

The Poisson boundary of CAT(0) cube complex groups

Amos Nevo and Michah Sageev*

Abstract. We consider a finite dimensional, locally finite CAT(0) cube complex X admitting a co-compact properly discontinuous countable group of automorphisms G . We construct a natural compact metric space $B(X)$ on which G acts by homeomorphisms, the action being minimal and strongly proximal. Furthermore, for any generating probability measure on G , $B(X)$ admits a unique stationary measure, and when the measure has finite logarithmic moment, it constitutes a compact metric mean-proximal model of the Poisson boundary. We identify a dense invariant G_δ subset $\mathcal{U}_{\text{NT}}(X)$ of $B(X)$ which supports every stationary measure, and on which the action of G is Borel-amenable. We describe the relation of $\mathcal{U}_{\text{NT}}(X)$ and $B(X)$ to the Roller boundary. Our construction can be used to give a simple geometric proof of property A for the complex. Our methods are based on direct geometric arguments regarding the asymptotic behavior of halfspaces and their limiting ultrafilters, which are of considerable independent interest. In particular we analyze the notions of median and interval in the complex, and use the latter in the proof that $B(X)$ is the Poisson boundary via the strip criterion developed by V. Kaimanovich.

Mathematics Subject Classification (2010). 20P05, 22F10, 37B05, 57S30, 60J50.

Keywords. CAT(0) cube complex, ultrafilters, co-compact action, cube complex group, boundary space, strong proximality, Poisson boundary.

Contents

1	Introduction	654
2	Basics on CAT(0) cube complexes	657
3	The definition of the boundary $B(X)$	664
4	General boundary theory	667
5	Establishing that $B(X)$ is a boundary	669
6	Intervals and property A	676
7	Amenability of the boundary action	685
8	Maximality of the boundary	688
9	Entropy and the Poisson boundary	692
	References	693

*The first author was supported in part by ISF grant. The second author was supported in part by ISF Grant.

1. Introduction

1.1. Motivation. Groups appearing as lattices in the automorphism groups of CAT(0) cube complexes have been a subject of considerable interest in recent years. On the one hand, the class of such groups include a broad spectrum of groups from across geometric group theory, including Coxeter groups, right-angled Artin groups, certain arithmetic lattices in real hyperbolic space, as well as small cancellation groups. On the other hand, the combinatorial nature of CAT(0) cube complexes allows one to develop techniques and results that are sharper than those that hold in the general context of CAT(0) spaces. The present paper is devoted to developing aspects of boundary theory for lattices in CAT(0) cube complexes, including a direct geometric construction of a boundary space. This construction exhibits useful analogies with boundary theory for lattices in semisimple algebraic groups over local fields, with the space constructed playing the role of a maximal boundary.

To motivate this analogy, note that the study of an irreducible lattice subgroup of semisimple algebraic groups is greatly facilitated by the study of its actions on various compact homogeneous spaces of the algebraic group, called boundary spaces. Taking $\mathrm{SL}_n(F)$ as an example, where $F = \mathbb{R}, \mathbb{C}$ or \mathbb{Q}_p say, one considers the action on the maximal boundary, namely the space \mathcal{F}_n of full flags of F^n , as well as on its equivariant factor spaces, the spaces of partial flags. Generally, for any semisimple algebraic group H , the maximal boundary of H is defined as the compact homogeneous space H/P , where P is a minimal parabolic subgroup, and its factors are given by H/Q , where Q is a parabolic subgroup containing P . The number of such factors is equal to the split rank of the group over the field F .

For a lattice subgroup of $\mathrm{SL}_n(F)$, the action on the maximal boundary is minimal and strongly proximal, and every generating probability measure on the lattice has a unique stationary measure. When the measure has finite logarithmic moment, the maximal boundary with the associated stationary measure constitutes a compact metric model of the Poisson boundary. In addition the maximal boundary is a mean-proximal and universally amenable action of the lattice. Furthermore, for large classes of discrete subgroups of $\mathrm{SL}_n(F)$ one can establish several of these properties (or natural modifications thereof) in considerable generality. As is well-known, the dynamical properties of the boundary actions have played a significant role in a wide array of applications in the study of discrete subgroups of semisimple algebraic groups, of which we mention the following:

- (1) The Tits alternative.
- (2) Simplicity and exactness of the reduced C^* -algebra.
- (3) Property A , a -T-menability, Baum–Connes and Novikov conjectures.
- (4) Patterson–Sullivan theory.
- (5) Super-rigidity of measure-preserving actions of higher rank lattices.
- (6) Classification of boundary factors and normal subgroup theorem for higher rank lattices.

It is natural to attempt to develop boundary theory for CAT(0) cube complex groups, to the point where it can be exploited in a manner analogous to the case of discrete subgroups of semisimple algebraic groups. This applies, in particular, to the search for a natural notion of split rank for the complex, generalizing the notion of split rank in the algebraic context.

1.2. Main results. A systematic development of boundary theory and its applications for CAT(0) cube complexes and their uniform lattices must begin with the construction of the right notion of maximal boundary, and this is our goal in the present paper. We will give a direct geometric construction of a boundary space, which in effect singles out a compact invariant subset $B(X)$ of the Roller boundary, which is the space of ultrafilters of a cube complex X . The compact metric space $B(X)$ plays a role analogous to the unique compact orbit H/P of a semisimple algebraic group in the Satake compactification. Namely, $B(X)$ is a limit set for the action of the lattice on the larger Roller compactification, and gives rise to a minimal and strongly proximal action. Furthermore, the action admits an equicontinuous decomposition, an important dynamical property we will describe below, which generalizes the irreducibility and proximality properties of the action of discrete subgroups of a semisimple algebraic group on its maximal boundary.

We will prove that for any generating measure on the lattice, the action on $B(X)$ is uniquely stationary, and in addition $B(X)$ realizes the Poisson boundary when the measure has finite logarithmic moment. $B(X)$ is also a mean proximal space for the group. We will identify a dense invariant G_δ subset denoted by $\mathcal{U}_{\text{NT}}(X)$ contained in $B(X)$, on which the action is Borel-amenable, and hence universal (or measure-wise) amenable. Furthermore, $\mathcal{U}_{\text{NT}}(X)$ has measure one with respect to any stationary measure as above. We will use $\mathcal{U}_{\text{NT}}(X)$ to give a simple geometric proof that cube complexes satisfy Yu's property A in our context (see [BC+] for the general case).

We note that $B(X)$ possesses, for general CAT(0) cube complexes, several further important structural features analogous to those of the maximal boundary of semisimple groups. In particular, utilizing the recent product decomposition theorem for cube complexes established in [CS] (see Theorem 2.5 below), the boundary $B(X)$ can be represented as a canonical product $\prod_{i=1}^r B(X_i)$ where each X_i is an essential, irreducible non-Euclidean cube complex, with the action on $B(X)$ being the direct product of the action of G on $B(X_i)$. The invariant r appearing here will be called the split rank of the cube complex. It is the natural generalization of the split rank associated with a semisimple algebraic group over a local field, coinciding with it in the case that where the Bruhat–Tits building is a product of trees, or equivalently, in the case the group has no simple factor groups of split rank at least two. In particular our construction yields 2^r continuous equivariant boundary factors of the Poisson boundary $B(X)$ (including $B(X)$ and the trivial factor), in analogy with the 2^r boundary factors of a semisimple algebraic group of split-rank r .

1.3. Other boundaries. Let us compare our results with previous ones existing in the literature.

In general, there are many compact metric realizations of the Poisson boundary of a given group, and the main challenge is to establish the dynamical properties of a given realization. These can vary widely, and the search is for a model that has as many desirable dynamical properties as possible, with the properties holding in the case of semisimple algebraic groups serving as a bench mark.

For general CAT(0) cube complexes there already exist two natural compactifications. First, one can consider the usual visual boundary defined for any CAT(0) complex, and second, specifically for cube complexes, one can consider the space consisting of all the ultrafilters on the partially ordered set of halfspaces, an important compactification introduced by Roller [Ro].

Thus for a CAT(0) cube complex one can consider

- (1) the boundary space $B(X)$ given by the closure of non-terminating ultrafilters, constructed in the present paper,
- (2) the Roller boundary,
- (3) the visual boundary,
- (4) the horofunction boundary (associated say with the ℓ^1 -metric),
- (5) the Gromov boundary of the Cayley graph with respect to a word metric (when the group is hyperbolic),
- (6) the space of ends (when non-trivial).

A systematic comparison of the relationships between these boundary spaces is a challenge yet to be met. However, it is worth pointing out that each of these boundaries constitute a compact metric model of the Poisson boundary of a properly discontinuous co-compact irreducible CAT(0) cube complex group (when applicable). For the boundary $B(X)$ consisting of the closure of non-terminating ultrafilters we prove this fact in the present paper, together with the fact that the action is minimal. It follows that the Roller boundary, which contains $B(X)$, also constitutes a compact metric model, but it is not a minimal action, in general. The horofunction boundary with respect to the ℓ^1 metric has been identified with the Roller boundary of the complex, by recent unpublished work of U. Bader and D. Guralnik. The fact that Gromov boundary and the space of ends (when applicable) constitute compact metric minimal models of the Poisson boundary is well-known.

As to the visual boundary, the fact that it constitute a compact metric model follows from a result of Karlsson–Margulis [KM], Cor. 6.2, which implies that for any co-compact discrete group acting properly on a proper CAT(0) complex, the Poisson boundary can be realized as a closed invariant subset of the visual boundary, using also Kaimanovich’s ray criterion [K]. When the cube complex is irreducible (and more generally when the group contains a rank-1 isometry) the action on the visual boundary is in addition minimal. Thus by the recent result of Caprace–Sageev [CS] on the rank rigidity of CAT(0) cube complexes (see Theorem 2.5 below) it follows

that the product of the visual boundaries of the irreducible non-Euclidean factors of the complex constitutes a compact metric minimal model of the Poisson boundary. The authors are indebted to the referee for these remarks on the visual boundary.

We note further two other boundary constructions discussed by Karlsson [Ka] and Caprace [Ca07], which typically also provide compact metric models of the Poisson boundary. However, these constructions do not give rise to minimal or strongly proximal actions, in general.

As far as we are aware, for uniform lattices in general CAT(0) complex, the existence of a compact, metric, minimal, strongly-proximal, uniquely-stationary, mean-proximal, universally amenable and equicontinuously decomposable realization of the Poisson boundary is an open problem. We establish these properties for the boundary space $B(X)$, the closure of the non-terminating ultrafilters of CAT(0) cube complex, thus providing a close analog of the maximal boundary given by the homogeneous projective variety H/P of a semisimple group H , with P being a minimal parabolic group.

We expect this construction to provide a unified approach to the applications of boundary theory to discrete subgroups mentioned above, but we note that some of them have been established for groups acting on CAT(0) cube complexes using a variety of other considerations. Thus the Tits alternative was established in [SW], property A was established in [BC+], and a-T-menability was established in [NR]. For a discussion of the Baum–Connes conjecture for some cube complex groups we refer to [CR], and for a discussion of exactness in this context see [GuNi].

2. Basics on CAT(0) cube complexes

2.1. Hyperplanes and halfspaces. We recall basic terminology and facts about CAT(0) cube complexes, referring for more details to [ChN], [Gu], [N], [Ro], [S1].

Definition 2.1. A CAT(0) *cube complex* is a simply-connected combinatorial cell complex whose closed cells are Euclidean n -dimensional cubes $[0, 1]^n$ of various dimensions such that:

- (1) Any two cubes either have empty intersection or intersect in a single face of each.
- (2) The link of each 0-cell is a *flag complex*, a simplicial complex such that any $(n + 1)$ adjacent vertices belong to an n -simplex.

Since an n -cube is a product of n unit intervals, each n -cube comes equipped with n natural projection maps to the unit interval. A *hypercube* is the preimage of $\{\frac{1}{2}\}$ under one of these projections; each n -cube contains n hypercubes. A *hyperplane* in a CAT(0) cube complex X is a subspace intersecting each cube in a hypercube or trivially. Hyperplanes are said to *cross* if they intersect non-trivially; otherwise

they are said to be *disjoint*. The *carrier* of a hyperplane is the union of all cubes intersecting it.

Here are some basic facts about hyperplanes in CAT(0) cube complexes which we will use throughout our arguments.

Basic properties. (1) Each hyperplane is embedded (i.e., it intersects a given cube in a single hypercube).

(2) Each hyperplane separates the complex into precisely two components, called *halfspaces*.

(3) Every collection of pairwise crossing hyperplanes has a non-empty intersection.

(4) Each hyperplane is itself a CAT(0) cube complex.

We shall use the ℓ_1 metric on X , namely simply the metric on the vertices which assigns to two vertices the number of hyperplanes separating them. This metric is equivalent to the path metric on the 1-skeleton of X .

2.2. Ultrafilters and the Roller boundary. Let \mathcal{H} denote the collection of halfspaces and $\hat{\mathcal{H}}$ denote the collection of hyperplanes. The collection of halfspaces comes equipped with a natural involution $\mathfrak{h} \mapsto \mathfrak{h}^*$, where \mathfrak{h} and \mathfrak{h}^* are the two complementary components of a given hyperplane. We denote by $\hat{\mathfrak{h}}$ the hyperplane associated to the halfspace \mathfrak{h} .

Recall that an *ultrafilter* on \mathcal{H} is a subset α of \mathcal{H} satisfying the following two conditions:

- (1) Choice. For every hyperplane $\hat{\mathfrak{h}}$, either $\mathfrak{h} \in \alpha$ or $\mathfrak{h}^* \in \alpha$ but not both.
- (2) Consistency. If $\mathfrak{h} \in \alpha$ and $\mathfrak{h} \subset \mathfrak{h}'$, then $\mathfrak{h}' \in \alpha$.

Sometimes we will want to construct an ultrafilter and this will be done by making a consistent choice of halfspaces. This means choosing halfspaces of \mathcal{H} according to (1) and (2) above.

Given two ultrafilters α and β and a hyperplane $\hat{\mathfrak{h}}$, we say that α and β are *separated* by $\hat{\mathfrak{h}}$ if $\mathfrak{h} \in \alpha$ and $\mathfrak{h}^* \in \beta$ or $\mathfrak{h} \in \beta$ and $\mathfrak{h}^* \in \alpha$.

We denote by $\mathcal{U} = \mathcal{U}(X)$ the collection of all ultrafilters on the collection of halfspaces \mathcal{H} of X . There is a natural embedding of the vertex set $X^{(0)}$ of X into \mathcal{U} , namely:

$$X^{(0)} \rightarrow \mathcal{U}, \quad v \mapsto \alpha_v = \{\mathfrak{h} \in \mathcal{H} \mid v \in \mathfrak{h}\}.$$

Every vertex of $v \in X$ may be viewed as ultrafilter, namely the collection of all those hyperplanes \mathfrak{h} such that $v \in \mathfrak{h}$. We will use α_v to denote the ultrafilter associated to v . As noted by [Gu], when X is finite dimensional, such ultrafilters are characterized by the Descending Chain Condition (DCC), namely every descending chain of halfspaces $\mathfrak{h}_1 \supset \mathfrak{h}_2 \supset \dots$ terminates. Such ultrafilters are called *principal* ultrafilters.

Following Roller, we may view \mathcal{U} as a compactification of $X^{(0)}$ in the following simple manner. We consider the product

$$\mathcal{P} = \prod_{\mathfrak{h} \in \widehat{\mathcal{H}}} \{\mathfrak{h}, \mathfrak{h}^*\}.$$

Since an ultrafilter is a choice for each pair $(\mathfrak{h}, \mathfrak{h}^*)$, we have $\mathcal{U} \subset \mathcal{P}$. The space \mathcal{P} endowed with the Tychonoff topology is a compact space. It is not difficult to see then that \mathcal{U} is a closed subset of \mathcal{P} and is therefore compact. Moreover, Roller shows that $X^{(0)}$ is open and dense in \mathcal{U} . We thus conclude that $\mathcal{U} \setminus X^{(0)}$ constitutes a compactification of $X^{(0)}$, which we refer to as the Roller boundary.

To develop a better understanding of what $\mathcal{U}(X)$ looks like, it is useful to recall how one metrizes the product topology. Recall that if $Y = \prod_{n=1}^\infty X_n$ is a countable product of metric spaces of uniformly bounded diameter, we may metrize the product topology as follows. Let $f : \mathbb{N} \rightarrow \mathbb{R}_+$ be any decreasing positive function such that $\lim_{n \rightarrow \infty} f(n) = 0$. Then given $\mathbf{x} = (x_n), \mathbf{y} = (y_n) \in Y$, we set

$$d(\mathbf{x}, \mathbf{y}) = \sup\{f(n)d(x_n, y_n) \mid n > 0\}.$$

In our setting, the space $Y = 2^{\mathcal{H}}$ is a countable product of two element spaces, each containing the two halfspaces associated to a given hyperplane. We simply need to describe a function $f : \widehat{\mathcal{H}} \rightarrow \mathbb{R}$ as above. When X is proper, one can order the hyperplanes using the metric on X . Pick a base vertex $\mathfrak{o} \in X$. For any hyperplane $\hat{\mathfrak{h}}$, we let $f(\hat{\mathfrak{h}}) = 1/d(\hat{\mathfrak{h}}, \mathfrak{o})$, where $d(\hat{\mathfrak{h}}, \mathfrak{o})$ denotes the 1-skeleton distance between \mathfrak{o} and $\hat{\mathfrak{h}}$. That is,

$$d(\hat{\mathfrak{h}}, \mathfrak{o}) = |\{\text{hyperplanes separating } \hat{\mathfrak{h}} \text{ from } \mathfrak{o}\}| + 1.$$

Since there are only finitely many hyperplanes a given distance from \mathfrak{o} , we have that f is a decreasing function approaching 0, as required. Explicitly, for two distinct ultrafilters $\alpha \neq \beta$ we have

$$d(\alpha, \beta) = \sup\{1/d(\hat{\mathfrak{h}}, \mathfrak{o}) \mid \hat{\mathfrak{h}} \text{ separates } \alpha \text{ and } \beta\}.$$

Note also that for each halfspace \mathfrak{h} we can define the following subset of \mathcal{U} :

$$U_{\mathfrak{h}} = \{\alpha \in \mathcal{U}(X) \mid \mathfrak{h} \in \alpha\}.$$

We refer to such a subset as an \mathcal{U} -halfspace. The collection of all \mathcal{U} -halfspaces forms a sub-basis for the Tychonoff topology on \mathcal{U} . Thus a basic open set consists of the intersection of finitely many \mathcal{U} -halfspaces.

2.3. Quotients. One can abstract the above construction in the following way. A pocset is a poset Σ with an order reversing involution $*$: $\Sigma \rightarrow \Sigma$. The pocset Σ is said to have the *finite interval condition* if for every $A < B$, there exist finitely many

C such that $A < C < B$. A pair of elements of $A, B \in \Sigma$ are said to be transverse if A and A^* are incomparable with B and B^* . The pocset Σ is said to have *finite width* if there is a universal bound on the size of a collection of incomparable elements. If one starts with a finite dimensional cube complex, the collection of halfspaces forms a pocset which satisfies both the finite width condition and the finite interval condition.

Given any poset Σ with an order-reversing involution, one can consider the collection of all ultrafilters $\mathcal{U}(\Sigma)$ on Σ as in Section 2.2. If Σ satisfies the finite interval condition and has finite width, then the collection of all principal ultrafilters is the vertex set of a finite dimensional CAT(0) cube complex $X(\Sigma)$. As noted by Roller, this construction is natural in that if one starts with a CAT(0) cube complex X , considers its pocset of halfspaces $\Sigma = \mathcal{H}(X)$ and then considers the cube complex whose vertices are the principal ultrafilters on Σ , then $X(\Sigma) = X$.

Let $\mathcal{H}(X)$ denote the halfspaces of X . Suppose that $\mathcal{K} \subset \mathcal{H}$ is a subset of \mathcal{H} closed under involution. we can consider the ultrafilters on the halfspace system \mathcal{K} , which we denote $\mathcal{U}_{\mathcal{K}}(X)$.

The collection of principal ultrafilters of $\mathcal{U}_{\mathcal{K}}(X)$ is the vertex set of a CAT(0) cube complex $X_{\mathcal{K}}(X)$. There is then a natural projection map.

$$\mathcal{U}(X) \rightarrow \mathcal{U}_{\mathcal{K}}(X), \quad \alpha \mapsto \alpha \cap \mathcal{K}.$$

This projection restricts to a projection on the principal ultrafilters, so that one has a projection map $X \rightarrow X_{\mathcal{K}}$.

Lemma 2.2. *Let X be finite dimensional CAT(0) cube complex. Then the natural projection $X \rightarrow X_{\mathcal{K}}$ is surjective.*

Proof. We need to show that given a consistent choice of half spaces for \mathcal{K} satisfying DCC, it can be extended to a consistent choice of halfspaces for half spaces in \mathcal{H} satisfying DCC. We do this by induction on the dimension on the dimension of X .

So let α be a principal ultrafilter on \mathcal{K} . We wish to extend it to an ultrafilter on \mathcal{H} .

When X is 1-dimensional, let \mathfrak{k} be a minimal element in α . The halfspace \mathfrak{k} is associated to some edge e . Let v be a vertex which is the endpoint e contained in \mathfrak{k} . It is now clear that the ultrafilter α_v is the desired ultrafilter.

We now proceed by induction. Again letting α be a principal ultrafilter on \mathcal{K} , we choose a minimal halfspace $\mathfrak{h}_0 \in \alpha$. Now the hyperplanes of $\hat{\mathcal{H}}$ are divided into those that meet $\hat{\mathfrak{h}}_0$ and those that do not. We focus first on those that meet $\hat{\mathfrak{h}}_0$. These correspond to hyperplanes in $\hat{\mathfrak{h}}_0$, viewed as a CAT(0) cube complex in its own right. By induction, for the collection of halfspaces associated to these hyperplanes, we have that there exists some vertex $v \in \hat{\mathfrak{h}}_0$ for which

$$\alpha_v = \{\mathfrak{h} \mid v \in \mathfrak{h} \text{ and } \hat{\mathfrak{h}} \cap \hat{\mathfrak{h}}_0 \neq \emptyset\}.$$

We now let e be the edge of X whose midpoint is v and let w be the endpoint of e which is contained in \mathfrak{h}_0 . Then α_w is our desired ultrafilter. □

We then have the following useful corollary.

Corollary 2.3. *Let X be a finite dimensional CAT(0) cube complex and let $\alpha \subset \mathcal{H}$ be a subset satisfying the choice and consistency conditions and satisfying DCC, then*

$$\bigcap_{\mathfrak{h} \in \alpha} \mathfrak{h} \neq \emptyset.$$

Proof. Let $\mathcal{K} = \{\mathfrak{h} \in \mathcal{H} \mid \mathfrak{h} \in \alpha \text{ or } \mathfrak{h}^* \in \alpha\}$. Then \mathcal{K} is an involution invariant subset and we may construct the space $X_{\mathcal{K}}$. The subset α is now an ultrafilter on \mathcal{K} and since it satisfies DCC, it corresponds to a vertex in $X_{\mathcal{K}}$. Now by Lemma 2.2 the map $X \rightarrow X_{\mathcal{K}}$ is surjective, so there exist vertices of X mapping to α . These vertices lie in $\bigcap_{\mathfrak{h} \in \alpha} \mathfrak{h}$, as required. \square

A particular example of this type of projection occurs when one eliminates a single hyperplane: $\mathcal{K} = \mathcal{H} - \{\mathfrak{h}, \mathfrak{h}^*\}$. In X , the collection of all cubes intersecting $\hat{\mathfrak{h}}$ is called the *carrier* of $\hat{\mathfrak{h}}$, which we denote $C(\hat{\mathfrak{h}})$, and naturally has a product structure $C(\hat{\mathfrak{h}}) \cong \hat{\mathfrak{h}} \times [0, 1]$. One then has a natural collapsing map

$$\rho: C(\hat{\mathfrak{h}}) \cong \hat{\mathfrak{h}} \times [0, 1] \rightarrow \hat{\mathfrak{h}} \times \{0\}.$$

One can apply this to the carrier of $\hat{\mathfrak{h}}$ to obtain a quotient \bar{X} of X , whose hyperplanes are $\mathcal{H} - \{\hat{\mathfrak{h}}\}$. More precisely, we define an equivalence relation \sim on X which identifies distinct points x, y if and only if $x, y \in C(\hat{\mathfrak{h}})$ and $\rho(x) = \rho(y)$. The resulting quotient space $\bar{X} = X/\sim$ is the cube complex associated to the pocset $\mathcal{H} - \{\mathfrak{h}, \mathfrak{h}^*\}$. In this instance we call the map $X \rightarrow \bar{X}$ a *collapsing* map. We also use the term collapsing for the quotient obtained by removing finitely many halfspaces.

2.4. Pruning. When a group acts on a tree, it is useful to pass to a minimal invariant subtree. For a CAT(0) cube complex, there is a similar process, described in [CS], which we now describe. A halfspace is called *deep* if it contains arbitrarily large balls. A hyperplane is called *essential* if both of its associated halfspaces are deep. A cube complex is called *essential* if all of its hyperplanes are essential.

We then have a the following result (described in more detail in [CS]).

Theorem 2.4. *Let X be a CAT(0) cube complex with cocompact automorphism group. Then there exists a canonical essential CAT(0) cube complex X_{ess} and a G -equivariant map $f: X \rightarrow X_{\text{ess}}$ such that the preimage under f of the collection of hyperplanes of X_{ess} is the collection of essential hyperplanes of X and f is bounded-to-one.*

Remark. A natural subclass of CAT(0) cube complexes is the collection of those which have *extendible geodesics*, which means that every finite geodesic path can be extended to a bi-infinite geodesic. It is easy to see that such complexes are essential.

We call Y the essential quotient of X . The map f above is easily seen to extend to a bounded-to-one map $\mathcal{U}(X) \rightarrow \mathcal{U}(X_{\text{ess}})$. In the rest of the paper, we will pass to this essential quotient and work with it.

2.5. Products. As was discussed in [CS], a decomposition of X as a product $X = X_1 \times X_2$ corresponds to a decomposition of $\hat{\mathcal{H}}$ as a disjoint union $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \hat{\mathcal{H}}_2$ where every hyperplane in $\hat{\mathcal{H}}_1$ intersects every hyperplane in $\hat{\mathcal{H}}_2$. A CAT(0) cube complex is called *irreducible* if it does not decompose as a product. We recall the following theorem, proved in [CS].

Theorem 2.5 (Product Decomposition Theorem). *Every finite dimensional CAT(0) cube complex admits a canonical (up to permutation of factors) decomposition as a finite product of irreducible CAT(0) cube complexes.*

We note that the canonical property of this decomposition is that up to passing to a subgroup of finite index, $\text{Aut}(X)$ preserves the decomposition. We now simply observe that an ultrafilter on X gives rise (by restriction) to an ultrafilter on each of the irreducible factors. Conversely, a choice of an ultrafilter on each factor gives us an ultrafilter on X . Thus, if $X \cong \prod_{i=1}^n X_i$ is the canonical product decomposition of X , we have an identification $\mathcal{U}(X) \cong \prod_{i=1}^n \mathcal{U}(X_i)$.

An unbounded cocompact CAT(0) cube complex is called *Euclidean* if it contains an $\text{Aut}(X)$ -invariant flat, otherwise it is called *non-Euclidean*. Only the non-Euclidean factors will play an important role in the description of the boundary. It is thus useful to separate all the Euclidean factors from all the non-Euclidean ones. We summarize this section as follows.

Corollary 2.6. *Let X be an unbounded, proper cocompact CAT(0) cube complex. Then*

- (1) X admits a bounded-to-one $\text{Aut}(X)$ -equivariant essential quotient X_{ess} ;
- (2) X_{ess} admits an $\text{Aut}(X)$ -invariant decomposition $X_{\text{ess}} = X_P \times X_E$, where X_E is Euclidean and X_P is a product of irreducible non-Euclidean complexes.

Part (1) of the corollary tells us that if we are considering proper cocompact actions on CAT(0) cube complexes, we may pass to actions on essential cube complexes. Note that then each of the factors described in part (2) of the corollary are also essential. An unbounded essential CAT(0) cube complex whose irreducible factors are all non-Euclidean will be called a *strictly non-Euclidean complex*.

2.6. Splitting off the Euclidean factor. Consider now a group G acting properly and cocompactly on an essential complex $X = X_P \times X_E$. We wish to obtain a corresponding splitting of G ; that is, we aim to show that uniform lattices in $\text{Aut}(X)$ are reducible. Such a result is true in a much more general setting, due to work of [CM]. In our setting, the matter is simplified by the following fact. By definition,

the Euclidean factor X_E has an $\text{Aut}(X_E)$ -invariant flat $\mathbb{E} \subset X_E$. Note that since the action of $\text{Aut}(X_E)$ on X_E is cellular, the action of $\text{Aut}(X_E)$ on \mathbb{E} is discrete. Thus the group G acting on $X_P \times X_E$ acts on $X_P \times \mathbb{E}$ and the projected action on \mathbb{E} is discrete.

However, a theorem of Caprace and Monod [CM] tells us that an irreducible lattice in a product of two infinite, proper CAT(0) spaces projects to an indiscrete action on each of its factors. We thus obtain as a corollary the following.

Corollary 2.7. *Let $X = X_P \times X_E$ be a CAT(0) cube complex and G a group acting properly and cocompactly on X . Then there exists a finite index subgroup $H < G$ such that $H = H_P \times H_E$, where H_P acts properly and cocompactly on X_P and H_E acts properly and cocompactly on X_E .*

2.7. Flipping, skewering and facing hyperplanes. We will be using some notions and results from [CS] which we record here for convenience. An automorphism $g \in \text{Aut}(X)$ of a CAT(0) cube complex is said to *skewer* a hyperplane \hat{h} if for some $n > 0$ and a halfspace \mathfrak{h} bounded by \hat{h} we have $g^n \mathfrak{h} \not\subset \mathfrak{h}$. We say that g *flips* a halfspace space \mathfrak{h} if $g\mathfrak{h}^* \subset \mathfrak{h}$. A halfspace for which there exists no $g \in \text{Aut}(X)$ which flips it is said to be *unflippable*. We then have the following results.

Lemma 2.8 (Single skewering). *Let that X is a finite dimensional CAT(0) cube complex and let G a group acting properly and cocompactly on X . Then for every essential hyperplane \hat{h} in X , there exists $g \in G$ such that g skewers \hat{h} .*

Theorem 2.9 (Flipping Lemma). *Let X be an unbounded CAT(0) cube complex and let G be a group acting properly and cocompactly on X . Let also \mathfrak{h} be a halfspace which is unflippable by the action of G .*

Then X has a decomposition $X = X_1 \times X_2$ into a product of subcomplexes, corresponding to a transverse hyperplane decomposition $\mathcal{H}(X) = \mathcal{H}_1 \cup \mathcal{H}_2$, which satisfies the following properties.

- (1) X_1 is irreducible and all of its hyperplanes are compact.
- (2) Some finite index subgroup $G' \leq G$ preserves the decomposition $X = X_1 \times X_2$.
- (3) The G' -orbit of \hat{h} is in \mathcal{H}_1 .
- (4) X_1 is \mathbb{R} -like, namely quasi-isometric to the real line.

Corollary 2.10 (Double skewering). *Let X be an essential CAT(0) cube complex and G a group acting properly and cocompactly on X . Then for every pair of disjoint hyperplanes, there exists a group element skewering both.*

Remark. The statement in [CS] is somewhat more general than what is stated here, but this is sufficient for our needs.

The next two results involve the existence of hyperplanes that lie in a particular configuration with respect to each other.

Proposition 2.11 (Corner Lemma). *Let X be an essential, non-Euclidean, irreducible cocompact CAT(0) cube complex. Let \hat{h}_1 and \hat{h}_2 be two intersecting hyperplanes of X . Then there exists a pair of hyperplanes which lie in diagonally opposite components of $X \setminus \{\hat{h}_1, \hat{h}_2\}$. (That is, the hyperplanes are separated both by \hat{h}_1 and \hat{h}_2).*

Proposition 2.12 (Facing Triple Lemma). *Let X be an essential non-Euclidean, cocompact complex. Then there exists a facing triple of hyperplanes in X . That is, there exists a triple of disjoint hyperplanes no one of which separates the other two.*

3. The definition of the boundary $B(X)$

3.1. Non terminating ultrafilters. For the rest of this paper we shall usually have as a standing assumption that X is an unbounded, locally finite, finite dimensional CAT(0) cube complex. The focus of this paper will be such complexes whose automorphism group acts cocompactly. A complex which has a cocompact automorphism group will be called a *cocompact complex*. By Corollary 2.6, it suffices to study such complexes which are essential.

Moreover, in light of the fact the Poisson boundary of an abelian group is a point, the Euclidean factor will play no role in the construction of the boundary. Thus given an arbitrary group G acting properly and cocompactly on an essential complex $X = X_P \times X_E$, we will consider a finite index subgroup H of G which preserves the decomposition into the strictly non-Euclidean and Euclidean factors, such that $H = H_P \times H_E$, and focus on the action of H_P on X_P . The boundary constructed will be the boundary of G as well. This follows from the fact that H_E has an Abelian subgroup of finite index which is central in H , and the center of any group acts trivially on its Poisson boundaries.

Thus, we will now let X be a cocompact, essential, strictly non-Euclidean CAT(0) cube complex. We shall be interested in a particular subset of the Roller boundary $\mathcal{U}(X)$. An ultrafilter which has the property that no descending collection of half-spaces terminates is called *non-terminating*. Let $\mathcal{U}_{\text{NT}}(X)$ denote the collection of non-terminating ultrafilters, namely:

$$\mathcal{U}_{\text{NT}}(X) = \{\alpha \in \mathcal{U}(X) \mid \mathfrak{h} \in \alpha \implies \text{there exists } \mathfrak{h}' \in \alpha \text{ with } \mathfrak{h}' \subsetneq \mathfrak{h}\}.$$

We let $B(X)$ denote the closure of $\mathcal{U}_{\text{NT}}(X)$ in the Tychonoff topology on $\mathcal{U}(X)$ defined above.

We note that if X decomposes as a product $X = \prod_i X_i$, then $\mathcal{U}_{\text{NT}}(X) = \prod_i \mathcal{U}_{\text{NT}}(X_i)$. Consequently, we obtain that $B(X) = \prod_i B(X_i)$.

Our aim will be to show that the action on the boundary $B(X)$ of X enjoys the dynamical properties we are interested in. The very first thing we need is that the boundary is not empty.

Theorem 3.1 (Essential complexes have non-empty boundary). *Let X be an essential, cocompact CAT(0) cube complex. Then $\mathcal{U}_{\text{NT}}(X) \neq \emptyset$.*

Remark. The conclusion that $\mathcal{U}_{\text{NT}}(X) \neq \emptyset$ is false without some assumption on the action of the automorphism group, see Figure 3.1. By Theorem 3.1, co-compactness is a sufficient condition.

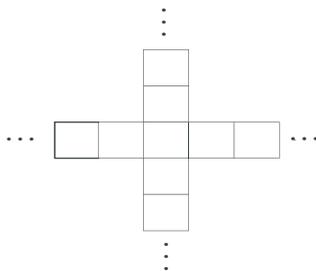


Figure 3.1. An essential complex with no non-terminating ultrafilter.

Proof. We consider the decomposition of $X = X_P \times X_E$ into factors, where X_P are strictly non-Euclidean and X_E is Euclidean. Note that by invariance of the decomposition under a finite index subgroup of $\text{Aut}(X)$, since X is cocompact so is each of the factors X_P and X_E . Since $\mathcal{U}_{\text{NT}}(X) = \mathcal{U}_{\text{NT}}(X_P) \times \mathcal{U}_{\text{NT}}(X_E)$, it suffices to treat separately the cases that X is Euclidean and X is strictly non-Euclidean.

First, let us treat the case that X is Euclidean. The complex X then contains an $\text{Aut}(X)$ -invariant flat $\mathbb{E} \subset X$. Since X is cocompact and each hyperplane is essential, it follows that every hyperplane meets \mathbb{E} . The intersection of the hyperplanes of X with \mathbb{E} gives rise to a collection of lines, which by the finite dimensionality of X fall into finitely many parallelism classes. Since X is essential, each such line has infinitely many parallel lines on either side of it. Choose a ray $r \subset \mathbb{E}$ which is not parallel to any of these lines. For each hyperplane, choose the halfspace containing the infinite part of r . This is then a non-terminating ultrafilter.

For the strictly non-Euclidean factor X_P , we first note that X_P is itself of a product of non-Euclidean irreducible factors $X_P = \prod X_i$. Since then $\mathcal{U}_{\text{NT}}(X_P) = \prod_i \mathcal{U}_{\text{NT}}(X_i)$, it thus suffices to consider the case that is that X is an essential, irreducible, non-Euclidean complex. We then proceed as follows.

We let $\mathcal{H} = \{\hat{h}_1, \hat{h}_2, \dots\}$ be some ordering of the hyperplanes of X . Recall that an ultrafilter α is a choice of a halfspace \mathfrak{h} for each hyperplane \hat{h} satisfying the consistency condition: if $\mathfrak{h} \in \alpha$ and $\mathfrak{h} \subset \mathfrak{k}$, then $\mathfrak{k} \in \alpha$. We will construct an ultrafilter α by describing the halfspaces $\mathfrak{h}_i \in \alpha$ associated to \hat{h}_i in order.

Choose $\mathfrak{h}_1 \in \alpha$ arbitrarily as a halfspace bounded by \hat{h}_1 . For every hyperplane $\hat{h} \subset \mathfrak{h}_1^*$, we choose the halfspace \mathfrak{h} such that $\mathfrak{h}_1 \subset \mathfrak{h}$ (this is dictated by the consistency condition).

Now we consider the hyperplane \hat{h}_n such that n is the smallest number for which $\hat{h}_n \not\subset h_1^*$. That is, the first n for which the choice of the halfspace h_n has yet to be made.

If $\hat{h}_n \cap \hat{h}_1 = \emptyset$, then we choose h_n so that $h_n \subset h_1$. If $\hat{h}_n \cap \hat{h}_1 \neq \emptyset$, then by Lemma 2.11, there exists a hyperplane \hat{h}_m such that $\hat{h}_m \subset h_1$ and $\hat{h}_m \cap \hat{h}_n = \emptyset$. Let h_m denote the halfspace bounded by \hat{h}_m not containing \hat{h}_1 and \hat{h}_n and let h_n denote the halfspace bounded by \hat{h}_n which contains \hat{h}_m . Now for every hyperplane $\hat{h} \subset h_m^*$, we choose the halfspace h such that $h_m \subset h$. Note that these choices do not change whatever choices were made previously. Note that we have now made choices for \hat{h}_1 and \hat{h}_n . We have thus arranged that there exists a halfspace $h \in \alpha$ such that $h \subset h_k$, for $k = 1, \dots, n$.

We continue in this manner choosing a halfspace for each hyperplane \hat{h} . Note that for any two hyperplanes \hat{h} and \hat{k} , the decisions for both are made at some finite stage so that the consistency condition for an ultrafilter are satisfied (i.e., if $h \in \alpha$ and $h \subset k$ then $k \in \alpha$.) Also, by construction, for each $h \in \alpha$, there exists $k \in \alpha$ such that $k \subset h$. Thus, α is non-terminating, as required. \square

Remark 3.2. Note that the above argument shows more, namely that there exists an element of $\mathcal{U}_{NT}(X)$ in each halfspace of X . This fact will be used below.

We now show that the non-terminating ultrafilters make up “most” of $B(X)$.

Proposition 3.3. *Let X be an essential, strictly non-Euclidean, cocompact CAT(0) cube complex. Then $\mathcal{U}_{NT}(X)$ is a dense G_δ in $B(X)$.*

Proof. For a given halfspace h in X we let

$$B_h(X) = \{\alpha \in B(X) \mid h \text{ is minimal in } \alpha\}.$$

Note that

$$B(X) - \mathcal{U}_{NT}(X) = \bigcup_h B_h(X).$$

Thus, to show that $\mathcal{U}_{NT}(X)$ is a dense G_δ in $B(X)$, it suffices to show that each $B_h(X)$ is closed and has empty interior.

Note that $B_h(X) \cap \mathcal{U}_{NT}(X)$ is empty and $B(X)$ is defined to be the closure of $\mathcal{U}_{NT}(X)$ so that $B_h(X)$ has empty interior. To see that $B_h(X)$ is closed, simply observe that

$$B(X) - B_h(X) = (U_{h^*} \cup \bigcup_{f \subsetneq h} U_f) \cap B(X).$$

Thus, the complement of $B_h(X)$ is open, as required. \square

4. General boundary theory

Let B be a compact metrizable space, and let G be a group of homeomorphisms of B . We recall some general definitions related to the dynamics of the G -action on B , and then present a general approach to boundary theory. Our approach is motivated by that of Margulis [M], Ch. VI, and Guivarc'h–Le Page [GLP], and we will show that it can handle not only linear groups acting on projective space, but is in fact well suited to handle the boundary theory of cube complexes as well.

Definition 4.1 (Minimality, proximality, contractibility). (1) The G -action on B is called *minimal* if for every $b \in B$ and for any non-empty open set $U \subset B$, there exist $g \in G$ such that $gb \in U$.

(2) The G -action on B is called *proximal* if for any two points $b_1, b_2 \in B$, there exists a point $c \in B$, such that for every neighbourhood U of c , there exists $g \in G$ such that $gb_1 \in U$ and $gb_2 \in U$.

(3) A neighbourhood V of a point $b \in B$ is called *contractible* [M], Ch. VI, §1, if there exists a point $c \in B$ such that for every neighbourhood U of c , there exists $g \in G$ such that $gV \subset U$.

We shall also make use of a notion which is stronger than the above three conditions. Roughly speaking, it says that the “attracting” conditions stated above for points of B hold with respect to the action on some larger space in which B is contained.

Definition 4.2. Let Y be a G -space. A compact subset $B \subset Y$ is said to be a *boundary limit set* for the action of G on Y if B is G -invariant and

- (1) for every point $y \in Y$, and every non-empty open set $U \subset Y$ such that $U \cap B \neq \emptyset$, there exists $g \in G$ such that $gy \in U$;
- (2) for any two points $y_1, y_2 \in Y$, there exists a point $c \in B$, such that for every neighbourhood U of c , there exists $g \in G$ such that $gy_1, gy_2 \in U$;
- (3) for every $y \in Y$ there exists a neighborhood V of y and a point $c \in B$, such that for every neighbourhood U of c , there exists $g \in G$ with $gV \subset U$.

We now recall the following definition (see [M], Ch. VI, 2.13).

Definition 4.3 (Equicontinuous decomposition). The G -action on B is said to admit an *equicontinuous decomposition* if we can write G as a finite union of subsets $G = \bigcup_{i=1}^N G_i$, and find non-empty open proper subsets $B_i \subset B$ satisfying $G \cdot B_i = B$, such that the set of functions G_i (from B_i to B) is equicontinuous on B_i . Equivalently, fixing a metric d on B , for every $\varepsilon > 0$ there exists $\delta > 0$ such that if $b, b' \in B_i$ satisfy $d(b, b') < \delta$, then $d(gb, gb') < \varepsilon$ for all $g \in G_i$ simultaneously.

We now consider the space $P(B)$ of Borel probability measures on B , which we take with the (metrizable, separable) w^* -topology, namely $\eta_n \rightarrow \eta$ if and only if for

every continuous function f on B , $\int_B f d\eta_n \rightarrow \int_B f d\eta$. For each $b \in B$, let δ_b denote the point measure supported on b . The map $\beta \mapsto \delta_b$ is continuous and embeds B as a compact subset of $P(B)$, the subset of point measures.

Given a probability measure μ on G , we can consider the convolution $\mu * \eta = \int_G g\eta d\mu(g)$, which is another probability measure on B . The operator $\eta \mapsto \mu * \eta$ is continuous and admits a fixed point ν satisfying $\mu * \nu = \nu$. Such a measure is called a μ -stationary measure on B . If ν is a μ -stationary measure, the pair (B, ν) is called a (G, μ) -space.

The following two definitions were introduced in [F1].

Definition 4.4 (Strong proximality, boundary space). (1) The G -action on B is called *strongly proximal* if given any probability measure $\eta \in P(B)$, there exists $b \in B$ and a sequence $g_n \in G$ such that $g_n\eta \rightarrow \delta_b$.

(2) A minimal strongly proximal G -space B is called a *boundary*.

A statistical version of proximality is given by the following.

Definition 4.5 (μ -boundary, μ -proximality, mean proximality). (1) The (G, μ) -space (B, ν) is called a μ -*boundary* (or a (G, μ) -*boundary*) if given a sequence $\omega = (\omega_n)_{n \in \mathbb{N}}$ of independent random variables with values in G and common distribution μ , the sequence of probability measures $\omega_1\omega_2 \dots \omega_n\nu$ converges with probability one to a limit point measure $\delta_{Z(\omega)}$.

(2) The G -space B is called μ -*proximal* if for every μ -stationary measure $\nu \in P(B)$, the (G, μ) -space (B, ν) is a μ -boundary.

We recall the following results.

Theorem 4.6 ([M], Ch. VI., Prop. 1.6). (1) *If B is a proximal G -space and every point has a contractible neighborhood, then B is strongly proximal.*

(2) *If B is proximal, minimal and contains a non-empty contractible open set, then it is strongly proximal, namely a boundary of G .*

As to the uniqueness of stationary measures, we note the following.

Theorem 4.7. *If the G -action on B is strongly proximal and admits an equicontinuous decomposition, then for every probability measure μ on G whose support generates G as a semigroup, the μ -stationary measure ν is unique, and (B, ν) is a (G, μ) -boundary.*

Proof. This result is a corollary of the discussion in [GLP], Ch. 5, as follows. Definition 5.2 in [GLP] introduces Markov systems and Markov kernels on $X \times Y$, with Y a compact metric space endowed with a (semi)group action. Eq. 5.1 introduces a condition of stationarity for the Markov kernel associated with a system of probability

measures on Y . When X reduces to a point, the stationarity condition reduces to the usual notion of stationarity under convolution with a measure μ on the (semi)group (that we are using in our discussion) as noted in the end of §5.1 in [GLP]. The issue of uniqueness of stationary kernels is discussed in §5.3 of [GLP], where a condition denoted by (I-P) is introduced in Definition 5.4 (generalizing condition (i-p) of irreducibility and proximality for linear groups on the space of flags). It is clear that the condition of equicontinuous decomposition stated above is stronger than condition (I-P), and thus the uniqueness of the stationary measure follows from Theorem 5.9 in [GLP], which asserts that the stationary kernel is unique for the more general Markov systems under discussion, provided they satisfy property (I-P). \square

We summarize the previous discussion as follows.

Corollary 4.8. *If the G -action on B is minimal, proximal, contains a non-empty contractible open set and admits an equicontinuous decomposition, then (B, ν) is a boundary. In addition, it is a (G, μ) -boundary for the unique μ -stationary measure ν on B , where μ is any probability measure on G whose support generates G as a semigroup.*

Corollary 4.8 will be our basic tool: we will prove below that the action of a subgroup G of the automorphism group of a CAT(0) cube complex satisfies the hypotheses required by the corollary, provided certain natural assumptions on the cube complex and the group G are satisfied.

5. Establishing that $B(X)$ is a boundary

The aim of this section is to show that for an essential, strictly non-Euclidean CAT(0) cube complex $B(X)$ is a boundary. More precisely, we establish the following theorem.

Theorem 5.1. *Let X be an essential, strictly non-Euclidean CAT(0) cube complex admitting a proper co-compact action of $G \subset \text{Aut}(X)$. Then the G -action on $B(X)$ is minimal and strongly proximal, and every probability measure μ whose support generates G as a semigroup has a unique stationary measure ν . The (G, μ) -space $(B(X), \nu)$ is a μ -boundary.*

In order to prove Theorem 5.1, by Corollary 4.8 it suffices to show that the action satisfies conditions laid out there. In particular, we will show that $B(X)$ is a boundary limit set and that it admits an equicontinuous decomposition.

5.1. $B(X)$ is a boundary limit set. In order to show that $B(X)$ is a boundary limit set, we need to explore the topology on $\mathcal{U}(X)$ and $B(X)$ a bit more closely. Recall

that the collection of subsets of the form $U_{\mathfrak{h}} = \{\alpha \mid \mathfrak{h} \in \alpha\}$, called the collection of $\mathcal{U}(X)$ -halfspaces, forms a sub-basis for the Tychonoff topology on $\mathcal{U}(X)$, so that a basic open set consists of the intersection of finitely many $\mathcal{U}(X)$ -halfspaces. Since $B(X)$ is a subset of $\mathcal{U}(X)$, we obtain a basis of open sets for $B(X)$.

We now observe that for points in $\mathcal{U}_{NT}(X)$ the local neighborhood basis has additional structure, and this fact will play a crucial role below. First, we introduce the following definition.

Definition 5.2 (Sectors). (1) A sector \mathfrak{s} in a cube complex X is a finite intersection of halfspaces $\bigcap_{i=1}^n \mathfrak{h}_i$, for which the corresponding hyperplanes $\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n$ meet, namely $\bigcap_{i=1}^n \hat{\mathfrak{h}}_i \neq \emptyset$.

(2) A sector S of $\mathcal{U}(X)$ is a finite intersection of $\mathcal{U}(X)$ -halfspaces of the form $\bigcap_{i=1}^n U_{\mathfrak{h}_i}$, for which the corresponding hyperplanes $\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n$ meet, namely for which $\bigcap_{i=1}^n \mathfrak{h}_i$ constitute a sector in X .

When S is a sector $\mathcal{U}(X)$, we will use \mathfrak{s} to denote the corresponding sector in X and visa versa.

We then have the following result.

Lemma 5.3. *Let X be a CAT(0) cube complex. Then every neighborhood of a point $\alpha \in \mathcal{U}_{NT}(X)$ contains a neighborhood of α which is a sector in $\mathcal{U}(X)$.*

Proof. First we define a certain notion of complexity for basic open sets. Given a basic open set $U = \bigcap_{i=1}^n U_{\mathfrak{h}_i}$, we define two numbers, $D(U)$ and $N(U)$. The number $N(U) = n$ is simply the number of halfspaces defining U . The number $D(U)$ is the sum of the distances between the hyperplanes (minimized over all presentations of U as an intersection of halfspaces),

$$D(U) = \sum d(\hat{\mathfrak{h}}_i, \hat{\mathfrak{h}}_j),$$

where the distance d between non-intersecting hyperplanes is the minimal length of a 1-skeleton path crossing both hyperplanes, minus one. When $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{h}}'$ intersect, $d(\hat{\mathfrak{h}}, \hat{\mathfrak{h}}') = 0$. We set $C(U) = (N(U), D(U))$, and get a partial ordering on the basic open sets by ordering $C(U)$ lexicographically.

Suppose that $U = \bigcap_{i=1}^n U_{\mathfrak{h}_i}$ is a basic open set containing $\alpha \in \mathcal{U}_{NT}(X)$ so that α is non-terminating. Now we may assume that $C(U)$ is minimal amongst all basic open sets contained in U and containing α .

If $N(U) = 1$ or $D(U) = 0$, then U is a sector and we are done. So suppose that $N(U) > 1$ and $D(U) > 0$.

First suppose that for some i, j , we have $\mathfrak{h}_i \subset \mathfrak{h}_j$. If this were the case, then U is the intersection of $\{\mathfrak{h}_1, \dots, \mathfrak{h}_n\} \setminus \{\mathfrak{h}_j\}$ and so $N(U)$ can be reduced. So we may assume that for all $i \neq j$, $\mathfrak{h}_i \not\subset \mathfrak{h}_j$. We call the process just described “removing extraneous halfspaces”.

Now after a possible renumbering of the h_i 's, the associated hyperplanes \hat{h}_1, \hat{h}_2 are disjoint and $d(\hat{h}_1, \hat{h}_2) > 0$. Now since α is non-terminating and $h_1 \in \alpha$, there exists a halfspace $h'_1 \in \alpha$ such that $h'_1 \subset h_1$. We will now let

$$U' = U'_{h'_1} \cap (\bigcap_{i=2}^n U_{h_i}).$$

By construction, we have $U' \subset U$ and $\alpha \in U'$. Thus, by the minimality of $C(U)$, we may assume that $C(U) \leq C(U')$.

For each $j = 2, \dots, n$, suppose that \hat{h}'_1 is disjoint from \hat{h}_j and does not separate \hat{h}_1 from \hat{h}_j , in which case we obtain that $h'_1 \subset h_j$. Thus U' has extraneous halfspaces that can be removed, reducing $N(U')$ and hence $C(U') < C(U)$, a contradiction.

We thus have two possibilities: \hat{h}'_1 intersects every $\hat{h}_j, j = 2, \dots, n$. In this case $D(U') < D(U)$, a contradiction.

There exists \hat{h}_j , disjoint from \hat{h}'_1 . In this case, for each such j , \hat{h}'_1 must separate \hat{h}_1 and \hat{h}_j , so that $d(\hat{h}'_1, \hat{h}_j) < d(\hat{h}_1, \hat{h}_j)$. We thus obtain $D(U') < D(U)$, again a contradiction. \square

We will now need a variant of the Corner Lemma which ensures that certain types of halfspace intersections contain hyperplanes. Let h_1 and h_2 be two halfspaces associated to intersecting hyperplanes \hat{h}_1 and \hat{h}_2 . Note that the hyperplane \hat{h}_1 separates the hyperplane \hat{h}_2 into two *half-hyperplanes*, $\hat{h}_2^+ = \hat{h}_2 \cap h_1$ and $\hat{h}_2^- = \hat{h}_2 \cap h_1^*$. Similarly, \hat{h}_1 is subdivided into two half-hyperplanes by \hat{h}_2 , where $\hat{h}_1^+ = \hat{h}_1 \cap h_2$. A half-hyperplane is called *R-shallow* if it contained in an R -neighborhood of $\hat{h}_1 \cap \hat{h}_2$. Otherwise it is called *deep*. Note that $h_1 \cap h_2$ is bounded by the union of the two half-hyperplanes \hat{h}_1^+ and \hat{h}_2^+ . Note also that since there are only finitely many orbits of half-hyperplanes, there exists some global constant $R > 0$ such that if \hat{h}_1^+ is not R -shallow, then \hat{h}_1^+ is deep. We record this fact.

Remark 5.4. If X is a cocompact cube complex, then there exists a number $R > 0$ such that any half-hyperplane which is not R -shallow is deep.

Lemma 5.5. *Let X be a cocompact, essential, irreducible, non-Euclidean CAT(0) cube complex and let \hat{h}_1 and \hat{h}_2 be hyperplanes in X . Let $h_1 \cap h_2$ be one of the four sectors defined by intersecting hyperplanes \hat{h}_1 and \hat{h}_2 . Suppose that one of the half-hyperplanes \hat{h}_1^+ and \hat{h}_2^+ is deep. Then the sector $h_1 \cap h_2$ contains a hyperplane.*

Proof. There are several cases, which we handle by order of difficulty.

Case 1: One of \hat{h}_1^+, \hat{h}_2^+ is shallow.

Suppose that, without loss of generality, the \hat{h}_1^+ is shallow, so that \hat{h}_2^+ is deep. Thus, there exists $R > 0$ such that \hat{h}_1^+ is contained in the R -neighborhood of \hat{h}_2 .

Since \hat{h}_2 is essential, there exists a hyperplane $\hat{f} \subset \hat{h}_2$ such that $d(\hat{f}, \hat{h}_2) > R$. It follows that $\hat{f} \cap \hat{h}_1 = \emptyset$. If $\hat{f} \subset \hat{h}_1$, we are done, so suppose that $\hat{f} \subset \hat{h}_1^*$. Consider a geodesic path γ joining \hat{f} and \hat{h}_2 and let $p = \gamma \cap \hat{h}_2$. Since the action of $\text{stab}(\hat{h}_2)$ on \hat{h}_2 is cocompact and \hat{h}_2^+ is deep, there exists some $g \in \text{stab}(\hat{h}_2)$ such that $gp \in \hat{h}_1$. We then have $g\gamma \in \hat{h}_1$, so that $g\hat{f} \in \hat{h}_1$, as required.

Case 2: Both \hat{h}_1^+, \hat{h}_2^+ are deep and one of \hat{h}_1^-, \hat{h}_2^- is shallow.

Suppose without loss of generality that \hat{h}_1^- is shallow. By applying case 1, we have that there exists a hyperplane $\hat{f} \subset \hat{h}_1 \cap \hat{h}_2^*$. Since the action of $\text{stab}(\hat{h}_1)$ on \hat{h}_1 is cocompact, there exists some $g \in \text{stab}(\hat{h}_1)$ such that $g\hat{h}_2 \subset \hat{h}_2$. Since \hat{h}_1^- is shallow, we know that g does not skewer \hat{h}_2 . It follows that $g\hat{h}_2^* \subset \hat{h}_2$. Thus $\hat{f} \subset \hat{h}_1 \cap \hat{h}_2$, as required.

Case 3: All half hyperplanes $\hat{h}_1^\pm, \hat{h}_2^\pm$ are deep.

By Lemma 2.11, there exists a pair of diagonally opposite sectors of \hat{h}_1 and \hat{h}_2 which contain hyperplanes. If that pair is $\hat{h}_1 \cap \hat{h}_2$ and $\hat{h}_1^* \cap \hat{h}_2^*$, then we are done, so we may assume that $\hat{h}_1 \cap \hat{h}_2^*$ and $\hat{h}_1^* \cap \hat{h}_2$ both contain hyperplanes. If $\hat{h}_2 \cap \hat{h}_1$ is flippable in \hat{h}_1 (i.e., there exists $g \in \text{stab}(\hat{h}_1)$ such that $g\hat{h}_2^* \subset \hat{h}_2$), then letting \hat{f} denote a hyperplane contained in $\hat{h}_1 \cap \hat{h}_2^*$, we would have $g\hat{f} \subset \hat{h}_1 \cap \hat{h}_2$, as required.

So let us assume that $\hat{h}_2 \cap \hat{h}_1$ is unflippable in \hat{h}_1 . We let \hat{f} denote a hyperplane such that $\hat{f} \subset \hat{h}_1^* \cap \hat{h}_2$. Since the hyperplane $\hat{h}_1 \cap \hat{h}_2$ is essential in \hat{h}_2 , there exists $g \in \text{stab}(\hat{h}_2)$ such that $g\hat{h}_1 \subset \hat{h}_1$. We then have that for some n , either $g^n\hat{f} \subset \hat{h}_1$ or $g^n\hat{f} \cap \hat{h}_1 \neq \emptyset$. By passing to possibly some larger power of g , we obtain that the half-hyperplane $g^n\hat{f}^+ = g\hat{f} \cap \hat{h}_1$ is deep.

Let $\hat{m} = g^n\hat{f}$. We now observe that for some power of $n > 0$, $g^n\hat{m} \cap \hat{m} = \emptyset$. This is clear since the axis of g does not lie in neighborhood of \hat{m} . There are three hyperplanes in \hat{h}_1 : $L_1 = \hat{h}_1 \cap \hat{h}_2$, $L_2 = \hat{h}_1 \cap \hat{m}$ and $L_3 = \hat{h}_1 \cap g^n\hat{m}$. We already are assuming that L_1 is essential in \hat{h}_1 . If either of L_2 or L_3 is inessential in \hat{h}_1 , we appeal to Case 1 above to produce the desired hyperplane. It follows that both L_2 and L_3 are essential in \hat{h}_1 . Thus, by the Flipping Lemma 2.9, the hyperplanes L_1, L_2 and L_3 do not form a facing triple. This means that one of them, without loss of generality L_2 separates L_1 from L_3 . But this contradicts the fact that \hat{h}_2, \hat{m} and $g^n\hat{m}$ form a facing triple in X . \square

Now we can prove our central technical result.

Proposition 5.6. *Let X be a cocompact, essential, irreducible, non-Euclidean CAT(0) cube complex admitting a proper cocompact group action. Let $S = \bigcap_{i=1}^k U_{h_i}$ be a sector neighborhood of a non-terminating ultrafilter. Then there exists a hyperplane $\hat{h} \in \mathfrak{s}$.*

Proof. The proof will be by induction on n . We start with $n = 2$. Let $\mathfrak{s} = \hat{h}_1 \cap \hat{h}_2$.

If either of the half-hyperplanes bounding \mathfrak{s} are deep, then we may apply Lemma 5.5 to conclude that \mathfrak{s} contains a hyperplane. Otherwise, suppose that \hat{h}_1^+ is shallow, then we may apply the fact that α is non-terminating to construct a sequence of halfspaces $h_2 \supset h_3 \supset \dots$ such that $\alpha \in U_{h_i}$ for all $i > 1$. Now since \hat{h}_1^+ is shallow there exists some n such that $\hat{h}_n \cap \hat{h}_1 = \emptyset$. We then have $\hat{h}_n \subset h_1 \cap h_2$, as required.

We now suppose that $n > 3$. Let $L = \bigcap_{i=2}^n \hat{h}_i$. Let $L^+ = L \cap h_1$ and $L^- = L \cap h_1^*$. We first claim that L^+ is not contained in any neighborhood of \hat{h}_1 . Since $\alpha \in U_{h_1}$ and α is non-terminating, there exists a sequence of halfspaces $h_1 \supset m_2 \supset m_3 \supset \dots$ such that $\alpha \in U_{m_i}$ for all i . Suppose that for some i we have $\hat{m}_i \cap L = \emptyset$. Then there exists some \hat{h}_i , $2 \leq i \leq n$ such that $\hat{h}_i \cap \hat{m} = \emptyset$. Let $F \subset \{2, \dots, n\}$ such that $\hat{h}_i \cap \hat{m} \neq \emptyset$.

By assumption, we have that F is a proper subset of $\{2, \dots, n\}$. If $F = \emptyset$, then we would have $\hat{m} \subset \mathfrak{s}$, and we are done. So now let $\mathfrak{s}' = m \cap \bigcap_{i \in F} h_i$. By construction $\mathfrak{s}' \subset \mathfrak{s}$ and $\alpha \in \mathfrak{s}'$. Thus, by induction, \mathfrak{s}' contains a hyperplane and we are done. We thus have that all the hyperplanes \hat{m}_i intersect L^+ . Since the hyperplanes \hat{m}_i correspond to a nested collection of halfspaces, we have shown our claim.

Next we apply the inductive hypothesis to the sector $\bigcap_{i=2}^n h_i$ to conclude that there exists a hyperplane $\hat{f} \in \bigcap_{i=2}^n h_i$. If $\hat{f} \subset h_1$, we are done, so suppose not, so that $\hat{f} \subset h_1^*$ or $\hat{f} \cap h_1 = \emptyset$.

Let R denote the number described in Remark 5.4, so that any half-hyperplane which is not R -shallow is deep. We claim that there exists $g \in \text{stab}(L)$ such that $g\hat{f} \cap h_1$ contains a point at distance greater than R from \hat{h}_1 . To see this, let p be some point of $\bigcap_{i=1}^n \hat{h}_i$. Let q be a point of \hat{f} and let $D = d(p, q)$. Now since L^+ is deep, there exists $g \in \text{stab}(L)$ such that $d(gp, \hat{h}_1) > R + D$. It follows from the triangle inequality that $d(gq, \hat{h}_1) > R$.

From this claim it follows that either $g\hat{f} \subset h_1$, in which case we are done, or $g\hat{f} \cap h_1 \neq \emptyset$ and the half-hyperplane $g\hat{f}^+ = g\hat{f} \cap h_1$ is deep. In the latter case, we apply Lemma 5.5 to conclude that both sectors bounded by $g\hat{f}^+$ and \hat{h}_1 contain hyperplanes. One of these sectors is contained in \mathfrak{s} , so that we have a hyperplane contained in \mathfrak{s} as required. \square

Putting together the previous two lemmas, we obtain the following.

Corollary 5.7. *Let X be a cocompact, essential, irreducible, non-Euclidean CAT(0) cube complex. If U is an open set in $\mathcal{U}(X)$ such that $U \cap B \neq \emptyset$, then U contains a sub-basic open set $U_{\hat{h}} \subset U$.*

We are now ready to prove that $B(X)$ is indeed a boundary limit set.

Theorem 5.8. *Let X be an essential, strictly non-Euclidean CAT(0) cube complex admitting a proper co-compact action of $G \subset \text{Aut}(X)$. Then $B(X)$ is a boundary limit set for the G action on $\mathcal{U}(X)$.*

Proof of Theorem 5.8. We need to establish the three properties in Definition 4.2 where $Y = \mathcal{U}(X)$ and $B = B(X)$.

We will prove the following stronger claim which immediately implies all three properties.

Claim. For any two ultrafilters α and β there exists an open neighborhood V of α and β such that for any open set $U \subset \mathcal{U}(X)$ such that $U \cap B(X) \neq \emptyset$, there exists $g \in G$ such that $gV \subset U$.

First, we establish the claim in the event that X is irreducible. Let $\alpha, \beta \in \mathcal{U}(X)$. By the Facing Triple Lemma 2.12, there exists in X a facing triple of hyperplanes T . It follows that there exists a halfspace \mathfrak{h} , associated to a hyperplane in T , such that $\alpha, \beta \in U_{\mathfrak{h}}$. We set $V = U_{\mathfrak{h}}$. Now given any open set U such that $U \cap B(X) \neq \emptyset$, we know by Corollary 5.7 that there exists a hyperplane $\hat{\mathfrak{f}} \subset U$. Thus, for one of the halfspaces bounded by $\hat{\mathfrak{f}}$, we have $U_{\hat{\mathfrak{f}}} \subset U$. Now by appealing to the Flipping Lemma or Double Skewering Lemma of [CS], we have $g \in G$ such that $gV \subset U_{\hat{\mathfrak{f}}} \subset U$, as required.

We now need to establish the claim when X is not irreducible. We shall establish the claim when X is a product of two irreducible factors. The general case is established by induction.

Suppose that $X = X_1 \times X_2$. After possible passing to a subgroup of finite index in G , we may assume that G preserves the decomposition. Let $\alpha = (\alpha_1, \alpha_2)$ and $\beta = (\beta_1, \beta_2)$ be two ultrafilters in $\mathcal{U}(X)$. As above, in each factor X_i , we may find a halfspace \mathfrak{h}_i such that $\alpha_i, \beta_i \in U_{\mathfrak{h}_i}$. We let $V = U_{\mathfrak{h}_1} \times U_{\mathfrak{h}_2}$.

Let U be some open set intersecting $B(X)$, so that U contains a non-terminating ultrafilter $\gamma = (\gamma_1, \gamma_2)$ in $\mathcal{U}_{\text{NT}}(X)$. It follows that there exists an elementary neighborhood $U_1 \times U_2$ such that $\gamma \in U_1 \times U_2 \subset U$. Note that γ_1 and γ_2 are non-terminating so that U_1 and U_2 intersect $B(X_1)$ and $B(X_2)$ respectively. By Corollary 5.7, we have sub-basic open sets $U_{\mathfrak{f}_1} \subset U_1$ and $U_{\mathfrak{f}_2} \subset U_2$.

We now choose a vertex $v \in X_2^{(0)}$ and let $X_1^v \equiv X_1 \times \{v\}$. Note that there is a natural isometric identification of X_1 with X_1^v ; we thus identify subsets of X_1 with subsets of X_1^v . Note that $\text{stab}(X_1^v)$ acts cocompactly on X_1^v . (This is because G acts preserving the product structure so that every element $g \in G$, either $g \in \text{stab}(X_1^v)$ or $g(X_1^v) \cap X_1^v = \emptyset$.) Thus we may apply the claim in the irreducible case to find some $g_1 \in \text{stab}(X_1^v)$ such that $g_1 U_{\mathfrak{h}_1} \subset U_{\mathfrak{f}_1}$.

Next we consider the $\hat{\mathfrak{m}} \equiv \hat{\mathfrak{f}}_1 \times X_2$ which is a hyperplane in X . Note again that $\text{stab}(\hat{\mathfrak{m}})$ acts on $\hat{\mathfrak{m}}$ cocompactly. Choose some vertex in $w \in \hat{\mathfrak{m}}$ and let $X_2^w \equiv \{w\} \times X_2$. Now $\text{stab}(X_2^w) < \text{stab}(\hat{\mathfrak{m}})$ acts cocompactly on X_2^w . Once again, we may apply the claim in the irreducible case, to find a $g_2 \in \text{stab}(X_2^w)$ such that $g_2(U_{\mathfrak{h}_2}) \subset U_{\mathfrak{f}_2}$. Note that $g_2 \in \text{stab}(\hat{\mathfrak{m}})$ so that it preserves the halfspace $\mathfrak{f}_1 \times X_2$. We then have $g_2 g_1 V \subset U$, as required. □

5.2. $B(X)$ admits an equicontinuous decomposition. In this section we will prove the following theorem.

Theorem 5.9. *Let X be an essential, strictly non-Euclidean CAT(0) cube complex and let $G < \text{Aut}(X)$ be a proper, cocompact action on X . Then the G -action admits an equicontinuous decomposition.*

The proof of the theorem rests on the lemma below which uses strong separation to produce equicontinuity. We recall the definition of Behrstock and Charney [BC] of strongly separated hyperplanes: two disjoint hyperplanes \hat{h} and \hat{k} are said to be *strongly separated* if there are no hyperplanes which intersect both \hat{h} and \hat{k} . We say that \hat{h} and \hat{k} are *super-strongly separated* if they are strongly separated and every hyperplane meeting \hat{h} is disjoint from every hyperplane meeting \hat{k} .

Lemma 5.10. *Let G be a group acting on a CAT(0) cube complex X . Let \hat{h} and \hat{k} be two super-strongly separated hyperplanes such that $\hat{h} \cap \hat{k} = \emptyset$. Let $G_{\hat{h}} \subset G$ denote the collection of elements which have an axis in \hat{h} or have a fixed point in \hat{h} . Then $G_{\hat{h}}$ acts equicontinuously on $B_{\hat{k}} \equiv U_{\hat{k}} \cap B$.*

Proof. Suppose that $G_{\hat{h}}$ does not act equicontinuously on $U_{\hat{k}}$. Then there exists some $\varepsilon > 0$, a sequence of elements $\alpha_n, \beta_n \in U_{\hat{k}}$, and a sequence of group elements $g_n \in G_{\hat{h}}$, such that $d(\alpha_n, \beta_n) \rightarrow 0$ and $d(g_n\alpha_n, g_n\beta_n) \geq \varepsilon$. Recall the definition of the metric on $\mathcal{U}(X)$ (and consequently on B) is a metric which is compatible with the Tychonoff topology on the product

$$\prod_{\hat{h}} \{\hat{h}, \hat{h}^*\}.$$

This means that the hyperplanes are ordered in some manner and the distance between points is $1/n$ where n is the first hyperplane on which they differ. We record now what the lack of equicontinuity above means explicitly in these terms: after possibly passing to a subsequence $\{g_n\}$ above, there exists a sequence of hyperplanes $\{\hat{h}_n\}$ such that

- (1) there exists some fixed hyperplane \hat{h}_0 so that $g_n\hat{h}_n = \hat{h}_0$,
- (2) $d(\hat{h}_n, w) \rightarrow \infty$, where w is some chosen vertex of X ,
- (3) $\hat{h}_n \cap \hat{k} \neq \emptyset$.

For convenience it is useful to work with $\gamma_n \equiv g_n^{-1}$. Note that since $g_n \in G_{\hat{h}}$, so is γ_n . We also choose $w \in \hat{k}^*$. By definition, $\gamma_n \in G_{\hat{h}}$ and $\gamma_n\hat{h}_0 = \hat{h}_n$. We will choose n to be sufficiently large; how large will become clear in due course. First note that because γ_n either has an axis in \hat{h} or a fixed point in \hat{h} , we know that either $\gamma_n\hat{h} \cap \hat{h} \neq \emptyset$ or $\gamma_n\hat{h} \subset \hat{h}$. In either case, strong separation of \hat{h} and \hat{k} tells us that $\gamma_n\hat{h} \subset \hat{k}^*$. If $\hat{h} \cap \hat{h}_0 \neq \emptyset$, then $\gamma_n\hat{h} \cap \gamma_n\hat{h}_0 \neq \emptyset$, so by super-strong separation $\gamma_n\hat{h}_0$ is disjoint from \hat{k} . This tells us that $\hat{h}_n = \gamma_n\hat{h}_0 \subset \hat{k}^*$ as well, contradicting (3) above. We can thus assume that $\hat{h} \cap \hat{h}_0 = \emptyset$.

Let $D_1 = d(\hat{h}, \hat{h}_0)$. Then $d(\gamma_n \hat{h}, \hat{h}_n) = D_1$, so that either $\hat{h}_n \cap \hat{f} = \emptyset$ or $d(\hat{h}_n, \hat{f}) \leq D_1$. Let $D_2 = d(w, \hat{f})$ and let $\hat{f}_1, \dots, \hat{f}_m$ denote the hyperplanes separating w and \hat{h}_n . We can choose n so that m is as large as we like, so we choose it so that $m > 2D_1 + D_2$. It follows that $m - D_1 - D_2$ of the \hat{f}_i 's intersect \hat{f} non-trivially. By strong separation of \hat{h} and \hat{f} , we then know that these \hat{f}_i 's do not intersect \hat{h} and hence separate \hat{h}_n from \hat{h} . If $\gamma_n \hat{h} \cap \hat{h} = \emptyset$, and hence $\gamma_n \hat{h} \subset \mathfrak{h}$. Thus the \hat{f}_i 's above separate $\hat{h}_n = \gamma_n \hat{h}_0$ and $\gamma_n \hat{h}$. But then $d(\hat{h}, \hat{h}_0) > D_1$, a contradiction. We thus have $\gamma_n \hat{h} \cap \hat{h} \neq \emptyset$. By super-strong separation, we then have that all the \hat{f}_i 's are disjoint from $\gamma_n \hat{h}$. Once again, we obtain that $d(\hat{h}_n, \gamma_n \hat{h}) > D_1$, a contradiction. \square

Proof of Theorem 5.9. Let $X = \prod_{i=1}^n X_i$ be the decomposition of X into irreducible factors. First we prove the theorem for the action of G on each of the irreducible factors. Since X has no Euclidean factor, by the Flipping Lemma 2.9 every halfspace in X is flippable. It follows that every halfspace in each factor X_i is flippable. This fact in turn implies that the action of G on the visual boundary of X_i has no fixed point at infinity. Since X_i is non-Euclidean, Theorem 2.12 tells us there exists a facing triple of hyperplanes $\hat{h}_1, \hat{h}_2, \hat{h}_3$. Since no element can skewer all such hyperplanes, it follows that $G = \bigcup_{i=1}^3 G_{\mathfrak{h}_i} \cup G_{\mathfrak{h}_i^*}$. We now need to show that for each halfspace \mathfrak{h} of X_i , there exists $\hat{f} \in \mathfrak{h}$ such that \hat{f} and \hat{h} are super-strongly separated. By Proposition 5.1 of [CS], there exist hyperplanes \hat{h}_1, \hat{h}_2 such that $\mathfrak{h}_1 \in \mathfrak{h}$ and $\mathfrak{h}_2 \in \mathfrak{h}^*$ such that \hat{h}_1 and \hat{h}_2 are strongly separated. By the Double Skewering Corollary 2.10, there exists a group element $g \in G$ double skewering \hat{h}_1 and \hat{h}_2 . That is $g\mathfrak{h}_2^* \subset \mathfrak{h}_1$. It follows that $g\hat{h}_1$ and \hat{h} are strongly separated. By applying a group element that double skewers $g\hat{h}_1$ and \hat{h} we obtain a hyperplane in \mathfrak{h} which is super-strongly separated from \hat{h} .

We can now apply Lemma 5.10 to produce a non-empty B_i for each of the finite subsets $G_{\mathfrak{h}_i^{(*)}}$ of the decomposition above.

Having established that equicontinuous decomposition for each of the factors, we then have $B(X) = \prod_{i=1}^n B(X_i)$. For each factor X_i , we now have equicontinuous decompositions $G = \bigcup_j G_j^i$, so that for each j , G_j^i , there exists $B_j^i \neq \emptyset$ such that G_j^i acts equicontinuously on B_j^i . We then observe that given any $J = \{j_1, \dots, j_n\}$ (where the j_i 's are within the requisite finite range), we have that $G_J \equiv G_{j_1}^1 \cap \dots \cap G_{j_n}^n$ acts equicontinuously on $B_J \equiv B_{j_1}^1 \times \dots \times B_{j_n}^n$. Note further that since G acts minimally on B and the B_J 's are all non-empty, we have $G \cdot B_J = B$. Thus $G = \bigcup_J G_J$ is an equicontinuous decomposition for the action. \square

6. Intervals and property A

We now describe the properties of intervals connecting a vertex in the cube complex to a boundary point or two boundary points. The existence of intervals will prove to

be an indispensable tool in analyzing the properties of the boundary, as we will see below. The discussion in the present section pertains to general finite dimensional locally-finite cube complexes, and the group action will not come into play at all.

Given three ultrafilters α, β, γ , recall that the *median* of the three ultrafilters is defined by

$$\text{med}(\alpha, \beta, \gamma) = (\alpha \cap \beta) \cup (\beta \cap \gamma) \cup (\alpha \cap \gamma).$$

It is easy to check that $\text{med}(\alpha, \beta, \gamma)$ is itself an ultrafilter. Given two ultrafilters, α and β , we now define the interval between α and β to be

$$[\alpha, \beta] = \{\gamma \mid \text{med}(\alpha, \beta, \gamma) = \gamma\}.$$

In the case that α and β are vertices of X , the interval $[\alpha, \beta]$ is simply the collection of vertices that lie on some 1-skeleton geodesic between α and β (see [ChN]).

The goal of the present section is to prove the following Følner-type property for such intervals. This will be used later on to show, among other things, that point stabilizers are amenable. Let $B(r)$ denote the ball of radius r about some base vertex and $|\cdot|$ the size of the intersection with the 0-skeleton.

Theorem 6.1. *Let v, w be a vertices in a cube complex X and $\alpha \in \mathcal{U}_{\text{NT}}(X)$ be a non-terminating ultrafilter. Then (for any fixed choice of base vertex as the center of the balls)*

$$\lim_{r \rightarrow \infty} \frac{|([v, \alpha] \Delta [w, \alpha]) \cap B(r)|}{|([v, \alpha] \cup [w, \alpha]) \cap B(r)|} = 0.$$

Remark. The property stated in the theorem may not hold for α not in $\mathcal{U}_{\text{NT}}(X)$. For example, consider the standard squaring of the Euclidean plane. There are vertical hyperplanes and horizontal ones. Let $\alpha = (\infty, 0)$ be the ultrafilter that contains the right halfspace of every vertical hyperplane, the upper halfspace of every horizontal hyperplane below the x -axis and the lower halfspace for every horizontal hyperplane above the x -axis. Let $v = (0, 0)$ and $w = (1, 0)$. Then it is easy to see that

$$\lim_{r \rightarrow \infty} \frac{|([v, \alpha] \Delta [w, \alpha]) \cap B(r)|}{|([v, \alpha] \cup [w, \alpha]) \cap B(r)|} = 1/2.$$

Before beginning the proof, we recall that a theorem of [BC+] tells us that intervals embeds (ℓ_1 -isometrically) in the standard Euclidean cube complex in \mathbb{R}^n . This property is the key to analyzing the properties of intervals, which will occupy us in the next five subsections. The proof of Theorem 6.1 will conclude in §6.6.

6.1. Basics of intervals. We will first need a few basic facts about intervals.

Lemma 6.2. *For any $\alpha, \beta \in \mathcal{U}(X)$, we have $[\alpha, \beta] = \{\gamma \mid \alpha \cap \beta \subset \gamma\}$.*

Proof. Observe that

$$\text{med}(\alpha, \beta, \gamma) = \gamma \Leftrightarrow (\alpha \cap \beta) \cup (\beta \cap \gamma) \cup (\alpha \cap \gamma) = \gamma \Leftrightarrow \alpha \cap \beta \subset \gamma. \quad \square$$

Lemma 6.3. *For any $\alpha, \beta, \gamma \in \mathcal{U}(X)$ and $m = \text{med}(\alpha, \beta, \gamma)$, we have $[\alpha, \gamma] \cap [\beta, \gamma] = [m, \gamma]$.*

Proof. We have $\alpha \cap m = \alpha \cap ((\alpha \cap \beta) \cup (\alpha \cap \gamma) \cup (\beta \cap \gamma)) = (\alpha \cap \beta) \cup (\alpha \cap \gamma)$. Hence by Lemma 6.2,

$$[\alpha, m] = \{\mu \mid \alpha \cap m \subset \mu\} = \{\mu \mid (\alpha \cap \beta) \cup (\alpha \cap \gamma) \subset \mu\} = [\alpha, \beta] \cap [\alpha, \gamma]. \quad \square$$

Recall that the carrier $C(\hat{h})$ of a hyperplane \hat{h} is the union of all cubes intersecting \hat{h} . Thus the vertices of $C(\hat{h})$ may be viewed as the ultrafilters satisfying DCC which contain \mathfrak{h} or \mathfrak{h}^* as minimal elements. We may then extend the notion of the carrier to all of $\mathcal{U}(X)$: an ultrafilter α is in $C(\hat{h})$ if it contains \mathfrak{h} or \mathfrak{h}^* as a minimal element.

We now have the following:

Lemma 6.4. *Suppose that $\alpha \Delta \beta < \infty$. Then for any $\gamma \in \mathcal{U}(X)$ we have*

$$[\alpha, \gamma] - [\beta, \gamma] \subset \bigcup_{\substack{\hat{h} \text{ separating} \\ \alpha \text{ and } \beta}} C(\hat{h}).$$

Proof. Consider $\mu \in [\alpha, \gamma] - [\beta, \gamma]$. Since $\mu \notin [\beta, \gamma]$, there exists \mathfrak{h} such that $\mathfrak{h} \in \beta, \gamma$ but $\mathfrak{h}^* \in \mu$. Since $\mu \in [\alpha, \gamma]$, we cannot have that \hat{h} separates μ from both α and γ . Thus, we have $\mathfrak{h}^* \in \alpha$. Notice now that by the same reasoning, for any hyperplane \hat{f} separating μ and \hat{h} , we also have that \hat{f} separates β and γ from μ and α .

Now there are only finitely many hyperplanes separating α and β , so choose \hat{f} to be such a hyperplane that is closest to μ . We then have $\mu \in C(\hat{f})$, as required. \square

6.2. Spheres, balls and invariance of basepoint. From now on in this section, we will let I denote an interval. We use the following notation for balls and spheres in I :

$$\begin{aligned} B_v(I, r) &= \text{the ball of radius } r \text{ about } v \text{ in } I, \\ S_v(I, r) &= \text{the sphere of radius } r \text{ about } v \text{ in } I. \end{aligned}$$

We also use $|\cdot|$ to denote the number of vertices in subset of I .

Proposition 6.5. *Suppose that I is an interval and $I = [v, \alpha]$, where v is a vertex of I . Then*

$$\lim_{r \rightarrow \infty} \frac{|S_v(I, r)|}{|B_v(I, r)|} = 0.$$

In fact, we shall prove the following more general and sharper statement.

Lemma 6.6. *Let D be a natural number and $\varepsilon > 0$ a real number. Then there exists a number $R(\varepsilon, D)$ such that for any interval of dimension less than or equal to D , we have $\frac{|S_v(I,r)|}{|B_v(I,r)|} < \varepsilon$ for all $r > R(D, \varepsilon)$.*

The difference between the proposition and the lemma is that the R chosen for the ε does not depend on the complex, just on the dimension and ε .

Proof. Since we will always be working with the same vertex v in this lemma, we shall write $S(r) = S_v(I, r)$ and $B(r) = B_v(I, r)$.

We proceed by induction on the dimension of the complex. For dimension 1, the theorem is clear since $|S(r)| = 1$ for all r , so $\frac{|S(r)|}{|B(r)|}$ is either 0 or $1/r$. So we assume that the theorem is true for complexes of dimension less than D . In particular we assume the proposition holds for hyperplanes in I .

We assume that ε is given. We make some choices of numbers now that we need for later. By induction, there exists an $R_1 = R(D - 1, \varepsilon/2)$, so that, for any complex of dimension less than D and any $r > R_1$, $\frac{|S_r|}{|B_r|} < \varepsilon/2$. Let S_{\max} denote the maximal number of points in the sphere of radius R_1 in a complex of dimension less than D , and let $M = R_1 \cdot S_{\max}$. (The number S_{\max} exists because there are only finitely many such possible spheres.)

We choose $R > \max\{R_1, 2DM/\varepsilon\}$. We claim that R is our required $R(D, \varepsilon)$.

We note that \mathbb{R}^D can be naturally factorized as a product $\mathbb{R}^D = \mathbb{R}^{D-1} \times \mathbb{R}$ in D different ways. Now we consider our embedding of I into \mathbb{R}^D . We assume that the sphere of radius R is non-empty, for otherwise, the proposition is clearly true. If we take a vertex $w \in I$ such that $d(v, w) = R$, then for one of the projections $\pi : \mathbb{R}^D \rightarrow \mathbb{R}$ has the property that $d(\pi(v), \pi(w)) > R/D > 2M/\varepsilon$ (recall that we are using the ℓ_1 metric, which is the same as the 1-skeleton metric on the vertices of I). We factorize $\mathbb{R}^D = \mathbb{R}^{D-1} \times \mathbb{R}$ so that the projection onto the second factor satisfies the above: some point in S_R projects onto a point at distance at least $2M/\varepsilon$ from v . We call the second factor the *vertical* direction and the hyperplanes transverse to them *horizontal* hyperplanes, so that we have at least $2M/\varepsilon$ horizontal hyperplanes intersecting the ball of radius R . Let N denote the number of hyperplanes meeting the ball of radius R .

We let L_n denote the n -th horizontal hyperplane (counting from the bottom). We let $\pi_n : I \rightarrow L_n$ denote the projection of I onto L_n . Note that $L_n = \pi_n(I)$ is an interval, namely $L_n = [\pi_n(v), \pi_n(\alpha)]$. Consequently, $B_{r,n} = B_r \cap L_n$ is a ball; we let $d_{n,r}$ denote the radius of this ball. We write $S_{r,n} = S_r \cap L_n$.

Observation. Note that for any n, r and $m > n$, we have $d_{m,r} \leq d(n, r) - (m - n)$.

Let L_k be the first hyperplane such that $L_k \cap S_R \neq \emptyset$ and $d_{k,R} \leq R_1$. By the observation above, we have at most R_1 hyperplanes above L_k which intersect the

ball of radius R . Moreover, for each of these, we have that the radius $d_{n,r} \leq R_1$. Thus there are at most M points in the spheres $S_{n,R}$ for $n \geq k$. For all $n < k$, we have that $S_{n,R}$ is empty or $d(k, R) > R_1$. In either case, for all $n < k$, we have $|S_{n,R}| < \frac{\varepsilon}{2} |B_{n,R}|$. We now look at the ratio we are considering and break it up into layers:

$$\frac{|S_R|}{|B_R|} = \frac{\sum_{i=1}^k |S_{i,R}| + \sum_{i=k+1}^N |S_{i,R}|}{\sum_{i=1}^N |B_{i,R}|} < \frac{\frac{\varepsilon}{2} \sum_{i=1}^k |B_{i,R}|}{\sum_{i=1}^N |B_{i,R}|} + \frac{M}{N} < \varepsilon/2 + \varepsilon/2 = \varepsilon.$$

□

Corollary 6.7. *For any two vertices v and w and for any $n \in \mathbb{Z}$, we have*

$$\lim_{r \rightarrow \infty} \frac{B_v(I, r+n)}{B_w(I, r)} \rightarrow 1.$$

Proof. First, we need to see that additive constants do not affect the limit. This follows immediately from Proposition 6.5, since it tells us that

$$\lim_{r \rightarrow \infty} \frac{|B_v(I, r)|}{|B_v(I, r+1)|} = 1.$$

Secondly, we show independence of basepoint. Let $d = d(v, w)$. Then for any $r > d$, we have

$$B_w(I, r-d) \subset B_v(I, r) \subset B_w(I, r+d).$$

We thus have

$$\frac{|B_w(I, r-d)|}{|B_w(I, r)|} \leq \frac{|B_v(I, r)|}{|B_w(I, r)|} \leq \frac{|B_w(I, r+d)|}{|B_w(I, r)|}.$$

By the lemma, we have that the left- and right-hand sides approach 1 as $r \rightarrow \infty$, so this gives

$$\lim_{r \rightarrow \infty} \frac{|B_v(I, r)|}{|B_w(I, r)|} \rightarrow 1.$$

The corollary follows. □

6.3. Volume and collapsing. We will need to understand something about the size of the portion of the interval taken up by its intersection with a hyperplane. Given an interval I and \hat{h} a hyperplane in I , we say that \hat{h} is *meager* if (for any fixed choice of base vertex as the center of the balls)

$$\lim_{r \rightarrow \infty} \frac{|C(\hat{h}) \cap B_v(I, r)|}{|B_v(I, r)|} = 0.$$

Remark 6.8. Let us note that by Corollary 6.7 the property of a hyperplane being meager is independent of the base point.

We now recall the construction called *collapsing*, discussed in §2.3. Since $C(\hat{h}) \cong \hat{h} \times I$, we may form a quotient of I by collapsing the interval direction of $C(\hat{h})$ to a point (i.e., via projection onto the \hat{h} factor). The resulting complex \bar{I} is called the complex obtained by *eliminating* \hat{h} . The quotient map $I \rightarrow \bar{I}$ is called the *collapsing map* for \hat{h} . One may also eliminate a finite number of hyperplanes. It is easy to check that such a collapsing map is well defined in that it does not depend on the order in which the hyperplanes are collapsed.

We will need to understand what collapsing does to growth of balls.

Lemma 6.9. *Let $\rho: I \rightarrow J$ be a collapsing for a single hyperplane and let v be a vertex in I . Then for all $r > 0$, we have*

$$\frac{|B_{\rho(I)}(J, r)|}{|B_v(I, r)|} \geq 1/2.$$

Proof. First we observe that since the distance between vertices simply counts the number of hyperplanes which separate them, we have that ρ is distance non-increasing, so that

$$\rho(B_v(I, r)) \subset B_{\rho(v)}(J, r).$$

Now we note that the collapsing map is at most 2-to-1, so that

$$2|B_{\rho(v)}(J, r)| \geq |B_v(I, r)|.$$

It follows that

$$\frac{|B_{\rho(v)}(J, r)|}{|B_v(I, r)|} \geq 1/2. \quad \square$$

Lemma 6.10. *Let $\rho: I \rightarrow J$ denote the collapsing map eliminating a meager hyperplane \hat{h} . Then*

$$\lim_{r \rightarrow \infty} \frac{|B_{\rho(v)}(J, r)|}{|B_v(I, r)|} = 1.$$

Proof. As in the previous lemma, we have $\rho(B_v(I, r)) \subset B_{\rho(v)}(J, r)$. Also, note that ρ is 1-1 on the complement of $C(\hat{h})$. Thus, we obtain

$$\frac{|B_{\rho(v)}(J, r)|}{|B_v(I, r)|} \geq \frac{|\rho(B_v(I, r))|}{|B_v(I, r)|} \geq \frac{|\rho(B_v(I, r) - C(\hat{h}))|}{|B_v(I, r)|} = \frac{|B_v(I, r) - C(\hat{h})|}{|B_v(I, r)|}.$$

By the definition of meager, we have

$$\lim_{r \rightarrow \infty} \frac{|B_v(I, r) - C(\hat{h})|}{|B_v(I, r)|} = 1.$$

The lemma follows. □

Corollary 6.11. *Let $I = [v, \alpha]$ be an interval with no more than n non-meager hyperplanes. Then for any collapsing map of finitely many hyperplanes $\rho: I \twoheadrightarrow J$, we have*

$$\liminf_{r \rightarrow \infty} \frac{|B_{\rho(v)}(J, r)|}{|B_v(I, r)|} \geq 1/2^n.$$

6.4. Projections of hyperplanes onto each other. As above, we let $I = [v, \alpha]$ be an interval embedded in \mathbb{R}^D , with v the origin and α an non-terminating ultrafilter. As before, we write $\mathbb{R}^D = \mathbb{R}^{D-1} \times \mathbb{R}$, with the last factor being the vertical direction and the hyperplanes transverse to this direction called horizontal hyperplanes. As before, the horizontal hyperplanes are ordered from their distance from the origin and are denoted by L_n .

As discussed previously, there exists a projection of I to L_n . Restricting, this we get a projection map $\pi_n: L_1 \rightarrow L_n$. We want to get a better handle on this map.

As usual, let $\hat{\mathcal{H}}(L_n)$ denote the collection of hyperplanes of the complex L_n : these are simply the hyperplanes of I which cross L_n . Note that any hyperplane \hat{h} which crosses L_1 and does not cross L_n must separate v from L_n . This is because both \hat{h} and L_n must separate v from α . Thus, there are finitely many hyperplanes which cross L_1 and do not cross L_n . We let $\rho_n: L_1 \twoheadrightarrow Y_n$ denote the resulting collapsing map.

Now for any hyperplane \hat{h} in I crossing L_n which does not cross L_1 , we let \mathfrak{h} denote the halfspace containing L_1 . We let

$$C_n = \bigcap_{\hat{h} \cap L_1 = \emptyset} \mathfrak{h} \cap L_n.$$

The subcomplex C_n is simply the image $\pi_n(L_1)$. Note that C_n is itself a convex subcomplex of L_n whose hyperplanes consist of those hyperplanes in I which intersect L_1 . Thus Y_n and C_n have the same halfspace system, namely the one coming from all hyperplanes which cross both L_1 and L_n . We thus have a natural isomorphism $i_n: Y_n \rightarrow C_n$. We then see that $\pi_n = i_n \circ \rho_n$.

We summarize the above discussion in the following lemma.

Lemma 6.12. *For each L_n , the projection map $L_1 \rightarrow L_n$ factors through a collapsing map $\rho: L_1 \twoheadrightarrow Y_n$ and an embedding $Y_n \rightarrow L_n$.*

6.5. Non-meager hyperplanes are inessential

Theorem 6.13. *Suppose that I is an interval. Then*

- (1) *I contains only finitely many non-meager hyperplanes.*
- (2) *For every non-meager hyperplane \hat{h} in I , there exists some R such that the R -neighborhood of \hat{h} contains $I \cap X$.*

Proof. We prove (1) by induction on D , the dimension of I . For $D = 0$, the statement is trivial.

Suppose that I has infinitely many non-meager hyperplanes. Since I is finite dimensional we may choose these hyperplanes L_1, L_2, \dots so that they are disjoint. We further choose the L_i 's so that their carriers are disjoint. As in Section 6.4, we view I as being embedded in R^D and we may pass to a subsequence of the $\{L_n\}$ so that they are all transverse to some vertical direction of a factorization $\mathbb{R}^D = \mathbb{R}^{D-1} \times \mathbb{R}$. We also order $\{L_n\}$ by distance to the base vertex v and let v_n denote the projection of v onto v_n , for each n . We let $d_n = d(v, v_n)$

Since L_1 is not meager, there exists some number $\varepsilon > 0$ and a sequence $\{r_k\}$ such that

$$\frac{|B_v(I, r_k) \cap C(L_1)|}{|B_v(I, r_k)|} > \varepsilon.$$

Since we are using the 1-skeleton metric, we have

$$B_v(I, r) \cap L_1 = B_{v_n}(L_n, r - d_n)$$

for all n . Furthermore, we know that the natural map $C(L_n) \rightarrow L_n$ is 2-1 and thus between 2-1 and 1-1 in the intersection with $B(I, r_k)$, we obtain

$$\frac{|B_{v_1}(L_1, r_k - d_1)|}{|B_v(I, r_k)|} = \frac{|B_v(I, r_k) \cap L_1|}{|B_v(I, r_k)|} \geq \frac{\frac{1}{2}|B_v(I, r_k) \cap C(L_1)|}{|B_v(I, r_k)|} > \varepsilon/2. \tag{6.1}$$

By induction we know that L_1 has only finitely many non-meager hyperplanes. So let N denote the number of non-meager hyperplanes in L_1 . By Lemma 6.12, we have that the projection $L_1 \rightarrow L_n$ factors through a collapse of finitely many hyperplanes and an embedding. We may thus apply Lemma 6.11 to obtain

$$\frac{|B_{v_n}(L_n, r_k - d_1)|}{|B_{v_1}(L_1, r_k - d_1)|} \geq 1/2^N.$$

By Lemma 6.7, for r_k sufficiently large, we get

$$\frac{|B_{v_n}(L_n, r_k - d_k)|}{|B_{v_n}(L_n, r_k - d_1)|} > 1/2, \tag{6.2}$$

which, for r_k sufficiently large, implies that

$$\frac{|B_{v_n}(L_n, r_k - d_k)|}{|B_{v_1}(L_1, r_k - d_1)|} > 1/2^{N+1}.$$

We choose $M > 2^{N+1}/\varepsilon$ and choose r_k sufficiently large so that inequality (6.2) holds for all $1 \leq n \leq M$. Putting this together with inequality (6.1), we obtain for all $1 \leq n \leq M$,

$$\frac{|B_v(I, r_k) \cap C(L_n)|}{|B_v(I, r)|} > \frac{|B_{v_n}(L_n, r_k - d_k)|}{|B_v(I, r)|} > \varepsilon/2^{N+1}.$$

Summing over the first M levels, we obtain

$$\sum_{n=1}^M \frac{|B_v(I, r_k) \cap C(L_n)|}{|B_v(I, r_k)|} > M\varepsilon/2^{N+1} > 1,$$

a contradiction.

To prove (2), we proceed in the same manner. Suppose that there is a non-meager hyperplane in I which does have a neighborhood containing $I \cap X$. Then, calling that hyperplane L_1 , we would have an infinite sequence of disjoint hyperplanes L_1, L_2, \dots in I . Now the same projection arguments as above show that each L_n is non-meager, contradicting part (1). \square

Corollary 6.14. *If $E = [p, \alpha]$, where $p \in I$ and α is a non-terminating ultrafilter, then all the hyperplanes in E are meager.*

Proof. By the definition of non-terminating, for every hyperplane \hat{h} , there exists another hyperplanes \hat{h}' separating \hat{h} from α . Therefore no neighborhood of \hat{h} contains all of I . \square

6.6. Proof of Theorem 6.1. Finally, we are now ready to prove the main theorem of this section.

Proof of Theorem 6.1. Let $m = \text{med}(v, w, \alpha)$. Then we have

$$[v, \alpha] \cap [w, \alpha] = [m, \alpha].$$

Also, $[v, \alpha] \Delta [w, \alpha] = [v, \alpha] \cup [w, \alpha] - [m, \alpha]$, so that in order to show our result, it suffices to show that

$$\lim_{r \rightarrow \infty} \frac{|([v, \alpha] - [m, \alpha]) \cap B(r)|}{|[m, \alpha] \cap B(r)|} = 0.$$

Let $\hat{h}_1, \dots, \hat{h}_k$ denote the hyperplanes separating v and m . Then

$$[v, \alpha] - [m, \alpha] \subset \bigcup_{i=1}^k C(\hat{h}_i).$$

Thus, we conclude that independently of the base point defining $B(r)$

$$\lim_{r \rightarrow \infty} \frac{|([v, \alpha] - [m, \alpha]) \cap B(r)|}{|[m, \alpha] \cap B(r)|} \leq \lim_{r \rightarrow \infty} \frac{|(\bigcup_{i=1}^k C(\hat{h}_i) \cap B(r))|}{|[m, \alpha] \cap B(r)|} = 0,$$

by Corollary 6.14. \square

6.7. Property A for groups acting on finite dimensional cube complexes. A discrete metric space X has property A if there exists a sequence of families of finitely supported probability measures $A_{n,x} \in \ell^1(X)$, indexed by $x \in X$, and a sequence of constants $d_n > 0$ such that:

- (1) For every n and x , the measure $A_{n,x}$ is supported in $B_{d_n}(x)$.
- (2) For every $D > 0$, we have $\|A_{n,x} - A_{n,x'}\| \rightarrow 0$ uniformly on the set $\{(x, x') \mid d(x, x') \geq D\}$ as $n \rightarrow \infty$.

A discrete group has property A if its underlying proper metric space does (this is independent of the choice of proper metric). In this case the definition amounts to a variant of Reiter’s condition for amenability.

It was shown in [BC+] that a finite dimensional cube complex satisfies Yu’s property A . We give here an alternative proof in the case the cube complex is in addition locally finite. Indeed, let us first assume that the cube complex is irreducible non-Euclidean. In that case, $\mathcal{U}_{NT}(X)$ is non-empty, namely there is a non-terminating ultrafilter α . For any vertex $v \in X$, consider the sets $B_n(v) \cap [v, \alpha] = A_{v,n}$, namely the intersection of the interval from v to α with a ball of radius n and center v in X . Then by Theorem 6.1 for any $\varepsilon > 0$ the sets satisfy

$$\frac{|A_{v,n} \Delta A_{w,m}|}{|A_{v,n}|} < \varepsilon,$$

provided that $m \geq n \geq n(\varepsilon)$ and of course $\text{diam } A_{v,n} \leq 2n$. This sequence of sets constitutes a direct generalization of the sequence usually used to show that a tree has property A , namely the sequence consisting of intervals of length n on the unique geodesic from a vertex v to a boundary point α . The existence of such a sequence implies property A for the cube complex, by definition. Now using the product decomposition theorem of [CS] the result follows for every finite dimensional locally finite CAT(0) cube complex, recalling also that Euclidean complexes have property A since they admit a Følner sequence.

We now turn to another ingredient that plays an important role in the proof of Theorem 5.1.

7. Amenability of the boundary action

In the present section we will discuss the amenability properties of the action of discrete groups on the boundary of a cube complex. The key to our analysis is the existence of intervals connecting vertices in the complex to boundary points, with balls in the intervals satisfying the Følner property established in Theorem 6.1 and Corollary 6.7. Later on, in our discussion of the Poisson boundary, we will also use the fact that intervals embed in \mathbb{R}^n and hence have polynomial growth.

Let us first recall the following definitions of amenability of an action. Let $P(G)$ denote the space of probability measures on the countable group G taken with the

total variation norm ($\ell^1(G)$ -metric). G acts on $P(G)$ by translations, denoted by $\mu \mapsto g\mu$, and the action is continuous. We denote by $P_c(G)$ the subspace of finitely supported measures.

Definition 7.1 (Topological, Borel and universal amenability). (1) The G -action on a locally compact metric space B is called *topologically amenable* if we can find a sequence of continuous functions $\beta_n : B \rightarrow P_c(G)$ such that

$$\lim_{n \rightarrow \infty} \|\beta_n(gb) - g\beta_n(b)\| = 0$$

for every $b \in B$, and the convergence is uniform on compact subsets in B .

(2) The G -action on a standard Borel space B is called *Borel-amenable* if we can find a sequence of Borel measurable functions $\beta_n : B \rightarrow P_c(G)$ which satisfies $\lim_{n \rightarrow \infty} \|\beta_n(gb) - g\beta_n(b)\| = 0$ for every $b \in B$.

(3) The G -action on a standard Borel space B is called *universally* (or *measure-wise*) *amenable* if for every G -quasi-invariant probability measure η on B , we can find a sequence of measurable functions β_n which are defined and satisfy $\lim_{n \rightarrow \infty} \|\beta_n(gb) - g\beta_n(b)\| = 0$ for η -almost all points $b \in B$.

The space B we will consider to begin with is $\mathcal{U}_{NT}(X)$, which by Proposition 3.3 is a dense G_δ of the boundary $B(X)$.

Theorem 7.2. *Let X be an irreducible, non-Euclidean CAT(0) cube complex. Then for any discrete subgroup $G \subset \text{Aut}(X)$ acting properly, the action on $\mathcal{U}_{NT}(X)$ is Borel-amenable, and hence also a universally (or measure-wise) amenable action.*

Proof. We will first construct a sequence of Borel maps $\tilde{\beta}_n : \mathcal{U}_{NT}(X) \rightarrow P(V)$, where $V = V(X)$ is the set of vertices of the complex X , taking values in finitely supported probability measures on the vertex set $V(X)$.

To define $\tilde{\beta}_n(b)$ when b is a non-terminating ultrafilter, fix a reference vertex $o \in V(X)$. Given any point $b \in \mathcal{U}_{NT}(X)$, draw the interval from o to b , denoted by $[o, b)$. Define $\tilde{\beta}_n : \mathcal{U}_{NT}(X) \rightarrow P(V)$, where $\tilde{\beta}_n(b)$ is the probability measure uniformly distributed on the finite set of vertices obtained as the intersection of a ball of radius n with center o and the interval $[o, b)$ from o to b .

Now note that $\tilde{\beta}_n(gb)$ is the measure on the complex uniformly distributed on $B(n, o) \cap [o, gb)$. On the other hand, $g\tilde{\beta}_n(b)$ is the measure uniformly distributed on $B(n, go) \cap [go, gb)$.

In view of Theorem 6.1, given $\varepsilon > 0$, there exists an n sufficiently large such that outside the ball $B(n, o)$, the symmetric difference between the intervals $[o, gb)$ and $[go, gb)$ has size bounded by ε times the sum of their sizes.

It follows the difference between the measures namely $\|\tilde{\beta}_n(gb) - g\tilde{\beta}_n(b)\|$ (in $\ell^1(V(X))$ -norm), does indeed converge to zero for any given g in G .

Let us also note that the convergence is in fact uniform on $\mathcal{U}_{NT}(X)$ by Lemma 6.6, namely for a fixed $\varepsilon > 0$, n can be chosen independent of $b \in \mathcal{U}_{NT}(X)$.

We now pass to a discrete subgroup $G \subset \text{Aut}(X)$ acting properly by restricting each measure $\tilde{\beta}_n$ to the various G -orbits in the complex, renormalizing and viewing the resulting measure as an element of $P(G)$. More explicitly, let T be a transversal to the G -orbits in $V(X)$, and for $v \in T$ let G_v be the stability group of v in G , which is finite since the action of G is proper. Let us view the probability measure $\tilde{\beta}_n(b)$ as function on the vertices in $V(X)$, and define

$$\beta_n(b, g) = \sum_{v \in T} \frac{\tilde{\beta}_n(b, gv)}{|G_v|}.$$

Then the $\ell^1(G)$ -norm of $\beta_n(b, \cdot)$ is

$$\begin{aligned} \|\beta_n(b, \cdot)\|_{\ell^1(G)} &= \sum_{g \in G} \beta_n(g, b) \\ &= \sum_{g \in G, v \in T} \frac{\tilde{\beta}_n(b, gv)}{|G_v|} \\ &= \sum_{v \in T} \sum_{w \in G \cdot v} \tilde{\beta}_n(b, w) \sum_{g \in G: gv=w} \frac{1}{|G_v|} \\ &= \sum_{v \in T} \sum_{w \in G \cdot v} \tilde{\beta}_n(b, w) \\ &= \|\tilde{\beta}_n(b, \cdot)\|_{\ell^1(V)}. \end{aligned}$$

Clearly a similar computation shows that $\|g\beta_n(b) - \beta_n(gb)\|_{\ell^1(G)} \rightarrow 0$ and thus $\mathcal{U}_{\text{NT}}(X)$ is a Borel-amenable action of G . □

A simple consequence of Theorem 7.2 is the following.

Corollary 7.3. *The stability group of a point in $\mathcal{U}_{\text{NT}}(X)$, namely of a non-terminating ultrafilter in a (finite dimensional) CAT(0) cube complex is an amenable group.*

Proof. Let $g \in S = St_G(b)$ be a group element which stabilized $b \in \mathcal{U}_{\text{NT}}(X)$. Then the sequence of finitely supported probability measures $\beta_n(b) \in P_c(G)$ is clearly asymptotically invariant under left translation by g . The sequence defines an mean on $\ell^\infty(G)$ in the usual way, and this mean is invariant under the subgroup S . Extending a bounded function on S to a bounded function on G by transferring it to the other cosets in the obvious way, we get an S -invariant mean on $\ell^\infty(S)$, so S is amenable. □

We note that Corollary 7.3 is in fact a consequence of more general results, considered in various formulations in [CN], [Ca07] and [BC+]. Of those, we quote the result of [BC+] most pertinent to us.

Theorem 7.4. *Let X be a (finite dimensional) CAT(0) cube complex. The stability group of every ultrafilter, namely every point in the Roller boundary and thus in particular every point in $B(X)$, is an amenable group.*

Remark 7.5. (1) In [Ca07] a structure theorem is proved for amenable closed subgroups of a group that acts properly, co-compactly and discontinuously on Hadamard spaces, and in particular on finite dimensional locally finite CAT(0) cube complexes. It is shown that an amenable group virtually admits a homomorphism into \mathbb{R}^d , with the kernel being a topologically locally finite group.

(2) It is not clear whether the action of G on $B(X)$ is *topologically* amenable. For that, one needs to define a sequence of functions $\beta_n : B(X) \rightarrow P_c(G)$, defined on all of $B(X)$, asymptotically invariant and *continuous*. The remark following Theorem 6.1 casts doubt as to whether the functions β_n we have defined on $\mathcal{U}_{NT}(X)$ can be extended *continuously* to $B(X)$. However, this observation is not conclusive, since there may in principle be other sequences that could demonstrate topological amenability.

8. Maximality of the boundary

In the present section we will show that the boundary $B(X)$ with its unique stationary measure ν is a compact metric model of the Poisson boundary $\mathcal{B}(G, \mu)$. We will use in our analysis an important criterion for boundary maximality developed by V. Kaimanovich, called the strip criterion. This criterion is applicable in the context of cube complexes, since intervals provide us with a natural notion of strips in the complex. In fact, any two distinct ultrafilters $\alpha \neq \beta$ determine a unique interval of ultrafilters between them, defined by

$$[\alpha, \beta] = \{\gamma \mid m(\alpha, \beta, \gamma) = \gamma\}.$$

By Lemma 6.2, we have

$$[\alpha, \beta] = \bigcap_{\mathfrak{h} \in \alpha \cap \beta} U_{\mathfrak{h}}. \tag{8.1}$$

The map $B(X) \times B(X) \rightarrow 2^{\mathcal{U}(X)}$ defined by $(\alpha, \beta) \mapsto [\alpha, \beta]$ is of course $\text{Aut}(X)$ -equivariant.

8.1. Antipodal pairs of ultrafilters. It may happen, however, that the interval consists of non-principal ultrafilters, i.e., the strip between two non-principal ultrafilters α and β may lie itself “at infinity”. This arises even in simple and natural examples, such as \mathbb{Z}^2 and $T_3 \times T_3$.

Thus not all pairs of boundary points are alike, and we must find the right notion an “antipodal pair”. A pair $(\alpha, \beta) \in B(X) \times B(X)$ is called *antipodal* if $\mathcal{S}(\alpha, \beta) \equiv [\alpha \cap \beta] \cap X \neq \emptyset$. The set $\mathcal{S}(\alpha, \beta)$ is called the *strip* between α and β . We let

$\mathcal{G} \subset B(X) \times B(X)$ denote the collection of antipodal pairs. We then have the following.

Proposition 8.1. *Let X be a strictly non-Euclidean, cocompact CAT(0) cube complex. Then the set \mathcal{G} is generic: it is a non-empty open invariant subset in $B(X) \times B(X)$.*

Proof. First, we need to show that \mathcal{G} is non-empty. First observe that it suffices to show this for the case that X is irreducible, because the antipodal pairs in the products appear as products of antipodal pairs in the factors.

So now assume that X is irreducible and suppose that $\mathcal{G} = \emptyset$. Let α, β be a pair of distinct non-terminating ultrafilters. By assumption we have $S(\alpha, \beta) \cap X = \emptyset$ for every such pair.

Since α and β are non-terminating ultrafilters, there exist infinitely many hyperplanes separating them. In particular there exists a pair of disjoint hyperplanes separating them. We thus see that given any two such α and β , we have that

- (1) there exists a collection of intersecting hyperplanes $\hat{h}_1, \dots, \hat{h}_n$ such that $\alpha \in \mathfrak{h}_i$ and $\beta \in \mathfrak{h}_i^*$,
- (2) there exists \hat{f} such that $\hat{f} \in \mathfrak{h}_i^*$,
- (3) $\alpha \in \mathfrak{f}$ and $\beta \in \mathfrak{f}^*$.

The above remark tells us that (1)–(3) holds for $n = 1$. The number n is bounded by the dimension of X , so that we may choose α and β such that n is maximal.

Note that by (8.1), we have

$$S(\alpha, \beta) = \bigcap_{\mathfrak{h} \in \alpha \cap \beta} \mathfrak{h}.$$

Note that if $\alpha \cap \beta = \emptyset$, then all hyperplanes of X separate α and β . In this case clearly every vertex of X is in $S(\alpha, \beta)$ and we are done.

So suppose that the above intersection is indeed an intersection of a non-empty collection of halfspaces. By Corollary 2.3 if $\mathcal{K} = \{\mathfrak{h} \mid \mathfrak{h} \in \alpha \cap \beta\}$ satisfies DCC, then we would have $S(\alpha, \beta) \neq \emptyset$. So by our assumption, we have that \mathcal{K} does not satisfy DCC. Let $\mathfrak{m}_1 \supset \mathfrak{m}_2 \supset \dots$ denote a non-terminating sequence of halfspaces in \mathcal{K} .

We claim that there exists some l such that \hat{m}_l intersects both \hat{f} and \hat{h}_i for all i .

Since each $\hat{m}_n \in \alpha \cap \beta$ and \hat{f} separates α and β , it follows that if $\hat{f} \cap \hat{m}_l = \emptyset$ then $\hat{f} \subset \mathfrak{m}_l$. But since there is a finite distance between \hat{m}_1 and \hat{f} , we must have $\hat{m}_l \cap \hat{f} \neq \emptyset$ for some l . Similarly, we can choose l large enough so that both $\hat{m}_l \cap \hat{f} \neq \emptyset$ and $\hat{m}_l \cap \hat{h}_i \neq \emptyset$ for all i .

We can further choose l such that the half-hyperplane $\hat{f} \cap \mathfrak{m}_l^*$ is deep. We then consider the sector $\mathfrak{m}_l^* \cap \mathfrak{f}^*$. By Lemma 5.5, there exists some hyperplane $\hat{f}_1 \subset$

$\pi_l^* \cap \mathfrak{F}^*$. We let \mathfrak{F}_1 denote the halfspace of $\hat{\mathfrak{F}}_1$ contained in α . By Remark 3.2, there exists a non-terminating ultrafilter $\gamma \in \mathfrak{F}_1^*$.

Now observe that for the pair α, γ the collection of hyperplanes $\hat{h}_1, \dots, \hat{h}_n, \hat{m}_l$ and the hyperplane $\hat{\mathfrak{F}}_1$ all satisfy properties (1)–(3) above. But this contradicts the maximality of n as chosen.

This completes the proof that $\mathcal{S} \neq \emptyset$. To prove that it is open we will use that the median map $m : \mathcal{U}(X) \times \mathcal{U}(X) \times \mathcal{U}(X) \rightarrow \mathcal{U}(X)$ is continuous in the Tychonoff topology. Suppose that (α, β) is a antipodal pair. Then there exists $v \in X^{(0)}$ such that $\text{med}(\alpha, \beta, v) = v$. Since vertex singletons in $X^{(0)}$ are open, we may find open sets U of α and V of β such that for any $\alpha' \in U$ and $\beta' \in V$, we have $\text{med}(\alpha', \beta', v) = v$. Thus $S(\alpha', \beta') \neq \emptyset$, as required. Finally, \mathcal{S} is clearly an invariant set under the product action of $\text{Aut}(X)$ on $B(X) \times B(X)$. \square

Remark 8.2. Since \mathcal{S} is $\text{Aut}(X)$ -invariant, any σ -finite measure η on $B(X) \times B(X)$ which is quasi-invariant and ergodic under a discrete subgroup G and charges every open set, must assign \mathcal{S} full measure.

8.2. Boundary maximality via the strip criterion. We will now use our construction of strips in the complex, for every antipodal pair of points in $B(X)$, in order to show that $B(X)$ is a maximal boundary, namely realizes the Poisson boundary.

First, recall that a probability measure μ on (a countable group) G is called of finite logarithmic moment if there exists a distance function $|g|$ which is quasi isometric to a word metric, and satisfies $\sum_{g \in G} \log |g| \mu(g) < \infty$. This definition is independent of the word metric chosen.

Recall also that the Avez entropy of a probability measure μ on (a countable group) G is defined by

$$H(\mu) = \lim_{n \rightarrow \infty} \frac{-1}{n} \sum_{g \in G} \mu^{*n}(g) \log \mu^{*n}(g).$$

This quantity is equal to the μ -entropy of the stationary measure ν on the Poisson boundary, see below.

The measure $\hat{\mu}$ is defined by $\hat{\mu}(g) = \mu(g^{-1})$, and it has finite logarithmic moment and finite entropy if μ does.

Now recall the following criterion for boundary maximality, due to V. Kaimanovich [K], Thm. 6.5.

Theorem 8.3 (Strip criterion, [K]). *Let μ be a probability measure of finite first logarithmic moment and finite entropy on a group G . Assume that (B, ν) is a (G, μ) -boundary, and $(\hat{B}, \hat{\nu})$ is a $(G, \hat{\mu})$ -boundary. Assume there exists a measurable G -equivariant map defined on $(B \times \hat{B}, \nu \times \hat{\nu})$, denoted by $(\alpha, \beta) \mapsto \mathcal{S}(\alpha, \beta) \subset G$ (viewed as assigning non-empty strips to pairs of boundary points). If for $\nu \times \hat{\nu}$ -almost all pairs (α, β) the strip $\mathcal{S}(\alpha, \beta)$ has polynomial growth (with respect to the distance*

function above), then both (B, ν) and $(\hat{B}, \hat{\nu})$ are maximal boundaries, namely they realize the Poisson boundaries of (G, μ) and $(G, \hat{\mu})$.

We will use this criterion to prove:

Theorem 8.4. *Let X be an irreducible non-Euclidean CAT(0) cube complex. Let G be a discrete subgroup of $\text{Aut}(X)$ acting properly and co-compactly on the complex. Let μ be a probability measure on G of finite logarithmic moment and finite entropy, whose support generates G as a semigroup. Denote the unique μ -stationary measure on $B(X)$ by ν . Then $(B(X), \nu)$ is a compact metric model of the Poisson boundary of (G, μ) .*

Proof. Any stationary measure is G -quasi-invariant, and hence its support is a closed non-empty G -invariant set. In Theorem 5.8 we have shown that the G -action on $B(X)$ is minimal, and hence the support of a quasi-invariant measure coincides with $B(X)$. In particular, the measure of every non-empty open set is strictly positive.

Consider now the $\text{Aut}(X)$ -equivariant map $(\alpha, \beta) \mapsto \mathcal{S}(\alpha, \beta)$ from $B(X) \times B(X)$ to strips in X . We first claim that for a set of pairs $(\alpha, \beta) \in B(X) \times B(X)$ of $\nu \times \hat{\nu}$ -measure 1, the strips are actually subsets of the complex, namely that the set of \mathcal{S} of antipodal pairs has full $\nu \times \hat{\nu}$ -measure.

The action of the product of the Poisson boundaries associated with μ and $\hat{\mu}$ is ergodic, in general (see [K]), and thus so is the action on any of its G -factor spaces. By Theorem 5.1 the space $B(X)$ is a (G, μ) -boundary and thus $(B(X), \nu) \times (B(X), \hat{\nu})$ is a factor of the product of the Poisson boundaries, so that $\nu \times \hat{\nu}$ on $B(X) \times B(X)$ is ergodic.

The set \mathcal{S} of antipodal pairs is invariant under the G -action on $B(X) \times B(X)$, and clearly has positive $\mu \times \hat{\mu}$ -measure, since it is non-empty and open, and the product measure charges every non-empty open set. By ergodicity, \mathcal{S} has full $\nu \times \hat{\nu}$ -measure.

We now pass to strips in the group G itself rather than in the complex. To that end, note that by co-compactness of G , it has only finitely many orbits of vertices in X , and let us call the different orbits “types” denoted by t_1, \dots, t_r , and choose a vertex in each orbit v_1, \dots, v_r . For each antipodal pair (α, β) , the strip $\mathcal{S}(\alpha, \beta) \subset X$ decomposes to a disjoint union of vertices belonging to these types. Clearly there exist at least one type (say t_1) such that for a positive measure subset of antipodal pairs, the associated strip contains vertices of type t_1 . But the latter set of antipodal pairs is clearly G -invariant, and so necessarily has full measure. We now define the map $(\alpha, \beta) \mapsto \mathcal{S}'(\alpha, \beta) \subset G$ to strips in the group as follows. $\mathcal{S}'(\alpha, \beta)$ is defined to be the union of all the cosets $gSt_G(v_1)$, as g ranges over all group elements with the property that gv_1 is in $\mathcal{S}(\alpha, \beta)$, namely gv_1 is a vertex of type t_1 in that strip. The map is clearly equivariant, and the stability group $St_G(v_1)$ is finite. Now $\mathcal{S}(\alpha, \beta) \subset X$ has polynomial growth since it is contained in an interval and thus embeds in \mathbb{R}^d [BC+]. The polynomial growth is with respect to the ℓ^1 -metric on vertices in X , and hence $\mathcal{S}'(\alpha, \beta)$ has polynomial growth with respect to a distance function quasi-isometric to a word metric on G . The desired result now follows from Theorem 8.3. \square

9. Entropy and the Poisson boundary

Theorem 8.4 establishes, in particular, that for measures μ on G with finite logarithmic moment, $B(X)$ gives rise to a compact metric uniquely-stationary model of the Poisson boundary. In the present section we would like to demonstrate that the existence of such a model for the Poisson boundary is a significant fact, which has important consequence for the boundary theory of a countable group G .

First let us recall the definition of μ -entropy of a standard Borel (G, μ) -space (see [NZ1] for a detailed discussion).

Definition 9.1. The μ -entropy of a (G, μ) -space (B, ν) is defined for a countable group G by

$$h_\mu(B, \nu) = \sum_{g \in G} \mu(g) \int_B -\log \frac{dg^{-1}\nu}{d\nu}(b) d\nu(b)$$

As noted in §8, the Avez entropy $H(\mu)$ coincides with the μ -entropy of the Poisson boundary. Furthermore, this value constitutes the largest value that the μ -entropy can assume, ranging over all (G, μ) -spaces (B, ν) . Recall that the Poisson boundary is the unique (up to ν -null sets) maximal standard Borel (G, μ) -space which is a μ -boundary. Here maximality means that any other (G, μ) -space is a factor of the Poisson boundary, with the factor map uniquely determined, up to ν -null sets. Recall also that the action on the Poisson boundary is amenable in the sense of Zimmer (see [Z2] for a detailed discussion).

Theorem 9.2 (Characterization of the Poisson boundary). *Let G be a countable group, μ a probability measure whose support generates G as a semigroup, and assume that there exists a compact metric G -space B which admits a unique μ -stationary measure ν such that (B, ν) realizes the Poisson boundary of (G, μ) . Then every amenable (G, μ) -space is a measure-preserving extension of the Poisson boundary of (G, μ) , and thus has maximal μ -entropy. In particular, an amenable (G, μ) -boundary space is (essentially) isomorphic to the Poisson boundary of (G, μ) . Thus the Poisson boundary is characterized as*

- (1) *the unique minimal amenable (G, μ) -space (i.e., it is a factor of every other amenable (G, μ) -space),*
- (2) *the unique maximal (G, μ) -boundary space (i.e., every other (G, μ) -boundary space is a factor of (B, ν)),*
- (3) *the unique (G, μ) -boundary space which is amenable,*
- (4) *the unique (G, μ) -boundary space of maximal μ -entropy.*

Proof. (1) To prove the first characterization, let us begin by recalling that if (Y, ν) is any amenable action of G with η a quasi-invariant probability measure, and B any compact metric G -space, then there exists a G -equivariant map $\phi: Y \rightarrow P(B)$,

where $P(B)$ is the space of probability measures on B and ϕ is defined η -almost everywhere [Z3], 4.3.9. If $\eta = \nu_Y$ is μ -stationary, its image under ϕ , denoted by $\phi_*(\nu_Y)$, is a μ -stationary measure on $P(B)$. Now if (B, ν) is a μ -proximal, then the measure $\phi_*(\nu_Y)$ must take values in δ -measures on B almost surely [M], Ch. VI, Cor. 2.10. Therefore ϕ arises from a measurable G -equivariant factor map $\phi'(Y, \nu_Y) \rightarrow B$. Clearly if the stationary measure ν on B is unique, then $\phi'_*(\nu_Y)$ and ν must coincide, so that (B, ν) is indeed a factor of (Y, ν_Y) and thus (B, ν) is indeed a minimal amenable space.

(2) The second characterization is well known to be valid for any Poisson boundary.

(3) The third characterization follows from the fact that an amenable (G, μ) -boundary space is both a cover and a factor of the Poisson boundary (B, ν) , and hence admits an equivariant endomorphism. But for a (G, μ) -boundary space such an endomorphism is necessarily the identity [M], Ch. VI, Cor. 2.10. Thus the space in question is isomorphic to (B, ν) .

(4) Any (G, μ) -boundary space is a factor of the Poisson boundary (B, ν) , and its μ -entropy is bounded by $h_\mu(B, \nu)$. Every proper factor of (B, ν) has strictly smaller μ -entropy, since otherwise the Poisson boundary would be a measure-preserving extension of this factor space (see [NZ1]). This is not possible, because every bounded function on (B, ν) is determined uniquely by its harmonic transform, so that (G, μ) -boundaries do not admit relatively-measure-preserving factors, as these produce distinct functions with the same harmonic transform. \square

References

- [BC+] J. Brodzki, S. J. Campbell, E. Guentner, G. A. Niblo, and N. J. Wright, Property A and CAT(0) cube complexes. *J. Funct. Anal.* **256** (2009), 1408–1431. [Zbl 1233.20036](#) [MR 2490224](#)
- [BC] J. Behrstock and R. Charney, Divergence and quasimorphisms of right-angled Artin groups. *Math. Ann.* **352** (2012), 339–356. [Zbl 1251.20036](#) [MR 2874959](#)
- [CN] S. Campbell and G. A. Niblo, Hilbert space compression and exactness of discrete groups. *J. Funct. Anal.* **222** (2005), 292–305. [Zbl 1121.20032](#) [MR 2132393](#)
- [Ca07] P.-E. Caprace, Amenable groups and Hadamard spaces with a totally disconnected isometry group. *Comment. Math. Helv.* **84** (2009), 437–455. [Zbl 1233.20037](#) [MR 2495801](#)
- [CM] P.-E. Caprace and N. Monod, Isometry groups of non-positively curved spaces: discrete subgroups. *J. Topol.* **2** (2009), 701–746. [Zbl 1187.53037](#) [MR 2574741](#)
- [CS] P.-E. Caprace and M. Sageev, Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.* **21** (2011), 851–891. [Zbl 1266.20054](#) [MR 2827012](#)
- [ChN] I. Chatterji and G. Niblo, From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.* **15** (2005), 875–885. [Zbl 1107.20027](#) [MR 2197811](#)

- [CR] I. Chatterji and K. Ruane, Some geometric groups with rapid decay. *Geom. Funct. Anal.* **15** (2005), 311–339. [Zbl 1134.22005](#) [MR 2153902](#)
- [F1] H. Furstenberg, A Poisson formula for semi-simple Lie groups. *Ann. of Math. (2)* **77** (1963), 335–386. [Zbl 0192.12704](#) [MR 0146298](#)
- [GuNi] E. Guentner and G. A. Niblo, Complexes and exactness of certain Artin groups. *Algebr. Geom. Topol.* **11** (2011), 1471–1495. [Zbl 1245.20040](#) [MR 2821432](#)
- [GJT] Y. Guivarc’h, L. Ji, and J. C. Taylor, *Compactifications of symmetric spaces*. Progr. Math. 156, Birkhäuser, Boston 1998. [Zbl 1053.31006](#) [MR 1633171](#)
- [GLP] Y. Guivarc’h and É. Le Page, Simplicité de spectres de Lyapounov et propriété d’isolation spectrale pour une famille d’opérateurs de transfert sur l’espace projectif. In *Random walks and geometry*, Walter de Gruyter, Berlin 2004, 181–259. [Zbl 1069.60005](#) [MR 2087783](#)
- [Gu] D. P. Guralnik, Coarse decompositions for boundaries of CAT(0) groups. Ph.D. thesis, Mathematics Department, Technion - Israel Institute of Technology, Haifa 2005. [arXiv:math/0611006](#)
- [HR] N. Higson and J. Roe, Amenable group actions and the Novikov conjecture. *J. Reine Angew. Math.* **519** (2000), 143–153. [Zbl 0964.55015](#) [MR 1739727](#)
- [K] V. A. Kaimanovich, The Poisson formula for groups with hyperbolic properties. *Ann. of Math. (2)* **152** (2000), 659–692. [Zbl 0984.60088](#) [MR 1815698](#)
- [Ka] A. Karlsson, On the dynamics of isometries. *Geom. Topol.* **9** (2005), 2359–2394. [Zbl 1120.53026](#) [MR 2209375](#)
- [KM] A. Karlsson and G. A. Margulis, A multiplicative ergodic theorem and nonpositively curved spaces. *Comm. Math. Phys.* **208** (1999), 107–123. [Zbl 0979.37006](#) [MR 1729880](#)
- [M] G. A. Margulis, *Discrete subgroups of semisimple Lie groups*. *Ergeb. Math. Grenzgeb.* (3) 17, Springer-Verlag, Berlin 1991. [Zbl 0732.22008](#) [MR 1090825](#)
- [NZ1] A. Nevo and R. J. Zimmer, Rigidity of Furstenberg entropy for semisimple Lie group actions. *Ann. Sci. École Norm. Sup. (4)* **33** (2000), 321–343. [Zbl 0956.22005](#) [MR 1775184](#)
- [NZ2] A. Nevo and R. J. Zimmer, A structure theorem for actions of semisimple Lie groups. *Ann. of Math. (2)* **156** (2002), 565–594. [Zbl 1012.22038](#) [MR 1933077](#)
- [NR] G. A. Niblo and M. A. Roller, Groups acting on cubes and Kazhdan’s property (T). *Proc. Amer. Math. Soc.* **126** (1998), 693–699. [Zbl 0906.20024](#) [MR 1459140](#)
- [N] B. Nica, Cubulating spaces with walls. *Algebr. Geom. Topol.* **4** (2004), 297–309. [Zbl 1131.20030](#) [MR 2059193](#)
- [Ro] M. Roller, Poc sets, median algebras and group actions. Southampton Preprint Archive, 1998.
- [S1] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3) **71** (1995), 585–617. [Zbl 0861.20041](#) [MR 1347406](#)
- [SW] M. Sageev and D. T. Wise, The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.* **37** (2005), 706–710. [Zbl 1081.20051](#) [MR 2164832](#)

- [Wo] W. Woess, Boundaries of random walks on graphs and groups with infinitely many ends. *Israel J. Math.* **68** (1989), 271–301. [Zbl 0723.60009](#) [MR 1039474](#)
- [Z1] R. J. Zimmer, Induced and amenable ergodic actions of Lie groups. *Ann. Sci. École Norm. Sup.* (4) **11** (1978), 407–428. [Zbl 0401.22009](#) [MR 521638](#)
- [Z2] R. J. Zimmer, Amenable ergodic group actions and an application to Poisson boundaries of random walks. *J. Funct. Anal.* **27** (1978), 350–372. [Zbl 0391.28011](#) [MR 0473096](#)
- [Z3] R. J. Zimmer, *Ergodic theory and semisimple groups*. Monogr. Math. 81, Birkhäuser Verlag, Basel 1984. [Zbl 0571.58015](#) [MR 776417](#)

Received May 10, 2011; revised December 4, 2011

A. Nevo, Mathematics Department, Technion - Israel Institute of Technology Haifa, 32000 Israel

E-mail: anevo@tx.technion.ac.il

M. Sageev, Mathematics Department, Technion - Israel Institute of Technology Haifa, 32000 Israel

E-mail: sageevm@tx.technion.ac.il