{C}$ is nonempty, there exists $c \in \hat{C}$. We have $c(h_1) \neq z(h_1)$ and $z(h_1) \neq w(h_1) = h_1^+$. Because c is an ultrafilter, we therefore have $c(h_2) = h_2^+$. If h_2 separates z from $\hat{C}$, then $c(h_2)$ disagrees with $z(h_2)$, which disagrees with $c(h_2)$, so $w(h_2) = h_2^+$. Otherwise, h_2 does not separate z from $\hat{C}$, and so $c(h_2)$ agrees with $z(h_2)$, which agrees with $w(h_2)$. Thus we have $w(h_2) = h_2^+$ in either case. $\square$

A subset A of a median algebra (M, m) is called *gated* if for every $x \in M$ there exists a unique $g \in A$ such that $m(a, g, x) = g$ for all $a \in A$. In view of the following lemma, we shall call the ultrafilter w constructed in Proposition 3.10 the *gate* of z to $\hat{C}$, and write $w = \mathfrak{g}_{\hat{C}}(z)$.

Lemma 3.11. *A subset $A \subset \hat{X}$ is gated if and only if it is P -convex.*

Proof. It is easy to check that if A is P -convex, then the map $\mathfrak{g}_A$ makes A gated. Conversely, suppose that $A \subset \hat{X}$ is gated. For each $z \in \hat{X} \setminus A$, let g_z be its gate to A , and set $P_z = \{h \in P : h \text{ separates } z \text{ from } g_z\}$. If $h \in P_z$ for some z , then orient it so that $z \in h^+$, $g_z \in h^-$. The set A must be contained in h^- , for otherwise we could find $a \in A \cap h^+$, and then $\mu(a, z, g_z) \in h^+$ cannot be g_z . Hence A is contained in an intersection of P -halfspaces in $\hat{X}$. But we showed that any wall separating a point from its gate must separate that point from A , so in fact A is equal to that intersection of P -halfspaces. $\square$

Remark 3.12 (Median algebras, convexity, and gatedness). Let (M, m) be a median algebra. Recall that a subset $B \subset M$ is *median-convex* if $m(b, b', x) \in B$ for all $b, b' \in B$, $x \in M$. Considering the family of maps $m(b, b', \cdot) : M \rightarrow M$ as b, b' vary over B , we can interpret median-convexity of B as saying that its two-point hulls are gated. Via Lemma 2.1, the four-point condition then implies that B is “finitely gated”: the hull of every finite subset is gated. In the discrete, finite-rank case, median-convexity and gatedness agree, but in general, this perspective shows that gatedness is perhaps more natural, because finitely-gated need not imply gated.

These notions give two families of walls that one can consider on a given median algebra M . The more common choice is the set of all *median-convex* walls: bipartitions $\{h^+, h^-\}$ where both h^+ and h^- are nonempty median-convex subsets [Rol98, Fio20]. Alternatively, one can consider *gated* walls, where h^+ and h^- are required to be nonempty and gated. In the case where M is a *Stone* median algebra (i.e. a compact and totally disconnected topological median algebra, see [Bow22, §12.5]), one can see that a wall is gated exactly when it is a *clopen* wall: a partition into two nonempty open halfspaces. The clopen walls are the ones that are used in duality statements.

A good example to keep in mind is the following. Let $S = \mathbf{Z}$, with P its usual cubical walls. The space $\hat{X}$ is the Roller compactification [Rol98]. It is a topological median algebra whose points can be thought of as living in $\mathbf{Z} \cup \{-\infty, \infty\}$. There are median-convex walls in $\hat{X}$ not coming from P . For instance, consider $h^+ = \{\infty\}$, $h^- = \hat{X} \setminus h^+$. This defines a median-convex wall of $\hat{X}$, but note that it is not clopen: the halfspace h^+ is not open. Correspondingly, the halfspace h^- is not gated. The clopen walls of $\hat{X}$ are exactly the gated walls, which are exactly those coming from P . This is part of what makes the P -convexity

in Definition 3.8 the appropriate notion here, even though S itself can sometimes fail to be P -convex in $\hat{X}$.

To clarify, if (S, P) is a set with walls, then the term “wall” will always mean an element of P , even though the median algebra $\hat{X}$ may have additional median-convex walls.

Lemma 3.13. *If $C = X_C \cap \hat{C}$ is nonempty, where $\hat{C}$ is some P -convex subset of $\hat{X}$, then $\mathbf{g}_{\hat{C}}(z) \in C$ for all $z \in X_C$.*

Proof. Let $x \in C$. Any $z \in X_C$ lies at finite distance from x . As $\mathbf{g}_{\hat{C}}(z)$ is obtained from z by switching the orientations of a subset of the walls separating z from x , every wall separating $\mathbf{g}_{\hat{C}}(z)$ from x must separate z from x . Hence $\mathbf{g}_{\hat{C}}(z)$ is at finite distance from x , and so lies in both X_C and $\hat{C}$. $\square$

In view of Proposition 3.10 and Lemmas 3.11 and 3.13, a subset $C \subset X_C$ is gated in the median algebra X_C exactly when there exists some P -convex set $\hat{C} \subset \hat{X}$ such that $C = X_C \cap \hat{C}$. We write $\mathbf{g}_C(z) = \mathbf{g}_{\hat{C}}(z)$ in this case when $z \in X_C$.

Both of the next two lemmas could equally well be stated for P -convex subsets of $\hat{X}$, but that is not to our purposes. The following is immediate from the construction of the gate and the definition of $\mathbf{d}_C$.

Lemma 3.14 (Closest-point). *Let $C \subset X_C$ be gated, and let $x \in X_C$. If $h \in P$ separates x from $\mathbf{g}_C(x)$, then h separates x from C . In particular, $\mathbf{d}_C(x, c) \geq \mathbf{d}_C(x, \mathbf{g}_C(x))$ for all $c \in C$.*

Lemma 3.15 (Lipschitz). *Let $C \subset X_C$ be gated and let $A, B \subset X_C$. If $h \in P$ separates $\mathbf{g}_C(A)$ from $\mathbf{g}_C(B)$, then h separates A from B . In particular, $\mathbf{g}_C$ is 1-Lipschitz.*

Proof. If h separates $\mathbf{g}_C(A)$ from $\mathbf{g}_C(B)$, then there are points of C on both sides of h , so h cannot separate any point of X_C from C . In particular, for all $z \in X_C$ the orientation of h determined by z is the same as that determined by $\mathbf{g}_C(z)$, so A and B are on different sides of h . $\square$

We finish this section with a simple observation about finite group actions. Given a dualisable system $\mathcal{C}$ for (S, P) , one could define a *first subdivision* of X_C by “doubling” each $h \in P$ into two identical partitions h_1, h_2 and leaving all other crossing relations the same. We refrain from making this formal, because we generally wish to work with P being a set of bipartitions, rather than a multiset, but it is completely analogous to the first cubical subdivision of a CAT(0) cube complex. After simply noting the following, subdivisions will not be mentioned again.

Proposition 3.16. *Let $\mathcal{C}$ be a dualisable system on (S, P) . If G is a finite group acting on S that preserves P , then G fixes a point in the first subdivision of X_C .*

Proof. Let A be a G -orbit in S . For each $h \in P$, let h^+ be the halfspace containing more than half of the elements of A , if it exists. Let Q be the set of such h . It is easy to see that $\phi(h) = h^+$ is a filter supported on Q , and any ultrafilter extending it lies in X_C , because it is separated from each element of A only by elements of P separating points of A . Also note that ϕ is fixed by G . If $P = Q$ then we are done.

Otherwise, $Q \neq P$. Let Q' be the set of all $h \in P \setminus Q$ such that there is some $g \in G$ for which $gh \neq h$ and gh does not cross h . Note that since h divides A into equal halves, both h and gh define the same partition of A . By definition, there are halfspaces h^- and $(gh)^-$ that are disjoint. We must have $(gh)^- = g(h^-)$, for otherwise we would have $g(h^+) \subsetneq h^+$, which

would imply that g had infinite order. One similarly argues that this labelling did not depend on the choice of g . In other words, the allocation $\phi(h) = h^+$ on Q' is G -invariant.

Let us show that this allocation ϕ is a filter on Q' . If not, then there are $h_1, h_2 \in Q'$ such that h_1^+ and h_2^+ are disjoint. Let $g \in G$ be such that $gh_1 \neq h_1$ and gh_1 does not cross h_1 . We have $h_2^+ \subset h_1^- \subset gh_1^+$, and so we must have $g^{-1}h_2^+ \subset h_1^+$. But this shows that h_2^+ is disjoint from its translate by g^{-1} , contradicting G -invariance of ϕ . Hence if $P = Q \cup Q'$, then ϕ is an ultrafilter that is fixed by G .

The final case is that there is some $h \in P \setminus Q$ such that for each $g \in G$, either h crosses gh or is equal to it. If $k \in P \setminus (Q \cup Q')$ does not cross h , then set $\phi(k)$ to be the halfspace containing a halfspace of h , and extend to $G \cdot \{k\}$ equivariantly. Let Q_h be the set of elements of $P \setminus (Q \cup Q')$ whose G -translates all cross h . We obtain an ultrafilter extending ϕ in the first subdivision of X_C by pointing the “doubled” copies of the elements of Q_h towards each other, and ϕ is fixed by G . $\square$

The corresponding statement does not always hold for *bounded* group actions, as can be seen from a transitive action of $\mathbf{Z}$ on an infinite clique.

4. SYSTEMS OF CHAINS

In the applications we have in mind for the constructions of Section 3, the dualisable system $\mathcal{C}$ will be a set of chains in P . Recall that by a chain in P , we mean a sequence $c = (h_i)_{i \in I}$, where I is some (finite or infinite) interval in $\mathbf{Z}$ and $h_i \in P$, such that $h_{i-1}^- \subset h_i^- \subset h_{i+1}^-$ for all i .

Definition 4.1 (System of chains). A dualisable system $\mathcal{C}$ is a *system of chains* if every element of $\mathcal{C}$ is a chain in P .

An immediate advantage of assuming that elements of $\mathcal{C}$ are chains is that every element of $\mathcal{C}$ then comes with a total order, so that one can speak of the *minimal* and *maximal* element of a finite chain $c \in \mathcal{C}$. In this section, we investigate some of the metric properties that one deduce about the dual space X_C when $\mathcal{C}$ is a system of chains.

Lemma 4.2. *If $\mathcal{C}$ is a system of chains on (S, P) , then every ball in X_C is gated.*

Proof. Let $x \in X_C$ and $r \geq 0$. Since d_C is integer-valued, we may assume that $r \in \mathbf{Z}$. Let Q be the set of all $h \in P$ such that there exists some $c \in \mathcal{C}$ with $|c| \geq r$ that separates x from h . For each $h \in Q$, let $\hat{C}(h) = x(h)$ be the halfspace containing x . This defines a P -convex set $\hat{C} \subset \hat{X}$. If $z \in \hat{C}$, then because $\mathcal{C}$ is a system of chains, we have $d_C(x, z) \leq r$. In particular, $\hat{C} \subset X$. Conversely, if $d_C(x, z) \leq r$, then no element of Q can separate z from x , so $z \in \hat{C}$. Thus the P -convex set $\hat{C}$ is the r -ball about x in X . But $\hat{C} \subset X_C$, so it is also the r -ball about x in X_C . $\square$

For systems of chains, we now describe a family of paths that are analogues of the *normal cube paths* introduced by Niblo–Reeves in [NR98]. In fact, in the case where 2^P is dualisable (see Example 3.5) and $\mathcal{C}$ is the set of all chains, X_C is a CAT(0) cube complex with the ℓ^∞ metric and the paths we shall define will exactly be normal cube paths.

Definition 4.3 (Normal wall path). Let $\mathcal{C}$ be a system of chains. Given $x, y \in X_C$, the *normal wall path* $\sigma(x, y) = \sigma_{xy}$ from x to y is the sequence

$$(x = \mathfrak{g}_{B(x,0)}(y), \mathfrak{g}_{B(x,1)}(y), \dots, \mathfrak{g}_{B(x,n)}(y) = y),$$

where $B(x, r)$ denotes the r -ball in X_C centred on x , which is gated by Lemma 4.2. See Figure 1.




FIGURE 1. $\sigma(x, y)$ is constructed by gating y to integer balls centred on x .

As with normal cube paths in CAT(0) cube complexes, $\sigma(x, y)$ need not agree with $\sigma(y, x)$.

By the definition of the gate map, $\sigma_{xy}(r)$ is obtained from y by switching the orientations of exactly the walls separating y from $B(x, r)$. Hence if $r_1 < r_2 < r_3$, then $\sigma_{xy}(r_2)$ is equal to its median with $\sigma_{xy}(r_1)$ and $\sigma_{xy}(r_3)$. In other words, σ_{xy} is a *median path* and cannot cross any wall twice. The following lemma can be viewed as a strengthening of this observation.

Lemma 4.4. *Let $x, y \in X_{\mathcal{C}}$. If $z \in [\sigma_{xy}(r), y]$, then $\sigma_{xz}(t) = \sigma_{xy}(t)$ for all $t \leq r$. In particular, $\sigma(x, \sigma_{xy}(r)) \subset \sigma(x, y)$.*

Proof. The point $\sigma_{xy}(r)$ is obtained from y by switching the orientations of exactly those walls that separate y from the ball $B(x, r)$. Thus the fact that $z \in [\sigma_{xy}(r), y]$ means both that the walls separating z from $B(x, r)$ are a subset of those separating y from $B(x, r)$, and that every wall separating x from $\sigma_{xy}(r)$ separates x from z . Hence no wall separates $\sigma_{xz}(r)$ from $\sigma_{xy}(r)$, so the two are equal. The rest follows. $\square$

An extremely useful property that a system of chains can have is the ability to combine certain elements of $\mathcal{C}$ to obtain a larger element of $\mathcal{C}$. This allows one to make certain constructions piecewise, which is often necessary when several points are involved. If c is a chain in P with maximal element h , then we write c^+ to mean the subset h^+ of S . If k is the minimal element of c , then we similarly write c^- for the subset k^- of S .

Definition 4.5 (Gluable). Let $\mathcal{C}$ be a system of chains. We say that $\mathcal{C}$ is *m-gluable* if the following holds. Suppose that $c_1, c_2 \in \mathcal{C}$ are such that $c_1 \cup c_2$ is a chain, $c_2 \subset c_1^+$, and $c_1 \subset c_2^-$. Write $c_1 = \{\dots, h_{-2}, h_{-1}\}$ and $c_2 = \{h_1, h_2, \dots\}$. There is a subset $b \subset \{h_{-m}, \dots, h_{-1}, h_1, \dots, h_m\}$ of consecutive halfspaces, of size at most m , such that $(c_1 \cup c_2) \setminus b$ is an element of $\mathcal{C}$.

We shall primarily be interested in *m-gluable* systems with $m \leq 3$. Note that the only 0-gluable system of chains in P containing all singletons is the set of all chains in P .

A *discrete geodesic* in a metric space (Y, d) is the image of an isometric embedding of an interval in $\mathbf{Z}$. A *k-rough geodesic* is the image of a $(1, k)$ -quasiisometric embedding of an interval in $\mathbf{Z}$. A *k-weak rough geodesic* from x to y is a sequence $(z_r)_{r=0}^n$ with $z_0 = x$ and $z_n = y$ such that $|d(x, z_r) - r| \leq k$ and $|d(z_r, y) - (n - r)| \leq k$, where $n = d(x, y)$.

Proposition 4.6. *If $\mathcal{C}$ is an m-gluable system of chains, then normal wall paths are m-weak rough geodesics and 3m-rough geodesics.*

More precisely, if $x, y \in X_{\mathcal{C}}$ have $d_{\mathcal{C}}(x, y) = n$, then for any $r \in [0, n]$ we have $r - m \leq d_{\mathcal{C}}(x, \sigma_{xy}(r)) \leq r$ and $n - r \leq d_{\mathcal{C}}(\sigma_{xy}(r), y) \leq n - r + m$.

Proof. Let us write $z_r = \sigma_{xy}(r)$. As in the proof of Lemma 4.2, let Q be the set of all $h \in P$ such that there exists some $c \in \mathcal{C}$ with $|c| \geq r$ that separates x from h .

First we control $d_{\mathcal{C}}(x, z_r)$. By definition, z_r lies in the r -ball about x , so we just need to establish the lower-bound. Recall that z_r is obtained from y by switching the orientations of exactly the elements of Q that separate x from y . Let $\{h_1, \dots, h_n\} \in \mathcal{C}$ realise $d_{\mathcal{C}}(x, y)$. If h_{r-m} does not separate z_r from x , then it separates z_r from y , so we must have $h_{r-m} \in Q$. Let $\{k_1, \dots, k_r\} \in \mathcal{C}$ separate x from h_{r-m} . Because $\mathcal{C}$ is m -gluable, there is then a subset of $\{k_1, \dots, k_r, h_{r-m}, \dots, h_n\}$ of size at least $n+1$ that is an element of $\mathcal{C}$. But this contradicts the fact that $d_{\mathcal{C}}(x, y) = n$. Hence h_{r-m} separates z_r from x , and in particular $d_{\mathcal{C}}(x, z_r) \geq r - m$.

Now we control $d_{\mathcal{C}}(z_r, y)$. Because $d_{\mathcal{C}}(x, z_r) \leq r$, the triangle inequality gives $d_{\mathcal{C}}(z_r, y) \geq n - r$. Now let $\{h'_1, \dots, h'_p\} \in \mathcal{C}$ realise $d_{\mathcal{C}}(z_r, y)$. Every element of P separating z_r from y is in Q , so in particular there is some $\{k'_1, \dots, k'_r\} \in \mathcal{C}$ separating x from h'_1 . By m -gluability of $\mathcal{C}$, we get $n = d_{\mathcal{C}}(x, y) \geq r + p - m$, and hence $d_{\mathcal{C}}(z_r, y) = p \leq (n - r) + m$. This in particular shows that (z_r) is an m -weak rough geodesic.

To complete the proof, let $r_1 < r_2$. By Lemma 4.4, $\sigma(x, z_{r_2}) \subset \sigma(x, y)$, so using the above inequalities, we compute

$$\begin{aligned} d_{\mathcal{C}}(x, y) &= r_2 + (n - r_2 + m) - m \\ &\geq d_{\mathcal{C}}(x, z_{r_2}) + d(z_{r_2}, y) - m \\ &\geq d_{\mathcal{C}}(x, z_{r_1}) + d_{\mathcal{C}}(z_{r_1}, z_{r_2}) + d_{\mathcal{C}}(z_{r_2}, y) - 2m \\ &\geq r_1 + d_{\mathcal{C}}(z_{r_1}, z_{r_2}) + n - r_2 - 3m, \end{aligned}$$

which shows that $d_{\mathcal{C}}(z_1, z_2) \leq r_2 - r_1 + 3m$. A similar computation shows that $d_{\mathcal{C}}(z_{r_1}, z_{r_2}) \geq r_2 - r_1 - 3m$, and hence $\sigma(x, y)$ is a $3m$ -rough geodesic. $\square$

A particular case of Proposition 4.6 is that if $\mathcal{C}$ is the set of all chains, then $X_{\mathcal{C}}$ is discretely geodesic. In other words, it is the vertex set of a graph.

A *bicombing* σ of a metric space Y is the choice of a path $\sigma(x, y)$ from x to y for each $x, y \in Y$. This notion is very general, so one often speaks of a bicombing by, say, geodesics, where the $\sigma(x, y)$ are required to be geodesics; or one imposes certain fellow-travelling conditions on the various paths in σ .

Proposition 4.7. *If $\mathcal{C}$ is an m -gluable system of chains, then normal wall paths form a bicombing of $X_{\mathcal{C}}$ by rough geodesics. Moreover, for every r we have*

$$d_{\mathcal{C}}(\sigma_{x_1 y_1}(r), \sigma_{x_2 y_2}(r)) \leq \max\{d_{\mathcal{C}}(x_1, x_2), d_{\mathcal{C}}(y_1, y_2)\} + 3m.$$

Proof. Let us write $d_{\mathcal{C}}(x_1, x_2) = R_x$ and $d_{\mathcal{C}}(y_1, y_2) = R_y$. After relabelling, we may assume that $d_{\mathcal{C}}(x_1, y_1) \leq d_{\mathcal{C}}(x_2, y_2)$, which in turn is at most $d_{\mathcal{C}}(x_1, y_1) + R_x + R_y$ by the triangle inequality. Let us write $z_r^i = \sigma_{x_i y_i}(r)$.

Given $r \leq d_{\mathcal{C}}(x_2, y_2)$, let $c \in \mathcal{C}$ realise $d_{\mathcal{C}}(z_r^1, z_r^2)$. Let $c_x \subset c$ be the subchain consisting of all elements that separate x_1 from x_2 , of which there are at most R_x . Define c_y similarly, and let $c' = c \setminus (c_x \cup c_y)$, which has $|c \setminus c'| \leq R_x + R_y$. Because $z_r^i \in [x_i, y_i]$, every $h \in P$ that separates z_r^i from either x_i or y_i also separates x_i from y_i . Thus every element of c' either separates $\{z_r^1, x_1, x_2\}$ from $\{z_r^2, y_2, y_1\}$ or separates $\{z_r^1, y_1, y_2\}$ from $\{z_r^2, x_2, x_1\}$. Because c is a chain, only one of these options can occur. The two cases are similar, so let us assume the former. See Figure 2.




FIGURE 2. The configuration considered in the proof of Proposition 4.7.

Suppose that $c' \neq \emptyset$. Because c is a chain, the sets c_x and c_y are disjoint. Let h be the minimal element of c' : it separates $c' \setminus \{h\}$ from x_2 . We know that h separates z_r^1 from y_1 . The set $H = \{v \in X_C : v(h) \neq z_r^1(h)\}$ is gated, so we can consider the point $w = \mathbf{g}_H(z_r^1) \in H$. Every wall separating z_r^1 from w separates z_r^1 from H , by Lemma 3.14. Also, since every element of c_x separates z_r^1 from w , we have $\mathbf{d}_C(z_r^1, w) \geq |c_x| + 1$.

As $w \in [z_r^1, y]$, Lemma 4.4 tells us that $\sigma_{x_1 w}(t) = z_t^1$ for $t \leq r$. By Proposition 4.6,

$$\mathbf{d}_C(x_1, w) \geq \mathbf{d}_C(x_1, z_r^1) + \mathbf{d}_C(z_r^1, w) - m \geq r + |c_x| + 1 - 2m.$$

From the fact that $w \in [z_r^1, y] \setminus \{z_r^1\}$, we also have that $\mathbf{d}_C(x_1, w) \geq r + 1$.

Let $b \in \mathcal{C}$ realise $\mathbf{d}_C(x_1, w)$. Because $|b| \geq r + 1$, at least one element of b separates z_r^1 from w . It follows that b separates x_1 from H , so $b \cup (c' \setminus \{h\}) \cup c_y$ is a chain separating x_1 from z_r^2 . By m -gluability of $\mathcal{C}$, there is some element of $\mathcal{C}$ that can be obtained from $b \cup (c' \setminus \{h\}) \cup c_y$ by removing at most m elements. Letting b_x be the subchain of b consisting of all elements separating x_1 from x_2 , which has cardinality at most R_x , we therefore have

$$\begin{aligned} r &\geq \mathbf{d}_C(x_2, z_r^2) \geq |b \cup (c' \setminus \{h\}) \cup c_y| - m - |b_x| \\ &= \mathbf{d}_C(x_1, w) + (|c| - |c_x| - 1) - m - |b_x| \\ &\geq (r + |c_x| + 1 - 2m) + |c| - |c_x| - 1 - m - R_x = r + |c| - 3m - R_x. \end{aligned}$$

This implies that $|c| \leq R_x + 3m$. We have shown that if $c' \neq \emptyset$, then $\mathbf{d}_C(z_r^1, z_r^2) \leq R_x + 3m$.

Now suppose that $c' = \emptyset$ but $|c| > m + \max\{R_x, R_y\}$. As $c = c_x \cup c_y$, each element of c either separates x_1 from x_2 or separates y_1 from y_2 (possibly both). Since $|c| > R_y$, at least one element of c lies in $c_x \setminus c_y$. Any such element h must either separate x_1 from $\{x_2, y_1, y_2\}$ or x_2 from $\{x_1, y_1, y_2\}$. Since the argument is similar in either case, let us assume that the former holds. Given this, let $h \in c$ be the minimal element of c : it separates z_r^1 from $c \setminus \{h\}$, and necessarily lies in $c_x \setminus c_y$. Similarly to the $c' \neq \emptyset$ case, let $H = \{v \in X : v(h) \neq z_r^1(h)\}$ and let $w = \mathbf{g}_H(z_r^1) \in H$.

Because h separates x_1 from y_1 , we have $w \in [z_r^1, y_1]$, and so $\mathbf{d}_C(x_1, w) > r$. Let $b \in \mathcal{C}$ realise $\mathbf{d}_C(x_1, w)$. Every element of b separates x_1 from H , and thus at least one separates z_r from H . In particular, $b \cup c \setminus \{h\}$ is a chain separating x_1 from z_r^2 . By m -gluability of $\mathcal{C}$, there is some element of $\mathcal{C}$ that can be obtained from $b \cup c \setminus \{h\}$ by removing at most m elements.

We therefore compute

$$\begin{aligned} d_{\mathcal{C}}(x_2, z_r^2) &\geq d_{\mathcal{C}}(x_1, z_r^2) - d_{\mathcal{C}}(x_1, x_2) \\ &\geq (|b| + |c| - 1 - m) - R_x \\ &\geq (r + 1) + (m + 1 + \max\{R_x, R_y\}) - 1 - m - R_x \geq r + 1, \end{aligned}$$

which is a contradiction. Hence $|c| \leq m + \max\{R_x, R_y\}$ if $c' = \emptyset$. $\square$

One cannot expect stronger properties of the bicombing: for normal cube paths in $\text{CAT}(0)$ cube complexes the inequality of Proposition 4.7 is optimal, as $m = 0$ in that case [NR98].

Definition 4.8. A metric space Y is k -coarsely injective if for each collection of balls $B(x_i, r_i)$ with $d(x_i, x_j) \geq r_i + r_j$ for all i, j , there is some point $z \in \bigcap B(x_i, r_i + k)$. A graph is *Helly* if this same property holds with $k = 0$ when one considers only balls of integer radius.

Theorem 4.9 (Coarsely injective). *Let $\mathcal{C}$ be an m -gluable system of chains. If $m > 0$, then $X_{\mathcal{C}}$ is $2m$ -coarsely injective. If $m = 0$, then $X_{\mathcal{C}}$ is a (not necessarily locally finite) Helly graph.*

Proof. Let $\{B_X(x_i, r_i)\}$ be a collection of balls with the property that $r_i + r_j \geq d_{\mathcal{C}}(x_i, x_j)$ for all i, j . If $m = 0$, assume that each r_i is an integer. (As noted above, if $m = 0$ then $\mathcal{C}$ is the set of all chains and $X_{\mathcal{C}}$ is the vertex set of a graph.) For each i , let

$$Q_i = \{w \in P : \text{there is some } c \in \mathcal{C} \text{ separating } x_i \text{ from } w \text{ with } |c| \geq r_i + 2m\}.$$

Define an orientation ϕ of $\bigcup Q_i$ by pointing each $w \in Q_i$ towards x_i .

We must first show that ϕ is well defined. For this, suppose that $w \in Q_i \cap Q_j$. Let $c_i, c_j \in \mathcal{C}$ be given by the definitions of Q_i and Q_j , so that $|c_i| \geq r_i + 2m$ and $|c_j| \geq r_j + 2m$. If w separates c_i from c_j , then $c_i \cup \{w\} \cup c_j$ is a chain. But then by twice applying m -gluability of $\mathcal{C}$, we find some $b \subset c_i \cup \{w\} \cup c_j$ such that $b \in \mathcal{C}$ and $|b| \geq |c_i| + 1 + |c_j| - 2m \geq r_i + r_j + 1$. As b separates x_i from x_j , this contradicts our assumption that $r_i + r_j \geq d_{\mathcal{C}}(x_i, x_j)$. Thus w cannot separate c_i from c_j , and so cannot separate x_i from x_j . This shows that ϕ is well defined.

Next we show that ϕ is a filter. Suppose that we have walls $w_i \in Q_i$ and $w_j \in Q_j$ such that $w_i^+ \subset w_j^+$ and $\phi(w_i) = w_i^+$. Let c_i and c_j be given by the definitions of Q_i and Q_j . If $\phi(w_j) = w_j^-$, then w_j separates c_j from w_i and c_i . With three applications of m -gluability of $\mathcal{C}$, we find some $b \subset c_i \cup \{w_i, w_j\} \cup c_j$ such that $b \in \mathcal{C}$ and $|b| \geq |c_i| + 2 + |c_j| - 3m \geq r_i + r_j + 2$. This contradicts the assumption that $r_i + r_j \geq d_{\mathcal{C}}(x_i, x_j)$. Thus $\phi(w_j) = w_j^+$, which shows that ϕ is a filter.

Let z be an ultrafilter extending ϕ . By construction, for each i there is no element of $\mathcal{C}$ of length greater than $r_i + 2m$ that separates z from x_i . In other words, $z \in \bigcap B_X(x_i, r_i + 2m)$, as desired. $\square$

Theorem 4.9 also implies the existence of a good bicombing on $X_{\mathcal{C}}$, thanks to work of Lang [Lan13]. While less related to the wall structure, it has the advantage of being symmetric and *roughly conical*, and so will not in general be the same as the bicombing by normal wall paths.

5. PRODUCING HYPERBOLIC SPACES

One use for the construction in Section 3 is to produce hyperbolic spaces that can help study a given space. For that we need a source of negative curvature, which is provided by *L-separation*. There may be many different natural hyperbolic spaces that one could produce in this way, and Section 5.2 gives a way to combine them into a single hyperbolic space subsuming them.

5.1. SEPARATED SYSTEMS

Recall that walls h_1 and h_2 are said to cross if all four quarterspaces $h_1^\pm \cap h_2^\pm$ are nonempty in S .

Definition 5.1. We say that a system of chains $\mathcal{C}$ is L -separated if for every $\{h_1, h_2\} \in \mathcal{C}$, if $c \in \mathcal{C}$ is such that every $h \in c$ crosses both h_1 and h_2 , then $|c| \leq L$.

Remark 5.2. In many situations, something stronger than L -separation holds. Indeed, there may be some larger set $\mathcal{D} \subset 2^P$ such that for every $\{h_1, h_2\} \in \mathcal{C}$, if $d \in \mathcal{D}$ has $h \cap h_i \neq \emptyset$ for $i = 1, 2$ and every $h \in d$, then $|d| \leq L$. It may even be that one can take $\mathcal{D}$ to be the set of all chains, as is the case in [Gen20] and [PSZ22]. Whilst these stronger properties may be useful, they are not needed for the arguments in this section.

We make a simple observation regarding elements of a gluable, separated system.

Lemma 5.3 (Gluing). *Suppose that $\mathcal{C}$ is an L -separated, m -gluable system of chains. If $c = \{\dots, h_{-1}\}$ and $c' = \{k_1, \dots\}$ are elements of $\mathcal{C}$ such that*

$$h_{-1}^+ \cap k_1^\pm \neq \emptyset \quad \text{and} \quad h_i^\pm \cap k_1^- \neq \emptyset,$$

for all i, j , then there is a subset $b \subset \{h_{-m-1}, \dots, h_{-1}, k_1, \dots, k_{L+m+1}\}$ of cardinality at most $L + m + 1$ such that $c \cup c' \setminus b \in \mathcal{C}$.

Proof. The wall h_{-2} cannot cross k_{L+1} , for then so would h_{-1} , contradicting L -separation of c . Hence $c \setminus \{h_{-1}\}$ and $c' \setminus \{k_1, \dots, k_{L+1}\}$ satisfy the hypotheses of the m -gluability assumption on $\mathcal{C}$. $\square$

With this in hand, we turn to showing that $X_{\mathcal{C}}$ is hyperbolic.

Definition 5.4 (Cross chain). Let $\mathcal{C}$ be a system of chains. A $\times$ -chain (cross chain) for $x_1, x_2, x_3, x_4 \in X_{\mathcal{C}}$ is a subset $\chi \subset P$ with a decomposition $\chi = \bigsqcup_{i=1}^4 \chi_i$ such that every element of χ_i separates x_i from $\{x_{i+1}, x_{i+2}, x_{i+3}\}$; and such that $\chi_i \cup \chi_j \in \mathcal{C}$ for all i, j .

Note that the notion of $\times$ -chain makes sense even when $\mathcal{C}$ is not separated. However, it is at its most useful in this setting, because maximal $\times$ -chains can be effectively compared with chains realising the distances between the defining points.


Lemma 5.5. *Let $\mathcal{C}$ be an L -separated, m -gluable system of chains. If $c \in \mathcal{C}$ realises $d_{\mathcal{C}}(x_1, x_2)$ and χ is a maximal $\times$ -chain for x_1, x_2, x_3, x_4 , then*

$$|\chi_1| + |\chi_2| - 2(L + m + 1) \leq |\{h \in c : h \text{ does not separate } x_3 \text{ from } x_4\}| \leq |\chi_1| + |\chi_2| + 4(L + m + 1).$$

Proof. Let $r = |\{h \in c : h \text{ does not separate } x_3 \text{ from } x_4\}|$. Let c_1 be the subset of c consisting of those elements that do not separate x_2 from x_3 or x_4 , and define c_2 similarly, so that $r = |c_1| + |c_2|$. By Lemma 5.3, after removing at most $L + m + 1$ elements of each of c_1, c_2, χ_3 , and χ_4 , we obtain a $\times$ -chain. Maximality of χ then implies that $r \leq |\chi_1| + |\chi_2| + 4(L + m + 1)$.

On the other hand, let $c' = c \setminus (c_1 \cup c_2)$ be the subset of c consisting of those elements that separate x_3 from x_4 . After relabelling, every element of c' separates $\{x_1, x_3\}$ from $\{x_2, x_4\}$. From Lemma 5.3 and the fact that $\chi_1 \cup \chi_2 \in \mathcal{C}$, there is some $b \subset \chi_1 \cup c' \cup \chi_2$ such that $b \in \mathcal{C}$ separates x_1 from x_2 and $|b| \geq |\chi_1| + |c'| + |\chi_2| - 2(L + m + 1)$. Since c realises $d_{\mathcal{C}}(x_1, x_2)$, we have $|b| \leq |c|$, so $r \geq |\chi_1| + |\chi_2| - 2(L + m + 1)$. $\square$

A metric space (Y, d) is said to be *four-point hyperbolic* with constant δ if for every $x_1, x_2, x_3, x_4 \in Y$ we have $d(x_1, x_2) + d(x_3, x_4) \leq \delta + \max\{d(x_1, x_3) + d(x_2, x_4), d(x_1, x_4) + d(x_2, x_3)\}$.

FIGURE 3. Lemma 5.5. χ is a union of “corners”.

Proposition 5.6. *If $\mathcal{C}$ is an L -separated, m -gluable set of chains, then the metric $d_{\mathcal{C}}$ is four-point hyperbolic with constant $22(L + m + 1)$.*

Proof. Given $x_1, x_2, x_3, x_4 \in X$, let χ be a maximal $\times$ -chain for them. If $d_{\mathcal{C}}(x_i, x_j)$ and $|\chi_1| + |\chi_2|$ differ by at most $4L + 4m + 5$ for all i, j then we are done. Otherwise we can relabel so that $d_{\mathcal{C}}(x_1, x_2) \geq |\chi_1| + |\chi_2| + 4L + 4m + 6$. Let $c_{12} \in \mathcal{C}$ realise $d_{\mathcal{C}}(x_1, x_2)$, and let c'_{12} be the subchain consisting of those elements that separate x_3 and x_4 . By Lemma 5.5 we have $|c'_{12}| \geq 2$. Perhaps after relabelling x_3 and x_4 , the chain c'_{12} separates $\{x_1, x_3\}$ from $\{x_2, x_4\}$.

Now let $c_{13} \in \mathcal{C}$ realise $d_L(x_1, x_3)$. If $|c_{13}| - |\chi_1| - |\chi_3| > 3L + 2m + 3$, then Lemma 5.5 implies the existence of a subchain c'_{13} of length $L + 1$ whose elements all separate $\{x_1, x_4\}$ from $\{x_2, x_3\}$, and hence cross every element of c'_{12} , which contradicts the assumption that $\mathcal{C}$ is L -separated. A similar argument shows that if $c_{24} \in \mathcal{C}$ realises $d_{\mathcal{C}}(x_2, x_4)$, then $|c_{24}| - |\chi_2| - |\chi_4| \leq 3L + 2m + 3$. In the other direction, we can apply Lemma 5.3 to χ_1 and χ_3 , and similarly to χ_2 and χ_4 . This shows that $d_{\mathcal{C}}(x_1, x_3) + d_{\mathcal{C}}(x_2, x_4)$ and $|\chi|$ differ by at most $2(3L + 2m + 3) + 2(L + m + 1)$.

Now let $c_{34} \in \mathcal{C}$ realise $d_{\mathcal{C}}(x_3, x_4)$, and let c'_{34} be the subchain consisting of those elements that separate x_1 from x_2 . By applying Lemma 5.3 twice to each of the pairs $(c'_{12}, c_{34} \setminus c'_{34})$ and $(c'_{34}, c_{12} \setminus c'_{12})$, one finds that $|c'_{12}|$ and $|c'_{34}|$ can differ by at most $2(L + m + 1)$. Combining with Lemma 5.5, one can then compute that $d_{\mathcal{C}}(x_1, x_2) + d_{\mathcal{C}}(x_3, x_4)$ differs from $|\chi| + 2|c'_{12}|$ by at most $10(L + m + 1)$.

Finally consider the analogously defined c_{14} , c'_{14} , c_{23} , and c'_{23} . Both of these are treated the same, so let us just consider the former. Because c'_{12} separates x_1 from x_4 , we can use Lemma 5.3 to see that $|c'_{14}| \geq |c'_{12}| - 2(L + m + 1)$. If $|c'_{14}| \leq L$, then certainly $|c'_{14}| \leq |c'_{12}| + L$, so suppose otherwise. If c'_{14} separates x_1 from x_3 then its elements all cross c'_{12} , which contradicts the fact that $c'_{12} \in \mathcal{C}$. Hence c'_{14} separates x_1 from x_2 . By Lemma 5.3 we thus have $|c'_{12}| \geq |c'_{14}| - 2(L + m + 1)$. We have shown that $|c'_{14}|$ differs from $|c'_{12}|$ by at most $2(L + m + 1)$.

By a similar argument to above, Lemma 5.5 now lets us show that $d_{\mathcal{C}}(x_1, x_4) + d_{\mathcal{C}}(x_2, x_3)$ differs from $|\chi| + 2|c'_{12}|$ by at most $12(L + m + 1)$. The result follows by combining the various estimates. $\square$

Together with Proposition 4.6, this shows that $X_{\mathcal{C}}$ is hyperbolic in the appropriate sense for non-geodesic spaces (cf. [PSZ22, Prop. A.2]).

Corollary 5.7. *If $\mathcal{C}$ is a separated, gluable system of chains, then $X_{\mathcal{C}}$ is a roughly geodesic hyperbolic space.*

Since four-point hyperbolicity passes to subsets, $S \subset X$ is four-point hyperbolic. However, four-point hyperbolicity is rather weak on its own, so it is interesting to know when (S, d_C) is roughly geodesic. One can also ask the *a priori* stronger question of when S is coarsely dense in X . The following proposition shows that these two properties are in fact equivalent.

Proposition 5.8 (Coarsely dense). *Let $\mathcal{C}$ be an L -separated, m -gluable system of chains. If (S, d_C) is k -weakly roughly geodesic, then S is $(3k + 4(L + m + 1))$ -coarsely dense in X_C .*

Proof. Let $x \in X_C$. Pick an arbitrary point $s \in S$, and let $c \in \mathcal{C}$ realise $d_C(s, x)$. If $|c| \leq 3k + 4(L + m + 1)$, then we are done, so assume otherwise. Because c is a finite chain, there is some $t \in S$ that is separated from s by every element of c . Let $b \in \mathcal{C}$ realise $d_C(t, x)$. Again, we can assume that $|b| > 3k + 4(L + m + 1)$. Because S is k -weakly roughly geodesic, there exists $s' \in S$ such that $d_C(s, s') \leq |c| + k$ and $d_C(s', t) \leq d_C(s, t) - |c| + k \leq |b| + k$. See Figure 4.

No element of c separates x from t , whereas every element of b separates x from t . Therefore, by Lemma 5.3 there is some $a \subset b \cup c$ with $a \in \mathcal{C}$ and $|a| \geq |b| + |c| - L - m - 1$. In particular, the fact that $d_C(s', t) \leq |b| + k$ implies that all but at most $k + L + m + 1$ elements of c must separate s' from s . Let $c_s \subset c$ consist of those elements that do separate s from s' .

Let $c' \in \mathcal{C}$ realise $d_C(s', x)$. Consider the subset $c'_s \subset c'$ consisting of those elements that do not separate s from s' . We can apply Lemma 5.3 to the pair (c_s, c'_s) to obtain an element of $\mathcal{C}$ that separates s from x and has cardinality at least $|c_s| + |c'_s| - L - m - 1$. Because c realises $d_C(s, x)$, this implies that $|c'_s| \leq k + 2(L + m + 1)$.




FIGURE 4. Proposition 5.8. The chain c realises $d_C(s, x)$, yielding t , and then b realises $d_C(t, x)$. Weak rough geodesicity gives s' , and c' realises $d_C(s', x)$.

Finally, consider $c' \setminus c'_s \in \mathcal{C}$. Every element of this chain separates $\{s, x\}$ from s' , so we can apply Lemma 5.3 to the pair $(c_s, c' \setminus c'_s)$ to find an element of $\mathcal{C}$ that separates s from s' and has cardinality at least $|c_s| + |c' \setminus c'_s| - L - m - 1$. From the fact that $d_C(s, s') \leq |c| + k$, we deduce that $|c' \setminus c'_s| \leq 2(k + L + m + 1)$. We therefore have $|c'| \leq 3k + 4(L + m + 1)$, completing the proof. $\square$

In practice it can often be verified from additional information about the pair (S, P) that (S, d_C) is weakly roughly geodesic. Proposition 5.8 then shows that working with (S, d_C) is largely the same as working with (X, d_C) , though the latter has the advantage that it has better fine properties, as illustrated by Section 4.

The following example shows how S can fail to be weakly roughly geodesic or coarsely dense in (X, d_C) .

Example 5.9. Let S be a copy of the real line, and identify S with a horocycle in the hyperbolic plane $\mathbf{H}^2$. Let P be the set of partitions obtained from the set of all $\mathbf{H}^2$ geodesics intersecting S by viewing each such geodesic as dividing $\mathbf{H}^2$ into two halves. Let $\mathcal{C}$ be the set of chains such that each pair of defining geodesics is at $\mathbf{H}^2$ -distance at least 1. With this system, the metric d_C on S is quasiisometric to the subspace metric from $\mathbf{H}^2$.

5.2. GRADED SYSTEMS

In [PSZ22], a key step in producing the curtain model of a CAT(0) space is the resolution of an infinite sequence of increasingly informative hyperbolic models into a single hyperbolic space. The goal of this section is to give that procedure a more general framing, which will be useful in Section 6.

Definition 5.10 (Graded system). A sequence $(\mathcal{C}_R)$ of dualisable systems on (S, P) is said to be a *graded system* on (S, P) if there are numbers $L_R \geq 1$, $m_R \geq 0$ and a sequence $(\kappa_R) \subset (0, \infty)$ such that the following hold.

- $\mathcal{C}_R \subset \mathcal{C}_{R+1}$ for all R .
- $\mathcal{C}_R$ is an L_R -separated, m_R -gluable system of chains for each R .
- For each $s, t \in S$ there exists M_{st} such that for every sequence (c^R) with $c^R \in \mathcal{C}_R$ separating s from t , we have $\sum_{R=1}^{\infty} \kappa_R |c^R| \leq M_{st}$.

Note that if we assume L_R to be minimal such that $\mathcal{C}_R$ is L_R -separated, then the condition that $\mathcal{C}_R \subset \mathcal{C}_{R+1}$ implies that $L_{R+1} \geq L_R$, but we could potentially have $m_{R+1} < m_R$.

Recall that $\hat{X}$ denotes the set of all ultrafilters defined by P . If $(\mathcal{C}_R)$ is a graded system, then the constructions of Section 3 produce a sequence of metric spaces $X_R = (X_{\mathcal{C}_R}, d_{\mathcal{C}_R})$, which will be hyperbolic by the results of Section 5.1. However, the sets X_R in general will be getting smaller (setwise, whilst the metric grows) as R increases, because as $\mathcal{C}_R$ gets larger the condition that $d_C(s, x) < \infty$ becomes more stringent. We unify these spaces as follows.

The graded dual. Given a graded system $(\mathcal{C}_R)$ on (S, P) , let (λ_R) be a sequence of positive numbers such that $\lambda_R \leq \kappa_R$ and $\sum_{R=1}^{\infty} \lambda_R (L_R + m_R + 1) = \Lambda < \infty$. Consider the extended metric on $\hat{X} \times \hat{X}$ given by

$$D(x, y) = \sum_{R=1}^{\infty} \lambda_R d_R(x, y).$$

The *graded dual* of S with respect to $(\mathcal{C}_R)$ is the space $X = \{x \in \hat{X} : D(s, x) < \infty \text{ for all } s \in S\}$, equipped with the metric D .

The final assumption in the definition of a graded system is what ensures that $D(s, t) < \infty$ for all $s, t \in S$.

Remark 5.11. If $x \in X$ then certainly $x \in X_R$ for every R , but it can happen that X is a strict subset of $\bigcap_{R=1}^{\infty} X_R$. For instance, if S is the region of the euclidean plane bounded between the x -axis and a sufficiently slowly growing sublinear function, P is the set of partitions induced by CAT(0) curtains, and $\mathcal{C}_R$ is as in [PSZ22], then any ultrafilter extending the filter defined by the unique CAT(0) boundary point of S lies in every X_R but not in X .

Definition 5.12 (Perichain). Given a graded system $(\mathcal{C}_R)$ on (S, P) , a *perichain* $\bar{c}$ is a choice of element $c^R \in \mathcal{C}_R$ for each R . A $\times$ -perichain is a choice of $\times$ -chain χ^R from $\mathcal{C}_R$ for each R .

For a perichain $\bar{c}$, write $\|\bar{c}\| = \sum \lambda_R |c^R|$. Note that if c^R realises $\mathbf{d}_R(x, y)$ for all R , then $\|\bar{c}\| = \mathbf{D}(x, y)$.

Lemma 5.13. *Let $(\mathcal{C}_R)$ be a graded system. If $\bar{c}$ is a perichain realising $\mathbf{D}(x_1, x_2)$ and $\bar{\chi}$ is a maximal $\times$ -perichain for x_1, x_2, x_3, x_4 , then*

$$\|\bar{\chi}_1\| + \|\bar{\chi}_2\| - 2\Lambda \leq \|\bar{c}'\| \leq \|\bar{\chi}_1\| + \|\bar{\chi}_2\| + 4\Lambda,$$

where $c'^R \subset c^R$ consists of all elements of c^R not separating x_3 from x_4 .

Proof. Apply Lemma 5.5 for each R . □

The proof of the following proposition is similar to that of Proposition 5.6, but one has to ensure some compatibility along the terms of the graded system.

Proposition 5.14. *If $(\mathcal{C}_R)$ is a graded system, then $(X, \mathbf{D})$ is four-point hyperbolic, with constant 16Λ .*

Proof. Fix $x_1, x_2, x_3, x_4 \in X$, and let $\bar{\chi}$ be a maximal $\times$ -perichain for them. For each distinct i, j, k , let $S^R(1i|jk)$ be the cardinality of a maximal element of $\mathcal{C}_R$ separating $\{x_1, x_i\}$ from $\{x_j, x_k\}$. Observe that since $\mathcal{C}_R$ is L_R -separated, at most one of the $S^R(1i|jk)$ can be greater than $2L_R$ for any given R .

Suppose that i is such that there is no value of R for which $S^R(1i|jk) > 2L_R$. Let $\bar{c}_{jk}$ be a perichain realising $\mathbf{D}(x_j, x_k)$, and let $\bar{b}_{jk}$ be the subperichain consisting of all elements of $\bar{c}_{jk}$ that either separate x_j from all three other x_l or separate x_k from all three other x_l . We have $\|\bar{b}_{jk}\| \geq \|\bar{c}_{jk}\| - 2\Lambda$. According to Lemma 5.13, this gives

$$(1) \quad \|\bar{\chi}_j\| + \|\bar{\chi}_k\| - 2\Lambda \leq \mathbf{D}(x_j, x_k) \leq \|\bar{\chi}_j\| + \|\bar{\chi}_k\| + 6\Lambda.$$

We similarly have the corresponding inequality for $\mathbf{D}(x_1, x_i)$.

If there is no i for which there is some value of R for which $S^R(1i|jk) > 2L_R$, then the inequality (1) holds for all pairs (x_j, x_k) . This implies the four-point inequality with constant 8Λ .

Otherwise, let R_0 be minimal such that some $S^{R_0}(1i|jk) > 2L_{R_0}$. After relabelling, we may assume that $i = 2$. In other words, there is an element of $\mathcal{C}_{R_0}$ of length greater than $2L_{R_0}$ that separates x_1 and x_2 from x_3 and x_4 . Because $\mathcal{C}_{R_0} \subset \mathcal{C}_R$ for all $R \geq R_0$, we can apply L_R -separation to a pair of walls in a chain realising $S^{R_0}(12|34)$ to see that if $R \geq R_0$ then $\max\{S^R(13|24), S^R(14|23)\} \leq L_R$. In particular, minimality of R_0 means that $\max\{S^R(13|24), S^R(14|23)\} \leq 2L_R$ for all R . Thus inequality (1) holds for the pairs (x_1, x_2) and (x_3, x_4) .

We now bound the distance between each of the other four pairs of points. The argument is the same for each, so for ease of notation we shall work with the pair (x_1, x_3) . By applying Lemma 5.3 twice, we see that $\mathbf{d}_L(x_1, x_3) \geq |\chi_1^R| + |\chi_3^R| + S^R(12|34) - 2(L_R + m_R + 1)$, and hence

$$\|\bar{\chi}_1\| + \|\bar{\chi}_3\| + \sum_{R=1}^{\infty} \lambda_R S^R(12|34) - 2\Lambda \leq \mathbf{D}(x_1, x_3).$$

On the other hand, let $c_{13}^R \in \mathcal{C}_R$ realise $\mathbf{d}_R(x_1, x_3)$, and let b_{13}^R be the subset consisting of those walls that do not separate x_2 from x_4 . By Lemma 5.3, after removing at most $L_R + m_R + 1$ elements from each of χ_2^R and χ_4^R and at most $2(L_R + m_R + 1)$ elements of b_{13}^R , we are left with

a $\times$ -chain. By maximality of $\bar{\chi}$, it follows that $|\chi^R| \geq |\chi_2^R| + |\chi_4^R| + |b_{13}^R| - 4(L_R + m_r + 1)$, or in other words that

$$|b_{13}^R| \leq |\chi_1^R| + |\chi_3^R| + 4(L_R + m_R + 1).$$

Every element of $c_{13}^R \setminus b_{13}^R$ either separates $\{x_1, x_2\}$ from $\{x_3, x_4\}$, or separates $\{x_1, x_4\}$ from $\{x_2, x_3\}$. But $S^R(14|23) \leq 2L_R$, and so $|c_{13}^R \setminus b_{13}^R| \leq S^R(12|34) + 2L_R$. Hence for every R we have

$$d_R(x_1, x_3) \leq |b_{13}^R| + S^R(12|34) + 2L_R \leq |\chi_1^R| + |\chi_3^R| + S^R(12|34) + 6(L_R + m_R + 1).$$

In tandem with the above lower bound on $D(x_1, x_3)$, summing this yields

$$\|\bar{\chi}_1\| + \|\bar{\chi}_3\| + \sum_{R=1}^{\infty} \lambda_R S^R(12|34) - 2\Lambda \leq D(x_1, x_3) \leq \|\bar{\chi}_1\| + \|\bar{\chi}_3\| + \sum_{R=1}^{\infty} \lambda_R S^R(12|34) + 6\Lambda.$$

As noted, the same argument holds for each of the pairs (x_1, x_4) , (x_2, x_3) , and (x_2, x_4) . The four-point inequality for x_1, x_2, x_3, x_4 follows, with constant 16Λ . $\square$

Proposition 5.15. *If $(\mathcal{C}_R)$ is a graded system such that the sequence (m_R) is bounded above by some number M , then (X, D) is $4M\Lambda$ -weakly roughly geodesic.*

Proof. Given $x, y \in X$, let R_0 be sufficiently large so that $\sum_{R>R_0} \lambda_R d_R(x, y) \leq \Lambda$. Let σ_{xy} be the normal wall path in (X_{R_0}, d_{R_0}) from x to y , as constructed in Section 4, and write $z_i = \sigma_{xy}(i)$. We do not know exactly how D compares with d_{R_0} on X , so this path could *a priori* fail to be a rough geodesic of X , but we shall nonetheless use it to show that X is weakly roughly geodesic. This will be a consequence of the following observations, which are built into the definition of the z_r .

- Every element of P separating x from z_r separates x from $\{z_r, y\}$.
- Every element of P separating z_r from z_{r+1} separates $\{x, z_r\}$ from $\{z_{r+1}, y\}$.
- Every element of P separating z_r from y separates $\{x, z_r\}$ from y .

First observe that these imply that z_r is indeed an element of X , because we must have $D(x, z_r) \leq D(x, y) < \infty$, and x lies at finite distance from S . In order to establish the proposition, it suffices to bound $D(z_r, z_{r+1})$ and $D(x, z_r) + D(z_r, y) - D(x, y)$.

For the latter, note that for each R we can apply Lemma 5.3 to find that $d_R(x, y) \geq d_R(x, z_r) + d_R(z_r, y) - L_R - m_R - 1$. Summing over R , we find that $D(x, y) \geq D(x, z_r) + D(z_r, y) - \Lambda$, as desired.

It remains to bound $D(z_r, z_{r+1})$. We know from Proposition 4.6 that $d_{R_0}(z_r, z_{r+1}) \leq 1 + 2m_{R_0} \leq 1 + 2M$. Because $\mathcal{C}_R \subset \mathcal{C}_{R+1}$ for all R , this implies that $d_R(z_r, z_{r+1}) \leq 1 + 2M$ for all $R \leq R_0$. On the other hand, if $R > R_0$, then by the above observations we have that $d_R(x, y) \geq d_R(z_r, z_{r+1})$. We can therefore compute

$$\begin{aligned} D(z_r, z_{r+1}) &= \sum_{R=1}^{R_0} \lambda_R d_R(z_r, z_{r+1}) + \sum_{R>R_0} \lambda_R d_R(z_r, z_{r+1}) \\ &\leq \sum_{R=1}^{R_0} \lambda_R (1 + 2M) + \sum_{R>R_0} \lambda_R d_R(x, y) \leq 4M\Lambda. \end{aligned} \quad \square$$

Remark 5.16. It follows that in the situation of Proposition 5.15, (X, D) is a coarsely injective, hence roughly geodesic, hyperbolic space [PSZ22, Prop. A.2]. Moreover, in the proof of Proposition 5.15, we showed that given $x, y \in X$, there is some R such that the normal wall path in X_R from x to y is a uniform weak rough geodesic in X . We therefore have *a posteriori* that these paths are uniform rough geodesics in X .

As in Section 5.1, it is also desirable to know that the original set S is a roughly geodesic hyperbolic space with respect to the metric D . Interestingly, in contrast to Proposition 5.8, it is uncertain whether (S, D) being roughly geodesic is equivalent to its being dense in (X, D) .

PART 2. APPLICATIONS

6. A UNIVERSAL HYPERBOLIC SPACE FOR CONTRACTING GEODESICS

The goal of this section is to construct, for a given geodesic space S , a hyperbolic space X with the property that every *strongly contracting* quasigeodesic in S is witnessed as a (parametrised) quasigeodesic in X . Any quasigeodesic in S can be perturbed to one whose image is a closed subset of S , so we shall always assume that quasigeodesics are closed. In particular, this means that every point in S has a nonempty set of closest points in each quasigeodesic.

Definition 6.1 (Strongly contracting). Let α be a quasigeodesic in a geodesic space S . Given $x \in S$, let $\pi_\alpha(x)$ be the (nonempty) set of closest points in α to x . We say that α is D -strongly contracting if for any ball $B \subset S$ disjoint from α we have $\text{diam}(\pi_\alpha(B)) < D$.

If α is strongly contracting, then although π_α may send points to sets of cardinality greater than one, the image of a point has bounded diameter. For a set $I \subset \alpha$, we write $\pi_\alpha^{-1}(I)$ to mean the set of all points $x \in S$ such that $\pi_\alpha(x) \subset I$.

For the remainder of this section, fix a geodesic space S . In order to apply the methods of the previous sections, we need two things: a good set of bipartitions of S , and a choice of which collections of bipartitions should be counted. As an intermediate step towards these goals, we consider a natural collection of subspaces of S , which we call *curtains*, in analogy with [PSZ22].

Definition 6.2 (Curtains). Suppose that α is a D -strongly contracting geodesic of length $20D$, with $D \geq 1$. A *curtain* dual to α is a set $\pi_\alpha^{-1}(I)$, where $I \subset \alpha$ is a subgeodesic of length $10D$ not containing any endpoint of α .

Let us write $\text{Cur } S$ for the set of all curtains in S . That is, $\text{Cur } S$ contains every curtain dual to every D -strongly contracting geodesic in S , for every $D \geq 1$. Note that $\text{Cur } S$ may very well be empty: this happens if S has no strongly contracting geodesics.

The fact that we only consider strongly contracting *geodesics* in the construction of curtains is justified by the following well-known lemma.

Lemma 6.3. *A q -quasigeodesic α is strongly contracting if and only if there is a constant D' such that the following hold.*

- For every $x \in S$ we have $\text{diam } \pi_\alpha(x) \leq D'$.
- For every $x, y \in S$ with $d(\pi_\alpha(x), \pi_\alpha(y)) > 20D'$, and for every geodesic β from x to y , every subpath of α from $\pi_\alpha(x)$ to $\pi_\alpha(y)$ lies in the $5D'$ -neighbourhood of β .

Proof. The forward direction is given by [CS15, Lem. 4.5]. For the reverse direction suppose that the two conditions hold for the q -quasigeodesic α , and let $B = B(x, r)$ be a ball in S that is disjoint from α . If $\text{diam } \pi_\alpha B > 50D'$, then by the first assumption there is some $y \in B$ such that $d(\pi_\alpha(x), \pi_\alpha(y)) > 20D'$. Let β be a geodesic from x to y , and let $z \in \pi_\alpha(x)$. By the second assumption, β comes $5D'$ -close to both z and $\pi_\alpha(y)$. But then the length of β must be at least $(d(x, z) - 5D') + 20D' - 5D' > d(x, z)$, which implies that $z \in B$, in conflict with the choice of B . Thus α is $50D'$ -strongly contracting. $\square$

Every curtain $\mathfrak{h}$ has two nonempty *halfspaces*, $\mathfrak{h}^+$ and $\mathfrak{h}^-$: if $\mathfrak{h}$ is dual to the strongly contracting geodesic α at an interval I , then they are the sets of points x such that $\pi_\alpha(x)$ intersects one of the two components of the complement of I in α . Note that the choice of length of I means that $x \in X$ cannot simultaneously be in $\mathfrak{h}^+$ and $\mathfrak{h}^-$, and in particular, $\{\mathfrak{h}^-, \mathfrak{h}, \mathfrak{h}^+\}$ is a tripartition of S . More strongly, we have $d(\mathfrak{h}^-, \mathfrak{h}^+) \geq 3D > 1$. Note that this “thickness” increases with the contracting constant.

We say that $\mathfrak{h}$ *separates* two points or subsets of S if they lie in opposite halfspaces of $\mathfrak{h}$. A *chain* of curtains is then a sequence $(\mathfrak{h}_i)$ such that $\mathfrak{h}_i$ separates $\mathfrak{h}_{i-1}$ from $\mathfrak{h}_{i+1}$ for all i .

Definition 6.4 (Ball-separation). For a natural number R , we say that disjoint curtains $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are R -ball-separated if there exists a ball $B \subset \mathfrak{h}_1^+ \cap \mathfrak{h}_2^-$ with radius at most R and such that every geodesic from $\mathfrak{h}_1^-$ to $\mathfrak{h}_2^+$ meets B .

An R -chain is a chain $(\mathfrak{h}_i)$ of curtains such that $\mathfrak{h}_i$ and $\mathfrak{h}_j$ are R -ball-separated for all i, j .

We say that an R -chain $(\mathfrak{h}_i)$ *crosses* a curtain $\mathfrak{k}$ if all four quarterspaces $\mathfrak{h}_i^\pm \cap \mathfrak{k}^\pm$ are nonempty for all i .

Lemma 6.5. *If curtains $\mathfrak{k}_1$ and $\mathfrak{k}_2$ are R -ball-separated, then every R -chain crossing both $\mathfrak{k}_1$ and $\mathfrak{k}_2$ has length at most $3R + 5$.*

Proof. Let B be a ball of radius at most R such that every geodesic from $\mathfrak{k}_1^-$ to $\mathfrak{k}_2^+$ passes through B . Let $\mathfrak{c} = (\mathfrak{h}_1, \dots, \mathfrak{h}_{2n+R+1})$ be an R -chain that crosses both $\mathfrak{k}_1$ and $\mathfrak{k}_2$. Because $d(\mathfrak{h}_i^-, \mathfrak{h}_i^+) \geq 1$ for each i , at most $R + 1$ elements of $\mathfrak{c}$ can intersect B . After switching the order of the $\mathfrak{h}_i$, we therefore have that $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ are all disjoint from B , and $B \subset \mathfrak{h}_n^+$.




FIGURE 5. Lemma 6.5. Geodesics from x to y have to pass through two balls of radius R .

Since $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are R -ball-separated, there is a ball B' of radius at most R such that every geodesic from $\mathfrak{h}_1^-$ to $\mathfrak{h}_2^+$ must pass through B' . Because $\mathfrak{h}_1$ crosses both $\mathfrak{k}_1$ and $\mathfrak{k}_2$, there are points $x \in \mathfrak{h}_1^- \cap \mathfrak{k}_1^-$ and $y \in \mathfrak{h}_1^- \cap \mathfrak{k}_2^+$. If α is a geodesic from x to y , then α must meet B . Let $z \in B \cap \alpha$, and let $\alpha_1 = \alpha[x, z]$, $\alpha_2 = \alpha[z, y]$. Both α_1 and α_2 must pass through B' . But $\text{diam } B' \leq 2R$, so the subsegment of α lying in $\mathfrak{h}_2^+$ must have length at most $2R$. Because $z \in \mathfrak{h}_n^+$ and $d(\mathfrak{h}_i^-, \mathfrak{h}_i^+) \geq 1$, it follows that $n \leq R + 2$. $\square$

The universal space. Each curtain $\mathfrak{h} \in \text{Cur } S$ induces two natural bipartitions of S , namely $(\mathfrak{h}^- \cup \mathfrak{h}, \mathfrak{h}^+)$ and $(\mathfrak{h}^-, \mathfrak{h} \cup \mathfrak{h}^+)$. We let P be the set of all bipartitions of S induced by curtains in this way (different curtains can induce a common bipartition). Let $\mathcal{C}_R$ be the set of all chains $\{h_i\} \subset P$ such that there exists an R -chain $(\mathfrak{h}_i) \subset \text{Cur } S$ with $\mathfrak{h}_i$ inducing h_i .

According to Lemma 6.6 below, the sequence $(\mathcal{C}_R)$ is a graded system. For each R , let X_R be the $\mathcal{C}_R$ -dual of S . Let X be the graded dual of S with respect to $(\mathcal{C}_R)$.

Lemma 6.6. *The sequence $(\mathcal{C}_R)$ defined above is a graded system on (S, P) , and each $\mathcal{C}_R$ is 3-gluable.*

Proof. Every element of $\mathcal{C}_R$ is induced by a chain of curtains, and is therefore itself a chain from P . If a pair of curtains is R -ball-separated then it is $(R + 1)$ -ball-separated, so $\mathcal{C}_R \subset \mathcal{C}_{R+1}$. Lemma 6.5 provides separation of the $\mathcal{C}_R$. It remains to check that $\mathcal{C}_R$ is 3-gluable and to find an appropriate sequence (κ_R) as in Definition 5.10.

For the latter, note that since $d(\mathfrak{h}^-, \mathfrak{h}^+) \geq 1$ for every $h \in \text{Cur } S$, if $c \in \mathcal{C}_R$ separates $s \in S$ from $t \in S$, then $|c| \leq d(s, t) + 2$. This means we can take $\kappa_R = \frac{1}{R^2}$, for instance.

To show that $\mathcal{C}_R$ is 3-gluable, suppose that $c_1, c_2 \in \mathcal{C}_R$ are such that $c_1 \cup c_2$ is a chain with $c_2 \subset c_1^+$ and $c_1 \subset c_2^-$. Let h_{-1}, h_{-2} be the two maximal elements of c_1 , and let k_1, k_2, k_3 be the three minimal elements of c_2 . Let $(\mathfrak{h}_i)$ be an R -chain inducing c_1 , and let $(\mathfrak{k}_i)$ be an R -chain inducing c_2 . Let B be a ball of radius at most R such that every geodesic from $\mathfrak{k}_2^-$ to $\mathfrak{k}_3^+$ meets B .

Since $(\mathfrak{k}_i)$ is a chain of curtains, $\mathfrak{k}_2 \subset \mathfrak{k}_1^+$. Also, the fact that $\{h_{-1}, k_1\}$ is a chain in P implies that $\mathfrak{h}_{-1}^- \subset \mathfrak{k}_1 \cup \mathfrak{k}_1^-$. Hence $\mathfrak{k}_2$ is disjoint from $\mathfrak{h}_{-1}^-$. In particular, $\mathfrak{k}_2$ is disjoint from $\mathfrak{h}_{-2}$. Moreover, any geodesic from $\mathfrak{h}_{-2}^-$ to $\mathfrak{k}_3^+$ is a geodesic from $\mathfrak{k}_2^-$ to $\mathfrak{k}_3^+$, so any such geodesic meets B . That is, $\mathfrak{h}_{-2}$ and $\mathfrak{k}_3$ are R -ball-separated curtains. We conclude that $(\mathfrak{h}_i)_{i \leq -2} \cup (\mathfrak{k}_j)_{j \geq 3}$ is an R -chain, and hence $c_1 \cup c_2 \setminus \{h_{-1}, k_1, k_2\} \in \mathcal{C}_R$. $\square$

In view of Propositions 5.14 and 5.15, we therefore have that X is a roughly geodesic hyperbolic space (see also Remark 5.16). According to the following proposition, one could also work in the roughly geodesic hyperbolic space (S, D) if preferred.

Proposition 6.7. *With the subspace metric D , the set $S \subset X$ is weakly roughly geodesic. Moreover, geodesics of (S, d) are uniform unparametrised rough geodesics of (X, D) .*

Proof. If $s_1, s_2 \in S$ have $d(s_1, s_2) \leq 2$, then $D(s_1, s_2) \leq \Lambda$, where Λ is as in Section 5.2. Hence it suffices, given $r, t \in S$ and s lying on an (S, d) -geodesic α from r to t , to upper bound $D(r, s) + D(s, t) - D(r, t)$. This will also imply the “moreover” statement. Let $\bar{c} = (c^R)$ be a perichain realising $D(r, s)$. Let c_s^R be the subset of c^R that does not separate r from t . We first show that $|c_s^R| \leq R + 3$.

For this, fix R , let $(\mathfrak{h}_i)$ be an R -chain of curtains inducing c^R , and let $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ be the curtains inducing c_s^R , with increasing index signifying increasing distance from s . Whilst we may have $s \in \mathfrak{h}_1$, we certainly have $s \in \mathfrak{h}_2^-$. Similarly, $r, t \in \mathfrak{h}_{n-1}^+$. Let B be a ball of radius at most R such that every geodesic from $\mathfrak{h}_{n-1}^+$ to $\mathfrak{h}_{n-2}^-$ must meet B . In particular, $\alpha[r, s]$ and $\alpha[s, t]$ both meet B . Because $\mathfrak{h}_2, \dots, \mathfrak{h}_{n-2}$ all separate B from s , and $d(\mathfrak{h}_i^-, \mathfrak{h}_i^+) \geq 1$, we therefore have $n \leq R + 3$, as desired.

Because of this, we can find perichains $\bar{b}_1$ and $\bar{b}_2$ such that $\bar{b}_1$ separates r from $\{s, t\}$ and $\bar{b}_2$ separates t from $\{s, r\}$, such that $D(r, s) \leq \|\bar{b}_1\| + \Lambda$ and $D(s, t) \leq \|\bar{b}_2\| + \Lambda$. Applying Lemma 5.3 in each $\mathcal{C}_R$, we have $D(r, t) \geq \|\bar{b}_1\| + \|\bar{b}_2\| - \Lambda$, and hence $D(r, t) \geq D(r, s) + D(s, t) - 3\Lambda$, which completes the proof. $\square$

The following theorem characterises strongly contracting quasigeodesics of S as those that quasiisometrically embed in X . In particular, it shows that the existence of a strongly contracting ray in (S, d) is sufficient for $(S, D) \subset X$ to be unbounded.

Theorem 6.8. *Given a geodesic space S , let X be the hyperbolic space constructed above.*

For each q, D there exists ν such that if α is a D -strongly contracting q -quasigeodesic in S , then $\alpha \rightarrow X$ is a ν -quasiisometric embedding.

Conversely, for each q, ν there exists D such that if α is a q -quasigeodesic in S with the property that $\alpha \rightarrow X$ is a ν -quasiisometric embedding, then α is D -strongly contracting.

Proof. Let $\alpha : [0, T) \rightarrow S$ be a D -strongly contracting q -quasigeodesic, with $T \in [0, \infty]$. By considering a geodesic from $\alpha(0)$ to $\alpha(n)$ for each possible n , Lemma 6.3 tells us that there is $R = R(D, q)$ such that there are R -chains of curtains separating points of α , the cardinality of which is uniformly linearly lower bounded in terms of R and n . Thus $\alpha \rightarrow X_R$ is a quasiisometric embedding. It follows from the construction of $\mathbf{D}$ that $\alpha \rightarrow X$ is also a quasiisometric embedding. This establishes the forward direction of the theorem.

Now let us consider the converse. Let Λ be as in Section 5.2. We may assume that q, ν , and Λ are positive integers. Write $K = 2q\nu\Lambda$. Let α be a q -quasigeodesic in S such that $\alpha \rightarrow X$ is a ν -quasiisometric embedding. After perturbing α , we can coarsely cover it by a sequence $(x_n) \subset \alpha$ such that $d(x_n, x_{n+1}) = 30K^2$.

As $\alpha \rightarrow X$ is a ν -quasiisometric embedding, we have $D(x_n, x_{n+1}) \geq 50K\Lambda$. By definition of Λ and the fact that the separation constant L_R of $\mathcal{C}_R$, coming from Lemma 6.5, is at least R , there must be some R such that $d_{\mathcal{C}_R}(x_n, x_{n+1}) \geq 50KR$. Hence there is an R -chain of curtains $(h_1^n, \dots, h_{50KR-2}^n)$ separating x_n from x_{n+1} . The fact that $d(h^-, h^+) \geq 1$ for every $h \in \text{Cur } S$ now implies that $50KR - 2 \leq d(x_n, x_{n+1}) = 30K^2$, and hence $R \leq K$.

By repeatedly applying Lemma 5.3, and recalling that $\mathcal{C}_K$ is 3-gluable, we obtain an element $c \in \mathcal{C}_K$ by taking the union of all (h_i^n) and removing at most $2(L_K + 4)$ from each. By Lemma 6.5, this leaves at least two elements of (h_i^n) in c for each n . Fix a choice of two, and label them $h_n, \mathfrak{t}_n$, with h_n separating x_n from $\mathfrak{t}_n$.

Because h_n and $\mathfrak{t}_n$ are K -ball-separated, there is some ball B_n of radius at most K such that every geodesic from h_n^- to $\mathfrak{t}_n^+$ must meet B_n . The subsegment of α from x_n to x_{n+1} has diameter at most $50K^2q^3$, because α is a q -quasigeodesic. Let γ_n be a geodesic from x_n to x_{n+1} . It meets B_n . Let $\gamma = \bigcup_{n \in \mathbb{Z}} \gamma_n$, which lies at Hausdorff-distance at most $100K^2q^3$ from α . We aim to apply Lemma 6.3 to find that γ , and hence α , is strongly contracting.




FIGURE 6. Theorem 6.8. x_n and x_{n+1} are distant points along α , separated by $\{h_n, \mathfrak{t}_n\} \in \mathcal{C}_K$. The path γ is a piecewise geodesic, and any geodesic δ from s to a projection $t \in \pi_\gamma(s) \cap h_n^-$ passes through B_n .

First we show that points of S have uniformly bounded projection to γ . Given $s \in S$, let n be maximal such that $s \in \mathfrak{t}_n^+$. If there is a point $t \in \pi_\gamma(s)$ lying in h_n^- , then any geodesic

δ from s to t must meet B_n , so δ comes $2K$ -close to γ at some point in B_n , and hence t is $2K$ -close to B_n . Similarly, if m is minimal such that $s \in \mathfrak{h}_m^-$, then any point of $\pi_\gamma(s)$ lying in $\mathfrak{k}_m^+$ is $2K$ -close to B_m . The choices of m and n ensure that $m \in \{n+1, n+2\}$. In particular, the piecewise geodesic $\gamma_n \cup \gamma_{n+1} \cup \gamma_{n+2}$ meets both B_n and B_m , so

$$\text{diam } \pi_\gamma(s) \leq 2K + \text{diam } B_n + |\gamma_n| + |\gamma_{n+1}| + |\gamma_{n+2}| + \text{diam } B_m + 2K \leq 100K^2.$$

Moreover, as $d(\mathfrak{h}^-, \mathfrak{h}^+) \geq 1$ for every $\mathfrak{h} \in \text{Cur } S$, we observe that no point of $\pi_\gamma(s)$ can be separated from s by more than $2K+2$ of the $\mathfrak{h}_i$, because points of $\pi_\gamma(s)$ that lie in $\mathfrak{h}_n^-$ are $2K$ -close to B_n , which meets $\mathfrak{h}_n^+$ and has diameter at most $2K$, and similarly points of $\pi_\gamma(s)$ that lie in $\mathfrak{h}_{m+1}^+ \subset \mathfrak{k}_m^+$ are $2K$ -close to B_m .

It remains to establish the second condition of Lemma 6.3. Suppose that $s_1, s_2 \in S$ have $d(\pi_\gamma(s_1), \pi_\gamma(s_2)) > (10K)^3 \geq 10(K+1)|\gamma_i|$, and let β be a geodesic in S from s_1 to s_2 . Let a be minimal such that $s_1 \in \mathfrak{h}_a^-$, and let b be maximal such that $s_2 \in \mathfrak{k}_b^+$. Perhaps after swapping s_1 and s_2 , the lower bound on $d(\pi_\gamma(s_1), \pi_\gamma(s_2))$ implies that $b-a > 5(K+1)$, because γ_i is a geodesic from $x_i \in \mathfrak{h}_i^-$ to $x_{i+1} \in \mathfrak{k}_i^+$. Thus the above observation yields that $\mathfrak{h}_{a+2K+2}, \mathfrak{k}_{a+2K+2}, \dots, \mathfrak{h}_{b-2K-2}, \mathfrak{k}_{b-2K-2}$ all separate $\{s_1\} \cup \pi_\gamma(s_1)$ from $\{s_2\} \cup \pi_\gamma(s_2)$.

We deduce from this that for any subpath δ of γ from $\pi_\gamma(s_1)$ to $\pi_\gamma(s_2)$, both β and δ meet all of $B_{a+2K+2}, \dots, B_{b-2K-2}$. In particular, the intersection of δ with the “middle” region $\mathfrak{h}_{a+2K+2}^+ \cap \mathfrak{k}_{b-2K-2}^-$ of S lies in a neighbourhood of β of radius at most $|\gamma_i| + \text{diam } B_i \leq 50K^2$.

The two “end” regions $S \setminus \mathfrak{h}_{a+2K+2}^+$ and $S \setminus \mathfrak{k}_{b-2K-2}^-$ are treated similarly, so let us just consider the former. By minimality of a , the above observation implies that every point of $\pi_\gamma(s_1)$ lies in $\mathfrak{h}_{a-2K-4}^+$. Since we know that β meets B_{a+2K+2} , this implies that the intersection of $\delta \subset \gamma$ with this region lies in a neighbourhood of β of radius at most $\text{diam } B_{a+2K+2} + \sum_{i=a-2K-4}^{a+2K+2} |\gamma_i| \leq (10K)^3$.

We have shown that the conditions of Lemma 6.3 are met by γ , so it is strongly contracting. Since γ and α lie at a bounded Hausdorff-distance, α is also strongly contracting. $\square$

Remark 6.9. The assumption that α is a quasigeodesic in S is essential for the converse direction of Theorem 6.8. For instance, if, as in Remark 5.11, S is the region of the euclidean plane bounded between the x -axis and a sufficiently slowly growing sublinear function, then X will be a quasiray even though no ray in S is strongly contracting. But in this example no geodesic ray in S is quasiisometrically embedded in X .

Next we show that, for isometries, being strongly contracting can be detected by just looking at pairs of curtains.

Definition 6.10. An isometry $g \in \text{Isom } S$ is *strongly contracting* if there exist D, q and $s \in S$ such that $\langle g \rangle \cdot s$ is a D -strongly contracting q -quasigeodesic in S .

We say that an isometry g *skewers* a pair of disjoint curtains $\mathfrak{h}_1, \mathfrak{h}_2$ if there exists n such that $g^n \mathfrak{h}_1^+ \subsetneq \mathfrak{h}_2^+ \subsetneq \mathfrak{h}_1^+$.

Corollary 6.11. *If $g \in \text{Isom } S$ has quasiisometrically embedded orbits in S , then the following are equivalent.*

- (1) g is strongly contracting.
- (2) g skewers a pair of ball-separated curtains.
- (3) g acts loxodromically on X .

Proof. Theorem 6.8 shows that (1) and (3) are equivalent.

Assume that g acts loxodromically on X . Let $s \in S$ and let n be sufficiently large that $D(s, g^n s) > 10\Lambda$. In particular, by definition $\Lambda \geq \sum \lambda_R(L_R + 1)$, there must be some R for

which $d_R(s, g^n s) \geq L_R + 9$. That is, there is some $c = \{h_1, \dots, h_k\} \in \mathcal{C}_R$ separating s from $g^n s$, with $k \geq L_R + 9$.

Consider $g^n c \in \mathcal{C}_R$, and recall from Lemma 6.6 that $\mathcal{C}_R$ is 3-gluable. According to Lemma 5.3, there is some $b \subset c \cup g^n c$, obtained by removing a subset of $\{h_{k-3}, h_{k-2}, h_{k-1}, h_k, g^n h_1, \dots, g^n h_{L_R+4}\}$ of cardinality at most $L_R + 4$, such that b separates s from $g^n s$. Repeating this inductively, we find that $\{g^{mn} h_{L_R+5}, \dots, g^{mn} h_{k-4}\} \in \mathcal{C}_R$ separates $g^m s$ from $g^{tn} s$ for all $r \leq m$ and all $t > m$. In particular, if (h_i) is a chain of curtains inducing $\{h_i\}$, then g^{2n} skewers the ball-separated curtains $h_{k-4}, g^n h_{k-4}$. Thus (3) implies (2).

Finally, suppose that g skewers a pair of R -ball-separated curtains $\mathfrak{h}$ and $\mathfrak{k}$. Because $\mathfrak{h}$ and $\mathfrak{k}$ are R -ball-separated, $(g^{nm} \mathfrak{h})_{m \in \mathbb{Z}}$ is an R -chain of curtains. In particular, for any $s \in S$ we have $d_{\mathcal{C}_R}(s, g^{nm} s) \geq m - 2$. This shows that g act loxodromically on X , so (2) implies (3). $\square$

More strongly, for any group G acting properly on S , we show that the action of G on X is weakly proper along the axis of each strongly contracting element.

Definition 6.12 (WPD). Let G be a group acting on a hyperbolic space Y . An element $g \in G$ is *WPD* if there is a point $x \in Y$ such that for each $\varepsilon > 0$ there exists $m > 0$ for which only finitely many $h \in G$ satisfy both $d(x, hx) < \varepsilon$ and $d(g^m x, hg^m x) < \varepsilon$.

Proposition 6.13. *Suppose that a group G acts properly on S . If $g \in G$ is strongly contracting, then g is a WPD loxodromic for the action of G on X .*

Proof. Fix $s \in S$ and let $\varepsilon > 0$. By definition, g being strongly contracting implies that it has quasiisometrically embedded orbits in S , so by Corollary 6.11, g acts loxodromically on X . In particular, there exists m such that $D(s, g^m s) \geq 2\varepsilon + 2\Lambda$. Let $\bar{c} = (c^R)$ be a perichain realising $D(s, g^m s)$.

If $h \in G$ has the property that at most $2R + 3$ elements of c^R separate hs from $hg^m s$ for every R , then for every R all but at most $2R + 3$ elements of c^R either separate s from hs or separate $g^m s$ from $hg^m s$ (or both). For such an element h , we therefore have

$$D(s, hs) + D(g^m s, hg^m s) \geq \sum_{R=1}^{\infty} \lambda_R(|c^R| - 2R - 3) \geq \|\bar{c}\| - 2\Lambda \geq 2\varepsilon.$$

This shows that if $h \in G$ satisfies $D(s, hs) \leq \varepsilon$ and $D(g^m s, hg^m s) \leq \varepsilon$, then there is some R for which at least $2R + 4$ elements of c^R separate hs from $hg^m s$. Letting (h_i) be an R -chain of curtains inducing the subset of c^R separating hs from $hg^m s$, we see that there is a ball $B \subset S$ of radius at most R such that every geodesic from h_{R+2}^- to h_{R+3}^+ must meet B , where $s, hs \in h_{R+2}^-$ and $g^m s, hg^m s \in h_{R+3}^+$.

Let $\gamma \subset S$ be a geodesic from s to $g^m s$. We have shown that $h\gamma$ comes $2R$ -close to γ . Moreover, because every $\mathfrak{h} \in \text{Cur } S$ has $d(\mathfrak{h}^-, \mathfrak{h}^+) \geq 1$, we must have $d(s, g^m s) \geq 2R + 4$. Hence

$$d(s, hs) \leq |\gamma| + 2R + |g\gamma| \leq 3d(s, g^m s).$$

But the action of G on S is proper, so there are only finitely many such elements h . $\square$

Next we extend Proposition 6.13 by showing that, in the terminology of [BCK⁺23], for many S the hyperbolic space X is a *universal recognising space* for stable subgroups of groups acting geometrically on S .

Definition 6.14 (Morse, stable). Given a function $M : \mathbb{R}^2 \rightarrow \mathbb{R}$, we say that a geodesic α in a metric space is *M-Morse* if every (λ, μ) -quasigeodesic with endpoints on α lies in the $M(\lambda, \mu)$ -neighbourhood of α .

A finitely generated subgroup H of a finitely generated group G is *stable* if its inclusion map is a quasiisometric embedding and there exists $M : \mathbf{R}^2 \rightarrow \mathbf{R}$ such that every geodesic in G between points of H is M -Morse.

For example, strongly contracting geodesics are always Morse, but the reverse can fail [RV21]. In many spaces, though, the two are equivalent [CS15, CW21b, SZ22]. Stability was introduced by Durham–Taylor [DT15], who showed that it reformulates *convex cocompactness* for mapping class groups [FM02]. The curve graph is a universal recognising space for stable subgroups of mapping class groups [KL08].

Theorem 6.15. *Let S be a geodesic space with the property that for each M there exists D such that every M -Morse geodesic in S is D -strongly contracting. Let X be the hyperbolic space constructed above. For any group G acting properly coboundedly on S , a finitely generated subgroup $H < G$ is stable if and only if its orbit maps on X are quasiisometric embeddings.*

Proof. Suppose that H is stable in G . Given $s \in S$ and $g, h \in H$, each geodesic from gs to hs is uniformly Morse. By assumption, such a geodesic is therefore uniformly strongly contracting, so by the construction of $\text{Cur } S$ and Lemma 6.3 we obtain a sequence of uniformly ball-separated curtains separating gs from hs at a uniform rate. This shows that there is some R such that orbit maps of H on X_R are uniform quasiisometric embeddings, and it follows that orbit maps of H on X are uniform quasiisometric embeddings.

Conversely, suppose that orbit maps of H on X are quasiisometric embeddings. Because H is finitely generated, its orbit maps on S are coarsely Lipschitz. Moreover, the fact that $d(\mathfrak{h}^-, \mathfrak{h}^+) \geq 1$ for every $\mathfrak{h} \in \text{Cur } S$ implies that the identity map $(S, d) \rightarrow (X, D)$ is coarsely Lipschitz, so orbit maps of H on S are in fact quasiisometric embeddings. In particular, the inclusion of H in G is a quasiisometric embedding. Furthermore, if $s \in S$ then for any $g, h \in H$, each geodesic α from gs to hs uniformly quasiisometrically embeds in X , so is uniformly strongly contracting by Theorem 6.8. This shows that H is stable in G . $\square$

7. STRONG COARSE MEDIAN SPACES

Coarse median spaces were introduced by Bowditch in [Bow13a], providing a general framework for studying groups that display features of median geometry up to controlled error, such as toral relatively hyperbolic groups [Bow13b] and mapping class groups [BM11], among many others. As discussed in the introduction, the idea is to take the tree-approximation lemma for hyperbolic spaces [Gro87], and use a higher-rank version of it as an axiomatisation. We do this slightly differently to [Bow13a].

Given a metric space S with a ternary operator μ , a subset $Y \subset S$ is said to be *k -coarsely convex* if $\mu(y_1, y_2, s)$ lies in the k -neighbourhood of Y for all $y_1, y_2 \in Y$, $s \in S$. A map $f : (S, \mu) \rightarrow (T, \nu)$ of spaces with ternary operators is called *k -quasimedial* if $f\mu(s_1, s_2, s_3)$ lies at distance at most k from $\nu(fs_1, fs_2, fs_3)$ for all $s_1, s_2, s_3 \in S$.

Definition 7.1 (Strong coarse median). Let S be a geodesic space with a ternary operator μ . We say that (S, μ) is a *strong coarse median space* if there exist an n and a nondecreasing function κ such that the following hold.

- μ is $\kappa(1)$ -coarsely Lipschitz in each parameter.
- For each finite subset $A \subset S$ there is a CAT(0) cube complex Q of dimension at most n , and maps $f : A \rightarrow Q$, $g : Q \rightarrow S$ such that:
 - g is a $\kappa(|A|)$ -quasimedial $\kappa(|A|)$ -quasiisometric embedding.
 - Q is the (combinatorial) convex hull of $f(A)$, and $g(Q)$ is $\kappa(|A|)$ -coarsely convex.

$$-d(a, gf(a)) \leq \kappa(|A|) \text{ for all } a \in A.$$

We say that S has *rank* at most n .

The difference between this and the definition of a coarse median space of finite rank is twofold. Firstly, in coarse median spaces, the map g is only required to be quasimedial, and need not include any metric information; in the terminology of [Bow19], we are replacing “quasimorphisms” by “strong quasimorphisms”. Secondly, the complex Q here is approximating the entire *coarse median hull* of A , whereas in coarse median spaces it only has to approximate the *coarse subalgebra* generated by A . For example, let S be $\mathbf{R}^2$ with the ℓ^1 median structure. If $|A| = 2$, then in a coarse median structure we can always take Q to be a single (unit) 1-cell, but in a strong coarse median structure, Q will be the (metric) rectangle spanned by A .

In the rank-one case the two notions define the same objects, namely hyperbolic spaces, by [Bow13a, Thm 2.1]. Both *hierarchically hyperbolic* spaces [BHS21] and the spaces considered in [Bow18a] are strong coarse median spaces.

The following construction appears in [Bow18a, §6], where it is considered for coarse median spaces more generally.

Definition 7.2 (Median hull). Let S be a strong coarse median space of rank n , and let $A \subset S$. The *median hull* of A , denoted $\text{Hull } A$, is the subset $J^n(A) \subset S$, where $J^0(A) = A$ and $J^{k+1}(A) = \{\mu(a, b, s) : a, b \in J^k(A), s \in S\}$. There is a constant $k_0 = k_0(\kappa, n)$ such that $\text{Hull } A$ is k_0 -coarsely convex for every $A \subset S$.

It follows from the definition of a strong coarse median space that if $A \subset S$ is finite, then $g(Q)$ lies at a uniform Hausdorff distance from $\text{Hull } A$, in terms of k_0 and $|A|$. Hence, after increasing κ by a controlled amount, we can extend f to a $\kappa(|A|)$ -quasimedial $\kappa(|A|)$ -quasiisometry $\hat{f} : \text{Hull } A \rightarrow Q$ such that $\hat{f}$ and g are $\kappa(|A|)$ -quasiinverse.

We shall sometimes write $\mu_{abc} = \mu(a, b, c)$ in order to simplify some expressions. For instance, the equality $\mu(a, b, \mu_{cde}) = \mu(\mu_{abc}, \mu_{abd}, e)$ holds in all median algebras, and therefore the same holds up to a uniform error in terms of $\kappa(5)$ in strong coarse median spaces; see [Bow18b, §6], [NWZ19, Lem 2.18]. After increasing κ by a controlled amount, we shall therefore assume that the two above expressions differ by at most $\kappa(5)$ inside S .

7.1. CURTAINS

Just as in Section 6, we shall go via a set of geometrically defined *curtains* in order to define sets P and $\mathcal{C}$ of partitions and chains on a strong coarse median space. Throughout this section, (S, d, μ) will be a strong coarse median space of rank n with associated function κ . For each finite subset $A \subset S$, fix a choice of g, Q satisfying the assumption of Definition 7.1, and, as discussed after Definition 7.2, fix a choice of $\hat{f} : \text{Hull } A \rightarrow Q$. Let $k_1 = \max\{k_0, \kappa(2^{n+1}), \kappa(5)\}$.

Definition 7.3 (Curtains). Let $a, b \in S$, and let $A = \{a, b\}$. Let c be a chain of hyperplanes in Q of length $20nk_1^5$ such that $d_Q^\infty(c^-, c^+) = |c|$. Let

$$\mathfrak{h}^- = \{s \in S : \hat{f}\mu(a, b, s) \in c^-\} \quad \text{and} \quad \mathfrak{h}^+ = \{s \in S : \hat{f}\mu(a, b, s) \in c^+\}.$$

The *curtain* defined by a, b, c is the set $\mathfrak{h} = S \setminus (\mathfrak{h}^- \cup \mathfrak{h}^+)$, and $\mathfrak{h}^-, \mathfrak{h}^+$ are the *halfspaces* of $\mathfrak{h}$.

A *chain of curtains* is a sequence $(\mathfrak{h}_i)$ of curtains such that $\mathfrak{h}_i$ separates $\mathfrak{h}_{i-1}$ from $\mathfrak{h}_{i+1}$ for all i , in the sense that (up to orientation) $\mathfrak{h}_{i-1} \subset \mathfrak{h}_i^-$ and $\mathfrak{h}_{i+1} \subset \mathfrak{h}_i^+$.

Remark 7.4. The set of curtains we have constructed depends on the choices of cube complex approximations of pairs of points in S . In particular, it need not be preserved by $\text{Isom } S$. We

chose this set of curtains because it is in a sense the most concrete option, and the underlying arguments of this section do not hinge on the specifics. Here are three natural alternative choices that would yield $\text{Isom } S$ -equivariance with essentially no modification to any of our proofs.

- (1) Declare the set of curtains to be the set of all translates of those defined above.
- (2) Fix a constant δ , and consider all cubical approximations (of pairs of points) that have constant at most δ . Define curtains as above but using all such approximations.
- (3) Fix a constant δ , and let K be sufficiently large in terms of δ and the parameters of X . Define a curtain to be any δ -coarsely convex subset $\mathfrak{h} \subset S$ with the property that $S \setminus \mathfrak{h}$ can be written as a disjoint union of two nonempty subsets $\mathfrak{h}^-$ and $\mathfrak{h}^+$ such that $d(\text{Hull } \mathfrak{h}^-, \text{Hull } \mathfrak{h}^+) > K$.

To keep matters simple, we shall proceed with the curtains of Definition 7.3, but in applications to groups one should use one of these above options.

The next lemma shows that, although the two halfspaces of a curtain could fail to be coarsely convex, their hulls are well controlled by the cube complex used to define them.

Lemma 7.5. *Let $\mathfrak{h}$ be the curtain defined by points a, b and a chain c of hyperplanes. If $x \in J^m(\mathfrak{h}^-)$, then at most $3m\kappa(5)^2$ elements of c can separate $\hat{f}\mu(a, b, x)$ from c^- . A similar statement holds for $J^m(\mathfrak{h}^+)$. In particular, $\text{Hull } \mathfrak{h}^- \subset S \setminus \mathfrak{h}^+$.*

Proof. We proceed by induction. If $x \in J^0(\mathfrak{h}^-)$, then by definition no element of c separates $\hat{f}\mu(a, b, x)$ from c^- . Suppose that we have established the lemma for some value of m , and let $x \in J^{m+1}(\mathfrak{h}^-)$. We can write $x = \mu(y, z, s)$ for some $y, z \in J^m(\mathfrak{h}^-)$ and some $s \in S$. We then have $d(\mu(a, b, x), \mu(\mu_{aby}, \mu_{abz}, s)) \leq \kappa(5)$. Since $\hat{f}$ is a $\kappa(2)$ -quasiisometry, we get

$$d(\hat{f}\mu(a, b, x), \hat{f}\mu(\mu_{aby}, \mu_{abz}, s)) \leq \kappa(2)\kappa(5) + \kappa(2),$$

and as $\hat{f}$ is $\kappa(2)$ -quasimedial, the latter point lies at distance at most $\kappa(2)$ from $\mu_Q(\hat{f}\mu_{aby}, \hat{f}\mu_{abz}, \hat{f}s)$. By the inductive hypothesis, at most $3m\kappa(5)^2$ elements of c separate this point from c^- , and so at most $3(m+1)\kappa(5)^2$ elements of c separate $\hat{f}\mu(a, b, x)$ from c^- . $\square$

The following technical lemma can be viewed as a kind of weak Helly property.

Lemma 7.6. *Let $\mathfrak{h}_1, \dots, \mathfrak{h}_m$ be curtains. If the halfspaces $\mathfrak{h}_i^-$ intersect pairwise, then there is a point $z \in \bigcap_{i=1}^m J^{m-1}(\mathfrak{h}_i^-)$.*

Proof. For each $j \neq i$, let $x_{ij} \in \mathfrak{h}_i^- \cap \mathfrak{h}_j^-$. We have $\mu(x_{ij}, x_{ik}, x_{jk}) \in J^1(\mathfrak{h}_i^-) \cap J^1(\mathfrak{h}_j^-) \cap J^1(\mathfrak{h}_k^-)$. By repeating this type of argument (see [HHP23, Lem. 2.18], for instance), one can find a point $z \in \bigcap_{i=1}^m J^{m-1}(\mathfrak{h}_i^-)$ as desired. $\square$

Definition 7.7 (Strong crossing). We say that two curtains $\mathfrak{h}_1$ and $\mathfrak{h}_2$ *strongly cross* if all four quarterspaces $\mathfrak{h}_1^\pm \cap \mathfrak{h}_2^\pm$ are nonempty.

Proposition 7.8. *Let S be a strong coarse median space of rank n . If $\mathfrak{h}_1, \dots, \mathfrak{h}_k$ are pairwise strongly crossing curtains, then $k \leq n$.*

Proof. Suppose that $k = n + 1$. Let $p \subset \{1, \dots, n + 1\}$, and write $p(i) \in \{+, -\}$ according to whether or not $i \in p$. By Lemma 7.6, for each such p there is a point $z_p \in \bigcap_{i=1}^{n+1} J^n(\mathfrak{h}_i^{p(i)})$. Let $A = \{z_p\}$, which has cardinality 2^{n+1} . Let Q be the CAT(0) cube complex approximating A , with corresponding map $\hat{f} : \text{Hull } A \rightarrow Q$.

The sets $J^n(\mathfrak{h}_i^\pm) = \text{Hull } \mathfrak{h}_i^\pm$ are k_0 -coarsely convex, as is $\text{Hull } A$. Hence for each i and each choice of sign, say $+$, the set $B_i^+ = \text{Hull } A \cap \text{Hull } \mathfrak{h}_i^+$ is k_0 -coarsely convex. From this, one

can compute that $\hat{f}(B_i^+) \subset Q$ is $(k_0\kappa(2^{n+1}) + 3\kappa(2^{n+1}))$ -coarsely convex. Because S has coarse median rank n , we have $\dim Q \leq n$, so $\text{Hull}_Q \hat{f}(B_i^+)$ lies in the $4nk_1^2$ -neighbourhood of $\hat{f}(B_i^+)$. Similarly, $\text{Hull}_Q \hat{f}(B_i^-)$ lies in the $4nk_1^2$ -neighbourhood of $\hat{f}(B_i^-)$. The map $\hat{f}$ is a $\kappa(2^{n+1})$ -quasiisometry, so

$$d_Q(\hat{f}(B_i^-), \hat{f}(B_i^+)) \geq \frac{1}{k_1} d(\text{Hull } \mathfrak{h}_i^-, \text{Hull } \mathfrak{h}_i^+) - k_1.$$

Let a_i, b_i be the points, and $c_i \subset Q_i$ the chain used to define $\mathfrak{h}_i$, with corresponding map $\hat{f}_i$. According to Lemma 7.5, all but at most $6n\kappa(5)^2$ elements of c_i can fail to separate $\hat{f}_i\mu(a_i, b_i, \text{Hull } \mathfrak{h}_i^-)$ from $\hat{f}_i\mu(a_i, b_i, \text{Hull } \mathfrak{h}_i^+)$. Combining this with the fact that μ is $\kappa(1)$ -coarsely Lipschitz in each parameter and $\hat{f}_i$ is $\kappa(2)$ -coarsely Lipschitz, it follows that

$$\kappa(1)(\kappa(2) d(\text{Hull } \mathfrak{h}_i^-, \text{Hull } \mathfrak{h}_i^+) + \kappa(2)) + \kappa(1) \geq |c_i| - 6n\kappa(5)^2,$$

and hence $d(\text{Hull } \mathfrak{h}_i^-, \text{Hull } \mathfrak{h}_i^+) \geq \frac{1}{k_1^2} |c_i| - 6n - 2 \geq 12nk_1^3$. Together with the above lower bounds, we find that

$$d_Q(\text{Hull}_Q \hat{f}(B_i^-), \text{Hull}_Q \hat{f}(B_i^+)) \geq \frac{1}{k_1} (12nk_1^3) - k_1 - 2(4nk_1^2)$$

is positive. As the sets $\text{Hull}_Q \hat{f}(B_i^\pm)$ are convex in the CAT(0) cube complex Q , there is a hyperplane w_i of Q separating the two. In particular, w_i separates $\hat{f}(z_p)$ from $\hat{f}(z_q)$ whenever $p(i) \neq q(i)$. But then the hyperplanes $w_1, \dots, w_{n+1}$ must pairwise cross, which is impossible because $\dim Q \leq n$. $\square$

7.2. THE INJECTIVE DUAL

We go from curtains to a set of partitions of S as in Section 6. More precisely, each curtain $\mathfrak{h}$ induces two natural bipartitions of S , namely $(\mathfrak{h}^- \cup \mathfrak{h}, \mathfrak{h}^+)$ and $(\mathfrak{h}^-, \mathfrak{h} \cup \mathfrak{h}^+)$. Let P be the set of all bipartitions induced in this way.

There are multiple dualisable systems that one can define using P . Here we shall consider what is perhaps the largest reasonable choice; we shall see more in Section 7.3. That is, we let $\mathcal{C}$ be the set of all chains $\{h_i\} \subset P$ such that there is a chain of curtains $(\mathfrak{h}_i)$ with $\mathfrak{h}_i$ inducing h_i . Throughout this section, $(X, d_{\mathcal{C}})$ will denote the $\mathcal{C}$ -dual of S . Because we are allowing all chains of curtains, it is easy to see that $\mathcal{C}$ is gluable.

Lemma 7.9. *$\mathcal{C}$ is a 2-gluable system of chains.*

Proof. Let $c_1 = \{\dots, h_{-2}, h_{-1}\}$ and $c_2 = \{k_1, k_2, \dots\}$ be elements of $\mathcal{C}$ such that $\{\dots, h_{-1}, k_1, \dots\}$ is a chain, $c_2 \subset h_{-1}^+$, and $c_1 \subset k_1^-$. Let $(\mathfrak{h}_i)$ induce c_1 and let $(\mathfrak{k}_i)$ induce c_2 . We have $\mathfrak{k}_2 \subset \mathfrak{k}_1^+ \subset k_1^+ \subset h_{-1}^+ \subset \mathfrak{h}_{-2}^+$, and similarly $\mathfrak{h}_{-2} \subset \mathfrak{k}_2^-$. Hence $c_1 \cup c_2 \setminus \{h_{-1}, k_1\} \in \mathcal{C}$. $\square$

The results of Section 4 therefore show that X is a coarsely injective space with a good bicombing by normal wall paths. Normal wall paths are median paths, so it is desirable to relate the median on X given by Lemma 3.6 to the coarse median on S .

Lemma 7.10. *The map $S \rightarrow X$ is a quasiisometric embedding that is 3-quasimedial.*

Proof. Let $s, t \in S$. From the CAT(0) cube complex Q approximating $\text{Hull}\{s, t\}$, we obtain a chain of curtains separating s from t whose length is linearly lower bounded by $d(s, t)$. Conversely, because $d(\mathfrak{h}^-, \mathfrak{h}^+) \geq 1$ for every curtain $\mathfrak{h}$, any element of $\mathcal{C}$ separating s from t can have length at most $d(s, t) + 2$. Thus $S \rightarrow X$ is a quasiisometric embedding.

To show that $S \rightarrow X$ is quasimedial, let $s_1, s_2, s_3 \in S$, and suppose that $\mathfrak{h}$ is a curtain with $s_1, s_2 \in \mathfrak{h}^-$. Because $\mu_S(s_1, s_2, s_3) \in J^1(\mathfrak{h}^-)$, Lemma 7.5 tells us that $\mu_S(s_1, s_2, s_3) \notin \mathfrak{h}^+$. This shows that any chain of curtains separating $\mu_S(s_1, s_2, s_3)$ from the majority of $\{s_1, s_2, s_3\}$ has length at most one. Hence $S \rightarrow X$ is 3-quasimedial. $\square$

Proposition 7.11. *Let S be a strong coarse median space of rank n , and let X be the dual space described above. The image of S in X is coarsely dense.*

The proof of this proposition is similar in spirit to that of Proposition 5.8, but is more complicated for two reasons. Firstly, we only have a coarse median, rather than points that roughly realise distances. Secondly, the separation property provided by Proposition 7.8 is weaker than L -separation.

Here is an outline of the argument. Given $x \in X$, we find a sequence of points $s_i \in S$ with the property that for any $c \in \mathcal{C}$ separating s_i from x , only a uniform number of elements of c can fail to separate x from all s_j with $j \leq i$. From this it will follow that if $c_i \in \mathcal{C}$ realises $\mathbf{d}_{\mathcal{C}}(s_i, x)$, then the tail of c_i either is empty or crosses the tails of all c_j with $j < i$. Proposition 7.8 then shows that the tail of c_{n+1} must be empty, which bounds $\mathbf{d}_{\mathcal{C}}(s_{n+1}, x)$.

The most difficult part is the construction of the s_i . The reason for requiring the above property involving all $j \leq i$ is that we wish to avoid a situation where the chains c_i “face” each other, as in that case there would be no way to make the process terminate. Informally, we need to ensure that s_{i+1} is not “on the opposite side” of x to s_i . To do this, we use an auxiliary point $t_i \in S$ and “project” x to the hull of $\{s_i, t_i\}$ to obtain s_{i+1} . This is reminiscent of the fact that in a CAT(0) cube complex, every wall separating a point z from its gate to a convex set C actually separates z from all of C . Again, we have to be slightly careful with this step, because x is not an element of S , so its gate in the sense of Section 3.2 might not be.

Proof of Proposition 7.11. Fix a sufficiently large constant C , which could be explicitly determined from the below arguments in terms of n and the parameters of strong coarse median space S .

Let $x \in X$, and choose an arbitrary point $s_1 \in S$. Let $c_1 \in \mathcal{C}$ realise $\mathbf{d}_{\mathcal{C}}(s_1, x)$. If $|c_1| \leq C$ then we are done. Otherwise there exists some point $t_1 \in S$ such that $t_1 \in c_1^+$, along with x . Write $A_1 = \{s_1, t_1\}$ and consider the corresponding finite CAT(0) cube complex $Q_1 = Q(A_1)$.

Given a curtain $\mathfrak{h}$ separating s_1 from t_1 (not necessarily arising from Q_1), there are two convex subcomplexes $H^- = \text{Hull}_{Q_1}(\hat{f}_1\mu(s_1, \mathfrak{h}^-, t_1))$ and $H^+ = \text{Hull}_{Q_1}(\hat{f}_1\mu(s_1, \mathfrak{h}^+, t_1))$, which may overlap. Let $h_1, h_2 \in P$ be walls coming from curtains $\mathfrak{h}_1$ and $\mathfrak{h}_2$ that separate s_1 from t_1 , and suppose that the halfspace $x(h_i)$ corresponds to either $\mathfrak{h}_i^+$ or $\mathfrak{h}_i^-$ for both i (as opposed to, say, $\mathfrak{h}_i^+ \cup \mathfrak{h}_i^-$). After relabelling the orientations of the h_i , let us assume for concreteness that $x(h_i) = \mathfrak{h}_i^-$. Since x is a filter, there must be some $p \in S$ such that $p \in \mathfrak{h}_1^- \cap \mathfrak{h}_2^-$. We then have $\hat{f}_1\mu(p, s_1, t_1) \in H_1^- \cap H_2^-$. That is, the convex subsets of Q_1 corresponding to $\mathfrak{h}_1^-$ and $\mathfrak{h}_2^-$ intersect.

Now consider the set $P_1 \subset P$ consisting of all walls h coming from curtains $\mathfrak{h}$ that separate s_1 from t_1 and such that $x(h)$ corresponds to either $\mathfrak{h}^-$ or $\mathfrak{h}^+$. By the previous paragraph, the pair (x, P_1) determines a set of pairwise intersecting convex subcomplexes of the finite CAT(0) cube complex Q_1 . By Helly’s theorem, there is a point σ_2 in the total intersection of those subcomplexes. Let $s_2 = g_1(\sigma_2) \in \text{Hull}(A_1)$.

Consider a chain $c \in \mathcal{C}$ separating s_2 from x . We claim that only a uniformly bounded number of elements of c can separate s_2 from either s_1 or t_1 . Firstly, boundedly many can

separate s_2 from both s_1 and t_1 . Indeed, $\sigma_2 = \mu_{Q_1}(\hat{f}_1(s_1), \hat{f}_1(t_1), \sigma_2)$ and g_1 is quasimedial, so s_2 is uniformly close to $\mu(s_1, s_2, t_1)$, and Lemma 7.10 shows that $S \rightarrow X$ is 3-quasimedial. Secondly, a long sequence in c separating s_2 from, say, s_1 would lead to a long sequence of curtains separating t_1 from x , from which we could find a long such sequence in P_1 , which is impossible.

Thus, if we let $c_2 \in \mathcal{C}$ realise $d_{\mathcal{C}}(s_2, x)$, then all but uniformly many elements of c_2 separate $\{s_1, t_1, s_2\}$ from x . Moreover, those elements have to occur at the end of c_2 that is closest to s_2 . By Lemma 7.9, the final eight elements of c_2 cannot be contained in the positive halfspace of the fifth-last element of c_1 , for then we would be able to elongate c_1 . Thus, either $|c_2| \leq C$, in which case we are done, or the final five elements of c_2 cross the final five elements of c_1 .

From here we proceed inductively. Let $k \leq n$ and suppose that we have constructed $s_1, \dots, s_k, t_1, \dots, t_{k-1}$, such that

- the final five elements of c_i cross the final five elements of c_j , where $c_l \in \mathcal{C}$ realises $d_{\mathcal{C}}(s_l, x)$;
- all but a uniformly bounded number of elements of c_i separate $\{s_1, \dots, s_{i-1}, t_1, \dots, t_{i-1}\}$ from x .

Let $c_k \in \mathcal{C}$ realise $d_{\mathcal{C}}(s_k, x)$. There exists $t_k \in S$ such that $t_k \in c_k^+$. Write $A_k = \{s_1, \dots, s_k, t_1, \dots, t_k\}$ and let $Q_k = Q(A_k)$ be the corresponding finite CAT(0) cube complex, with $g_k : Q_k \rightarrow \text{Hull } A_k$. Referring to Lemma 2.1, write $\mathbf{g}_k$ for the map

$$z \mapsto \mu(s_k, z, \mu(t_k, z, \mu(s_{k-1}, z, \mu(t_{k-1}, z, \mu(s_{k-2}, \dots, \mu(s_1, z, t_1))) \dots)).$$

If $z \in S$, then up to a uniform error we have $\mathbf{g}_k(z) \in \text{Hull}(A_k)$ for all $z \in X$. Moreover, in view of Lemma 2.1, if $z \in S$, then by considering $Q' = Q(A_k \cup \{z\})$ we see that $\mathbf{g}_k(z)$ is uniformly close to the image in $\text{Hull}(A_k \cup \{z\})$ of the gate of z to $\text{Hull}_{Q'}(f'(A_k))$ in Q' , where the constant depends on $k \leq n$.

Given a curtain $\mathfrak{h}$ that separates some pair of elements of A_k , there are two convex subcomplexes $H^- = \text{Hull}_{Q_k}(\hat{f}_k \mathbf{g}_k(\mathfrak{h}^-))$ and $H^+ = \text{Hull}_{Q_k}(\hat{f}_k \mathbf{g}_k(\mathfrak{h}^+))$. Let $P_k \subset P$ be the set of walls h coming from curtains $\mathfrak{h}$ that separate a pair of elements of A_k and such that $x(h)$ corresponds to either $\mathfrak{h}^+$ or $\mathfrak{h}^-$. For each $h \in P_k$, let $\text{sign}_x(h) \in \{-, +\}$ be such that $x(h) = h^{\text{sign}_x(h)}$. Given $h_1, h_2 \in P_k$, there is a point $p \in S$ lying in $\mathfrak{h}_1^{\text{sign}_x(h_1)} \cap \mathfrak{h}_2^{\text{sign}_x(h_2)}$, and hence the convex subcomplexes $H_1^{\text{sign}_x(h_1)}$ and $H^{\text{sign}_x(h_2)}$ of the finite CAT(0) cube complex Q_k intersect. From this pairwise intersection, Helly's theorem provides a point $\sigma_{k+1} \in Q_k$ that lies in $H^{\text{sign}_x(h)}$ for all $h \in P_k$. Let $s_{k+1} = g_k(\sigma_{k+1})$.

Consider an arbitrary chain $c \in \mathcal{C}$ separating s_{k+1} from x . As $\sigma_{k+1} \in Q_k$, Lemma 2.1 tells us that s_{k+1} lies at a uniform distance from $\mathbf{g}_k(s_{k+1})$, and hence only uniformly many elements of c can separate s_{k+1} from A_k . Furthermore, any subchain of c separating s_{k+1} from a subset $B \subset A_k$ yields a slightly shorter chain of elements of P_k separating s_{k+1} from B . But g_k is quasimedial, and Lemma 7.10 states that $S \rightarrow X$ is 3-quasimedial, so chains in P_k that separate s_{k+1} from x have uniformly bounded length. This shows that, for any chain in $\mathcal{C}$ separating s_{k+1} from x , all but a uniformly finite number (depending on $k \leq n$) of elements must separate x from $s_{k+1} \cup A_k$.

Let $c_{k+1} \in \mathcal{C}$ realise $d_{\mathcal{C}}(s_{k+1}, x)$. If $|c_{k+1}| \leq C$ then we are done, so suppose otherwise. If the final eight elements of c_{k+1} were contained in the positive halfspace of the fifth-last element of c_i for some $i \leq k$, then Lemma 7.9 would tell us that we could have made a longer choice of c_i , which is impossible. Thus the final five elements of c_{k+1} cross the final five elements of c_i for all $i \leq k$.

Suppose that the above process has not terminated before the final step $k = n$, because $d_C(s_i, x) > C$ for all $i \leq k$. If it is also the case that $d_C(s_{n+1}, x) > C$, then for each $i \leq k+1$ let $h_1^i, \dots, h_5^i$ be the final five elements of c_i . If $i \neq j$, then every h_l^i crosses every h_l^j . This implies that the curtains $h_3^1, \dots, h_3^{k+1}$ pairwise strongly cross, which contradicts Proposition 7.8. So we must have $d_C(s_i, x) \leq C$ for some $i \leq n+1$. $\square$

Remark 7.12. One can tweak the definition of $\mathbf{g}_k$ in the above proof to provide a notion of gate map to hulls of finite subsets of general coarse median spaces of finite rank, with constants independent of the cardinality of the finite subset.

The following is the main result of this section. It includes the notion of (n, δ) -hyperbolicity, which was introduced in [JL22] as a higher-rank form of negative curvature. We shall not discuss this notion in detail here, but it replaces the four-point condition for hyperbolicity by a $(2n+2)$ -point condition.

Theorem 7.13. *Let S be a strong coarse median space of rank n , and let X be as constructed above. The map $S \rightarrow X$ is a quasimedial quasiisometry to a coarsely injective space. Moreover, X is (n, δ) -hyperbolic.*

Proof. X is coarsely injective by Theorem 4.9, and Lemma 7.10 together with Proposition 7.11 shows that $S \rightarrow X$ is a quasimedial quasiisometry. The fact that X is (n, δ) -hyperbolic is a combination of [Bow13a, Thms 2.2, 2.3], which control the asymptotic rank of X , and then [JL22, Thm 1.4]. $\square$

By a *strong-coarse-median group*, we mean a finitely generated group G with a G -equivariant ternary operator making it a strong coarse median space. We refer to [JL22] for the notion of having *slim simplices*, which generalises that of having slim triangles as in a hyperbolic space.

Corollary 7.14. *If G is a strong-coarse-median group of rank n , then G acts properly coboundedly on an injective metric space with slim n -simplices.*

Proof. Referring to Remark 7.4, the construction of X can be done in a G -equivariant way. In view of Theorem 7.13, it then follows from [HHP23, Prop. 1.1] and [Lan13, Prop. 3.7] that G acts properly coboundedly on the injective hull of X , which has the slim simplex property by [JL22, Thm 1.3]. $\square$

Theorem 7.13 should be compared with [HHP23, Cor. 3.6]. The construction of the metric in that paper works for any strong coarse median space, but the argument that it is coarsely injective requires the space to be hierarchically hyperbolic, and relies on the hierarchy structure. Additionally, there is no equivalent of normal wall paths in that setting. In any case, since hierarchically hyperbolic groups are strong-coarse-median groups [BHS21], Corollary 7.14 recovers the main result of [HHP23], that hierarchically hyperbolic groups act properly coboundedly on injective spaces. Corollary 7.14 generalises [JL22, Cor. 1.6].

7.3. HYPERBOLIC MODELS

One can also construct hyperbolic spaces from the set P of partitions constructed at the beginning of Section 7.2, in a similar manner to Section 6. Let us say that two disjoint curtains h_1 and h_2 are R -separated if there is no chain c of curtains of cardinality greater than R whose elements all strongly cross both h_1 and h_2 . By an R -chain of curtains, we mean a chain of curtains such that each pair is R -separated.

Let $\mathcal{C}$ be as in Section 7.2. For each natural number R , we define a dualisable system $\mathcal{C}_R$ on S by letting $\mathcal{C}_R \subset \mathcal{C}$ consist of all elements that are induced by R -chains of curtains. It is routine to check the following (cf. Lemma 6.6).

Lemma 7.15. *The sequence $(\mathcal{C}_R)$ is a graded system on (S, P) , and each $\mathcal{C}_R$ is 3-gluable.*

Let Y be the graded dual of S with respect to $(\mathcal{C}_R)$. By Propositions 5.14 and 5.15, Y is a roughly geodesic hyperbolic space. We already described a way to construct a hyperbolic space associated to S in Section 6, but Y has the advantage that it can more easily be compared with existing hyperbolic spaces in known examples, as we now briefly discuss.

In [ABD21], Abbott–Behrstock–Durham constructed, for each hierarchically hyperbolic group G , a hyperbolic space Z witnessing a largest acylindrical action of G . A detailed discussion is beyond our scope, so we refer the reader to [BHS19] for background on hierarchical hyperbolicity, and to [ABD21, §3] for details about the construction of Z .

Proposition 7.16. *Let G be a hierarchically hyperbolic group. Let Y be the hyperbolic space constructed above, and let Z be the hyperbolic space constructed in [ABD21]. The two are G -equivariantly quasiisometric.*

Outline of proof. Z is a cone-off of G , and Y is the metric quotient of an equivariant pseudometric on G . We show that the natural maps $G \rightarrow Z$ and $G \rightarrow Y$ yield a quasiisometry between Y and Z , and this map is automatically G -equivariant. Let $s, t \in G$ and let c be a chain of curtains obtained from the cube complex approximating $\text{Hull}\{s, t\}$. If two elements of c are such that their images under the projection map $\pi_Z : G \rightarrow Z$ are far apart, then those curtains are R -separated for sufficiently large R . This shows that $Y \rightarrow Z$ is coarsely Lipschitz.

For the other direction, consider a subpath σ' of the normal wall path σ from s to t such that $\pi_Z \sigma'$ has small diameter. There are two cases. The first is that there is no *relevant domain* U for $\{s, t\}$ with ρ_Z^U near $\pi_Z \sigma'$ (see [BHS19]). In this case, σ' itself has small diameter in G . Otherwise there is such a U , and then [ABD21, Thm 3.7] shows that U is one factor of a nontrivial product. This product structure prevents U from contributing any separated curtains. Thus separated curtains can only arise when σ makes progress in Z , so $Z \rightarrow Y$ is coarsely Lipschitz. $\square$

In particular, Proposition 7.16 shows that if S is the mapping class group of a surface, then Y is quasiisometric to the curve graph of that surface.

8. OTHER DIRECTIONS

We believe that there should be many situations where the construction of Section 3 can be applied. Here we suggest a few possible avenues.

CAT(0) spaces and cube complexes. Much of this article stems from constructions in CAT(0) spaces [PSZ22], though the combinatorial perspective in terms of ultrafilters used here is rather different. That said, if one uses the curtains constructed in [PSZ22] to induce a set of partitions as in Sections 6 and 7, then the graded dual with respect to the systems of L -separated curtains should be essentially equivalent to the *curtain model*.

If one starts with S a CAT(0) cube complex and P the set of hyperplanes, then, as discussed in Example 3.5, taking $\mathcal{C}$ to be the set of all chains simply makes X the Helly thickening of S . Letting $\mathcal{C}_L$ be the set of chains of pairwise L -separated hyperplanes, one obtains a space very similar to that of [Gen20, §6.6] (the difference being in the precise definition of separation). The sequence $(\mathcal{C}_L)$ is a graded system on S , so one could investigate the graded dual, which is a more natural hyperbolic space than the curtain model in this setting.

The Bestvina–Bromberg–Fujiwara construction. In influential work, Bestvina–Bromberg–Fujiwara gave a general method for assembling a collection of metric spaces in a quasitree-like

fashion [BBF15], by gluing them at bounded sets. This is often used when the spaces in question are hyperbolic, quasitrees, or even quasilines [Sis18, BBF21, BHS17].

Since geodesic hyperbolic spaces are quasiisometric to the dual space of some dualisable system, in many cases one can interpret the assumptions in the BBF construction in terms of extensions of the partitions of the component spaces. Whilst this alternative axiomatisation may not be directly useful, it seems plausible that it could open the door to generalisations that allow for gluings along larger subspaces: a coarse point could perhaps be replaced by a subset that is approximately gated in the sense of Section 3.2.

Curve graph of a surface. Let $\text{MCG}(\Sigma)$ be the mapping class group of a finite-type surface Σ . One corollary of Proposition 7.16 is that the hyperbolic space constructed for $\text{MCG}(\Sigma)$ in Section 7.3 is equivariantly quasiisometric to the curve graph of Σ . The mapping class group also admits a proper cobounded actions on an injective space S , either via [HHP23] or Theorem 7.13. According to [SZ22, Thm A], Morse subsets of injective metric spaces are strongly contracting, so orbits of pseudo-Anosovs on S are strongly contracting. The results of Section 6 therefore produce a hyperbolic space X dual to S on which $\text{MCG}(\Sigma)$ acts. Is X also quasiisometric to the curve graph of Σ ? This seems especially likely in view of Theorem 6.15 and [DT15].

Hyperbolic models for other groups. There are various “nonpositively curved” groups for which hyperbolic models have been constructed, such as $\text{Out } F_n$ [HV98, KL09, HM13] and various Artin–Tits groups [KK13a, CW21a, Mor21, MP22]. It would be interesting to know whether some of these models can be (coarsely) reconstructed from a suitable set of walls. For instance, can one find a natural set of curtains in Culler–Vogtmann outer space, and does this reproduce, say, the free factor complex? How about for the Deligne complex?

Moreover, in view of Theorems C and D, it would also be desirable to have more examples of metric spaces where all Morse geodesics are strongly contracting.

Higher-rank hyperbolicity. In Section 5.1, we saw that if a dualisable system is separated and gluable then fairly straightforward combinatorial arguments show four-point hyperbolicity of the dual. In [JL22], Jørgensen–Lang introduced a family of higher-rank generalisations of the four-point inequality, and a space satisfying their $(2n + 2)$ -point inequality is said to be (n, δ) -hyperbolic.

For a higher-rank version of the L -separation condition, let us say that a dualisable system of chains is (n, L) -separated if, whenever $c_1, \dots, c_{n+1}$ are elements of $\mathcal{C}$ such that every element of c_i crosses every element of c_j for all i, j , we necessarily have $|c_k| \leq L$ for some k . Is it true that if $\mathcal{C}$ is gluable and (n, L) -separated then X is (n, δ) -hyperbolic? Does an analogue of Proposition 5.8 hold? Note that this (n, L) -separation assumption holds in the setting of strong coarse median spaces, by Proposition 7.8. In that setting, though, (n, δ) -hyperbolicity follows *a posteriori* from coarse injectivity of X (Theorem 4.9) and coarse density of S (Proposition 7.11).

Metric quotients. One very useful consequence of the duality between CAT(0) cube complexes and discrete wallspaces [CN05, Nic04] is a simple trick for producing quotients of a given CAT(0) cube complex Q . Namely, one takes a subset P' of the walls of Q , and lets Q' be the cube complex dual to P' . This procedure, known as a *restriction quotient*, was introduced to CAT(0) cube complexes in [CS11], though it appeared earlier for median graphs [Mul78].

It is easy to see that restriction quotients can be taken in the generality of Section 3. More precisely, let $\mathcal{C}$ be a dualisable system for a set with walls (S, P) , and let X be the $\mathcal{C}$ -dual.

Given a subset $P' \subset P$, the set $\mathcal{C}'$ consisting of all elements of $\mathcal{C}$ supported on P' (i.e., elements $\{h_1 \cdots, h_n\} \in \mathcal{C}$ with each $h_i \in P'$) is a dualisable system for (S, P') , and there is a natural quotient map from X to the $\mathcal{C}'$ -dual of S .

Many desirable properties of a dualisable system, such as gluability and the property of being a system of chains, are preserved by this restriction, so often the quotient X' will have similar properties to X . For instance, if $\mathcal{C}$ is the set of all chains, so that X is a Helly graph, then X' will also be a Helly graph.

Random walks. A common application for producing actions of a finitely generated group G on hyperbolic spaces is that it can yield information about random walks on G [Kai00, MT18, QRT20]. In the very general setting where G acts properly with a strongly contracting element on some geodesic space, a combination of Theorem 6.8 and Proposition 6.13 with [CFFT22, Thm 1.2] shows that if μ is a generating probability measure on G with finite entropy, then the Poisson boundary of G is modelled by $(\partial X, \nu)$, where X is the hyperbolic space constructed in Section 6 and ν is the hitting measure on the Gromov boundary ∂X of the random walk driven by μ .

For more precise information about the limiting behaviour of a random walk, one can ask whether a central limit theorem holds [Bjö10, Hor18, FLM21]. A new approach to this type of problem was introduced by Benoist–Quint [BQ16a, BQ16b], and this was used together with [PSZ22] by Le Bars to establish a central limit theorem for random walks on groups acting on CAT(0) spaces [Le 22, Le 23]. Since the construction of Section 6 bears many geometric similarities to that of [PSZ22], it is natural to ask whether one can work along similar lines to prove a central limit theorem for random walks on groups with a contracting element.

A continuous variant. Spaces with walls can be generalised to spaces with *measured walls* [CMV04], and these still exhibit a duality with median metric spaces [CDH10, Fio20]. One could consider a continuous generalisation of the constructions of this article. For this, measures cease to be appropriate, because, for instance, the union of two chains need not be a chain, and so the set $\mathcal{C}$ will generally not be a σ -algebra. However, the property of being a measure is not really an essential feature for the construction, and one can just request a function $\nu : \mathcal{C} \rightarrow [0, \infty]$ satisfying certain compatibility criteria.

We believe that many of the statements in the present article would then admit continuous formulations. One concrete question would be: is there an alternative construction of the injective hull of a metric space S that can be achieved by letting P be the set of all balls in S and taking an appropriate function ν ?

APPENDIX. QUASIMORPHISMS (WITH DAVIDE SPRIANO)

The purpose of this appendix is to study the vector space $\widetilde{\text{QM}}(\Gamma)$ of (nontrivial, homogeneous) quasimorphisms of groups Γ acting coboundedly on spaces with strongly contracting geodesics. We use the construction of Section 6 together with the Bestvina–Fujiwara criterion [BF02]. The first step is the following, which may be of independent interest.

Proposition A.1. *Let S be a non-hyperbolic geodesic space with the property that for each Morse gauge M there exists D such that every M -Morse geodesic in S is D -strongly contracting. Suppose that S contains some biinfinite strongly contracting geodesic. For every group $\Gamma < \text{Isom } S$ acting coboundedly on S there exists q such that for each D , there is a strongly contracting element $g \in \Gamma$ that is not D -strongly contracting, and g has a q -quasiaxis.*

Proof. Let X be the hyperbolic space constructed in Section 6. The change-of-metric map $(S, d) \rightarrow (X, D)$ is coarsely Lipschitz and $\text{Isom } S$ -equivariant, and Theorem 6.8 shows that

every contracting geodesic in S is quantitatively quasiisometrically embedded in X . Let D_0 be such that there is a D_0 -strongly contracting biinfinite geodesic $A \subset S$. Let $s_0 \in A$, and fix any number $D > D_0$.

According to [GHP⁺23, Cor. 3.6], for each Morse gauge M there is a Morse gauge M' , a constant r , and an M' -Morse ray $\beta \subset S$ emanating from s_0 , such that for any M -Morse ray α emanating from s_0 , we have

$$\text{diam}(\pi_\alpha \beta) \leq r \quad \text{and} \quad \text{diam}(\pi_\beta \alpha) \leq r,$$

where $\pi_\alpha : S \rightarrow \alpha$ is a map such that $D(s, \pi_\alpha(s)) = D(s, \alpha)$, and π_β satisfies a similar property. Note that although the statement in [GHP⁺23] assumes that S is a group, that property is not used in the proof: all that is required is that $S \rightarrow X$ is coarsely Lipschitz.

Let M be a sufficiently large Morse gauge, defined in terms of D and the correspondence between strong contraction and Morseness. Let M' , r , and β be obtained from M as above, and let R be sufficiently large in terms of M' , r , and the coboundedness of Γ .

Let s_1 be a point in β with $d(s_1, s_0) > R$, and let β' be the subsegment of β from s_0 to s_1 . By coboundedness of Γ , there is a translate A' of A passing uniformly close to s_1 . Because A is a geodesic, only one direction of A' can fellow-travel β' for more than a uniformly bounded distance. Let s_2 be a point of A' in the other direction that has $d(s_2, s_1) > R$. By coboundedness of Γ , there exists $h \in \Gamma$ translating s_0 into a uniform neighbourhood of s_2 . We claim that if M and R were chosen large enough, then h is strongly contracting, but not D -strongly contracting.

Consider the points $\langle h \rangle \cdot s_0$. We can connect them by the $\langle h \rangle$ -translates of the union of β' with a perturbation α of a subsegment of A' with endpoints near s_1 and hs_0 . Let γ be the $\langle h \rangle$ -invariant path obtained in this way. We emphasise that, by construction, β' and α do not fellow-travel, nor do α and $h\beta'$ by [GHP⁺23, Cor. 3.6].

By [SZ22, Prop. 4.7], the space S is *Morse local-to-global* in the sense of [RST22]. The construction of $h^n \beta'$ ensures that it fellow-travels with each of $h^{n-1} \alpha$ and $h^n \alpha$ for a uniformly bounded amount of time. Hence γ is a uniform quasixis for h . Hence γ is obtained as the concatenations of sufficiently long Morse segments such that consecutive ones do not fellow-travel. Thus γ is locally a Morse quasigeodesic, and so a (global) Morse quasigeodesic by the Morse local-to-global property, for some sufficiently large Morse gauge. However, γ fellow-travels with a long initial subsegment of β , and so is not M -Morse. The result follows from the quantitative relation between Morseness and strong contraction that we are assuming. $\square$

Theorem A.2. *Let S be a non-hyperbolic geodesic space with the property that for each M there exists D such that every M -Morse geodesic in S is D -strongly contracting. Suppose that S contains a biinfinite strongly contracting geodesic. For every group $\Gamma < \text{Isom } S$ acting coboundedly on S , the space $\widetilde{\text{QM}}(\Gamma)$ is infinite-dimensional.*

Proof. Let X be the δ -hyperbolic space constructed in Section 6, with natural map $S \rightarrow X$. According to Proposition 6.7, there is a constant q such that every geodesic in S defines an unparametrised q -quasigeodesic in X . Let $B = B(q, \delta)$ be the constant of [BF02, p.72], and let n be a sufficiently large constant, defined in terms of B .

By Proposition A.1, there is some strongly contracting $g \in \Gamma$. Let D be such that g is D -strongly contracting. By the same lemma, there is some strongly contracting $h \in \Gamma$ that is not nD -strongly contracting. Moreover, g and h have uniform-quality quasixes.

By Theorem 6.8, both g and h act loxodromically on X . Let A_g and A_h be projections to X of uniform quasixes in S , which have uniform constants in terms of q . Given a constant

R , suppose that there is some $k \in \Gamma$ such that some subsegment $I \subset A_g$ of length at least R lies in the B -neighbourhood of kA_h in X .

Let $s, t \in S$ be the endpoints of I , and let $s', t' \in kA_h$ have $D(s, s'), D(t, t') \leq B$. Let α be a geodesic in S from s to t , which is uniformly strongly contracting in terms of D . Moreover, since $S \rightarrow X$ is coarsely Lipschitz, the length of α is lower-bounded in terms of R . It follows that there is a chain of curtains dual to α of length lower-bounded in terms of $\frac{R}{D}$ that separate s' from t' . If R is chosen to be sufficiently large in terms of the S -translation length of h , then we can apply Lemma 6.6 to the khk^{-1} -translates of that chain to obtain a lower bound on the X -translation length of h that is a uniform multiple of D . According to Theorem 6.8, that gives an upper-bound on the strong-contracting constant of h that is a uniform multiple of D . The choice of n ensures that this is a contradiction of the fact that h is not nD -strongly contracting.

Thus there is an upper bound on the length of subsegments of A_g that can lie in the B -neighbourhoods of Γ -translates of A_h . This shows that g and h are not equivalent in the sense of [BF02, p.72]. The result is given by [BF02, Thm 1]. $\square$

REFERENCES

- [AB90] Stephanie B. Alexander and Richard L. Bishop. The Hadamard-Cartan theorem in locally convex metric spaces. *Enseign. Math.* (2), 36(3-4):309–320, 1990. (Cited p. 1).
- [AB95] Juan M. Alonso and Martin R. Bridson. Semihyperbolic groups. *Proc. London Math. Soc.* (3), 70(1):56–114, 1995. (Cited p. 5).
- [ABD21] Carolyn Abbott, Jason Behrstock, and Matthew G. 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. (Cited pp. 3, 6, and 38).
- [ACT15] Goulmira N. Arzhantseva, Christopher H. Cashen, and Jing Tao. Growth tight actions. *Pacific J. Math.*, 278(1):1–49, 2015. (Cited p. 2).
- [Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning. (Cited p. 7).
- [Alg11] Yael Algom-Kfir. Strongly contracting geodesics in outer space. *Geom. Topol.*, 15(4):2181–2233, 2011. (Cited p. 2).
- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.*, 122:1–64, 2015. (Cited pp. 2 and 39).
- [BBF21] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Proper actions on finite products of quasi-trees. *Ann. H. Lebesgue*, 4:685–709, 2021. (Cited p. 39).
- [BBFS20] Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. Acylindrical actions on projection complexes. *Enseign. Math.*, 65(1-2):1–32, 2020. (Cited p. 2).
- [BCK⁺23] Sahana Balasubramanya, Marissa Chesser, Alice Kerr, Johanna Mangahas, and Marie Trin. (non-)recognizing spaces for stable subgroups. *arXiv:2311.15187*, 2023. (Cited pp. 3 and 30).
- [BF92] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992. (Cited p. 1).
- [BF02] Mladen Bestvina and Koji Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geom. Topol.*, 6:69–89, 2002. (Cited pp. 2, 40, 41, and 42).
- [BF14] Mladen Bestvina and Mark Feighn. Hyperbolicity of the complex of free factors. *Adv. Math.*, 256:104–155, 2014. (Cited p. 3).
- [BH83] Hans-J. Bandelt and Jarmila Hedlíková. Median algebras. *Discrete Math.*, 45(1):1–30, 1983. (Cited p. 8).
- [BHS17] 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. (Cited p. 39).
- [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. (Cited pp. 5 and 38).

- [BHS21] Jason Behrstock, Mark Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):909–996, 2021. (Cited pp. 4, 5, 32, and 37).
- [Bjö10] Michael Björklund. Central limit theorems for Gromov hyperbolic groups. *J. Theoret. Probab.*, 23(3):871–887, 2010. (Cited p. 40).
- [BM11] Jason Behrstock and Yair N. Minsky. Centroids and the rapid decay property in mapping class groups. *J. Lond. Math. Soc. (2)*, 84(3):765–784, 2011. (Cited pp. 4 and 31).
- [Bow98] Brian H. Bowditch. A topological characterisation of hyperbolic groups. *J. Amer. Math. Soc.*, 11(3):643–667, 1998. (Cited p. 7).
- [Bow08] Brian H. Bowditch. Tight geodesics in the curve complex. *Invent. Math.*, 171(2):281–300, 2008. (Cited p. 2).
- [Bow13a] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013. (Cited pp. 4, 31, 32, and 37).
- [Bow13b] Brian H. Bowditch. Invariance of coarse median spaces under relative hyperbolicity. *Math. Proc. Cambridge Philos. Soc.*, 154(1):85–95, 2013. (Cited p. 31).
- [Bow18a] Brian H. Bowditch. Convex hulls in coarse median spaces. *Preprint available at homepages.warwick.ac.uk/~masgak/papers/hulls-cms.pdf*, 2018. (Cited pp. 4 and 32).
- [Bow18b] Brian H. Bowditch. Large-scale rigidity properties of the mapping class groups. *Pacific J. Math.*, 293(1):1–73, 2018. (Cited p. 32).
- [Bow19] Brian H. Bowditch. Quasiflats in coarse median spaces. *Preprint available at homepages.warwick.ac.uk/~masgak/papers/quasiflats.pdf*, 2019. (Cited pp. 4 and 32).
- [Bow21] Brian H. Bowditch. A Cartan–Hadamard theorem for median metric spaces. *Preprint available at homepages.warwick.ac.uk/~masgak/papers/ch-median.pdf*, 2021. (Cited p. 1).
- [Bow22] Brian H. Bowditch. Median algebras. *Preprint available at homepages.warwick.ac.uk/~masgak/papers/median-algebras.pdf*, 2022. (Cited pp. 8 and 12).
- [BQ16a] Yves Benoist and Jean-François Quint. Central limit theorem for linear groups. *Ann. Probab.*, 44(2):1308–1340, 2016. (Cited p. 40).
- [BQ16b] Yves Benoist and Jean-François Quint. Central limit theorem on hyperbolic groups. *Izv. Ross. Akad. Nauk Ser. Mat.*, 80(1):5–26, 2016. (Cited p. 40).
- [BV91] H.-J. Bandelt and M. van de Vel. Superextensions and the depth of median graphs. *J. Combin. Theory Ser. A*, 57(2):187–202, 1991. (Cited p. 10).
- [BW12] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012. (Cited p. 4).
- [Can87] James W. Cannon. Almost convex groups. *Geom. Dedicata*, 22(2):197–210, 1987. (Cited p. 6).
- [CD95] Ruth M. Charney and Michael W. Davis. Strict hyperbolization. *Topology*, 34(2):329–350, 1995. (Cited p. 1).
- [CDH10] Indira Chatterji, Cornelia Druţu, and Frédéric Haglund. Kazhdan and Haagerup properties from the median viewpoint. *Adv. Math.*, 225(2):882–921, 2010. (Cited pp. 2, 7, and 40).
- [CF10] Pierre-Emmanuel Caprace and Koji Fujiwara. Rank-one isometries of buildings and quasi-morphisms of Kac-Moody groups. *Geom. Funct. Anal.*, 19(5):1296–1319, 2010. (Cited p. 3).
- [CFFT22] Kunal Chawla, Behrang Forghani, Joshua Frisch, and Giulio Tiozzo. The Poisson boundary of hyperbolic groups without moment conditions. *arXiv:2209.02114*, 2022. (Cited p. 40).
- [Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000. (Cited p. 8).
- [CMV04] Pierre-Alain Cherix, Florian Martin, and Alain Valette. Spaces with measured walls, the Haagerup property and property (T). *Ergodic Theory Dynam. Systems*, 24(6):1895–1908, 2004. (Cited p. 40).
- [CN05] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005. (Cited pp. 8 and 39).
- [Cou23] Rémi Coulon. Ergodicity of the geodesic flow for groups with a contracting element. *arXiv:2303.01390*, 2023. (Cited p. 2).
- [CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21(4):851–891, 2011. (Cited p. 39).
- [CS15] Ruth Charney and Harold Sultan. Contracting boundaries of CAT(0) spaces. *J. Topol.*, 8(1):93–117, 2015. (Cited pp. 3, 25, and 31).
- [CW21a] Matthieu Calvez and Bert Wiest. Hyperbolic structures for Artin-Tits groups of spherical type. In *Geometry at the frontier—symmetries and moduli spaces of algebraic varieties*, volume 766 of *Contemp. Math.*, pages 83–98. Amer. Math. Soc., Providence, RI, 2021. (Cited p. 39).

- [CW21b] Matthieu Calvez and Bert Wiest. Morse elements in Garside groups are strongly contracting. *arXiv:2106.14826*, 2021. (Cited pp. 3 and 31).
- [Dah03] François Dahmani. Combination of convergence groups. *Geom. Topol.*, 7:933–963, 2003. (Cited p. 1).
- [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. (Cited p. 2).
- [DHS17] Matthew G. Durham, Mark Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017. (Cited p. 8).
- [DJ91] Michael W. Davis and Tadeusz Januszkiewicz. Hyperbolization of polyhedra. *J. Differential Geom.*, 34(2):347–388, 1991. (Cited p. 1).
- [DMS23] Matthew G. Durham, Yair N. Minsky, and Alessandro Sisto. Stable cubulations, bicomings, and barycenters. *Geom. Topol.*, 27(6):2383–2478, 2023. (Cited pp. 4, 5, 6, and 7).
- [DT15] Matthew G. Durham and Samuel J. Taylor. Convex cocompactness and stability in mapping class groups. *Algebr. Geom. Topol.*, 15(5):2839–2859, 2015. (Cited pp. 3, 31, and 39).
- [DT19] Spencer Dowdall and Samuel J. Taylor. Contracting orbits in outer space. *Math. Z.*, 293(1-2):767–787, 2019. (Cited p. 3).
- [Dur16] Matthew Gentry Durham. The augmented marking complex of a surface. *J. Lond. Math. Soc. (2)*, 94(3):933–969, 2016. (Cited p. 4).
- [Dur23] Matthew Gentry Durham. Cubulating infinity in hierarchically hyperbolic spaces. *arXiv preprint arXiv:2308.13689*, 2023. (Cited p. 7).
- [DZ22] Matthew G. Durham and Abdul Zalloum. The geometry of genericity in mapping class groups and Teichmüller spaces via $CAT(0)$ cube complexes. *arXiv:2207.06516*, 2022. (Cited p. 4).
- [EMR17] Alex Eskin, Howard Masur, and Kasra Rafi. Large-scale rank of Teichmüller space. *Duke Math. J.*, 166(8):1517–1572, 2017. (Cited p. 4).
- [Far06] Benson Farb. Some problems on mapping class groups and moduli space. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 11–55. Amer. Math. Soc., Providence, RI, 2006. (Cited p. 6).
- [Fio20] Elia Fioravanti. Roller boundaries for median spaces and algebras. *Algebr. Geom. Topol.*, 20(3):1325–1370, 2020. (Cited pp. 2, 7, 12, and 40).
- [FLM01] Benson Farb, Alexander Lubotzky, and Yair Minsky. Rank-1 phenomena for mapping class groups. *Duke Math. J.*, 106(3):581–597, 2001. (Cited p. 5).
- [FLM21] Talia Fernós, Jean Lécureux, and Frédéric Mathéus. Contact graphs, boundaries, and a central limit theorem for $CAT(0)$ cubical complexes. *arXiv:2112.10141*, 2021. (Cited p. 40).
- [FM02] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152, 2002. (Cited p. 31).
- [Gen20] Anthony Genevois. Hyperbolicities in $CAT(0)$ cube complexes. *Enseign. Math.*, 65(1-2):33–100, 2020. (Cited pp. 19 and 38).
- [Gen23] Anthony Genevois. Algebraic properties of groups acting on median graphs. *Preprint available at sites.google.com/view/agenevois/books*, 2023. (Cited p. 8).
- [GHP⁺23] Antoine Goldsborough, Mark Hagen, Harry Petyt, Jacob Russell, and Alessandro Sisto. Induced quasi-isometries of hyperbolic spaces, Markov chains, and acylindrical hyperbolicity. *arXiv:2309.07013*, 2023. (Cited p. 41).
- [Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987. (Cited pp. 1, 4, and 31).
- [GY22] Ilya Gekhtman and Wen-yuan Yang. Counting conjugacy classes in groups with contracting elements. *J. Topol.*, 15(2):620–665, 2022. (Cited p. 2).
- [Hae21] Thomas Haettel. Lattices, injective metrics and the $k(\pi, 1)$ conjecture. *arXiv:2109.07891*, 2021. (Cited p. 4).
- [Hae22] Thomas Haettel. A link condition for simplicial complexes, and cub spaces. *arXiv:2211.07857*, 2022. (Cited p. 1).
- [Hae23] Thomas Haettel. Group actions on injective spaces and Helly graphs. *arXiv:2307.00414*, 2023. (Cited pp. 4 and 5).
- [HHP23] Thomas Haettel, Nima Hoda, and Harry Petyt. Coarse injectivity, hierarchical hyperbolicity and semihyperbolicity. *Geom. Topol.*, 27(4):1587–1633, 2023. (Cited pp. 3, 4, 5, 6, 7, 33, 37, and 39).
- [HM13] Michael Handel and Lee Mosher. The free splitting complex of a free group, I: hyperbolicity. *Geom. Topol.*, 17(3):1581–1672, 2013. (Cited p. 39).

- [HMS21] Mark Hagen, Alexandre Martin, and Alessandro Sisto. Extra-large type Artin groups are hierarchically hyperbolic. *arXiv:2109.04387*, 2021. (Cited p. 5).
- [HO21] Jingyin Huang and Damian Osajda. Helly meets Garside and Artin. *Invent. Math.*, 225(2):395–426, 2021. (Cited p. 4).
- [Hor18] Camille Horbez. Central limit theorems for mapping class groups and $\text{Out}(F_N)$. *Geom. Topol.*, 22(1):105–156, 2018. (Cited p. 40).
- [HV98] Allen Hatcher and Karen Vogtmann. The complex of free factors of a free group. *Quart. J. Math. Oxford Ser. (2)*, 49(196):459–468, 1998. (Cited p. 39).
- [HW15a] Mark Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the general case. *Geom. Funct. Anal.*, 25(1):134–179, 2015. (Cited p. 4).
- [HW15b] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. *Invent. Math.*, 199(2):293–331, 2015. (Cited pp. 1 and 4).
- [IPS21] Alessandra Iozzi, Cristina Pagliantini, and Alessandro Sisto. Characterising actions on trees yielding non-trivial quasimorphisms. *Ann. Math. Qué.*, 45(1):185–202, 2021. (Cited p. 3).
- [JL22] Martina Jørgensen and Urs Lang. A combinatorial higher-rank hyperbolicity condition. *arXiv:2206.08153*, 2022. (Cited pp. 37 and 39).
- [Kai00] Vadim A. Kaimanovich. The Poisson formula for groups with hyperbolic properties. *Ann. of Math. (2)*, 152(3):659–692, 2000. (Cited p. 40).
- [KK13a] Sang-Hyun Kim and Thomas Koberda. Embedability between right-angled Artin groups. *Geom. Topol.*, 17(1):493–530, 2013. (Cited p. 39).
- [KK13b] Marcin Kotowski and Michał Kotowski. Random groups and property (T) : Żuk’s theorem revisited. *J. Lond. Math. Soc. (2)*, 88(2):396–416, 2013. (Cited p. 7).
- [KL08] Autumn E. Kent and Christopher J. Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008. (Cited pp. 3 and 31).
- [KL09] Ilya Kapovich and Martin Lustig. Geometric intersection number and analogues of the curve complex for free groups. *Geom. Topol.*, 13(3):1805–1833, 2009. (Cited p. 39).
- [KMV22] Daniel Keppeler, Philip Möller, and Olga Varghese. Automatic continuity for groups whose torsion subgroups are small. *J. Group Theory*, 25(6):1017–1043, 2022. (Cited p. 5).
- [Lan13] Urs Lang. Injective hulls of certain discrete metric spaces and groups. *J. Topol. Anal.*, 5(3):297–331, 2013. (Cited pp. 5, 7, 18, and 37).
- [Le 22] Corentin Le Bars. Central limit theorem on $\text{CAT}(0)$ spaces with contracting isometries. *arXiv:2209.11648*, 2022. (Cited p. 40).
- [Le 23] Corentin Le Bars. Marches aléatoires et éléments contractants sur des espaces $\text{CAT}(0)$. *arXiv:2310.18710*, 2023. (Cited p. 40).
- [Lea13] Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013. (Cited p. 1).
- [Mie14] Benjamin Miesch. Injective metrics on cube complexes. *arXiv:1411.7234*, 2014. (Cited p. 4).
- [MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999. (Cited p. 2).
- [MM00] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000. (Cited p. 5).
- [Mor21] Rose Morris-Wright. Parabolic subgroups in FC-type Artin groups. *J. Pure Appl. Algebra*, 225(1):1–13, 2021. (Cited p. 39).
- [MP22] Alexandre Martin and Piotr Przytycki. Acylindrical actions for two-dimensional Artin groups of hyperbolic type. *Int. Math. Res. Not. IMRN*, 17:13099–13127, 2022. (Cited p. 39).
- [MR08] Mahan Mj and Lawrence Reeves. A combination theorem for strong relative hyperbolicity. *Geom. Topol.*, 12(3):1777–1798, 2008. (Cited p. 1).
- [MT91] Jie Hua Mai and Yun Tang. On the injective metrization for infinite collapsible polyhedra. *Acta Math. Sinica (N.S.)*, 7(1):13–18, 1991. (Cited p. 4).
- [MT18] Joseph Maher and Giulio Tiozzo. Random walks on weakly hyperbolic groups. *J. Reine Angew. Math.*, 742:187–239, 2018. (Cited p. 40).
- [Mul78] Martyn Mulder. The structure of median graphs. *Discrete Math.*, 24(2):197–204, 1978. (Cited p. 39).
- [Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004. (Cited pp. 8 and 39).
- [NR97] Graham A. Niblo and Lawrence Reeves. Groups acting on $\text{CAT}(0)$ cube complexes. *Geom. Topol.*, 1:1–7, 1997. (Cited p. 7).

- [NR98] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998. (Cited pp. 5, 14, and 18).
- [NWZ19] Graham A. Niblo, Nick Wright, and Jiawen Zhang. A four point characterisation for coarse median spaces. *Groups Geom. Dyn.*, 13(3):939–980, 2019. (Cited p. 32).
- [Ont20] Pedro Ontaneda. Riemannian hyperbolization. *Publ. Math. Inst. Hautes Études Sci.*, 131:1–72, 2020. (Cited p. 1).
- [Osi16] D. Osin. Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.*, 368(2):851–888, 2016. (Cited p. 2).
- [Pap22] Chloé Papin. Strongly contracting axes for fully irreducible automorphisms of generalized Baumslag–Solitar groups. *arXiv:2206.04010*, 2022. (Cited p. 2).
- [Pet21] Harry Petyt. Mapping class groups are quasicubical. *arXiv:2112.10681*, 2021. (Cited p. 5).
- [PS23] Harry Petyt and Davide Spriano. Unbounded domains in hierarchically hyperbolic groups. *Groups Geom. Dyn.*, 17(2):479–500, 2023. (Cited p. 8).
- [PSZ22] Harry Petyt, Davide Spriano, and Abdul Zalloum. Hyperbolic models for CAT(0) spaces. *arXiv:2207.14127*, 2022. (Cited pp. 3, 19, 20, 22, 24, 25, 38, and 40).
- [QRT20] Yulan Qing, Kasra Rafi, and Giulio Tiozzo. Sublinearly morse boundary II: Proper geodesic spaces. *arXiv:2011.03481*, 2020. (Cited p. 40).
- [Rol98] Martin Roller. Poc sets, median algebras and group actions. *Habilitationsschrift, Universität Regensburg*, 1998. Available on the arXiv at arXiv:1607.07747. (Cited pp. 8 and 12).
- [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. (Cited p. 41).
- [RV21] Kasra Rafi and Yvon Verberne. Geodesics in the mapping class group. *Algebr. Geom. Topol.*, 21(6):2995–3017, 2021. (Cited p. 31).
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3), 71(3):585–617, 1995. (Cited pp. 1 and 8).
- [Sel97] Z. Sela. Acylindrical accessibility for groups. *Invent. Math.*, 129(3):527–565, 1997. (Cited p. 2).
- [She22] Sam Shepherd. A cubulation with no factor system. *arXiv:2208.10421*, 2022. (Cited p. 5).
- [Sis18] Alessandro Sisto. Contracting elements and random walks. *J. Reine Angew. Math.*, 742:79–114, 2018. (Cited p. 39).
- [SZ] Mireille Soergel and Abdul Zalloum. Morse geodesics in weakly modular graphs. *In preparation*. (Cited p. 3).
- [SZ22] Alessandro Sisto and Abdul Zalloum. Morse subsets of injective spaces are strongly contracting. *arXiv:2208.13859*, 2022. (Cited pp. 2, 3, 31, 39, and 41).
- [Wis21] Daniel T. Wise. *The structure of groups with a quasiconvex hierarchy*, volume 209 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2021. (Cited pp. 7 and 8).
- [Ż03] A. Żuk. Property (T) and Kazhdan constants for discrete groups. *Geom. Funct. Anal.*, 13(3):643–670, 2003. (Cited p. 7).
- [Zal23] Abdul Zalloum. Injectivity, cubical approximations and equivariant wall structures beyond CAT(0) cube complexes. *arXiv:2305.02951*, 2023. (Cited p. 4).
- [Zbi] Stefanie Zbinden. Hyperbolic spaces that detect all strongly-contracting directions. *In preparation*. (Cited p. 3).
- [Zbi23] Stefanie Zbinden. Small cancellation groups with and without sigma-compact Morse boundary. *arXiv preprint arXiv:2307.13325*, 2023. (Cited p. 3).

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, UK
 Email address: `petyt@maths.ox.ac.uk`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TORONTO, CANADA
 Email address: `abdul.zalloum@utoronto.ca`

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, UK
 Email address: `spriano@maths.ox.ac.uk`