

The median class and superrigidity of actions on CAT(0) cube complexes

Indira Chatterji, Talia Fernós and Alessandra Iozzi

With Appendix B by Pierre-Emmanuel Caprace

ABSTRACT

We define a bounded cohomology class, called the *median class*, in the second bounded cohomology, with appropriate coefficients, of the automorphism group of a finite-dimensional CAT(0) cube complex X . The median class of X behaves naturally with respect to taking products and appropriate subcomplexes and defines in turn the *median class of an action* by automorphisms of X . We show that the median class of a non-elementary action by automorphisms does not vanish and we show to what extent it does vanish if the action is elementary. We obtain as a corollary a superrigidity result and show, for example, that any irreducible lattice in the product of at least two locally compact connected groups acts on a finite-dimensional CAT(0) cube complex X with a finite orbit in the Roller compactification of X . In the case of a product of Lie groups, the appendix by Caprace allows us to deduce that the fixed point is in fact inside the complex X . In the course of the proof, we construct a Γ -equivariant measurable map from a Poisson boundary of Γ with values in the non-terminating ultrafilters on the Roller boundary of X .

1. Introduction

The goal of this paper is to define a cohomological invariant of some non-positively curved metric spaces X for a non-elementary action of a group $\Gamma \rightarrow \text{Aut}(X)$ and to use this invariant to establish rigidity phenomena.

The paradigm is that bounded cohomology with non-trivial coefficients is the appropriate framework to study negative curvature. The first instance of this fact is the Gromov–Sela cocycle on the real hyperbolic n -space X (in fact, on any simply connected space with pinched negative curvature) with values into the L^2 differential one-forms on X (see [27, 7.E₁, 49]).

The same philosophy has been promoted by Monod [39], Monod–Shalom [40, 41] and Mineyev–Monod–Shalom [37]. They prove that a non-elementary isometric action on a negatively curved space (belonging to a very rich class) yields the non-vanishing of second bounded cohomology with appropriately defined coefficients of a geometric nature. Such negatively curved spaces include proper CAT(−1) spaces, Gromov hyperbolic graphs of bounded valency, Gromov hyperbolic proper cocompact geodesic metric spaces or simplicial trees.

On the other hand, if G is a simple Lie group with rank at least 2 and \mathcal{H} is any unitary representation with no invariant vectors, then $H_{\text{cb}}^2(G, \mathcal{H}) = 0$, [16, 17], thus showing that in non-positive curvature the situation cannot be expected to be completely analogous.

Received 5 January 2013; revised 7 August 2015.

2010 *Mathematics Subject Classification* 20XX, 20J06, 22FXX, 22F50, 57XX.

Indira Chatterji was partially supported by ANR QuantiT JC08-318197; Alessandra Iozzi was partially supported by the Swiss National Science Foundation projects 2000021-127016/2 and 200020-144373 and by the Simons Foundation. Pierre-Emmanuel Caprace is an F.R.S.-FNRS Research Associate, supported in part by the European Research Council (grant #278469).

In this paper, we move away from the negative curvature case and look at actions on CAT(0) cube complexes.

CAT(0) cube complexes are simply connected combinatorial objects introduced by Gromov [26]. They have been used in several important contexts, such as Moussong’s characterization of word hyperbolic Coxeter groups in terms of their natural presentation [42]. A prominent use of CAT(0) cube complexes was made by Sageev in his thesis [47]: generalizing Stallings’s theorem on the equivalence between splittings of groups and actions on trees, [48, 50, 51], he proved an equivalence between the existence of an action of a group Γ on a CAT(0) cube complex and the existence of a subgroup $\Lambda < \Gamma$ such that the pair (Γ, Λ) has more than one end. More recently, Agol’s proof of the last standing conjecture in 3-manifolds, the virtual Haken conjecture, uses (special) cube complexes in a crucial way, thus indisputably asserting their relevance in the mathematical scenery.

The first example of a CAT(0) cube complex X is a simplicial tree; the midpoint of a vertex is the analogue of a *hyperplane* for a general CAT(0) cube complex. Hyperplanes separate X into two connected components, called *halfspaces*, the collection of which is denoted by $\mathfrak{H}(X)$. If the vertex set of X is locally countable, then $\mathfrak{H}(X)$ is countable as well.

A CAT(0) cube complex is in particular a *median space*; that is, given any three vertices, there is a unique vertex, the *median*, that is on the combinatorial geodesics joining any two of the three points. For $n \geq 2$, let $\mathfrak{H}(X)^n$ denote the set of n -tuples of halfspaces in X . If $1 \leq p < \infty$, then we define a 1-parameter family of $\text{Aut}(X)$ -invariant cocycles

$$c_{(n,R)} : X \times X \times X \longrightarrow \ell^p(\mathfrak{H}(X)^n)$$

as the sum of the characteristic functions of some appropriate finite subsets of nested halfspaces (called *über-parallel*; see Definitions 2.12 and 2.15) ‘around’ the median of three points and at distance less than R . Incidentally, this is not a distance but just a pseudo-distance on the set of hyperplanes, and will be discussed more in Subsection 2.G. Choosing a basepoint $v_0 \in X$ and evaluating c on an $\text{Aut}(X)$ -orbit, we get what we call a *median cocycle* on $\text{Aut}(X) \times \text{Aut}(X) \times \text{Aut}(X)$. We then prove that, for every $n \geq 2$ and $R \geq 0$, the cocycle so defined is bounded and hence defines a bounded cohomology class $\mathfrak{m}_{(n,R)}(X)$ in degree 2, which we call a median class of X . Note that for any $n \geq 2$ there is a median class, but in the following we will not necessarily make a distinction of the various median classes for different n . (See (3.5), Proposition 3.4 and Lemma 3.11 for the precise definition and the proof of the above statements.)

If $\rho : \Gamma \rightarrow \text{Aut}(X)$ is an action of a group Γ by automorphisms on X , then the *median class of the Γ -action* is the pullback

$$\rho^*(\mathfrak{m}_{(n,R)}(X)) \in H_b^2(\Gamma, \ell^p(\mathfrak{H}(X)^n)).$$

THEOREM 1.1. *Let X be a finite-dimensional CAT(0) cube complex with a Γ -action. If the Γ -action is non-elementary, then there is an $R_\Gamma \geq 1$ so that the median class of the Γ -action $\rho^*(\mathfrak{m}_{(n,R)}(X))$ does not vanish for all $n \geq 2$ and all $R \geq R_\Gamma$.*

We call an action $\Gamma \rightarrow \text{Aut}(X)$ *non-elementary* if there is no finite orbit in $X \sqcup \partial_\triangleleft X$, where $\partial_\triangleleft X$ denotes the visual boundary of X as a CAT(0) space. We note that, owing to [18], the visual boundary of a finite-dimensional CAT(0) cube complex is both well defined, and indeed well-behaved.

Let us say a word about what it means for a Γ -action to be non-elementary in the context of CAT(0) cube complexes. First of all, the assumption implies, in particular that, by passing to a subgroup of finite index, there are no Γ -fixed points in $\partial_\triangleleft X$. Under this hypothesis, using the work of Caprace–Sageev [20, Proposition 3.5], one can pass to a non-empty convex subset of X , called the Γ -essential core (see Subsection 2.E), which will have rather nice dynamic

properties. Furthermore, the exclusion of a finite orbit on $\partial_{\triangleleft} X$ excludes the existence of a Euclidean factor in the essential core (see Corollary 2.35).

A key object in this paper is the *Roller boundary* ∂X of a CAT(0) cube complex, defined in Subsection 2.A. It arises naturally from considering the hyperplane (and hence halfspace) structure of X . The vertex set of X , together with its Roller boundary, can be thought of as a closed subset of a Bernoulli space (with $\mathfrak{H}(X)$ as the indexing set) and is hence compact and totally disconnected. Although, in the case of a tree, the Roller boundary and the visual boundary coincide, we remark that in general, there is no natural map between them. (There is a map from the CAT(0) boundary to a quotient of the Roller boundary [28], but we will not use it in this paper.) In Proposition 2.26, we prove nevertheless a result relating, to the extent to which it is possible, finite orbits in the Roller boundary to finite orbits in the CAT(0) boundary. The dichotomy that we obtain is analogous to the one in the case of a group Γ acting on a symmetric space \mathcal{X} of non-compact type; in this case, if Γ fixes a point at infinity in the CAT(0) boundary of \mathcal{X} , then the image of Γ is contained in a parabolic subgroup, via which it acts on the symmetric space of non-compact type associated to the semisimple component of the parabolic. The latter action may very well be non-elementary.

The dichotomy in Proposition 2.26 leads to the following converse of Theorem 1.1.

THEOREM 1.2. *Let X be a finite-dimensional CAT(0) cube complex with an elementary Γ -action. Then:*

- (1) *either there is a finite orbit in the Roller compactification $\overline{X} = X \cup \partial X$ of X and hence the median class $\rho^*_{(\mathfrak{m}_{(n,R)}(X))}$ of the Γ -action on X vanishes for all $n \geq 2$ and all $R \geq 1$;*
- (2) *or there exists a finite index subgroup $\Gamma' < \Gamma$ and a Γ' -invariant subcomplex $X' \subset \partial X$ (of lower dimension) on which the Γ' -action is non-elementary. In this case any median class $\rho^*_{(\mathfrak{m}_{(n,R)}(X))}$, $n \geq 2$ and $R \geq 1$, of the Γ -action on X restricts to a median class of the Γ' -action on X' . In particular, $R_{\Gamma} = R_{\Gamma'}$.*

We say that an action of a group Γ on X is *Roller elementary* if it has a finite orbit on the Roller compactification (that is, X union its Roller boundary). Combining the above theorem with Theorem 1.1 we get the following formulation.

THEOREM 1.3. *Let X be a finite-dimensional CAT(0) cube complex. A Γ -action on X is Roller elementary if and only if the median class $\rho^*_{(\mathfrak{m}_{(n,R)}(X))}$ vanishes for some (equivalently, any) $n \geq 2$ and all $R \geq 1$.*

One of the nice features of the Roller boundary is its robustness when considering products. Because of this, the median cocycle can be defined for each irreducible factor of the essential core of X . We refer the reader to Proposition 3.2 for a description of the cocycle in the case in which the CAT(0) cube complex is not irreducible and hence of the naturality of the behavior of the median class with respect to products. This, together with Theorem 1.1 yields immediately the following corollary.

COROLLARY 1.4. *Let X be a finite-dimensional CAT(0) cube complex with a non-elementary action $\Gamma \rightarrow \text{Aut}(X)$. Then for all $n \geq 2$ and $1 \leq p < \infty$,*

$$\dim H_{\mathbb{b}}^2(\Gamma, \ell^p(\mathfrak{H}(X)^n)) \geq m,$$

where $m \geq 1$ is the number of irreducible factors in the essential core of the Γ -action on X .

This result might not be sharp, in the sense that $H_b^2(\Gamma, \ell^p(\mathfrak{H}(X)^n))$ could be in some cases infinite-dimensional.

In a similar vein Bestvina–Bromberg–Fujiwara have proven the non-vanishing of the second bounded cohomology with general uniformly convex Banach spaces as coefficients and for weakly properly discontinuous actions on CAT(0) spaces in the presence of a rank 1 isometry (that is, an isometry whose axis does not bound a half-flat), see [4].

Our results are different in that in Theorem 1.1, we are neither assuming that the action of $\Gamma \rightarrow \text{Aut}(X)$ is proper or weakly properly discontinuous, nor that the CAT(0) cube complex is proper or has a cocompact group of automorphisms.

Moreover, if the CAT(0) cube complex is a product, then there are no rank 1 isometries. Caprace–Sageev proved [20] that there is always a decomposition of a CAT(0) cube complex analogous to the decomposition of symmetric spaces into ‘irreducible’ (or ‘rank 1’) factors. Our result is not sensitive to this decomposition and hence also applies to products.

But, more than anything else, we want to emphasize that the existence of a well-behaved and concrete bounded cohomological class goes well beyond the mere knowledge that the bounded cohomology group does not vanish and is the starting point of a wealth of rigidity results (see, for example, [7, 8, 10–15, 24, 25, 31, 33, 35–37, 40, 53]).

Furthermore, our coefficients reflect geometric properties of the CAT(0) cube complex, and this is essential to draw conclusions about the action. An example of this is the following superrigidity result.

THEOREM 1.5 (Superrigidity). *Let Y be an irreducible finite-dimensional CAT(0) cube complex and $\Gamma < G_1 \times \cdots \times G_\ell =: G$ be an irreducible lattice in the product of $\ell \geq 2$ locally compact groups. Let $\Gamma \rightarrow \text{Aut}(Y)$ be an essential and non-elementary action on Y . Then the action of Γ on Y extends continuously to an action of G , by factoring via one of the factors.*

Here the group $\text{Aut}(Y)$ is a topological group endowed with the topology of the pointwise convergence on vertices. This theorem is proved in Section 6, to which we refer the reader also for an analogous result that does not require Y to be irreducible and the action to be essential.

We remark that requiring that the action be essential is necessary if one wants an irreducible CAT(0) cube complex, as there is no guarantee that the essential core will be irreducible even when X is.

A result similar to Theorem 1.5 was proved by Monod [39, Theorems 6 and 7] (see also [18, Corollary 1.9]) in the case of an infinite-dimensional CAT(0) space, with conditions both on the action and on the lattice Γ . For example, if Γ is not uniform, then in order to apply Monod’s version of Theorem 1.5, Γ has to be square-integrable and weakly cocompact. Although these conditions are verified for a large class of groups (such as, for example, Kazhdan Kac–Moody lattices and lattices in connected semisimple Lie groups), they are in general rather intractable. To give a sense of this, let us only remark that already finite generation (needed for example for square integrability) is not known for a lattice $\Gamma < \text{Aut}(\mathcal{T}_1) \times \text{Aut}(\mathcal{T}_2)$, not even by imposing strong conditions on the closure of the projections on Γ in $\text{Aut}(\mathcal{T}_i)$ to ensure irreducibility. Furthermore, the more specific nature of a CAT(0) cube complex versus a CAT(0) space allows us to extend the action to the whole complex.

As an illustration we have the following corollary.

COROLLARY 1.6. *Let Γ be an irreducible lattice in the product $G := G_1 \times \cdots \times G_\ell$ of $\ell \geq 2$ locally compact groups with a finite number of connected components. Then any Γ -action on a finite-dimensional CAT(0) cube complex is elementary and has a finite orbit in the Roller compactification $\bar{X} = X \cup \partial X$.*

Proof. Indeed, since G has finitely many connected components and, for any CAT(0) cube complex X , the group $\text{Aut}(X)$ is totally disconnected, a continuous map from G to $\text{Aut}(X)$ must have finite image. In view of Theorem 1.5 (in fact, more precisely of Corollary 6.2), this implies that every Γ -action on a finite-dimensional CAT(0) cube complex is elementary. Note that every finite index subgroup $\Gamma' \leq \Gamma$ is itself a lattice in G . Moreover, the closure of the projection of Γ' to each factor G_i is a closed subgroup of finite index in G_i . It is thus open, and therefore contains the connected component of the identity of G_i . This shows that Γ' is itself an irreducible lattice in the product of $\ell \geq 2$ locally compact groups with a finite number of connected components. By Theorem 1.2, this implies that every Γ' -action, and thus also every Γ -action, has a finite orbit in the Roller compactification of X . \square

Combining Corollary 1.6 with a description of the structure of a point stabilizer in the Roller boundary, established by Pierre-Emmanuel Caprace in Appendix B, one obtains the following Fixed Point property for lattices in semisimple groups, which was pointed out to us by him.

COROLLARY 1.7. *Let Γ be an irreducible lattice in a semisimple Lie group of rank at least 2. Then every Γ -action on a finite-dimensional CAT(0) cube complex X has a fixed point.*

Proof. If the semisimple Lie group has only one simple factor, then Γ has property (T) and the desired conclusion is well known, see [44]. Otherwise, we apply Corollary B.2 from Appendix B: Condition (a) holds as a consequence of Margulis' Normal Subgroup Theorem, while Condition (b) holds by Corollary 1.6. \square

It is conjectured that the conclusion of Corollary 1.7 holds without the hypothesis that X is finite-dimensional; see [23]. In fact, Yves de Cornulier shows in [23] that this is indeed the case, provided the ambient semisimple Lie group has at least one simple factor of rank at least 2.

On a different tone, recall that the concept of measure equivalence was introduced by Gromov as a measure theoretical counterpart of quasi-isometries. The vanishing or non-vanishing of bounded cohomology is not invariant under quasi-isometries (see [16, Corollary 1.7]); on the other hand, Monod–Shalom proved that vanishing of bounded cohomology with coefficients in the regular representation is invariant under measure equivalence [41] and hence introduced a class of groups $\mathcal{C}_{\text{reg}} := \{\Gamma : H_b^2(\Gamma, \ell^2(\Gamma)) \neq 0\}$. They also proved, for example, that if $\Gamma \in \mathcal{C}_{\text{reg}}$ and $\Gamma \times \Gamma$ is measure equivalent to Λ , then $\Gamma \times \Gamma$ and Λ are commensurable. We can add to the groups in this list.

COROLLARY 1.8. *Let Γ be a group acting on a finite-dimensional irreducible CAT(0) cube complex. If the action is metrically proper, non-elementary and essential, then $H_b^2(\Gamma, \ell^p(\Gamma)) \neq 0$ for $1 \leq p < \infty$, and hence in particular $\Gamma \in \mathcal{C}_{\text{reg}}$.*

We remark that the same result does not hold if X is not irreducible. In fact, it can be easily seen, using [17, Theorem 16], that if $\Gamma < G_1 \times G_2$ is an irreducible lattice in the product of locally compact groups, then $H_b^2(\Gamma, \ell^p(\Gamma)) = 0$, provided G_1 and G_2 are not compact. An example of such a group is any irreducible lattice Γ in $\text{SL}(2, \mathbb{Q}_p) \times \text{SL}(2, \mathbb{Q}_q)$, while it is easy to see that it acts non-elementarily and essentially on the product of two regular trees $\mathcal{T}_{p+1} \times \mathcal{T}_{q+1}$.

A result similar to Corollary 1.8 has been proved by Hamenstädt in the case of a group Γ acting properly on a proper CAT(0) space, also under the assumption that there exists a rank 1 isometry and that the group Γ is closed in the isometry group of X [30]. Similarly, Hull and Osin proved that every group Γ with a sufficiently nice hyperbolic subgroup has infinite-dimensional $H_b^2(\Gamma, \ell^p(\Gamma))$ for $1 \leq p < \infty$ (see [32]). Examples of groups satisfying such a condition encompass, among others, groups Γ acting properly on a proper CAT(0) space with a rank 1 isometry and groups Γ acting on a hyperbolic space also with a

rank 1 isometry and containing a loxodromic element satisfying the Bestvina–Fujiwara ‘weakly properly discontinuous’ condition. We emphasize that our CAT(0) cube complexes are allowed to be locally countable and reducible (recall that irreducibility is equivalent to the existence of a rank 1 isometry [20]).

The proof of Theorem 1.1 uses the functorial approach to bounded cohomology developed in [9, 17, 38]; the main point here is to be able to realize bounded cohomology via essentially bounded alternating cocycles on a *strong* Γ -boundary. Recall from [34] that a strong Γ -boundary is a Lebesgue space (B, ϑ) endowed with a measure class preserving Γ -action that is in addition amenable and doubly ergodic ‘with coefficients’ (see [34] for the precise definition or Subsection 3.C). An example of a strong Γ -boundary is the Poisson boundary of any spread-out non-degenerate symmetric probability measure on Γ , see [34]. The advantage of the approach using a strong Γ -boundary is that the second bounded cohomology is not a quotient anymore (hence allowing one to determine easily when a cocycle defines a non-trivial class); the disadvantage is that the pullback via a representation has to be realized by a boundary map (with consequent technical difficulties [9]). The amenability of a strong Γ -boundary implies immediately the existence of a boundary map into probability measures on the Roller compactification of X , but going from probability measures to Dirac masses is often the sore point of many rigidity questions. In the case of a proper CAT(0) cube complex and a cocompact group of isometries in $\text{Aut}(X)$, Nevo–Sageev identified the closure of the set of non-terminating ultrafilters (see Subsection 2.A for the definition) as a metric model for a Poisson boundary of Γ (see [43]). In this case, the boundary map could have been taken simply to be the identity. In general, we have the following theorem.

THEOREM 1.9. *Let $\Gamma \rightarrow \text{Aut}(X)$ be a non-elementary group action on a finite-dimensional CAT(0) cube complex X . If (B, ϑ) is a strong Γ -boundary, then there exists a Γ -equivariant measurable map $\varphi : B \rightarrow \partial X$.*

In fact, one can obtain something a bit more precise, namely that the boundary map takes values in the non-terminating ultrafilters of the Γ -essential core of X (see Theorem 4.1 and Corollary 4.2). To prove Theorem 1.9, we develop some methods that take inspiration from [40, Proposition 3.3] in the case of a simplicial tree but are considerably more involved in the case of a CAT(0) cube complex due to the lack of hyperbolicity. For the sake of completeness we remark that, in the case in which X is a CAT(0) space of finite telescopic dimension, Bader–Duchesne–Lécureux proved the existence of such a boundary map into the visual boundary $\partial_{\triangleleft} X$ of X [1].

The first step in the identification of a Poisson boundary in [43] is the proof that the set of non-terminating ultrafilters is not empty, under the assumption that the action is essential and the CAT(0) cube complex is cocompact. The same assertion with the cocompactness of X replaced by the non-existence of $\text{Aut}(X)$ -fixed points in the CAT(0) boundary follows from our proof that the boundary map takes values into the set of the non-terminating ultrafilters.

COROLLARY 1.10. *Let Y be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(Y)$ acts essentially and without fixed points in $\partial_{\triangleleft} Y$. Then the set of non-terminating ultrafilters in ∂Y is not empty.*

The structure of the paper is as follows. In Section 2, we recall the appropriate definitions and fix the terminology of CAT(0) cube complexes; we establish moreover some basic results needed in the paper, by pushing a bit further than what was available in the literature; the knowledgeable reader should have no problem parsing through the subsections. In Section 3, we construct the cocycle on the Roller compactification of the CAT(0) cube complex X and show that it is bounded. We conclude the section with an outlook on the proof of the non-vanishing

of a median class. The boundary map and Theorem 1.9 are discussed in Section 4. We prove Theorems 1.1 and 1.2 in Subsection 5.B, Corollary 1.4 is a consequence of Theorem 1.1 and Proposition 3.2, while Theorem 1.5 and Corollary 1.8 are proved in Section 6.

2. Preliminaries and basic results

2.A. Generalities on CAT(0) cube complexes, hyperplanes, duality and boundaries

A cube complex X is a metric polyhedral complex with cells isomorphic to $[0, 1]^n$ and isometries $\varphi_j : [0, 1]^j \rightarrow X$ as gluing maps. The cube complex is CAT(0) if it is non-positively curved with the induced Euclidean metric and has finite dimension D if the m -dimensional skeleton X^m of X is empty for $m > D$ and non-empty for $m = D$. We always assume our cube complexes to be finite-dimensional. A cube complex X is CAT(0) if and only if it is both simply connected and the link of every vertex is a flag complex: recall that a flag complex is a simplicial complex such that any $(n + 1)$ -vertices that are pairwise connected by an edge actually span an n -simplex [5, Theorem II.5.20]. A combinatorial isometry between two CAT(0) cube complexes is a homeomorphism $f : X \rightarrow Y$ such that the composition $f \circ \varphi_j : [0, 1]^j \rightarrow Y$ is an isometry into a cube of Y . Note that any combinatorial isometry preserves also the CAT(0) metric. We denote by $\text{Aut}(X)$ the group of combinatorial isometries from X to itself.

Given a finite-dimensional cube complex X , we can define an equivalence relation on edges, generated by the condition that two edges are equivalent if they are opposite sides of the same square (that is, a 2-cube). A midcube of an n -cube σ with respect to the above equivalence relation is the convex hull of the set of midpoints of elements in the equivalence relation. A hyperplane is the union of the midcubes that intersect the edges in an equivalence class. So a hyperplane is a closed convex subspace and it defines uniquely two halfspaces, that is, the two complementary connected components. On the countable collection $\mathfrak{H}(X)$, or simply \mathfrak{H} , when no confusion arises, of halfspaces on X defined by the hyperplanes, one can define a fixed-point-free involution

$$\begin{aligned} * : \mathfrak{H} &\longrightarrow \mathfrak{H} \\ h &\longmapsto h^* := X \setminus h, \end{aligned} \tag{2.1}$$

so that a hyperplane is the geometric realization of a pair $\{h, h^*\}$. In the following, we identify the hyperplane \hat{h} with the pair of halfspaces $\{h, h^*\}$ that it defines. We denote by $\hat{\mathfrak{H}}(X)$ the set of hyperplanes.

We say that two halfspaces h, k are transverse, and we write $h \pitchfork k$ if all the intersections

$$h \cap k, \quad h \cap k^*, \quad h^* \cap k, \quad h^* \cap k^* \tag{2.2}$$

are not empty. Two halfspaces h, k are parallel, and we write $h \parallel k$, if they are not transverse, equivalently if (exactly) one of the following relations,

$$h \subset k^*, \quad h \subset k, \quad h^* \subset k^*, \quad h^* \subset k \tag{2.3}$$

holds; two parallel halfspaces h and k are said to be facing if $h \subset k^*$. We say that two hyperplanes \hat{h}, \hat{k} are transverse (respectively, parallel) if some (and hence any) choice of corresponding halfspaces h and k is transverse (respectively, parallel). Finally, we say that two points u and v are separated by a halfspace h (or a hyperplane $\hat{h} = \{h, h^*\}$) if $u \in h$ and $v \in h^*$ (or vice versa).

Two halfspaces h, k are said to be nested if either $h \subset k$ or $k \subset h$. A subset of hyperplanes is transverse (respectively, parallel) if all of its elements are pairwise transverse (respectively, parallel).

Recall that a family of pairwise transverse hyperplanes must have a common intersection ([47] or [46]). We can think of the dimension of a CAT(0) cube complex as the largest

cardinality of a family of pairwise transverse hyperplanes, because such a maximal intersection defines a cube of maximal dimension.

Given a subset $\alpha \subset \mathfrak{H}$ of halfspaces, we denote by α^* the set $\{h^* : h \in \alpha\}$.

DEFINITION 2.1. We say that a subset $\alpha \subset \mathfrak{H}$ of halfspaces satisfies:

- (i) the *partial choice* condition if $\alpha \cap \alpha^* = \emptyset$, that is if whenever $h \in \alpha$, then $h^* \notin \alpha$;
- (ii) the *choice* condition if $\alpha \cap \alpha^* = \emptyset$ and $\alpha \sqcup \alpha^* = \mathfrak{H}$;
- (iii) the *consistency* condition if whenever $h \in \alpha$ and $h \subset k$, then $k \in \alpha$.

Then an *ultrafilter* on \mathfrak{H} is a subset of \mathfrak{H} that satisfies the choice and consistency properties. In other words, an ultrafilter on \mathfrak{H} is a choice of a halfspace for each hyperplane in X with the condition that as soon as a halfspace is in the ultrafilter, any halfspace containing it must also be in the ultrafilter. We call *partially defined ultrafilter* a subset $\alpha \subset \mathfrak{H}$ that satisfies the partial choice and consistency properties.

REMARK. We point out that the notion of ultrafilter used in the theory of CAT(0) cube complexes is slightly off from the classical one in set theory and topology (see, for example, [21]). In fact, in the context of CAT(0) cube complexes, subsets of ultrafilters are never ultrafilters and thus, in particular, the intersection of two ultrafilters is never an ultrafilter.

We say that an ultrafilter satisfies the *Descending Chain Condition (DCC)* if every descending chain of halfspaces terminates. Such ultrafilters are called *principal* and are in one-to-one correspondence with the vertices of the CAT(0) cube complex X (see [28]). By abuse of notation, we do not usually make a distinction between X , its vertex set or the collection of principal ultrafilters.

The consideration of X as a collection of ultrafilters leads in a natural way to an inclusion of X into the Bernoulli space $2^{\mathfrak{H}}$, where $v \mapsto \{h \in \mathfrak{H} : v \in h\}$. This justifies a further (standard) abuse: thinking of $X \subset 2^{\mathfrak{H}}$, by duality we get that $h \in v$ if and only if $v \in h$ and we can hence write $v = \bigcap_{h \in v} h$. Let \overline{X} be the closure of X in $2^{\mathfrak{H}}$. One can check that the elements of \overline{X} , thought of as subsets of \mathfrak{H} , are ultrafilters.

The correspondence that associates to an ultrafilter a vertex in \overline{X} can be pushed further to give a duality between finite-dimensional CAT(0) cube complexes and those *pocsets* that satisfy both the *finite interval condition* and the *finite width condition*. Recall that a *pocset* Σ is a partially ordered set with an order-reversing involution. The pocset satisfies the *finite interval condition* if, for every pair $\alpha, \beta \in \Sigma$ with $\alpha \subset \beta$, there are only finitely many $\gamma \in \Sigma$ such that $\alpha \subset \gamma \subset \beta$; moreover, it satisfies the *finite width condition* if there is an upper bound on the size of a collection of incomparable elements. Given a pocset Σ , one can consider the space of ultrafilters on Σ . The CAT(0) cube complex $X(\Sigma)$ corresponding to the pocset Σ has the principal ultrafilters as vertices, edges joining ultrafilters that differ only in the assignment on one element in Σ and cubes attached to the 1-skeleton whenever it is possible.

The set of halfspaces $\mathfrak{H}(X)$ in a CAT(0) cube complex is a pocset with the above properties and the CAT(0) cube complex obtained with the above construction from the set of principal ultrafilters on $\mathfrak{H}(X)$ is exactly X .

The boundary $\partial X := \overline{X} \setminus X$ is called the *Roller boundary*, and consists of all ultrafilters that are not principal [45]. The compact set \overline{X} is called the *Roller compactification*.

An ultrafilter $v \in \partial X$ is said to be *non-terminating* if every finite descending chain can be extended, that is, if given a finite collection $\{h_0, \dots, h_N\} \subset v$ such that $h_0 \supset \dots \supset h_N$, then there is an $h_{N+1} \in v$ such that $h_{N+1} \subset h_N \subset \dots \subset h_0$.

While the Roller boundary ∂X is not empty if the CAT(0) cube complex is unbounded, it is unclear as to when the set of non-terminating ultrafilters is not empty. However, one can

impose reasonable conditions that do guarantee that these exist (see [43] and Section 4). In the case of a tree, the entire Roller boundary consists of non-terminating ultrafilters; in the case of a \mathbb{Z}^D , there are only 2^D -many non-terminating ultrafilters, while there are examples, such as the wedge of two strips [43, Remark 3.2], in which the set of non-terminating ultrafilters is empty.

A CAT(0) cube complex has of course also a visual boundary $\partial_{\leq} X$ with respect to its CAT(0) metric [5]. We recall that $\partial_{\leq} X$ is the set of endpoints of geodesic rays in X , where we identify two geodesic rays if they stay at bounded distance from each other.

2.B. *Intervals and median structure*

If $u, v \in X$, then their *combinatorial distance* $d(u, v)$ is the number of hyperplanes by which the two corresponding ultrafilters differ. We will call a sequence of points $u = x_0, \dots, x_n = v$ a *combinatorial geodesic* if $d(x_i, x_{i+1}) = 1$ for $i = 0, \dots, n - 1$ and

$$d(x_i, x_j) + d(x_j, x_k) = d(x_i, x_k)$$

for all $1 \leq i \leq j \leq k \leq n$. Hence the combinatorial distance corresponds to the graph metric on the 1-skeleton of X . The oriented *interval* of halfspaces

$$[u, v] = \{h \in \mathfrak{H} : h \in v \setminus u\}$$

is the (finite) set of halfspaces containing v and not u and the counting measure on \mathfrak{H} is consistent with the combinatorial metric d on X in that $|[u, v]| = d(u, v)$. Note that on $[u, v]$ there is a partial order given by the inclusion. It is immediate to check that

$$[u, v] = [v, u]^*,$$

where we recall that $[v, u]^* = \{h \in \mathfrak{H} : h^* \in [v, u]\}$.

Combinatorial geodesics and oriented intervals are related as follows.

LEMMA 2.2. *For each $u, v \in X$, combinatorial geodesics between u and v are in one to one correspondence with enumerations of the elements of $[u, v]$ that are compatible with the (reverse) partial order given by inclusions.*

Proof. Let $u = x_0, \dots, x_n = v$ be a combinatorial geodesic. One obtains an enumeration h_1, \dots, h_n of elements of $[u, v]$ by setting h_i to be the halfspace corresponding to the unique oriented edge between x_i and x_{i+1} . Let us now show that this order is consistent with the (reverse) partial order. Indeed, suppose that $i < j$. If $h_i \cap h_j$, then they are incomparable and there is nothing to check. Otherwise, observe that $x_i \in h_i \setminus h_j$ as x_i is obtained by starting at u and crossing each of the elements in $\{h_1, \dots, h_i\}$ (which does not contain h_j). Therefore, $h_j \subset h_i$.

Conversely, let $u, v \in X$ with $d(u, v) = n$ and assume we are given an enumeration h_1, \dots, h_n consistent with the inclusion, where $h_j \in [u, v]$. By consistency, if $h \neq h_j$ were a halfspace between u and h_1 , then $h \in [u, v]$, contradicting that $d(u, v) = n$. Therefore, there is a unique oriented edge starting at u corresponding to h_1 . Let x_1 be the terminal vertex. Inductively, this defines a sequence x_1, \dots, x_n where $x_n = v$. Since $|[u, v]| = d(u, v)$, it follows that this describes a combinatorial geodesic. □

We consider also the *vertex-interval*

$$\mathcal{I}(u, v) := \{w \in X : w \cap (u \cap v) = u \cap v\},$$

that is, the set of all vertices that are crossed by some combinatorial geodesic between u and v .

The following fact seems to be folklore and is essential for our result. We refer the reader to [6, Theorem 1.16] for a complete proof.

LEMMA 2.3 (Intervals embedding). *Let $u, v \in \overline{X}$. Then the vertex intervals $\mathcal{I}(u, v)$ isometrically embed into $\overline{\mathbb{Z}^D}$ (with the standard cubulation), where D is the dimension of X .*

The image $\mathcal{I}_{u,v}$ in \mathbb{Z}^D of the above embedding is exactly the CAT(0) cube complex associated to the halfspaces $\mathfrak{H}(u, v) := [u, v] \cup [v, u]$.

REMARK 2.4. In general, if $u \in X$ is an ultrafilter, its opposite u^* is not an ultrafilter. It is easy to see that if u^* is an ultrafilter, then $\mathfrak{H}(X) = [u, u^*] \cup [u^*, u]$ and hence X is an interval.

Also recall that the vertex set of a CAT(0) cube complex with the edge metric is a *median space* [45]; namely, for every triple of vertices $u, v, w \in X$, the intersection $\mathcal{I}(u, v) \cap \mathcal{I}(v, w) \cap \mathcal{I}(w, u)$ is exactly a singleton. This unique point is called the *median* of u , v and w and we denote it by $m(u, v, w)$. It is a standard fact that

$$m(u, v, w) = (u \cap v) \cup (v \cap w) \cup (w \cap u). \quad (2.4)$$

2.C. Isometric embeddings

If $\mathfrak{H}' \subset \mathfrak{H}(X)$ is an involution invariant subset of halfspaces, then \mathfrak{H}' is a pocset in its own right and hence one can consider the associated CAT(0) cube complex $X(\mathfrak{H}')$. A priori, the complex $X(\mathfrak{H}')$ that one obtains with this construction cannot be embedded as a subcomplex of X , but there is always a combinatorial quotient map $\pi_{\mathfrak{H}'} : \overline{X} \rightarrow \overline{X}(\mathfrak{H}')$, defined by $\alpha \mapsto \alpha \cap \mathfrak{H}'$, that restricts to $\pi_{\mathfrak{H}'} : X \rightarrow X(\mathfrak{H}')$. If the subset $\mathfrak{H}' \subset \mathfrak{H}$ is invariant for the action of a group $\Gamma \rightarrow \text{Aut}(X)$ of combinatorial automorphisms, then $\overline{X}(\mathfrak{H}')$ inherits a Γ -action with respect to which the map $\pi_{\mathfrak{H}'}$ is Γ -equivariant.

There are however conditions under which $\overline{X}(\mathfrak{H}')$ can be embedded in \overline{X} .

DEFINITION 2.5. Let $\mathfrak{H}' \subset \mathfrak{H}(X)$ be an involution invariant subset of halfspaces. A *lifting decomposition* of \mathfrak{H}' is a choice of a subset $W \subset \mathfrak{H}(X) \setminus \mathfrak{H}'$ satisfying the partial choice and consistency conditions (see Definition 2.1), and so that $\mathfrak{H}(X) = \mathfrak{H}' \sqcup (W \sqcup W^*)$.

We note that a lifting decomposition need not exist. A collection $\mathfrak{H}' \subset \mathfrak{H}$ is said to be *tight* if it satisfies the following: for every $h, k \in \mathfrak{H}'$, if $h \subset \ell \subset k$, then $\ell \in \mathfrak{H}'$. We remark that the existence of a lifting decomposition W of $\mathfrak{H}' \subset \mathfrak{H}$ implies that \mathfrak{H}' is tight. Indeed, suppose that $h, k \in \mathfrak{H}'$ and $h \subset \ell \subset k$. If $\ell \notin \mathfrak{H}'$, then $\ell \in W \sqcup W^*$. Since \mathfrak{H}' is involution invariant, we may assume that $\ell \in W$. But this means that $k \in W$, which contradicts the fact that $W \cap \mathfrak{H}' = \emptyset$. This shows that the condition that \mathfrak{H}' is tight is necessary for the existence of a lifting decomposition for it.

LEMMA 2.6. *Let $\mathfrak{H}' \subset \mathfrak{H}(X)$ be an involution invariant tight subset of halfspaces. Assume that \mathfrak{H}' admits a lifting decomposition $\mathfrak{H} = \mathfrak{H}' \sqcup (W \sqcup W^*)$. Then there is an isometric embedding $i : \overline{X}(\mathfrak{H}') \hookrightarrow \overline{X}(\mathfrak{H})$, defined by $i(\alpha) := \alpha \sqcup W$, whose image is $i(\overline{X}(\mathfrak{H}')) = \bigcap_{h \in W} h$.*

As particular cases, if $\mathfrak{H}' = \emptyset$, then $i(\overline{X}(\mathfrak{H}'))$ is a point; or if W contains an infinite descending chain, then $i(\overline{X}(\mathfrak{H}')) \subset \partial X$.

Proof. We first show that if α is an ultrafilter on \mathfrak{H}_W , then $\alpha \sqcup W$ is an ultrafilter on \mathfrak{H} . By construction, $\alpha \sqcup W$ satisfies the choice condition. We need to verify the consistency condition, that is, that if $h \in \alpha \sqcup W$, $k \in \mathfrak{H}$ with $h \subset k$, then $k \in \alpha \sqcup W$.

If $h \in W$ and $k \in \mathfrak{H}$ is such that $h \subset k$, then $k \in W$, since $W \subset \mathfrak{H}$ satisfies the consistency condition. If $h \in \alpha$ and $k \in \mathfrak{H}$ is such that $h \subset k$, then either $k \in \alpha \sqcup \alpha^*$ and hence $k \in \alpha$ because α is an ultrafilter on $\alpha \sqcup \alpha^*$, or $k \in W \sqcup W^*$. But k cannot be in W^* . In fact, if

$k \in W^*$, then $k^* \in W$; since $k^* \subset h^*$ and W satisfies the consistency condition, it follows that $h^* \in W$, contradicting that $\alpha \cap (W \sqcup W^*) = \emptyset$.

Now, assume that $\alpha, \beta \in \overline{X}(\mathfrak{H}')$. It is easy to check that

$$(\alpha \sqcup W)\Delta(\beta \sqcup W) = \alpha\Delta\beta,$$

since $\mathfrak{H}' \cap W = \emptyset$. This shows that the embedding is an isometry and hence extends to the cube structure on $\overline{X}(\mathfrak{H}')$.

By definition $\alpha \sqcup W = \bigcap_{h \in \alpha \sqcup W} h \subset \bigcap_{h \in W} h$, so that $\mathfrak{H} \subseteq (\alpha \sqcup \alpha^*) \sqcup (W \sqcup W^*)$. Moreover, $\bigcap_{h \in W} h$ consists of all partially defined ultrafilters on $W \sqcup W^*$: to complete an element $x \in \bigcap_{h \in W} h$ to an ultrafilter on \mathfrak{H} is exactly equivalent to choosing $\alpha \in \overline{X}(\mathfrak{H}')$. \square

REMARK 2.7. As alluded to above, the existence of a lifting decomposition is a very restrictive condition. However, the existence of a *strongly convex* set, that is a set $B \subset X$ such that for any $x, y \in B$ also $\mathcal{I}(x, y) \subset B$, implies the existence of a lifting decomposition. In fact, if

$$\mathfrak{H}' := \{h \in \mathfrak{H}(X) : h \text{ crosses } B\} \quad \text{and} \quad W := \{h \in \mathfrak{H}(X) : B \subset h\},$$

then $\mathfrak{H} = \mathfrak{H}' \sqcup (W \sqcup W^*)$ is a lifting decomposition of \mathfrak{H}' with which $\overline{X}(\mathfrak{H}')$ gets isometrically embedded in \overline{X} with image B .

DEFINITION 2.8. A map $\varrho : \overline{X} \rightarrow \overline{X}$ is a *projection* if there exists a lifting decomposition $\mathfrak{H}(X) = \mathfrak{H}_W \sqcup (W \sqcup W^*)$ and if $\varrho = i \circ \pi_{\mathfrak{H}_W}$, where $i : \overline{X}(\mathfrak{H}_W) \hookrightarrow \overline{X}$ is the isometric embedding in Lemma 2.6.

It is easy to verify that the composition of two projections is still a projection.

If $\Gamma \rightarrow \text{Aut}(X)$ is an action and \mathfrak{H}_W is Γ -invariant, then $\overline{X}(\mathfrak{H}_W)$ inherits a Γ -action. If in addition ϱ is a projection and the embedding is Γ -equivariant, then the image of the projection is a Γ -invariant subcomplex in \overline{X} . This happens exactly when the choice of subset $W \subset \mathfrak{H}(X)$ of the lifting decomposition is Γ -invariant, so that $i(\overline{X}(\mathfrak{H}_W))$ is a Γ -invariant subcomplex in $\overline{X}(\mathfrak{H})$.

2.D. Decomposition into products

The product of CAT(0) cube complexes is a CAT(0) cube complex in a natural way. If $X = Y \times Z$, then there is the following decomposition of the hyperplanes:

$$\hat{\mathfrak{H}}(X) = \{\hat{h}_Y \times Z : \hat{h}_Y \in \hat{\mathfrak{H}}(Y)\} \sqcup \{Y \times \hat{h}_Z : \hat{h}_Z \in \hat{\mathfrak{H}}(Z)\} \cong \hat{\mathfrak{H}}(Y) \sqcup \hat{\mathfrak{H}}(Z), \quad (2.5)$$

and $(\hat{h}_Y \times Z) \pitchfork (Y \times \hat{h}_Z)$ for any $h_Y \in \mathfrak{H}(Y)$ and $h_Z \in \mathfrak{H}(Z)$.

Conversely, any such partition of the hyperplanes into mutually transverse subsets corresponds to a decomposition of the CAT(0) cube complex into a product. In fact, by [20, Proposition 2.6] any CAT(0) cube complex decomposes as a product $X = X_1 \times \cdots \times X_m$ of irreducible factors, $m \geq 1$, which are unique up to permutations and are often referred to as the *rank 1 factors* of X .

The induced CAT(0) metric (respectively, the combinatorial metric) on the product is the ℓ^2 -product (respectively, ℓ^1 -product) of the factor metrics. We record the following standard fact.

LEMMA 2.9. *Let $X = X_1 \times \cdots \times X_k$ be the product of CAT(0) spaces X_j , $j = 1, \dots, k$ and let $G := G_1 \times \cdots \times G_k$, where $G_j \leq \text{Aut}(X_j)$ is a subgroup of the isometries of the j th factor X_j . Then any G_j -fixed point in $\partial_{\triangleleft} X_j$ defines a G -fixed point in $\partial_{\triangleleft} X$.*

Proof. Let us denote by δ_j and δ the CAT(0) metrics, respectively, on X_j and on X . Assume that, up to permuting the indices, there is a G_1 -fixed point $\xi_1 \in \partial_{\triangleleft} X_1$. Let $\ell_1 : [0, \infty) \rightarrow X_1$ be

a geodesic in X_1 representing ξ_1 , that is, $\xi_1 = \ell_1(\infty)$. Since ξ_1 is G_1 -invariant, it follows that $\sup_{t \in [0, \infty)} \delta_1(\gamma \ell_1(t), \ell_1(t)) < \infty$. If $x_j \in X_j$ for $2 \leq j \leq m$ is any point, then $\ell : [0, \infty) \rightarrow X$ defined by $\ell(t) := (\ell_1(t), x_2, \dots, x_m)$ is a geodesic in X . Then, for any $\gamma \in G$, we have

$$\sup_{t \in [0, \infty)} \delta(\gamma \ell(t), \ell(t))^2 := \sup_{t \in [0, \infty)} \left[\delta_1(\gamma \ell_1(t), \ell_1(t))^2 + \sum_{j=1}^m \delta_j(\gamma x_j, x_j)^2 \right] < \infty,$$

hence $\ell(\infty)$ is G -invariant. \square

In addition, there is a corresponding decomposition of the Roller boundary,

$$\partial X = \bigcup_{j=1}^m \overline{X_1} \times \cdots \times \overline{X_{j-1}} \times \partial X_j \times \overline{X_{j+1}} \times \cdots \times \overline{X_m},$$

and $\text{Aut}(X)$ contains $\text{Aut}(X_1) \times \cdots \times \text{Aut}(X_m)$ as a finite index subgroup ($\text{Aut}(X)$ is allowed to permute isomorphic factors). If $\Gamma \rightarrow \text{Aut}(X)$ is a group acting by automorphisms, then there is a subgroup $\Gamma_0 < \Gamma$ of finite index ($\leq m!$) that acts on X_j via the projection $\Gamma_0 \rightarrow \text{Aut}(X_j)$.

2.E. The essential core

A halfspace $h \in \mathfrak{H}$ is said to be Γ -essential if, for some (equivalently all) $x \in X$, the Γ -orbit of x inside h , that is, $\Gamma \cdot x \cap h$, is not at bounded distance from the hyperplane \hat{h} . A hyperplane $\hat{h} \in \hat{\mathfrak{H}}$ is called Γ -essential (or essential for short) if each of the corresponding halfspaces is Γ -essential, and half- Γ -essential (or half-essential) if only one of the two corresponding halfspaces is Γ -essential. The Γ -essential core (or essential core) Y of the Γ -action on X is a CAT(0) cube complex corresponding to the Γ -essential (or essential) hyperplanes. The Γ -action on the Γ -essential core Y is essential and any non-empty Γ -invariant convex subcomplex of Y is equal to Y . Following the notation of [20], we denote by $\text{Ess}(X, \Gamma)$ the set of Γ -essential hyperplanes in X , so that we can write

$$\hat{\mathfrak{H}}(X) = \text{Ess}(X, \Gamma) \sqcup \text{nEss}(X, \Gamma),$$

where the set of non-essential hyperplanes $\text{nEss}(X, \Gamma)$ includes both the half-essential and the trivial ones. Since both $\text{Ess}(X, \Gamma)$ and $\text{nEss}(X, \Gamma)$ are Γ -invariant subsets of $\hat{\mathfrak{H}}(X)$, the decompositions

$$\hat{\mathfrak{H}}(X) = \text{Ess}(X, \Gamma) \sqcup \text{nEss}(X, \Gamma) = \text{Ess}(Y, \Gamma) \sqcup \text{nEss}(X, \Gamma)$$

are Γ -invariant.

While in general the essential core of an action can be empty, it is proved in [20, Proposition 3.5] that if there are no Γ -fixed points in the visual boundary $\partial_{\triangleleft} X$ of X and no Γ -fixed points in X , then the essential core Y is a non-empty Γ -invariant convex subcomplex $Y \subset X$. As a consequence, one has both that $\partial_{\triangleleft} Y \subset \partial_{\triangleleft} X$ and $\partial Y \subset \partial X$. However, even if X is irreducible, its essential core Y need not be. Let $Y = Y_1 \times \cdots \times Y_m$ be the decomposition into irreducible factors. Using the decomposition of hyperplanes for products discussed above, we obtain

$$\hat{\mathfrak{H}}(X) = \text{Ess}(Y, \Gamma) \sqcup \text{nEss}(X, \Gamma) = \hat{\mathfrak{H}}(Y_1) \sqcup \cdots \sqcup \hat{\mathfrak{H}}(Y_m) \sqcup \text{nEss}(X, \Gamma), \quad (2.6)$$

where we used for simplicity the notation $\hat{\mathfrak{H}}(Y_j)$ to indicate $\text{Ess}(Y_j, \Gamma)$ (since by hypothesis they coincide because the induced action on Y_j is Γ -essential).

Let $\mathfrak{H}(X)^n$ denote the set of n -tuples of halfspaces in X . Since if n -tuple $s \in \mathfrak{H}(X)^n$ is essential, then any other halfspace containing the halfspaces in s is essential as well, the

decomposition in (2.6) induces a decomposition

$$\mathfrak{H}(X)^n = \mathfrak{H}(Y_1)^n \sqcup \cdots \sqcup \mathfrak{H}(Y_m)^n \sqcup \mathfrak{H}_{nEss}(X)^n, \tag{2.7}$$

where $\mathfrak{H}_{nEss}(X)^n$ consists of n -tuples such that at least one halfspace is non-essential.

2.F. *Skewering, flipping: über-separated and über-parallel pairs of halfspaces*

Flipping and double skewering are important tools introduced by Caprace–Sageev [20].

DEFINITION 2.10 ([20]). We say that $\gamma \in \text{Aut}(X)$ *flips* a halfspace $h \in \mathfrak{H}(X)$ if $\gamma h^* \subset h$. Moreover, we say that γ *skewers* \hat{h} if $\gamma h \subset h$ (or $h \subset \gamma h$).

Under reasonable hypotheses such combinatorial automorphisms can always be found. More precisely, if X is a finite-dimensional CAT(0) cube complex and $\Gamma \rightarrow \text{Aut}(X)$ acts essentially on X without fixing any point in the visual boundary $\partial_{\infty} X$, then, for every halfspace $h \in \mathfrak{H}(X)$, there exists $\gamma \in \Gamma$ that flips h (see [20, Flipping Lemma, Subsection 1.2]). As a simple consequence, we have also that, given any two halfspaces $k \subset h$, there exists $\gamma \in \Gamma$ such that $\gamma h \subset k \subset h$ (see [20, Double Skewering Lemma, Subsection 1.2]).

The following notion was first introduced by Behrstock–Charney [3].

DEFINITION 2.11 ([3]). We say that two parallel hyperplanes are *strongly separated* if there is no hyperplane that is transverse to both.

By the usual abuse of terminology we say that two halfspaces are strongly separated if the corresponding hyperplanes are.

The existence of strongly separated hyperplanes is definitively a rank 1 phenomenon. In fact, it is easy to see that if X is reducible, then there are no strongly separated hyperplanes. For non-elementary CAT(0) cube complexes, the fact that the existence of strongly separated hyperplanes is actually equivalent to the irreducibility of the CAT(0) cube complex was proved in [20], although the case of a right-angled Artin group can already be found in [3].

We will need a finer notion of strong separation, which is less standard but will be key to our work.

DEFINITION 2.12. Two strongly separated halfspaces h_1 and h_2 are said to be an *über-separated pair* if any two halfspaces k_1, k_2 with the property that $h_i \cap k_i$ for $i = 1, 2$ are parallel. We say that two strongly separated hyperplanes are *über-separated* if their halfspaces are.

REMARK 2.13. If $h \subset k \subset \ell$ are pairwise strongly separated halfspaces, then h and ℓ form an über-separated pair.

LEMMA 2.14. Let Y be a finite-dimensional irreducible CAT(0) cube complex and $\Gamma \rightarrow \text{Aut}(Y)$ be a group acting essentially and non-elementarily. Given any hyperplane \hat{h} , there exists $\gamma \in \Gamma$ such that \hat{h} and $\gamma \hat{h}$ form an über-separated pair and $h \subset \gamma h$ (or $\gamma h \subset h$).

Proof. By [20, Proposition 5.1] for any halfspace h there is a pair of halfspaces h_1, h_2 such that $h_1 \subset h \subset h_2$ and \hat{h}_1 and \hat{h}_2 are strongly separated. We apply now the Double Skewering Lemma in [20, Subsection 1.2] to the pair $h_1 \subset h_2$ to obtain that $h_1 \subset h_2 \subset \gamma_0 h_1$ for some $\gamma_0 \in \Gamma$. By construction, and since Γ acts by automorphisms of Y , we have the chain of inclusions

$$h_1 \subset h \subset h_2 \subset \gamma_0 h_1 \subset \gamma_0 h \subset \gamma_0 h_2 \subset \gamma_0^2 h_1 \subset \gamma_0^2 h \subset \gamma_0^2 h_2.$$

Since \hat{h}_1 and \hat{h}_2 are strongly separated, the same is true for $\gamma_0\hat{h}_1$ and $\gamma_0\hat{h}_2$, and for $\gamma_0^2\hat{h}_1$ and $\gamma_0^2\hat{h}_2$. Hence \hat{h} and $\gamma_0^3\hat{h}$ is an über-separated pair by Remark 2.13. \square

Note that über-separated pairs are in particular strongly separated and hence they do not exist in the reducible case [20, Proposition 5.1]. To deal also with the reducible case we will use the following generalization.

DEFINITION 2.15. Two parallel halfspaces h_1 and h_2 are said to be *über-parallel* if, for every pair of halfspaces k_1, k_2 such that $h_i \cap k_i$ for $i = 1, 2$, then either the halfspaces k_1 and k_2 are parallel, or they each cross both h_1 and h_2 . Two parallel hyperplanes are *über-parallel* if their halfspaces are.

According to the definition, an über-separated pair is in particular über-parallel. If X is a product, then über-separated pairs from an irreducible factor will not be über-separated in X but will be über-parallel. Even when X is irreducible, there may be reducible subcomplexes of X . In such a reducible subcomplex, such as, for example, a copy of \mathbb{Z}^2 inside $\mathbb{Z} * \mathbb{Z}^2$, there may be pairs of halfspaces that are über-separated in one of the factors of that subcomplex but not über-separated in X . The notion of über-parallel captures these types of pairs, as is the case, for example, in the Salvetti complex associated to $\mathbb{Z} * \mathbb{Z}^2$.

2.G. The bridge

The concept of *bridge* of two parallel hyperplanes was introduced by Behrstock–Charney [3].

DEFINITION 2.16. Let $h_1 \subset h_2$ be a nested pair of halfspaces. Consider the set of pairs of points in $h_1 \times h_2^*$ minimizing the distance between h_1 and h_2^* , that is,

$$M_{h_1, h_2} = \{(x, y) \in h_1 \times h_2^* : \text{if } (a, b) \in h_1 \times h_2^* \text{ then } d(x, y) \leq d(a, b)\}.$$

It will be convenient to denote by M_1 and M_2 the projections of M_{h_1, h_2} into h_1 and h_2^* , respectively.

The *combinatorial bridge* connecting h_1 and h_2^* is the union of intervals between such minimal distance pairs:

$$b(\hat{h}_1, \hat{h}_2) = \bigcup_{(x, y) \in M_{h_1, h_2}} \mathcal{I}(x, y).$$

In the following, we will drop the dependence on the hyperplanes whenever no confusion can arise.

We observe that if $(x_1, y_1), (x_2, y_2) \in M_{h_1, h_2}$, then $d(x_1, y_1) = d(x_2, y_2)$. The following lemma on the distance between a point and a halfspace, permeates several proofs to come. We denote by $\hat{\mathfrak{H}}(u, h)$ the hyperplanes separating u from h and define the distance of u from h to be

$$d(u, h) := \min\{d(u, v) \mid v \in h\}.$$

LEMMA 2.17. *Let $u \in X$ and let h be a halfspace so that $u \in h^*$. Then $d(u, h)$ equals the cardinality $|\hat{\mathfrak{H}}(u, h)|$.*

Proof. If a hyperplane separates u from h , it will have to be crossed by any combinatorial geodesic from u to v for any $v \in h$ and hence it will contribute to $d(u, v)$. It follows that $|\hat{\mathfrak{H}}(u, h)| \leq d(u, h)$.

Conversely, take $v \in h$ minimizing the distance to u and assume that a combinatorial geodesic from v to u crosses a hyperplane \hat{k} transverse to \hat{h} . Since \hat{k} and \hat{h} are not comparable, there

is a (perhaps different) combinatorial geodesic that crosses \hat{k} before crossing \hat{h} . Let v' be the point reached just after crossing \hat{k} , then $v' \in h$ because the geodesic has not crossed \hat{h} yet, and $d(u, v') < d(u, v)$. Since this is impossible by definition of v , the geodesic must cross at most all hyperplanes separating u from h . \square

The structure of the bridge is obtained in the following lemma.

LEMMA 2.18. *Let $h_1 \subset h_2$ be any pair of nested halfspaces.*

- (1) *If \hat{h} separates two points in M_i , $i = 1$ or 2 , then \hat{h} crosses \hat{h}_i .*
 - (2) *Let $(p_1, p_2) \in M_{h_1, h_2}$ and suppose that \hat{h} separates p_1 and p_2 . Then \hat{h} is parallel to both \hat{h}_1 and \hat{h}_2 .*
 - (3) *If a hyperplane \hat{h} separates any two points on the bridge and \hat{h} is transverse to either of the \hat{h}_i , with $i \in \{1, 2\}$, then \hat{h} is transverse to both the \hat{h}_i .*
 - (4) *For any $(p_1, p_2) \in M_{h_1, h_2}$, the distance $d(p_1, p_2)$ is exactly the number of hyperplanes separating h_1 from h_2^* , including \hat{h}_1 and \hat{h}_2 .*
 - (5) *The bridge $b(\hat{h}_1, \hat{h}_2)$ is isomorphic to a product and strongly convex.*
- More precisely, $b(\hat{h}_1, \hat{h}_2) \cong M_1 \times \mathcal{I}(p_1, p_2)$, where M_1 , the projection of M_{h_1, h_2} into h_1 , is strongly convex, and (p_1, p_2) is any pair in M_{h_1, h_2} .*

Before starting the proof, we make the general observation that if $p_1 \in M_1$, then no hyperplane \hat{h} can separate p_1 from \hat{h}_1 . In fact, if there is such a hyperplane, the geodesic joining p_1 to the point $p_2 \in M_2$ such that $(p_1, p_2) \in M_{h_1, h_2}$ would have to cross this hyperplane before crossing \hat{h}_1 , contradicting that (p_1, p_2) is a minimizing pair. The same argument holds of course for $p_2 \in M_2$.

We also establish the following easy claim.

CLAIM 2.19. *Let $p_1 \in M_1$ and $p_2 \in M_2$ be such that (p_1, p_2) minimizes the distance. Assume that there exists \hat{h} such that $\hat{h} \cap \hat{h}_1$, $p_1 \in h$ and $p_2 \in h^*$. Then either p_1 belongs to the cube identified by \hat{h} and \hat{h}_1 or there exists \hat{h}' such that $\hat{h}' \cap \hat{h}_1$, $p_1 \in h'$ and $\hat{h}' \subset h$.*

Proof of Claim. If p_1 does not belong to the cube determined by \hat{h} and \hat{h}_1 , then there exists a hyperplane \hat{h}' separating p_1 from \hat{h} . If \hat{h}' were not transverse to \hat{h}_1 , then \hat{h}' would be a hyperplane separating p_1 from \hat{h}_1 , which we observed is not possible. \square

Proof of Lemma 2.18. (1) For simplicity let us set $i = 1$ and let $p_1, p'_1 \in M_1$ be the points separated by \hat{h} . If \hat{h} is not transverse to \hat{h}_1 , then \hat{h} must separate, say, p_1 from \hat{h}_1 and we observed already that this is not possible.

(2) If \hat{h} were to cross \hat{h}_1 , then we could assume, by applying repeatedly the claim, that p_1 belongs to the cube identified by \hat{h} and \hat{h}_1 . Then, by crossing \hat{h} , one would still remain in \hat{h}_1 and reach a point closer to p_2 , contradicting the minimality of (p_1, p_2) .

(3) Let \hat{h} be a hyperplane that separates two points on the bridge and $\hat{h} \cap \hat{h}_1$. Let us assume that \hat{h} is parallel to \hat{h}_2 . Then, up to replacing h by h^* , $M_2 \subset h^*$. If it were also $M_1 \subset h^*$, then the interval between any element of M_1 and any element in M_2 would be contained in h and hence $b(\hat{h}_1, \hat{h}_2) \subset h$, which contradicts the assumption on \hat{h} . Therefore, $M_1 \cap h^* \neq \emptyset$. Let $p_1 \in M_1 \cap h^*$ and let $p_2 \in M_2$ such that $(p_1, p_2) \in M_{h_1, h_2}$. By construction, \hat{h} separates p_1 and p_2 and, hence, by (2), cannot be transverse to either \hat{h}_1 or \hat{h}_2 , contradicting the hypothesis.

(4) Let $(p_1, p_2) \in M_{h_1, h_2}$. Clearly, $d(p_1, p_2)$ is greater than or equal to the number of hyperplanes separating h_1 from h_2 . The other inequality is the assertion in (2).

DEFINITION 2.20. Let $b(\hat{h}_1, \hat{h}_2)$ be the bridge of the hyperplanes \hat{h}_1, \hat{h}_2 .

(1) The hyperplanes crossing both \hat{h}_1 and \hat{h}_2 are *horizontal hyperplanes* and are denoted by $\hat{\beta}_h$.

(2) The hyperplanes separating \hat{h}_1 and \hat{h}_2 are called *vertical hyperplanes* and are denoted by $\hat{\beta}_v$.

Continuation of the proof of Lemma 2.18. (5) From (3), we see that the hyperplanes of the bridge are either horizontal or vertical. Then any element of $\hat{\beta}_h$ crosses any element of $\hat{\beta}_v$ and vice versa; therefore, the bridge $b(\hat{h}_1, \hat{h}_2)$ is a product $X(\hat{\beta}_h) \times X(\hat{\beta}_v)$. Furthermore, by part (2) we have that $\mathcal{I}(p_1, p_2) \cong X(\hat{\beta}_v)$ for any $(p_1, p_2) \in M_{\hat{h}_1, \hat{h}_2}$.

To conclude, it remains to show that M_1 is strongly convex. First, we note that, by (2) and (3), if $\hat{\beta}_h$ is not empty, each element of $\hat{\beta}_h$ separates elements of M_1 and of M_2 . Now take $s_1, t_1 \in M_1$ and $u_1 \in \mathcal{I}(s_1, t_1)$. Let $s_2, t_2 \in M_2$ be the other end of the minimizing pairs for s_1, t_1 . Let $u_2 \in \mathcal{I}(s_2, t_2)$ be the element obtained by starting at s_2 and crossing the hyperplanes separating s_1 from t_1 . This is well defined because, by (1) and (3), the hyperplanes separating s_1 from t_1 are all in $\hat{\beta}_h$. Then only the hyperplanes separating h_1 from h_2^* can separate u_1 from u_2 . Hence the pair $(u_1, u_2) \in M_{\hat{h}_1, \hat{h}_2}$, so that $\mathcal{I}(s_1, t_1) \subset M_1$. Hence, M_1 is strongly convex and thus is $b(\hat{h}_1, \hat{h}_2) = M_1 \times \mathcal{I}(p_1, p_2)$. \square

Although we will not need it, we observe that, since the bridge is strongly convex, if $H := \{h \in \mathfrak{H} : b(\hat{h}_1, \hat{h}_2) \subset h\}$, then $\mathfrak{H} = (\beta_h \sqcup \beta_v) \sqcup (H \sqcup H^*)$ is a lifting decomposition of the hyperplanes associated to the bridge and hence the bridge is isometrically embedded in X as a product.

In view of Lemma 2.18 the following is well defined.

DEFINITION 2.21. Let $h_1 \subset h_2$ be a nested pair of halfspaces. The *length* $\ell(b(\hat{h}_1, \hat{h}_2))$ of the bridge of \hat{h}_1 and \hat{h}_2 is the cardinality of the set β_v .

We will adopt the usual abuse of terminology and refer to *horizontal halfspaces* (respectively, *vertical halfspaces*) as the halfspaces corresponding to the horizontal (respectively, vertical) hyperplanes, and denote them by β_h (respectively, β_v).

It is straightforward to see that if \hat{h}_1 and \hat{h}_2 are strongly separated, then the corresponding set of horizontal halfspaces is empty [3, Lemma 2.2].

The following lemma is probably well known, but we include it here because we could not find a reference for it.

LEMMA 2.22. Let X be a CAT(0) cube complex, $x \in X$ and $B \subseteq X$ be a strongly convex subset. There is a unique point $p_B(x) \in B$ minimizing the combinatorial distance between x and B .

Note that this lemma is standard in the case of the CAT(0) distance, see, for instance, [5]. In the case of a CAT(0) space, however, the proof of the existence of an orthogonal projection is a bit more difficult than the proof of its uniqueness.

Proof. Since the combinatorial distance takes discrete values, the existence of a point in B minimizing the distance is obvious. To prove uniqueness, let y and y' be two points in B minimizing the distance between B and x . Let us show that $\hat{\mathfrak{H}}(x, y) \subseteq \hat{\mathfrak{H}}(x, y')$, where $\hat{\mathfrak{H}}(x, y)$ is the collection of hyperplanes separating x from y : If a hyperplane $\hat{h}_0 \in \hat{\mathfrak{H}}(x, y)$ does not separate x from y' , then it has to belong to $\hat{\mathfrak{H}}(y, y')$, and so do the hyperplanes $\hat{h}_1, \dots, \hat{h}_s$ separating \hat{h}_0 from y . Let $p \in \mathcal{I}(x, y)$ obtained by starting at y and crossing the hyperplanes

$\hat{h}_s, \dots, \hat{h}_0$. Then p also belongs to $\mathcal{I}(y, y') \subseteq B$; but $d(x, p) < d(x, y)$ because p is also on a geodesic from x to y , contradicting that y was distance minimizing. So $\hat{\mathfrak{H}}(x, y) = \hat{\mathfrak{H}}(x, y')$. By Lemma 2.2, we deduce that $y = y'$. \square

We can hence give the following definition.

DEFINITION 2.23. For $x \in X$ and B a strongly convex subset in X , denote by $p_B(x) \in B$ the projection of x on B .

LEMMA 2.24. Let $h_1 \subset h_2$ be a pair of nested halfspaces, $x_1 \in h_1$ and $x_2 \in h_2^*$. Denote by b the bridge connecting h_1 and h_2^* . Then the following conditions are satisfied.

- (1) A horizontal hyperplane of the bridge cannot separate x_i from $p_b(x_i)$ for $i = 1, 2$.
- (2) The following holds true:

$$d(x_1, x_2) = d(x_1, p_b(x_1)) + d(p_b(x_1), p_b(x_2)) + d(p_b(x_2), x_2).$$

Proof. (1) Observe first of all that $p_b(x_1) \in b(\hat{h}_1, \hat{h}_2) \cap h_1$. Now let \hat{h} be a horizontal hyperplane of the bridge. If \hat{h} separates x_1 from $p_b(x_1)$, say $x_1 \in h$ and $p_b(x_1) \in h^*$, then there is a point in $b(\hat{h}_1, \hat{h}_2) \cap h_1$ different from $p_b(x_1)$ and at distance from x_1 smaller than $d(x_1, p_b(x_1))$, contradicting that $p_b(x_1)$ is the projection of x_1 on $b(\hat{h}_1, \hat{h}_2)$.

(2) That $d(x_1, x_2) \leq d(x_1, p_b(x_1)) + d(p_b(x_1), p_b(x_2)) + d(p_b(x_2), x_2)$ follows from the triangle inequality, so let us show the other inequality, which we do by showing that $\hat{\mathfrak{H}}(x_1, p_b(x_1)) \cup \hat{\mathfrak{H}}(p_b(x_1), p_b(x_2)) \cup \hat{\mathfrak{H}}(p_b(x_2), x_2) \subset \hat{\mathfrak{H}}(x_1, x_2)$.

A hyperplane \hat{h} separating $p_b(x_1)$ from $p_b(x_2)$ cuts the bridge and hence by Lemma 2.18(3) is either vertical or horizontal. If it is vertical, it separates h_1 from h_2^* and hence x_1 from x_2 as well. If \hat{h} is horizontal, then it cannot separate x_i from $p_i(x_i)$ by part (1) of this lemma. Since \hat{h} is separating $p_b(x_1)$ from $p_b(x_2)$, this forces \hat{h} to separate x_1 from x_2 .

By part (1) a hyperplane \hat{h} separating x_i from $p_b(x_i)$, for $i = 1$ or 2 , cannot be horizontal; by Lemma 2.18 it cannot cross the bridge, so it has to separate x_1 from x_2 . \square

2.H. Finite orbits in the CAT(0) boundary versus finite orbits in the Roller boundary

Non-elementarity of the action is defined in terms of the non-existence of a finite orbit in the CAT(0) boundary. We describe in this section to what extent this is equivalent to the same property with respect to the Roller boundary.

We start with one direction of the equivalence that is very easy and is here for completeness, since it will not be needed in the following.

PROPOSITION 2.25. Let Y be a finite-dimensional CAT(0) cube complex and let $\Gamma \rightarrow \text{Aut}(Y)$ be an action on Y . If the action is essential and there is a finite orbit in the Roller boundary, then there is a finite orbit in the CAT(0) boundary.

Proof. Let $\xi \in \partial Y$ be a point in the finite Γ -orbit and let Γ_ξ be its stabilizer, whose action is still essential since it is of finite index in Γ . Let $h \in \mathfrak{H}$ be a halfspace containing ξ . If there were no Γ_ξ -fixed point in $\partial_{\leq} Y$, then we could apply the Flipping Lemma (see Subsection 2.F); hence there would exist $\gamma \in \Gamma_\xi$ that flips h^* , so that $\gamma h \subset h^*$. But this would contradict the fact that $\xi = \gamma\xi \in \gamma h$. \square

The following proposition pins down to what extent an elementary action implies the existence of a finite orbit in the Roller boundary.

PROPOSITION 2.26. *Let X be a finite-dimensional CAT(0) cube complex and let $\Gamma \curvearrowright \text{Aut}(X)$ be an action on X . If there is a finite orbit in the CAT(0) boundary, then either*

- (1) *there is a finite Γ -orbit in the Roller boundary, or*
- (2) *there exists a subgroup of finite index $\Gamma' < \Gamma$ and a Γ' -invariant subcomplex $X' \hookrightarrow \partial X$ on which the Γ' -action is non-elementary. Moreover, X' corresponds to a lifting decomposition of halfspaces.*

The argument will depend on the following lemma, which we assume now, and whose verification we defer to right after the proof of the proposition. We start with the following natural construction, that can also be found in [28, Subsection 4.1]. Note that there is a similar construction in [19, Section 3].

Let $\xi \in \partial_{\triangleleft} X$, let $g : [0, \infty) \rightarrow X$ be a geodesic asymptotic to ξ and let us define T_ξ to be the following set of halfspaces:

$$T_\xi := \{h \in \mathfrak{H} : \text{for every } \epsilon > 0 \text{ there exists } t_\epsilon \geq 0 \text{ such that } N_\epsilon(g(t_\epsilon, \infty)) \subset h\}, \quad (2.8)$$

where $N_\epsilon(g(t_\epsilon, \infty))$ is the ϵ -neighborhood of the image of the geodesic ray $g|_{(t_\epsilon, \infty)}$.

LEMMA 2.27. *Let X be a CAT(0) cube complex with a Γ -action and let $\xi \in \partial_{\triangleleft} X$. The set T_ξ in (2.8) satisfies the following properties.*

- (1) *It is independent of the geodesic g and Γ_ξ -invariant, where $\Gamma_\xi < \Gamma$ is the stabilizer of ξ in Γ .*
- (2) *It is not empty.*
- (3) *It satisfies the partial choice and consistency conditions.*
- (4) *It contains an infinite descending chain.*

Proof of Proposition 2.26. Let $\xi \in \partial_{\triangleleft} X$ be one of the points in the finite Γ -orbit, and let $\Gamma_\xi < \Gamma$ be the stabilizer of ξ , $[\Gamma : \Gamma_\xi] < \infty$. It follows from the above Lemma 2.27 and from Subsection 2.C that T_ξ induces a lifting decomposition

$$\mathfrak{H}(X) = \mathfrak{H}_\xi \sqcup (T_\xi \sqcup T_\xi^*) \quad (2.9)$$

and there is a Γ_ξ -equivariant projection $\rho : \overline{X} \rightarrow \overline{X}$ whose image is the isometrically embedded Γ_ξ -invariant subcomplex $\overline{X}_\xi := i(\overline{X}(\mathfrak{H}_\xi))$. Observe that because of Lemmas 2.27(4) and 2.6, $\dim X_\xi < \dim X$. Moreover, if $\mathfrak{H}_\xi = \emptyset$, then $\overline{X}_\xi = X_\xi$ is a Γ_ξ -fixed point in ∂X .

Proceeding inductively, we can conclude the proof. In fact, if the Γ_ξ -action on X_ξ is non-elementary, then we are in case (2) of the proposition. If, on the other hand, there is a finite Γ_ξ -orbit in ∂X_ξ , using the fact that the composition of two projections is a projection, we can repeat the argument. The finite dimensionality of X ensures that the process terminates. \square

Proof of Lemma 2.27. (1) Only for this part of the proof we denote by $T_\xi(g)$ and by $T_\xi(g')$ the sets defined in (2.8) with respect to two asymptotic geodesics g and g' . Then g and g' are at bounded distance from each other, that is, there exists an $r > 0$ such that $g'([0, \infty)) \subset N_r(g[0, \infty))$. By the triangle inequality, if $\epsilon > 0$, then $N_\epsilon(g'[t, \infty)) \subset N_\epsilon(g[t, \infty))$ for all $t \geq 0$. But this implies that if $h \in \mathfrak{H}$ is such that there exists $t_{\epsilon+r}$ with $N_\epsilon(g'(t_{\epsilon+r}, \infty)) \subset h$, then also $N_\epsilon(g(t_{\epsilon+r}, \infty)) \subset h$. Thus $T_\xi(g) = T_\xi(g')$ and therefore $T_\xi(g)$ is independent of g .

Since ξ is Γ_ξ -invariant, the geodesics g and γg are asymptotic and hence $T_\xi(g) = T_\xi(\gamma g)$. The Γ_ξ -invariance of T_ξ then follows at once, since $\gamma N_\epsilon(g(t, \infty)) = N_\epsilon(\gamma g(t, \infty))$ for all $t \geq 0$.

- (2) We will show that if, for every, $h \in \mathfrak{H}$ there exists $\epsilon_h > 0$ such that

$$N_{\epsilon_h}(g(t, \infty)) \cap h \neq \emptyset \quad \text{and} \quad N_{\epsilon_h}(g(t, \infty)) \cap h^* \neq \emptyset \quad (2.10)$$

for all $t > 0$, then there is an infinite family of pairwise transverse hyperplanes.

We may assume that $h \in \mathfrak{H}$ is not compact; otherwise (2.10) is never verified for any ϵ . Moreover, observe that any geodesic γ crosses infinitely many hyperplanes. Order the cubes according to the order in which they are crossed by γ . This gives rise, up to choosing an order of hyperplanes on each of these cubes, to an order $\widehat{h}_1, \widehat{h}_2, \dots$ on the hyperplanes according to the order in which they are crossed by γ . If $\epsilon_i > 0$ is the smallest ϵ such that (2.10) is verified for \widehat{h}_i , then any $\gamma(t)$ is at CAT(0) distance at most ϵ_i from \widehat{h}_i .

We claim that, for every \widehat{h}_ℓ , there exists $n_\ell \in \mathbb{N}$ such that, for every $j \geq n_\ell$, the hyperplane \widehat{h}_j is transverse to \widehat{h}_ℓ . In fact, the CAT(0) distance is quasi-isometric to the combinatorial distance and the combinatorial distance between $\gamma(t)$ and \widehat{h}_ℓ is the number of hyperplanes parallel to \widehat{h}_ℓ that separate \widehat{h}_ℓ from $\gamma(t)$. On the other hand, for every $j > n_\ell$, any hyperplane \widehat{h}_j parallel to \widehat{h}_ℓ that intersects γ will contribute to the distance from $\gamma(t)$ and \widehat{h}_ℓ for t large enough. By the previous observation this is not possible and hence eventually \widehat{h}_j must intersect \widehat{h}_ℓ .

By setting $m_{\ell+1} := n_{n_\ell}$, for every $d \in \mathbb{N}$ the hyperplanes $\widehat{h}_1, \widehat{h}_{m_1}, \widehat{h}_{m_2}, \dots, \widehat{h}_{m_d}$ form a family of $d + 1$ pairwise transverse hyperplanes.

(3) This is obvious from the construction.

(4) If T_ξ does not contain an infinite descending chain, then $\overline{X}_\xi \cap X \neq \emptyset$, where \overline{X}_ξ is the complex associated to the lifting decomposition in (2.9). Because of Lemma 2.6, $\xi \in \partial_{\triangleleft} X_\xi$. We can hence apply the construction in the beginning of the proof of Proposition 2.26 to the complex X_ξ , whose halfspaces are now $\mathfrak{H} \setminus (T_\xi \sqcup T_\xi^*)$, thus contradicting (2). \square

2.1. From products to irreducible essential factors

The following lemma identifies the important properties that are passed down from a complex to the irreducible factors of the essential core. The content of the lemma is already in [20], but we recall it here in the form in which we will need it.

LEMMA 2.28. *Let X be a finite-dimensional CAT(0) cube complex and let $\Gamma \rightarrow \text{Aut}(X)$ be a non-elementary action. Then the Γ_0 -action on the irreducible factors of the essential core is also non-elementary and essential, where Γ_0 is the finite index subgroup preserving this decomposition.*

Proof. Let $Y \subset X$ be the essential core, $Y = Y_1 \times \dots \times Y_m$ be its decomposition into irreducible factors, and Γ_0 be the finite index subgroup preserving this decomposition. We need to show that the following hold.

- (1) The action of Γ_0 on the Y_j , $j = 1, \dots, m$ is essential as well.
- (2) If Γ_0 has no finite orbit on the visual boundary $\partial_{\triangleleft} X$, then the same holds for the action on $\partial_{\triangleleft} Y_j$, $j = 1, \dots, m$.

(1) By [20, Proposition 3.2], the Γ_0 -action on Y (respectively, on Y_i) is essential if and only if every hyperplane $\widehat{h} \in \widehat{\mathfrak{H}}$ (respectively, $\widehat{h}_i \in \widehat{\mathfrak{H}}_j$) can be skewed by some element in Γ_0 . If $\widehat{h}_j \in \widehat{\mathfrak{H}}(Y_j)$ is a hyperplane in Y_j , then $\widehat{h} := Y_1 \times \dots \times Y_{j-1} \times \widehat{h}_j \times Y_{j+1} \times \dots \times Y_m$ is a hyperplane in Y . Since the action on Y is essential, there exists $\gamma \in \Gamma_0$ that skews \widehat{h} and hence it skews \widehat{h}_j . Then the Γ_0 -action on Y_j is essential.

(2) We prove the contrapositive of the statement. Let $\Gamma_0 < \Gamma$ be the finite subgroup that preserves each of the factors Y_j and let us assume, by passing if necessary to a further subgroup of finite index, that there is a Γ_0 -fixed point in $\partial_{\triangleleft} Y_j$ for some $1 \leq j \leq m$. Then, by Lemma 2.9 there is a Γ_0 -fixed point in $\partial_{\triangleleft} Y$ and hence a finite Γ -orbit in $\partial_{\triangleleft} Y$. Since Y is a convex subset of X and hence $\partial_{\triangleleft} Y \subset \partial_{\triangleleft} X$, there is a finite Γ -orbit in $\partial_{\triangleleft} X$. \square

2.J. Euclidean (sub)complexes

DEFINITION 2.29. Let X be a CAT(0) cube complex. We say that X is *Euclidean* if the vertex set with the combinatorial metric embeds isometrically in \mathbb{R}^D with the ℓ^1 -metric for some $D < \infty$.

In [20, Theorem 7.2], under some natural conditions on the action of $\text{Aut}(X)$, the authors relate the existence of an $\text{Aut}(X)$ -invariant *Euclidean flat* with the non-existence of a facing triple of halfspaces, in the following sense.

DEFINITION 2.30. Let $n \in \mathbb{N}$. An n -tuple of halfspaces is called a *facing n -tuple* if they are pairwise disjoint. An n -tuple of hyperplanes is called a *facing n -tuple* if there is a choice of halfspaces forming a facing n -tuple.

As our setting differs slightly from the one used in [20], we discuss briefly in this section the notion of Euclidean complexes and subcomplexes. The following definition is from [20].

DEFINITION 2.31. A CAT(0) cube complex X is said to be \mathbb{R} -like if there is an $\text{Aut}(X)$ -invariant bi-infinite CAT(0) geodesic.

PROPOSITION 2.32. Let Y be a CAT(0) cube complex on which $\text{Aut}(Y)$ acts essentially. Consider the following statements:

- (1) Y is Euclidean;
- (2) Y is an interval;
- (3) Y is a product of \mathbb{R} -like factors.

Then (3) \Rightarrow (2) \Rightarrow (1).

Proof. Observe that conditions (1) and (2) are preserved under taking products. Also, the hypothesis of having an essential action is preserved by passing to the irreducible factors by Lemma 2.28. Therefore, it is sufficient to consider the case in which Y is irreducible.

(3) \Rightarrow (2). Assume that Y is \mathbb{R} -like. Let $\ell \subset Y$ be the $\text{Aut}(Y)$ -invariant CAT(0) geodesic. We claim that ℓ crosses every hyperplane of Y . In fact, otherwise there would be a halfspace h_0 containing ℓ and, since ℓ is $\text{Aut}(Y)$ -invariant, then \hat{h}_0 would not be essential.

Let $\ell : \mathbb{R} \rightarrow Y$ be a parameterization of ℓ . One can check that, because of the above claim, the collection of halfspaces

$$\alpha := \{h \in \mathfrak{H}(Y) : \text{there exists } t \in \mathbb{R} \text{ such that } h \supset \ell(t, \infty)\}$$

defines a non-terminating ultrafilter. Then Y is an interval on α and its opposite ultrafilter $\alpha^* = \mathfrak{H} \setminus \alpha$.

(2) \Rightarrow (1) This is Lemma 2.3. □

We prove next that, under the assumption that there are no fixed points in the visual boundary and the action is essential, being Euclidean is equivalent to the non-existence of facing triples of hyperplanes. (It is possible that a Euclidean CAT(0) cube complex Y on which $\text{Aut}(Y)$ acts essentially and without fixed points in the visual boundary, is a point (cf. [20, Theorem E]).) As a byproduct, using [20, Theorem 7.2] we can conclude that also (1) implies (3) under the above hypotheses. We start with the following easy lemma.

LEMMA 2.33. *If X is a Euclidean CAT(0) cube complex that isometrically embeds into \mathbb{R}^D , then any set of pairwise facing halfspaces has cardinality at most $2D$.*

Proof. Indeed, any collection of halfspaces can be arranged in at most D chains. Hence, for each dimension there can be at most one pair of facing halfspaces and the assertion follows from the fact that the ℓ^1 -metric on \mathbb{R}^D is the sum of the ℓ^1 -metrics on its factors. \square

More precisely, we have the following dichotomy that is compatible with the terminology in [20] but holds also in the case in which the CAT(0) cube complex does not have a cocompact group of automorphisms.

COROLLARY 2.34. *Let Y be a finite-dimensional irreducible CAT(0) cube complex and assume that $\text{Aut}(Y)$ acts essentially and without fixed points on $\partial_{\triangleleft} Y$. Then Y is Euclidean if and only if $\mathfrak{H}(Y)$ does not contain a facing triple of halfspaces.*

Proof. We first prove that if Y is Euclidean, then there are no facing triples of halfspaces. Since the action is essential and there are no fixed points in $\partial_{\triangleleft} Y$, if there is a facing triple of halfspaces, we can skewer several times two of the halfspaces into the third one to obtain a set of pairwise facing halfspaces of arbitrarily large cardinality. Then Lemma 2.33 implies that Y is not Euclidean.

Conversely, we assume that there are no facing triples of hyperplanes and prove that Y must be Euclidean. Since Y is irreducible, let $\{h_n\}$ be a descending sequence of strongly separated halfspaces, $h_{n+1} \subset h_n$. The strategy of the proof consists in showing that $\bigcap h_n$ consists of a single point $\alpha \in \partial Y$ and in using the non-existence of facing triples of hyperplanes to show that α^* is also an ultrafilter. Then Remark 2.4 will complete the proof.

To show that $\bigcap h_n$ is a single point, let us assume by contradiction that $\bigcap h_n$ contains at least two distinct points, $u, v \in \bigcap h_n$. Let \hat{h} be a hyperplane that separates them. Observe that, for every $n \in \mathbb{N}$,

$$\begin{aligned} u &\in h \cap h_n \neq \emptyset \quad \text{and} \\ v &\in h^* \cap h_n \neq \emptyset. \end{aligned} \tag{2.11}$$

From this and the fact that the h_n are a descending chain, one can check that if there exists $N \in \mathbb{N}$ such that $\hat{h} \pitchfork h_N$, then $\hat{h} \pitchfork h_n$ for all $n \geq N$, which is impossible since the $\{h_n\}$ are pairwise strongly separated. So $\hat{h} \parallel \hat{h}_n$ for every $n \in \mathbb{N}$.

Again from (2.11) it follows that $\hat{h} \subset h_n$ for all $n \in \mathbb{N}$. But this is also not possible since there exist finitely many hyperplanes between \hat{h} and \hat{h}_n . Hence $\alpha := \bigcap h_n$ is a single point.

To see that α^* is an ultrafilter, we need only to check the consistency condition, namely, that if $h \in \alpha^*$ and $h \subset k$, then $k \in \alpha^*$. Observe that this is equivalent to verifying that if $\alpha \in h^*$ and $h \subset k$, then $\alpha \in k^*$. Suppose that this is not the case, that is, that there exists $h, k \in \mathfrak{H}(Y)$ such that $h \subset k$ and $\alpha \in h^* \cap k$.

We first claim that

$$\text{there exists } n_0 \in \mathbb{N} \text{ such that } h_n \subset k \text{ for all } n > n_0. \tag{2.12}$$

In fact, suppose that there exists $n'_0 \in \mathbb{N}$ such that $\hat{h}_{n'_0} \pitchfork \hat{k}$. Since the $\{h_n\}$ are pairwise strongly separated, then $\hat{h}_n \parallel \hat{k}$ for all $n > n'_0$. Using (2.2), the fact that the $\{h_n\}$ are a descending chain and that $\alpha \in h_n$ for all $n \in \mathbb{N}$, it is easy to verify that $h_n \subset k$ for all $n > n'_0$.

On the other hand, if $\hat{h}_n \parallel \hat{k}$ for all $n \in \mathbb{N}$, using again that the $\{h_n\}$ are a descending chain and that there are only finitely many hyperplanes between any \hat{h}_n and \hat{k} , one can easily verify that there exists $n''_0 \in \mathbb{N}$ such that $h_n \subset k$ for all $n > n''_0$. Hence (2.12) is verified with $n_0 = \max\{n'_0, n''_0\}$.

Since $h \subset k$ and there are only finitely many hyperplanes between \hat{h} and \hat{k} , there exists $n_1 \geq n_0$ such that either $\hat{h}_{n_1} \pitchfork \hat{h}$ or $h_{n_1} \subset h$. But h_{n_1} cannot be contained in h since $\alpha \in h_{n_1} \cap h^*$, hence $\hat{h}_{n_1} \pitchfork \hat{h}$.

Again because the hyperplanes $\{\hat{h}_n\}$ are strongly separated, if $n > n_1$, then $\hat{h}_n \parallel \hat{h}$. This, the fact that $\hat{h}_{n_1} \pitchfork \hat{h}$ and that $h_n \subset h_{n_1}$ imply that $h_n \cap h = \emptyset$.

It follows that h_n^*, h^* and k is a facing triple of halfspaces, contradicting the hypothesis. Hence α^* is an ultrafilter and the proof is complete. \square

We conclude the section with the following corollary that will be paramount in the sequel.

COROLLARY 2.35. *Let X be a finite-dimensional CAT(0) cube complex and $\Gamma \rightarrow \text{Aut}(X)$ be a non-elementary action. Then there are no Euclidean factors in the essential core.*

Proof. Let $Y \subset X$ be the essential core of the Γ -action and let Y_0 be an irreducible factor of Y . By Lemma 2.28, the Γ -action on Y_0 is also essential and non-elementary. By Corollary 2.34, Y_0 cannot be Euclidean. \square

2.K. Facing triples of halfspaces

In this section, we show how the hypotheses of non-elementarity and essentiality of the action are used to construct a suitable facing triple of hyperplanes.

DEFINITION 2.36. A facing n -tuple of halfspaces is a *facing über-separated (or parallel) n -tuple* if all the pairs of halfspaces are über-separated (or parallel) pairs.

As usual we extend the above definition to hyperplanes in the obvious way. We will need the following lemma only in the case $n = 3$, but the extension to larger n is very easy.

LEMMA 2.37. *Let X be a CAT(0) cube complex with a non-elementary action $\Gamma \rightarrow \text{Aut}(X)$ and $n \in \mathbb{N}$. Then any essential halfspace $h \in \mathfrak{H}(X)$ belongs to a facing über-parallel n -tuple all of whose halfspaces can be taken to be in a single Γ -orbit.*

Proof. First, we assume that X is irreducible and essential. We show the existence of a facing über-separated n -tuple. According to Corollary 2.34, since X is non-Euclidean, it contains a facing triple of halfspaces, call it a, b, c . Using Lemma 2.14, we find $\gamma_1, \gamma_2, \gamma_3 \in \Gamma$ such that $\gamma_1 a \subset a$, $\gamma_2 b \subset b$ and $\gamma_3 c \subset c$ are über-separated pairs. Hence the triple $\gamma_1 a, \gamma_2 b$ and $\gamma_3 c$ is facing and über-separated. To get a facing n -tuple out of a facing $(n-1)$ -tuple h_1, \dots, h_{n-1} , we flip and skewer two elements of the $(n-1)$ -tuple into a third one; for example, we flip and skewer h_1 and h_{n-2} into h_{n-1} via $\gamma_1 \gamma_2 \in \text{Aut}(X)$, and now the n -tuple $h_1, h_2, \dots, \gamma_1 h_1, \gamma_2 h_2$ will be über-separated.

To get a facing über-separated n -tuple in an orbit, take h , and any facing über-separated $(n+1)$ -tuple of halfspaces. Then h crosses at most one element of this facing $(n+1)$ -tuple. Skewer and flip h into the n other halfspaces to get a facing über-separated n -tuple of halfspaces in the orbit of h .

Let $\Gamma_0 < \Gamma$ be a finite index subgroup preserving each irreducible factor of the essential core Y of X . Note that the hypotheses that the action is non-elementary and essential are preserved up to passing to Γ_0 . One then deduces the general case where X is not necessarily irreducible and essential by using that any essential halfspace belongs to an irreducible factor of Y . We find the über-separated n -tuple in that irreducible factor of the essential core, and use it to produce an über-parallel n -tuple on X . \square

3. Construction and boundedness of the median class

Let Γ be a group and E be a coefficient Γ -module, that is, the dual of a separable Banach space on which Γ acts by linear isometries. The bounded cohomology of Γ with coefficients in E is the cohomology of the subcomplex of Γ -invariants in $(C_b(\Gamma^{k+1}, E), d)$, where

$$C_b(\Gamma^k, E) := \left\{ f : \Gamma^k \rightarrow E : \sup_{g \in \Gamma^k} \|f(g)\|_E < \infty \right\}, \tag{3.1}$$

is endowed with the Γ -action

$$(gf)(g_1, \dots, g_k) := g \cdot f(g^{-1}g_1, \dots, g^{-1}g_k),$$

and

$$d : C_b(\Gamma^k, E) \longrightarrow C_b(\Gamma^{k+1}, E)$$

is the usual homogeneous coboundary operator defined by

$$df(g_0, \dots, g_k) := \sum_{j=0}^k (-1)^j f(g_0, \dots, g_{j-1}, g_{j+1}, \dots, g_k).$$

3.A. The median cocycle

Let X be an irreducible finite-dimensional CAT(0) cube complex. Recall that \overline{X} denotes the Roller compactification of X , that is, the set of ultrafilters on $\mathfrak{H}(X)$ (see Section 2). For $n \geq 2$, we denote by $\mathfrak{H}(X)^n$ the set of n -tuples of halfspaces of X .

If $1 \leq p < \infty$, then $\ell^p(\mathfrak{H}(X)^n)$ is the dual of a separable Banach space. In fact, if $1 < p < \infty$, then $\ell^p(\mathfrak{H}(X)^n)$ is the dual of $\ell^q(\mathfrak{H}(X)^n)$, where $1/p + 1/q = 1$. On the other hand, $\ell^1(\mathfrak{H}(X)^n)$ is the dual of the Banach space $C_0(\mathfrak{H}(X)^n)$ of functions on $\mathfrak{H}(X)^n$ that vanish at infinity, which is separable since $\mathfrak{H}(X)^n$ is countable. For further use, we set the notation

$$\mathcal{E}_p := \begin{cases} \ell^q(\mathfrak{H}(X)^n) & 1 < p < \infty \text{ and } 1/p + 1/q = 1, \\ C_0(\mathfrak{H}(X)^n) & p = 1. \end{cases} \tag{3.2}$$

For each $1 \leq p < \infty$ and each integer $n \geq 1$, we define in this section a 1-parameter family of cocycles

$$c_{(n,R)} : \overline{X} \times \overline{X} \times \overline{X} \longrightarrow \ell^p(\mathfrak{H}(X)^n),$$

that, by evaluation on a basepoint in \overline{X} , will give a cocycle on $\Gamma \times \Gamma \times \Gamma$. We define the *median cocycle* $c_{(n,R)}$ as the coboundary of an $\text{Aut}(X)$ -invariant map $\omega_{(n,R)}$ on $\overline{X} \times \overline{X}$ whose values are not in general p -summable and we will show that, on the other hand, $c_{(n,R)} = d\omega_{(n,R)}$ is bounded in the sense of (3.1) if $n \geq 2$. For $n \geq 2$, the *median class* $m_{(n,R)}$ will be defined as the cohomology class of $c_{(n,R)}$ (which is independent of the basepoint).

If X is irreducible with an essential and non-elementary Γ -action, then the collection of sequences of length n of über-separated pairs at consecutive distance less than R is non-empty for R sufficiently large. Indeed, according to Caprace–Sageev [20] since X is irreducible and non-elementary, it contains a strongly separated pair, and by essentiality we can repeatedly skewer this pair to get an über-separated and nested n -tuple for any $n \in \mathbb{N}$. In the general case one can always find über-parallel sequences in the essential core of the action, and extend those to the whole space. We hence define $[[u, v]]_n$ for $u, v \in \overline{X}$ to be the collection of pairwise über-parallel n -tuples $(h_1, \dots, h_n) \in \mathfrak{H}(X)^n$ such that $h_1 \subset \dots \subset h_n$ and $h_i \in v \setminus u$ for each i .

For $R > 0$, we also define

$$[[u, v]]_{(n,R)} = \{(h_1, \dots, h_n) \in [[u, v]]_n : d(h_i, h_{i+1}) \leq R\}. \tag{3.3}$$

So, $[[u, v]]_{(n, R)}$ is the collection of sequences of length n of nested über-parallel halfspaces containing v and not u and at consecutive distance less than or equal to R . We hope that the notation suggests that these are in some sense subintervals.

For $u, v \in \bar{X}$, let us define

$$\omega_{(n, R)}(u, v) := \mathbb{1}_{[[u, v]]_{(n, R)}} - \mathbb{1}_{[[v, u]]_{(n, R)}}. \quad (3.4)$$

We will simply write c , ω and $[[u, v]]$ for $c_{(n, R)}$, $\omega_{(n, R)}$ and $[[u, v]]_{(n, R)}$ when the context is clear.

Fixing $u, v \in \bar{X}$, the function has finitely many values

$$\omega(u, v) : \mathfrak{H}(X)^n \longrightarrow \{-1, 0, 1\}$$

and is finitely supported when $u, v \in X$.

Note that ω is not necessarily bounded when thought of as a function with values in $\ell^p(\mathfrak{H}(X)^n)$ and, in fact, its norm is proportional to the distance between u and v .

Let us now consider the $\text{Aut}(X)$ -equivariant cocycle taking values in the functions on $\mathfrak{H}(X)^n$, defined as

$$\begin{aligned} c(u_1, u_2, u_3) &:= (d\omega)(u_1, u_2, u_3) \\ &= \omega(u_2, u_3) - \omega(u_1, u_3) + \omega(u_1, u_2) = \omega(u_2, u_3) + \omega(u_3, u_1) + \omega(u_1, u_2) \\ &= \mathbb{1}_{[[u_2, u_3]]} + \mathbb{1}_{[[u_3, u_1]]} + \mathbb{1}_{[[u_1, u_2]]} - (\mathbb{1}_{[[u_3, u_2]]} + \mathbb{1}_{[[u_1, u_3]]} + \mathbb{1}_{[[u_2, u_1]]}). \end{aligned} \quad (3.5)$$

We will show that, contrary to ω , the cocycle c on \bar{X} actually takes values in $\ell^p(\mathfrak{H}(X)^n)$ (Proposition 3.4) and is bounded in the sense of (3.1).

REMARK 3.1. Let $Y \subset X$ be the essential core of the Γ -action on X and $Y = Y_1 \times \cdots \times Y_m$ be the decomposition of Y into irreducible CAT(0) cube complexes. From the decomposition in (2.7), we have a corresponding decomposition

$$\ell^p(\mathfrak{H}(X)^n) \cong \ell^p(\mathfrak{H}(Y_1)^n) \oplus \cdots \oplus \ell^p(\mathfrak{H}(Y_m)^n) \oplus \ell^p(\mathfrak{H}_{nEss}(X)^n) \quad (3.6)$$

given by $f \mapsto \mathbb{1}_{\mathfrak{H}(Y_1)^n} f + \cdots + \mathbb{1}_{\mathfrak{H}(Y_m)^n} f + \mathbb{1}_{\mathfrak{H}_{nEss}(X)^n} f$, where the direct sum is in the ℓ^p sense. The direct summand $\ell^p(\mathfrak{H}(Y_j)^n)$ is invariant for the action of a finite index subgroup $\Gamma' < \Gamma$.

PROPOSITION 3.2. *Let Y be an essential CAT(0) cube complex and consider the cocycle defined in (3.5)*

$$c_{(n, R)} : \bar{Y} \times \bar{Y} \times \bar{Y} \longrightarrow \ell^p(\mathfrak{H}(Y)^n),$$

where R is chosen to be large enough so that in each irreducible component Y_i of Y , the set of über-separated n -tuples at consecutive distance less than or equal to R is not empty. Then $c_{(n, R)}$ decomposes as

$$c_{(n, R)}(\xi, \eta, \zeta) = c_{(n, R)}^1(\pi_1(\xi), \pi_1(\eta), \pi_1(\zeta)) \oplus \cdots \oplus c_{(n, R)}^m(\pi_m(\xi), \pi_m(\eta), \pi_m(\zeta)),$$

where

$$c_{(n, R)}^j : \bar{Y}_j \times \bar{Y}_j \times \bar{Y}_j \longrightarrow \ell^p(\mathfrak{H}(Y_j)^n)$$

is the cocycle on the irreducible factors and $\pi_j : \bar{Y} \rightarrow \bar{Y}_j$ is the projection. Moreover, $c_{(n, R)}(\xi, \eta, \zeta) \neq 0$ if and only if $c_{(n, R)}^j(\pi_j(\xi), \pi_j(\eta), \pi_j(\zeta)) \neq 0$ for some $1 \leq j \leq m$.

Proof. Let ω and ω_j , for $j = 1, \dots, m$, be defined as in (3.4), respectively, on \bar{Y} and \bar{Y}_j . Since $c_{(n, R)}^j = d\omega_{(n, R)}^j$ for $1 \leq j \leq m$ and $c_{(n, R)} = d\omega_{(n, R)}$, it is enough to verify that

$$\omega_{(n, R)} = \omega_{(n, R)}^1 + \cdots + \omega_{(n, R)}^m.$$

Let $(\xi, \eta) \in \overline{Y} \times \overline{Y}$ and set $\xi_j := \pi_j(\xi)$ for $1 \leq j \leq m$. Since $\omega_{(n,R)}^j(\xi_j, \eta_j) = \mathbb{1}_{[[\xi_j, \eta_j]]_{(n,R)}^j} - \mathbb{1}_{[[\eta_j, \xi_j]]_{(n,R)}^j}$ and $\omega_{(n,R)}(\xi, \eta) = \mathbb{1}_{[[\xi, \eta]]_{(n,R)}} - \mathbb{1}_{[[\eta, \xi]]_{(n,R)}}$, it is enough to see that

$$[[\xi, \eta]]_{(n,R)} = [[\xi_1, \eta_1]]_{(n,R)}^1 \sqcup \cdots \sqcup [[\xi_m, \eta_m]]_{(n,R)}^m,$$

where $[[\xi_j, \eta_j]]_{(n,R)}^j \subset \mathfrak{H}(Y_j)^n$. But this follows immediately from the structure of the hyperplanes and halfspaces in a product. \square

Corollary 1.4 will then be a direct consequence of Proposition 3.2, once Theorem 1.1 is proved.

Another property of the median class of an action is that it behaves nicely with respect to subcomplexes in the following sense.

PROPOSITION 3.3. *Let X be a finite-dimensional CAT(0) cube complex, $\Gamma \rightarrow \text{Aut}(\overline{X})$ be an action and $\Gamma_0 < \Gamma$ be a finite index subgroup. Let $W \subset \mathfrak{H}(X)$ be a consistent and Γ_0 -invariant subset, so that $\mathfrak{H}(X) = \mathfrak{H}_W \sqcup (W \sqcup W^*)$ is a lifting decomposition. Let $X_W \subset \partial X$ be the corresponding subcomplex. Then the median class of the Γ -action on X restricts to the median class of the Γ_0 -action on \overline{X}_W .*

Proof. Since $\mathfrak{H}_W^n \subset \mathfrak{H}(X)^n$, there is a map $j : \ell^p(\mathfrak{H}(X)^n) \rightarrow \ell^p(\mathfrak{H}_W^n)$ obtained by restriction. If $c : \Gamma \times \Gamma \times \Gamma \rightarrow \ell^p(\mathfrak{H}(X)^n)$ is the median Γ -equivariant cocycle on X , then $j \circ c|_{\Gamma_0^3} : \Gamma_0 \times \Gamma_0 \times \Gamma_0 \rightarrow \ell^p(\mathfrak{H}_W^n)$ is the median Γ_0 -equivariant cocycle on X_W . \square

3.B. Boundedness of the median class

PROPOSITION 3.4. *Let X be a finite-dimensional CAT(0) cube complex and, for $1 \leq p < \infty$, let $c_{(n,R)}$ be the 1-parameter family of cocycles defined in (3.5). Then*

$$c_{(n,R)} : \overline{X} \times \overline{X} \times \overline{X} \longrightarrow \ell^p(\mathfrak{H}(X)^n) \tag{3.7}$$

and

$$\sup_{u_1, u_2, u_3 \in \overline{X}} \|c_{(n,R)}(u_1, u_2, u_3)\|_p < \infty.$$

More precisely, if D is the dimension of X , then, for any $u_1, u_2, u_3 \in \overline{X}$, the support of $c_{(n,R)}(u_1, u_2, u_3)$ has cardinality bounded above by

$$6(2(n-1)R)^{2D+n-2}.$$

To prove this proposition, we need a few preliminary results. For a set $S \subset X$ define $V_\ell(S)$ to be the ℓ -neighborhood of S , that is, the set of vertices at combinatorial distance less than or equal to ℓ from some element of S .

We start with the following key result, where über-parallel is needed:

LEMMA 3.5. *Let $h_1 \subset h_2$ be an über-parallel pair of halfspaces, $x \in h_1$ and $y \in h_2^*$. Let ℓ be the length of the corresponding bridge. Then*

$$\mathcal{I}(x, y) \subset V_\ell(h_1 \cup h_2^*),$$

that is, the interval between x and y stays within ℓ of $h_1 \cup h_2^*$.

Before proceeding with the proof, we record the following important remark, straightforward from the proof of Lemma 2.17 and the concepts used in Subsection 2.G.

REMARK 3.6. Let $u, v \in X$, let h be a halfspace so that $u, v \in h^*$ and assume that $\hat{\mathfrak{H}}(u, h) \subseteq \hat{\mathfrak{H}}(v, h)$. Let $u = x_0, \dots, x_n = v$ be a combinatorial geodesic from u to v , and let $d_i = d(x_i, h)$. Our assumptions on u and v force the sequence d_i to be increasing, and the proof of the above lemma shows that it remains constant as long as the hyperplanes crossed are transverse to \hat{h} , but increase when they are parallel. In other words, crossing a hyperplane parallel to \hat{h} will take the geodesic away from h .

Proof of Lemma 3.5. We will show that any geodesic from $x \in h_1$ to $y \in h_2^*$ stays within ℓ of h_1 , then goes to the bridge $b(\hat{h}_1, \hat{h}_2)$ and then stays within ℓ of h_2^* to reach y . Since by Lemma 2.18, $b(\hat{h}_1, \hat{h}_2) \subset V_\ell(h_1 \cup h_2^*)$, we will have shown that the geodesic never leaves $V_\ell(h_1 \cup h_2^*)$.

According to Lemma 2.2, a geodesic between x and y corresponds to an enumeration of all the hyperplanes separating x from y , and hence, by Lemma 2.24(2), it has to cross all hyperplanes separating x from $p_b(x)$, where $p_b(x) \in h_1$ is the projection of x on the bridge $b(\hat{h}_1, \hat{h}_2)$, all those separating $p_b(x)$ from $p_b(y)$ and all those separating $p_b(y)$ from y , not necessarily in this order. In fact, when two hyperplanes are parallel, the enumeration in the geodesic has to respect the order given by the inclusion of the corresponding halfspaces, but when two hyperplanes are transverse, the geodesic can cross either one first.

Thus, to understand how far away from $h_1 \cup h_2^*$ a combinatorial geodesic can possibly go, we have to study the possible intersections of elements belonging to the following disjoint sets:

$$\hat{\mathfrak{H}}(x, p_b(x)), \quad \hat{\mathfrak{H}}(p_b(x), p_b(y)) \subset \beta_h \sqcup \beta_v \quad \text{and} \quad \hat{\mathfrak{H}}(y, p_b(y)),$$

where $\beta_h \sqcup \beta_v$ is the decomposition of the halfspaces in the bridge into horizontal and vertical halfspaces according to Lemma 2.18.

By Lemmas 2.18 and 2.24, since h_1 and h_2 are über-parallel, none of the hyperplanes from $\hat{\mathfrak{H}}(x, p_b(x))$ can cross a hyperplane from $\hat{\mathfrak{H}}(y, p_b(y))$. Hence a geodesic from x to y must enumerate all the hyperplanes from $\hat{\mathfrak{H}}(x, p_b(x))$ before enumerating any hyperplane from $\hat{\mathfrak{H}}(y, p_b(y))$.

Now the hyperplanes from β_h all cross \hat{h}_1 and so, according to Remark 3.6, they will not allow the geodesic to travel away from h_1 . The only hyperplanes that can take a geodesic away from h_1 are the ones from β_v and from $\hat{\mathfrak{H}}(y, p_b(y))$. There are at most ℓ hyperplanes from β_v and the hyperplanes from $\hat{\mathfrak{H}}(y, p_b(y))$ will not matter as they will take the geodesic away from h_1 when it is already ℓ -close to h_2^* . Indeed, since the geodesic has to exhaust all the elements of $\hat{\mathfrak{H}}(x, p_b(x))$ before using a hyperplane from $\hat{\mathfrak{H}}(y, p_b(y))$, the same argument for h_2^* shows that it will be ℓ -close to h_2^* . \square

The above lemma says that, in case h_1 and h_2 are über-parallel, in order to go from h_1 to h_2^* , one needs to travel on the bridge. The relevance of the hypothesis of being über-parallel is exemplified in the following.

EXAMPLE 3.7. Take six quarter planes glued in a natural way around their boundaries.

Let h_1, h_2^* be the halfspaces corresponding to the hyperplanes \hat{h}_1 and \hat{h}_2 in the figure. In this case $\ell(b(\hat{h}_1, \hat{h}_2)) = 2$, but one can easily find $x \in h_1$ and $y \in h_2^*$ such that a geodesic joining x and y is not contained in a 2-neighborhood of $h_1 \cup h_2^*$. In fact, there are pairs x, y (that will be away from the bridge), with geodesics arbitrarily far away from the bridge joining them. In this case the pair \hat{h}_1, \hat{h}_2 is strongly separated but not über-separated.




FIGURE 1. The point m is not in the 2-neighborhood of $h_1 \cup h_2^*$.

COROLLARY 3.8. *Let $x, y \in X$ and $h_1 \subset h_2$ be an über-parallel pair of halfspaces such that $x \in h_1 \subset h_2$ and $y \in h_2^*$. Take $z \in h_1^* \cap h_2$. Then the median $m(x, y, z) \in V_\ell(h_1 \cup h_2^*)$, where ℓ is the length of the bridge between h_1 and h_2 .*

Proof. The proof follows directly from the fact that the median is contained in the interval between x and y , which in turn is contained in $V_\ell(h_1 \cup h_2^*)$. \square


COROLLARY 3.9. *Let $x, y \in X$, and $m \in \mathcal{I}(x, y)$. Let $h_1 \in [m, x]$ and $h_2 \in [y, m]$, with $h_1 \subset h_2$, be an über-parallel pair at distance less than or equal to R . Then $B_R(m) \cap (h_1 \cup h_2^*) \neq \emptyset$.*

Proof. This is just a reformulation of Lemma 3.5. In fact, if $d(\hat{h}_1, \hat{h}_2) \leq R$, then $\ell(b(\hat{h}_1, \hat{h}_2)) \leq R$. Then Lemma 3.5 implies that any point in an interval $\mathcal{I}(x, y)$, with $x \in h_1$ and $y \in h_2^*$, is at distance at most R from h_1 or h_2^* . \square

We remark that Example 3.7 shows that if \hat{h}_1 and \hat{h}_2 are only strongly separated, the assertion of Corollary 3.9 does not hold, as one can see in Figure 1 with $R = 2$.

EXAMPLE 3.10. The following ‘infinite staircase’ shows an irreducible CAT(0) cube complex with a pair of hyperplanes \hat{h}_1 and \hat{h}_2 that are parallel and strongly separated but not über-parallel. This example is elementary and there are no über-separated pairs. Furthermore, what captures the pathology of this example is the fact that the median can be

arbitrarily far from the bridge $b(\hat{h}_1, \hat{h}_2)$. The notion of über-separated precisely excludes this pathology.




We also need the following result on the structure of the support of the cocycle.

LEMMA 3.11. *The support of $c(u_1, u_2, u_3)$ is the disjoint union of the six sets obtained by permuting the indices of $[[u_1, u_3]] \setminus ([[u_1, u_2]] \cup [[u_2, u_3]])$. On each of these sets the cocycle is identically equal to 1 or -1 .*

Proof. Let us first examine the structure of the intersections of the six sets appearing in the definition of c . Observe that, for $a, b, i, j \in \{1, 2, 3\}$ and $a \neq b$ and $i \neq j$, we have that

$$\text{if } a \neq i \text{ and } b \neq j, \text{ then } [[u_a, u_b]] \cap [[u_i, u_j]] = \emptyset.$$

This is described more clearly by the following diagram:



Indeed, consider $s \in [[u_a, u_b]] \cap [[u_i, u_a]]$. Then every $h \in s$ must contain u_b and not u_a but must also contain u_a and not u_i , which shows that the intersection is empty. Likewise, $[[u_a, u_b]] \cap [[u_b, u_j]] = \emptyset$.

However, c vanishes on each of the pairwise non-empty intersections. Indeed, if $s \in [[u_1, u_2]] \cap [[u_1, u_3]]$, then

$$c(u_1, u_2, u_3)(s) = 0 + 0 + \mathbb{1}_{[[u_1, u_2]]}(s) - (0 + \mathbb{1}_{[[u_1, u_3]]}(s) + 0) = 0.$$

The other cases are computed similarly. □

We are now ready to prove Proposition 3.4.

Proof of Proposition 3.4. According to Lemma 3.11, an n -tuple $h_1 \subset h_2 \subset \dots \subset h_n$ contributing to the cocycle at a triple u_1, u_2, u_3 has to be in a set of the type

$$[[u_i, u_j]] \setminus ([[u_i, u_k]] \cup [[u_k, u_j]]),$$

where $i, j, k \in \{1, 2, 3\}$ are all different, so that there are six such sets. In other words, if $(h_1, h_2, \dots, h_n) \in [[u_i, u_j]] \setminus ([[u_i, u_k]] \cup [[u_k, u_j]])$, then there exists j , with $1 \leq j < n$, such that $m(u_1, u_2, u_3) \in h_{j+1}$ but $m(u_1, u_2, u_3) \notin h_j$. We hence need to count the number of such n -tuples that ‘hug’ the median $m(u_1, u_2, u_3)$.

We start with the case $n = 2$ and look at the contribution from the set $[[u_1, u_2]]_{(2,R)} \setminus ([[u_1, u_3]]_{(2,R)} \cup [[u_3, u_2]]_{(2,R)})$. According to Lemma 2.3, the interval $\mathcal{I}(u_1, u_2)$ embeds in Euclidean space. Let $m = m(u_1, u_2, u_3) \in \mathcal{I}(u_1, u_2) \subseteq \mathbb{R}^D$. Let $h_1 \subset h_2$, so that $u_2 \in h_1 \subset h_2$, $u_1 \in h_2^*$ and $m \in h_1^* \cap h_2$. According to Corollary 3.9, $B_R(m) \cap (h_1 \cup h_2^*) \neq \emptyset$. Without loss of generality, we may assume that $B_R(m) \cap h_1 \neq \emptyset$. Then, there are at most $(2R)^D$ choices for h_1 ; since h_2 is at distance R from h_1 , there are at most $(2R)^D$ choices for h_2 .

Hence, since there are six terms in the definition of the cocycle and each is a characteristic function on a set of at most $(2R)^{2D}$ elements, we get $6(2R)^{2D}$ for the bound of the cocycle.

For the general case, we count in how many ways we can construct a contributing n -tuple from $[[u_1, u_2]]_{(n,R)} \setminus ([[u_1, u_3]]_{(n,R)} \cup [[u_3, u_2]]_{(n,R)})$, call it $h_1 \subset \dots \subset h_n$. According to the case $n = 2$, there are at most $(2(n-1)R)^{2D}$ choices for h_1 and h_n , since h_1 and h_n are at distance less than or equal to $(n-1)R$ from each other. Therefore, we must count the possible ways of choosing h_2, \dots, h_{n-1} . To this end, we note that h_2, \dots, h_{n-1} must belong to the set β_v of the bridge $b(\hat{h}_1, \hat{h}_n)$ between h_1 and h_n . The bridge has length at most $(n-1)R$ and hence there are at most $((n-1)R)^{n-2}$ many choices for h_2, \dots, h_{n-1} . This means that there are at most $(2(n-1)R)^{2D}((n-1)R)^{n-2}$ contributing n -tuples from $[[u_1, u_2]]_{(n,R)} \setminus ([[u_1, u_3]]_{(n,R)} \cup [[u_3, u_2]]_{(n,R)})$.

Hence, since there are six terms in the definition of the cocycle and each is a characteristic function on a set of at most $2^{2D}((n-1)R)^{2D+n-2}$ elements, we get $6(2(n-1)R)^{2D+n-2}$ for the bound of the cocycle. □

3.C. Toward the Proof of Theorem 1.1

We defined at the beginning of this section the bounded cohomology of Γ with coefficients in $\ell^p(\mathfrak{H}(X)^n)$ as the cohomology of the complex of the Γ -equivariant bounded functions on the Cartesian product Γ^k with values in $\ell^p(\mathfrak{H}(X)^n)$. So far, for any $n \geq 2$ we constructed a 1-parameter family of Γ -equivariant cocycles $c_{(n,R)} : \overline{X} \times \overline{X} \times \overline{X} \rightarrow \ell^p(\mathfrak{H}(X)^n)$ and we remarked that a choice of a basepoint will give a cocycle in $C_b(\Gamma^3, \ell^p(\mathfrak{H}(X)^n))$. We still need to show that the cohomology class represented by this cocycle does not vanish if the action is not elementary. In order to do this, we recall from [17, 38] that if (B, ϑ) is a strong Γ -boundary, there is an isometric isomorphism

$$H_b^2(\Gamma, \ell^p(\mathfrak{H}(X)^n)) \cong \mathcal{Z}L_{\text{alt},*}^\infty(B^3, \ell^p(\mathfrak{H}(X)^n))^\Gamma, \tag{3.8}$$

where the space on the right-hand side is the space of L^∞ alternating Γ -equivariant cocycles on $B \times B \times B$, with the measurability intended with respect to the weak- $*$ topology on $\ell^p(\mathfrak{H}(X)^n)$, $1 \leq p < \infty$ as a dual of \mathcal{E}_p (see (3.2)).

We recall from the introduction that a strong Γ -boundary (B, ϑ) is a Lebesgue space endowed with a measure class preserving Γ -action that is in addition

- (1) amenable, and
- (2) doubly ergodic with coefficients; namely:

DEFINITION 3.12. Let Γ be a group and (B, ϑ) be a Lebesgue space endowed with a measure class preserving Γ -action. The action of Γ on B is *doubly ergodic with (Hilbert) coefficients* if any weak-* measurable Γ -equivariant map $B \times B \rightarrow \mathcal{E}$ into the dual \mathcal{E} of a separable Banach (Hilbert) space on which Γ acts by isometries is essentially constant.

One of the advantages of the realization (3.8) is that, because of (2) with $\mathcal{E} = \ell^p(\mathfrak{H}(X)^n)$ for $1 \leq p < \infty$, in degree 2 there are no coboundaries: hence showing that a cohomology class does not vanish is reduced to showing that the corresponding cocycle is non-zero. The disadvantage is that realizing the pullback via $\rho : \Gamma \rightarrow \text{Aut}(X)$ of a bounded cohomology class defined on the boundary is possible under the condition that there exists a Γ -equivariant measurable boundary map $\varphi : B \rightarrow \overline{X}$ and that the bounded cohomology class can be represented by a bounded Borel measurable alternating strict cocycle [9].

It can be immediately verified that the cocycle c defined in this section is alternating in (u_1, u_2, u_3) , that is to say that if σ is a permutation of $\{u_1, u_2, u_3\}$, then

$$c(\sigma(u_1, u_2, u_3)) = \text{sign}(\sigma)c(u_1, u_2, u_3). \quad (3.9)$$

The Borel measurability of c is proved in Lemma A.4.

Furthermore, a strong Γ -boundary with properties (1) and (2) exists for any locally compact and compactly generated group according to [17], and for arbitrary locally compact groups with respect to a spread-out non-degenerate symmetric measure according to [34]; the existence of the boundary map will take up the next section.

4. The boundary map

This section is devoted to the proof of the following theorem, with an eye to the implementation of the isomorphism in (3.8).

THEOREM 4.1. *Let $\Gamma \rightarrow \text{Aut}(Y)$ be a group action on an irreducible finite-dimensional CAT(0) cube complex Y . Assume that the action is essential and non-elementary. If B is a strong Γ -boundary, then there exists a Γ -equivariant measurable map $\varphi : B \rightarrow \partial Y$ taking values into the non-terminating ultrafilters in ∂Y .*

To realize the isomorphism in (3.8) in our generality, we will in fact need the following stronger statement, which guarantees the existence of some kind of boundary map when the action is not assumed to be essential and the complex is not necessarily irreducible.

COROLLARY 4.2. *Let $\Gamma \rightarrow \text{Aut}(X)$ be a group acting on a finite-dimensional CAT(0) cube complex X . Assume that there is no finite orbit in the visual boundary $\partial_{\triangleleft} X$ and denote by Y the essential core of X . Then there exists a Γ -equivariant measurable map $\varphi : B \rightarrow \partial Y \subseteq \partial X$.*

Proof. Since the action of Γ has no finite orbit in $\partial_{\triangleleft} X$, it has no fixed point. Therefore, the essential core Y is not empty, [20, Proposition 3.5], and Γ also has no finite orbit in $\partial_{\triangleleft} Y$. If $Y = Y_1 \times \cdots \times Y_m$ is the decomposition of Y into a product of irreducible subcomplexes, by Lemma 2.28, Γ also has no finite orbit in $\partial_{\triangleleft} Y_i$, for $i = 1, \dots, m$, and moreover the action on each Y_i is essential.

If $j = 1, \dots, q$, let $\varphi_j : B \rightarrow \partial Y_j$ be the Γ -equivariant measurable boundary map whose existence is proved in Theorem 4.1. Since $\prod_{j=1}^q \partial Y_j \subseteq \partial Y \subseteq \partial X$, the map $\varphi : B \rightarrow \partial Y$ defined by $\varphi(b) := (\varphi_1(b), \dots, \varphi_q(b))$ has the desired properties. \square

The idea of the proof of Theorem 4.1 is as follows. Since \overline{X} is a continuous compact metric G -space, the space $\mathcal{P}(\overline{X})$ of probability measures on \overline{X} endowed with the weak- $*$ topology is a subset of the (unit ball in the) dual of the continuous functions on \overline{X} . By amenability of the Γ -action on B , there exists a Γ -equivariant measurable map $\psi : B \rightarrow \mathcal{P}(\overline{X})$ into the probability measures on \overline{X} (see [54, Proposition 4.3.9]). Each probability measure μ on \overline{X} divides the set of halfspaces into ‘balanced’ (that is halfspaces such that $\mu(h) = \mu(h^*)$) and ‘unbalanced’ ones. If all halfspaces are unbalanced, this defines an ultrafilter, hence the map $\psi : B \rightarrow \mathcal{P}(\overline{X})$ gives a Γ -equivariant map $\psi : B \rightarrow \overline{X}$. Since the measure ϑ on B is ergodic, so is the pushforward measure on \overline{X} . Hence up to measure 0 the image of ψ is either in X or in ∂X . If it is in X , then it is essentially constant, so we get a Γ -fixed point, hence it had to land in ∂X . The whole work in the proof will be to exclude the presence of balanced halfspaces using non-elementary actions assumptions as well as essentiality.

4.A. General preliminary lemmas using ergodicity

The following lemma can be thought of as a weaker version of the statement that a strong Γ -boundary for a lattice is a strong Γ -boundary for its ambient group and vice versa.

LEMMA 4.3. *Let Γ be a group acting on a measure space (M, ϑ) . If Γ acts ergodically on $(M \times M, \vartheta \times \vartheta)$, then every finite index subgroup $\Gamma_0 \leq \Gamma$ acts ergodically on (M, ϑ) .*

Proof. We prove the contrapositive of the statement. Let $\Gamma_0 \leq \Gamma$ be a finite index subgroup that does not act ergodically.

Let (M_0, ϑ_0) be the Mackey’s point realization of the measure algebra generated by the Γ_0 -invariant sets. In other words, M_0 is a standard measure space equipped with a measurable map $p : M \rightarrow M_0$ such that $p_*(\vartheta) = \vartheta_0$. Since, by passing to a finite index subgroup if necessary (that will still act non-ergodically), Γ_0 can be taken to be normal in Γ , this measure algebra is Γ -invariant and hence it defines a Γ -action on M_0 with respect to which the map $p : M \rightarrow M_0$ is Γ -equivariant. Hence there is an ergodic action of the finite group Γ/Γ_0 on M_0 , which is therefore an atomic space, but cannot consist of one point (otherwise the Γ_0 -action would be ergodic).

Now take any point $m_0 \in M_0$ and define $A := p^{-1}(m_0) \subset M$. By construction A is neither null nor conull and Γ_0 -invariant. Consider the subset $\bigcup_{[\gamma] \in \Gamma/\Gamma_0} \gamma A \times \gamma A \subset M \times M$ (which is well defined by the Γ_0 -invariance of M_0). This set is Γ -invariant and not null. Furthermore, it is not conull. Indeed, let $A^c = M \setminus A$ denote the complement. We claim that $A \times A^c \subset (\bigcup_{[\gamma] \in \Gamma/\Gamma_0} \gamma A \times \gamma A)^c$. Indeed, if there is a $\gamma' \in \Gamma$ such that $A \times A^c \cap (\gamma' A \times \gamma' A)$ has positive measure, then $\vartheta(A \cap \gamma' A) > 0$ and $\vartheta(A^c \cap \gamma' A) > 0$, while, by construction, $\vartheta(A \cap \gamma A) = \vartheta(A)$ or $\vartheta(A \cap \gamma A) = 0$ for all $\gamma \in \Gamma$.

Therefore, the Γ -invariant set $\bigcup_{[\gamma] \in \Gamma/\Gamma_0} \gamma A \times \gamma A$ is neither null nor conull and hence the diagonal action of Γ on $M \times M$ is not ergodic. \square

LEMMA 4.4. *Let C be a countable set with a Γ action and (B, ϑ) be a Lebesgue space with a measure class preserving Γ -action that is in addition doubly ergodic with Hilbert coefficients. Let $\mathcal{B} := B$ or to $\mathcal{B} := B \times B$. If $\psi : \mathcal{B} \rightarrow C$ is a Γ -equivariant measurable map, then ψ is essentially constant.*

Proof. We prove the assertion for $\mathcal{B} := B \times B$. The assertion for $\mathcal{B} := B$ follows then from the first one applied to the precomposition with the projection $\pi_1 : B \times B \rightarrow B$ on the first component.

As the action of Γ on $B \times B$ is ergodic, so is the pushforward measure $\psi_*(\beta \times \beta)$ and hence the image of ψ is supported on an orbit. We now assume that the Γ -action on C is transitive.

If C is finite, then there is a finite index subgroup Γ_0 which acts trivially on C . But as the Γ_0 action on B is still ergodic, by Lemma 4.3 we conclude that the action of Γ_0 on C is still transitive and hence C is a single point.

Next, assume that C is infinite. This means that the corresponding generalized Bernoulli action of Γ on 2^C is ergodic (indeed, it is weakly mixing) and measure preserving with respect to the standard Bernoulli measure λ on 2^C generated by taking 0 and 1 with equal mass. By the double ergodicity with coefficients of B (see [17]; also [2, Lemma 2.2]), we conclude that the diagonal Γ -action on $B \times B \times 2^C$ is ergodic. Let $(x, y) \in B \times B$ and $S \subset C$. It is clear that the following evaluation function is essentially constant as it is invariant under the diagonal Γ -action

$$(x, y, S) \longmapsto \mathbf{1}_S(\psi(x, y)) \in \{0, 1\}.$$

By Fubini's Theorem, there is a point $(x_0, y_0) \in B \times B$ so that, for λ -almost every $\mathbf{1}_S \in 2^C$, the value of $\mathbf{1}_S(\psi(x_0, y_0))$ is identically 0, or 1. This gives a contradiction. Indeed, for any $c \in C$ we know that

$$\lambda(\{\mathbf{1}_S \in 2^C : \mathbf{1}_S(c) = 0\}) = \lambda(\{\mathbf{1}_S \in 2^C : \mathbf{1}_S(c) = 1\}) = 1/2,$$

in particular for $c_0 := \psi(x_0, y_0)$. □

We apply the previous lemma to the countable set $2_f^{\mathfrak{H}(X)}$ consisting of finite subsets of $\mathfrak{H}(X)$.

COROLLARY 4.5. *Let \mathcal{P} be equal to either $\mathcal{P}(\overline{X})$ or $\mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X})$. If there exists a Γ -equivariant measurable map $\mathcal{P} \rightarrow 2_f^{\mathfrak{H}(X)}$, then the Γ -action on X is not essential.*

Proof. By hypothesis there is a finite Γ -invariant subset of $\mathfrak{H}(X)$ and in particular, there is a finite Γ -orbit $\Gamma \cdot h$. Then, the corresponding CAT(0) cube complex $X(\Gamma \cdot h)$ is finite and by [20, Proposition 3.2, (i) \Rightarrow (iii)], the action is inessential. □

4.B. Heavy and balanced halfspaces, and properties of their associated complexes

Let $\mathcal{P}(\overline{X})$ denote the space of probability measures on \overline{X} . If $\mu \in \mathcal{P}(\overline{X})$, define

$$\begin{aligned} H_\mu &:= \{h \in \mathfrak{H}(X) : \mu(h) = \mu(h^*)\}, \\ H_\mu^+ &:= \{h \in \mathfrak{H}(X) : \mu(h) > 1/2\}, \\ H_\mu^- &:= \{h \in \mathfrak{H}(X) : \mu(h) < 1/2\}, \\ H_\mu^\pm &:= \{h \in \mathfrak{H}(X) : \mu(h) \neq 1/2\}. \end{aligned}$$

We refer to H_μ as the *balanced* halfspaces and to H_μ^+ as the *heavy* halfspaces. The terms *unbalanced* and *light* halfspaces are also self-explanatory.

We record a few easy consequences of the definition.

LEMMA 4.6. *Let $\mu, \nu \in \mathcal{P}(\overline{X})$ be any two measures.*

(1) *The family H_μ is closed under the involution $h \mapsto h^*$ and the involution is a bijection between H_μ^+ and H_μ^- .*

(2) *There is the following partition of halfspaces: $\mathfrak{H}(X) = H_\mu \sqcup H_\mu^\pm$, where $H_\mu^\pm = H_\mu^+ \sqcup H_\mu^-$.*

- (3) If h, k belong to H_μ (respectively, H_μ^+ or H_μ^-), then either $h \pitchfork k$ or all halfspaces between h, k are in H_μ (respectively, H_μ^+ or H_μ^-).
- (4) There are no facing triples of halfspaces in H_μ . If X is not Euclidean, it follows that $H_\mu^+ \neq \emptyset$.
- (5) If X is not Euclidean, H_μ and H_ν are not empty and $H_\mu \cap H_\nu = \emptyset$, then $H_\mu \cap H_\nu^\epsilon \neq \emptyset$ for $\epsilon \in \{+, -\}$.
- (6) If $h, k \in H_\mu$ are two parallel halfspaces with $h \subset k$, then $\mu(h^* \cap k) = 0$.
- (7) The assignments $\mu \mapsto H_\mu$ and $\mu \mapsto H_\mu^\epsilon$, for $\epsilon \in \{+, -\}$, are $\text{Aut}(X)$ -equivariant for the natural actions on $\mathcal{P}(\overline{X})$ and $2^{\mathfrak{H}(X)}$.

Proof. Assertions (1), (2) and (3) are obvious.

To see (4), assume that h_1, h_2, h_3 were a facing triple of halfspaces in H_μ , so that $h_2^* \subset h_1, h_3^* \subset h_1$ and $\hat{h}_2 \parallel \hat{h}_3$. This would imply that $1/2 = \mu(h_1) \geq \mu(h_2^*) + \mu(h_3^*) = 1$, which is a contradiction. Since X is not Euclidean, and hence there are facing triples of halfspaces, then $H_\mu^\pm \neq \emptyset$ and also $H_\mu^+ \neq \emptyset$.

Assertion (5) follows from the fact that if $H_\mu \cap H_\nu = \emptyset$, then $H_\mu \subset H_\mu^\pm$. But then, since H_μ is invariant under the involution $h \mapsto h^*$, both $H_\mu \cap H_\nu^+$ and $H_\mu \cap H_\nu^-$ must be non-empty.

Assertion (6) is immediate since $\mu(k) = \mu(h^*) + \mu(h \cap k)$ and $h^*, k \in H_\mu$ and (7) is immediate from the definitions. □

It follows from Lemma 2.6 with $W := H_\mu^+$ and $\mathfrak{H}_W := H_\mu$ that there is an isometric embedding $\overline{X}(H_\mu) \hookrightarrow \overline{X}$ and to simplify the notation, we denote by \overline{X}_μ its image in \overline{X} ($\overline{X}_{H_\mu^+}$ in the notation of Lemma 2.6).

We remark again that if $H_\mu = \emptyset$, then \overline{X}_μ is a single vertex in \overline{X} . Note, moreover, that the H_μ^+ are not Γ -invariant and the subcomplex $\overline{X}_\mu \subset \overline{X}$ is not Γ -invariant.

LEMMA 4.7. *The complex $\overline{X}(H_\mu)$ is an interval.*

Proof. Let us consider the projection $p : \overline{X} \rightarrow \overline{X}(H_\mu)$ and let $\alpha_0 \in \text{supp}(p_*\mu)$. Let α_0^* be the ‘opposite’ of α_0 (in H_μ). Observe that α_0^* is an ultrafilter on H_μ : indeed, the only non-trivial condition we must check is that if $h \in \alpha_0^*$ and $h \subset k$, then $k \in \alpha_0^*$. If instead $k \notin \alpha_0^*$, then $k \in \alpha_0$, which means that $h^* \cap k$ is an open neighborhood of α_0 , contradicting that α_0 is in the support of μ with Lemma 4.6(6). By construction, $H_\mu = [\alpha_0, \alpha_0^*] \cup [\alpha_0^*, \alpha_0]$, where the intervals are taken in $\overline{X}(H_\mu)$. □

DEFINITION 4.8. Let \mathfrak{H}' be a subset of $\mathfrak{H}(X)$. An element $h \in \mathfrak{H}'$ is called:

- (i) *minimal in \mathfrak{H}'* if, for every $k \in \mathfrak{H}'$, either $k \pitchfork h, h \subset k$, or $h \subset k^*$;
- (ii) *maximal in \mathfrak{H}'* if, for every $k \in \mathfrak{H}'$, either $k \pitchfork h, k \subset h$, or $k^* \subset h$, that is to say, h is maximal if h^* is minimal;
- (iii) *terminal in \mathfrak{H}'* if it is either maximal or minimal.

REMARK 4.9. The number of terminal elements is bounded above by $2d$ not just for H_μ but for any union of pairwise incomparable chains in H_μ .

4.C. Proof of Theorem 4.1

Proof of Theorem 4.1. Since the Γ -action on (B, ϑ) is amenable, there exists a Γ -equivariant measurable map $\psi : B \rightarrow \mathcal{P}(\overline{X})$ into the probability measures on \overline{X} . We consider $\mathcal{P}(\overline{X})$ endowed with the pushforward of the quasi-invariant, doubly ergodic measure ϑ on B , so that Γ acts ergodically on $\mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X})$. We will show that under the hypotheses of the theorem,

we may associate to every μ in the image of ψ a point in ∂X and the composition will be the required Γ -equivariant measurable boundary map $\varphi : B \rightarrow \partial X$.

The map $C_1 : \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$ defined by $\mu \mapsto |H_\mu|$ is measurable (Corollary A.2(1)) and Γ -equivariant, hence by ergodicity it is essentially constant.

I. $H_\mu = \emptyset$ for almost all μ

If the essential value of C_1 is 0, then, for almost every $\mu \in \mathcal{P}(\overline{X})$, $H_\mu = \emptyset$. This means that, up to measure 0, the image of ψ lies in the set $\mathcal{E} := \{\mu \in \mathcal{P}(\overline{X}) : H_\mu = \emptyset\}$. Thus we have a well-defined composition $\varphi : B \rightarrow \mathcal{E} \rightarrow \overline{X}$ defined by $x \mapsto \psi(x) \mapsto \overline{X}_{\psi(x)}$, whose image is the single point $\overline{X}_{\psi(x)} \in \overline{X}$. (Lemma 2.6). Measurability is guaranteed by Lemmas A.1 and A.3. The equivariance under Γ follows from Lemma 4.6(7). Proposition 4.11 will show that, in fact, φ takes values into the non-terminating ultrafilters of X .

The rest of the proof will show that all other cases cannot occur.

II. $0 < |H_\mu| < \infty$ for almost all μ

If the essential value of C_1 were to be finite, then Corollary 4.5 with $\mathcal{P} = \mathcal{P}(\overline{X})$ would imply that the action is not essential.

III. $|H_\mu| = \infty$ for almost all μ

To deal with this case, we consider the Γ -equivariant and measurable function $C_2 : \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by $(\mu, \nu) \mapsto |H_\mu \cap H_\nu|$ (Corollary A.2(2)). Again by ergodicity of the Γ -action on $\mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X})$, the function C_2 is essentially constant.

III.a $0 < |H_\mu \cap H_\nu| < \infty$ for almost all μ, ν

If the essential value of C_2 were finite and non-zero, then Corollary 4.5 with $\mathcal{P} = \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X})$ would again imply that the action is not essential.

III.b $|H_\mu \cap H_\nu| = 0$ for almost all μ, ν

Now suppose that the essential value of C_2 is 0, so that, for almost every $\mu, \nu \in \mathcal{P}(\overline{X})$, $H_\mu \cap H_\nu = \emptyset$. Let us consider the measurable (Corollary A.2(3)) Γ -equivariant function $T : \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by

$$T(\mu, \nu) := |\tau((H_\mu \cap H_\nu^+) \cup (H_\nu \cap H_\mu^+))|,$$

where

$$\tau : 2^{\mathfrak{H}(X)} \longrightarrow 2^{\mathfrak{H}(X)} \quad (4.1)$$

is the map that assigns to a subset of halfspaces its terminal elements. By double ergodicity T is essentially constant. Using the fact that both H_μ and H_ν are Euclidean, any subset of them must have finitely many terminal elements and therefore this essential value must be finite (see Remark 4.9). Once more, essentiality of the action, along with Corollary 4.5 assures us that the essential value is 0.

This leaves us with the case in which the essential value is zero, that is, $H_\mu \cap H_\nu^+$ has no terminal elements for almost every (μ, ν) . In this case the following proposition (whose proof we postpone to Subsection 4.D) allows us to conclude that this case cannot happen.

PROPOSITION 4.10. *Suppose that, for almost every $\mu, \nu \in \mathcal{P}(\overline{X})$, $|H_\mu| = |H_\nu| = \infty$, $H_\mu \cap H_\nu = \emptyset$ and $H_\mu \cap H_\nu^+$ has no minimal elements. Then X contains cubes of arbitrarily large dimension.*

III.c $|H_\mu \cap H_\nu| = \infty$ for almost all μ, ν

Finally, let us suppose that the essential value of C_2 is ∞ , namely $|H_\mu \cap H_\nu| = \infty$ for almost every $(\mu, \nu) \in \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X})$.

If $H_\mu = H_\nu$ for almost every $\mu, \nu \in \mathcal{P}(\overline{X})$, then applying Fubini, there is a $\mu_0 \in \mathcal{P}(\overline{X})$ such that, for every ν in a conull Γ -invariant subset, we have that $H_{\mu_0} = H_\nu$. Hence, $H_{\gamma_*\nu} = \gamma H_\nu = H_\nu$. Since the action is essential without fixed points on the visual boundary, we may flip

any $h \in H_\nu^+$. This means that $H_\nu^- \cap H_{\gamma_\nu}^+ \neq \emptyset$ and so H_ν^+ is not Γ -invariant. As a result, the corresponding embedded subcomplexes $\overline{X}_{\gamma_\nu}$ are not invariant. We will see in Proposition 4.17 that this implies that X is a product, which is a contradiction.

We are therefore left in the case in which $H_\mu \cap H_\nu$ is infinite but $H_\mu \neq H_\nu$, for almost all $\mu, \nu \in \mathcal{P}(\overline{X})$.

We now consider whether or not H_μ has strongly separated halfspaces. Observe that the set

$$\mathcal{S} = \{(h_1, h_2) \in \mathfrak{H}(X) \times \mathfrak{H}(X) : h_1, h_2 \text{ are strongly separated}\}$$

is Γ -invariant. Therefore, the map $\mu \rightarrow |(H_\mu \times H_\mu) \cap \mathcal{S}|$ is measurable (Corollary A.2(4)) and Γ -invariant, and hence essentially constant.

If H_μ contains pairs of strongly separated halfspaces, then H_μ^+ satisfies the DCC (Lemma 4.18). This implies that the action is again inessential by extracting the finitely many terminal elements of the set $(H_\mu^+ \cap (H_\nu \setminus H_\mu)) \cup (H_\nu^+ \cap (H_\mu \setminus H_\nu))$ (Corollary A.2(5)), and we proceed as before to conclude that the action is inessential.

If, on the other hand, H_μ does not contain pairs of strongly separated halfspaces, then by Corollary 4.21 the action is inessential. \square

4.D. Further properties and proofs

PROPOSITION 4.11. *Let X be a finite-dimensional CAT(0) cube complex, $\Gamma \rightarrow \text{Aut}(X)$ be an essential action on X , (B, ν) be a doubly ergodic Γ -space with quasi-invariant measure ν and $\varphi : B \rightarrow \overline{X}$ be a measurable Γ -equivariant map. Then φ takes values in the non-terminal ultrafilters of X .*

We start with a few easy observations. Recall that if α and β are two ultrafilters,

$$\mathfrak{H}(\alpha, \beta) := [\alpha, \beta] \cup [\beta, \alpha] = [\alpha, \beta] \cup [\alpha, \beta]^*.$$

Then it is easy to check that

$$\tau(\mathfrak{H}(\alpha, \beta)) = \tau([\alpha, \beta]) \cup \tau([\alpha, \beta]^*) \tag{4.2}$$

and hence $|\tau(\mathfrak{H}(\alpha, \beta))|$ is finite.

LEMMA 4.12. *Let α and β be two ultrafilters and $h \in \tau(\alpha)$. Then $\beta \notin h$ if and only if $h \in \tau(\mathfrak{H}(\alpha, \beta))$.*

Proof. If $\beta \in h$, then h does not separate α and β , so that $h \notin \mathfrak{H}(\alpha, \beta)$ and, even more so, $h \notin \tau(\mathfrak{H}(\alpha, \beta))$. The converse is equally easy and will not be needed. \square

Proof of Proposition 4.11. We may assume that X is irreducible. The general case will follow from this case as in the proof of Corollary 4.2, since the set of non-terminating ultrafilters in a product is the Cartesian product of the sets of non-terminating ultrafilters of each factor.

The composition of φ with the map τ defined in (4.1) that assigns to a set of halfspaces its terminal element gives a Γ -equivariant measurable map $B \rightarrow 2^{\mathfrak{H}(X)}$ defined by $x \mapsto \tau(\phi(x))$. The function $C_4 : B \rightarrow \mathbb{N} \cup \{\infty\}$ defined by $x \mapsto |\tau(\phi(x))|$ is hence essentially constant.

Therefore, we want to show that $|C_4(x)| = 0$ for almost every x , that is, that the set $\tau(\phi(x))$ is empty, thus showing that $\varphi(x)$ is non-terminating.

To this purpose let us consider the map $\theta : B \times B \rightarrow 2^{\mathfrak{H}(X)}$ that to a pair $(x, y) \in B \times B$ associates the set of terminal elements in $\mathfrak{H}(\varphi(x), \varphi(y))$. Again by ergodicity the function $C_5 : B \times B \rightarrow \mathbb{N} \cup \{\infty\}$, defined by $C_5(x, y) := |\tau(\mathfrak{H}(\varphi(x), \varphi(y)))|$, is essentially constant and, by (4.2), $0 \leq |\tau(\mathfrak{H}(\varphi(x), \varphi(y)))| < \infty$.

By Corollary 4.5 with $\mathcal{P} = B$, we deduce that, for almost every $x, y \in B$, $\tau(\mathfrak{H}(\varphi(x), \varphi(y))) = \emptyset$. We show now that this is incompatible with $|\tau(\varphi(x))| > 0$ for almost every $x \in B$, thus proving the proposition.

Let $x_0 \in B$ be such that $|\tau(\varphi(x_0))| > 0$ and let $B_0 \subset B$ be a set of full measure such that $\tau(\mathfrak{H}(\varphi(x_0), \varphi(y))) = \emptyset$ for all $y \in B_0$. Then by Lemma 4.12, if $h \in \tau(\varphi(x_0))$, then we must have that $\varphi(y) \in h$ for all $y \in B_0$. But B_0 contains a Γ -orbit and hence this contradicts the fact that the action is essential. \square

4.D.1. *Proof of Proposition 4.10 (in step III.b).* We will find arbitrarily a large family of pairwise intersecting halfspaces. To this purpose, choose a sequence $\{\mu_i\}_{i \in \mathbb{N}}$ of pairwise generic measures satisfying the hypotheses of Proposition 4.10. For each i , choose an infinite descending chain $h_n^i \in H_{\mu_0}^+ \cap H_{\mu_i}$.

Consider the following property of an ordered pair (μ_i, μ_j) of measures .

(*) There exists $C(i, j) \in \mathbb{N}$ such that, for every $n \geq C(i, j)$, there is an $M_n \geq C(i, j)$ such that if $m > M_n \geq C(i, j)$, then $\hat{h}_n^i \pitchfork \hat{h}_m^j$.

LEMMA 4.13. *Up to switching i and j , any pair of measures μ_i and μ_j satisfies (*).*

We postpone the proof of this lemma and show how to conclude the proof of Proposition 4.10.

Let us consider a graph $\mathcal{G} := \mathcal{G}(V, E)$, where $V := \{\mu_i\}$ and where two measures μ_i and μ_j are connected by an edge $e \in E$ with source μ_i and target μ_j if the ordered pair (μ_i, μ_j) satisfies (*). By Lemma A.8, given $D \in \mathbb{N}$, there exist (relabelled) measures $\mu_1, \dots, \mu_D \in \{\mu_i\}_{i \in \mathbb{N}}$ such that, for $1 \leq i < j \leq D$, each ordered pair (μ_i, μ_j) satisfies (*).

By choosing

$$C := \max\{C(i, j) : 1 \leq i < j \leq D\}$$

and

$$M := \max\{M_C(i, j) : 1 \leq i < j \leq D\}.$$

we obtain that, for all $n, m \geq C$ and $1 \leq i, j \leq D$, the corresponding hyperplanes are transverse, $\hat{h}_n^i \pitchfork \hat{h}_m^j$. This concludes the proof of Proposition 4.10.

Proof of Lemma 4.13. Fix two measures that we denote, for ease of notation, μ and ν . Let $h_n \in H_{\mu_0}^+ \cap H_\mu$ and $k_m \in H_{\mu_0}^+ \cap H_\nu$ be the corresponding infinite descending sequences. Since all the halfspaces in question belong to $H_{\mu_0}^+$, for each pair n, m we have the following decomposition:

$$\mathbb{N} \times \mathbb{N} = N_1 \sqcup N_2 \sqcup N_3 \sqcup N_4,$$

where

$$\begin{aligned} N_1 &= \{(n, m) : h_n \pitchfork k_m\}, \\ N_2 &= \{(n, m) : h_n \subset k_m\}, \\ N_3 &= \{(n, m) : h_n \subset k_m\}, \\ N_4 &= \{(n, m) : h_n \supset k_m\}. \end{aligned}$$

We claim that if we allow ourselves to throw away a finite number of pairs (n, m) if necessary, then the decomposition of $\mathbb{N} \times \mathbb{N}$ takes a simpler shape. Namely, we make the following claim.

CLAIM 4.14. *There exists a constant $C \in \mathbb{N}$ depending on μ and ν , such that*

$$\mathbb{N}_C := (\mathbb{N} \times \mathbb{N}) \cap ([C, \infty) \times [C, \infty)) = N_1 \sqcup N_j,$$

where $j = 2, 3$ or 4 .

In fact, let us suppose that $N_2 \neq \emptyset$ and $N_3 \neq \emptyset$ and let us take $(n_3, m_3) \in N_3$ and $(n, m) \in N_2$. Set $m' := \min\{m, m_3\}$, such that

$$h_n^* \subset k_{m'} \quad \text{and} \quad h_{n_3} \subset k_{m'}.$$

If $n \geq n_3$, then $h_n \subset h_{n_3} \subset k_{m'}$, which is impossible since also $h_n^* \subset k_{m'}$. Hence there is no pair $(n, m) \in N_2$ such that $n \geq \min\{n_3 : (n_3, m_3) \in N_3\} =: A_3$. It follows that

$$\{(n, m) \in N_2 : n \geq A_3\} \cap N_3 = \emptyset. \quad (4.3)$$

Now let us suppose that $N_3 \neq \emptyset$ and $N_4 \neq \emptyset$ and let us take $(n_3, m_3) \in N_3$ and $(n, m) \in N_4$. If $n \geq n_3$, then

$$k_n \subset h_n \subset h_{n_3} \subset k_{m_3},$$

which is impossible by Lemma 4.6 part (3). Hence, analogously to the previous case, we have that

$$\{(n, m) \in N_2 : n \geq A_3\} \cap N_4 = \emptyset. \quad (4.4)$$

Finally, let us suppose that $N_2 \neq \emptyset$ and $N_4 \neq \emptyset$ and let us take $(n, m) \in N_2$ and $(n_4, m_4) \in N_4$. Set $n' := \min\{n, n_4\}$, so that

$$h_{n'} \supset k_m \quad \text{and} \quad h_{n'} \supset k_{m_4}.$$

If $m \geq m_4$, then $h_{n'} \supset k_{m_4} \subset k_m$, which is impossible since also $h_{n'} \supset k_m^*$. Hence there is no pair $(n, m) \in N_2$ such that $m < \min\{m_4 : (n_4, m_4) \in N_4\} =: B_4$. It follows that

$$\{(n, m) \in N_2 : m \geq B_4\} \cap N_4 = \emptyset. \quad (4.5)$$

By setting $C := \max\{A_3, B_4\}$ we have proved the claim.

Let us suppose now that, for $n_0, m_0 \geq C$, the pair $(n_0, m_0) \in N_1 \sqcup N_3$ and, in fact, that $(n_0, m_0) \in N_3$ (otherwise there is nothing to prove). Choose $m_0 = m_0(n_0)$ to be the largest integer such that $(n_0, m_0) \in N_3$. Then, because $\mathbb{N}_C = N_1 \sqcup N_3$, for every $m \geq m_0(n_0) + 1$, the hyperplanes \hat{h}_{n_0} and \hat{k}_m are transverse. Hence the assertion of the lemma is proved in the case in which $\mathbb{N}_C = N_1 \sqcup N_3$.

Remark that the same identical argument shows the assertion if $\mathbb{N}_C = N_1 \sqcup N_2$, since we only used that there is a minimal element k_{m_0} in the sequence k_m that contains the hyperplane \hat{h}_{n_0} .

The argument if $\mathbb{N}_C = N_1 \sqcup N_4$ is analogous, but with the role of h_n and k_m reversed, as now there is a minimal element h_{n_0} in the sequence h_n that contains the hyperplane \hat{k}_{m_0} . Namely, let $(n_0, m_0) \in N_4$ be such a pair. Then, for all $n \geq n_0(m_0) + 1$, the hyperplanes \hat{h}_n and \hat{k}_{m_0} are transverse. \square

REMARK 4.15. The last assertion in the proof relative to the case $\mathbb{N}_C = N_1 \sqcup N_4$ holds also for the case $\mathbb{N}_C = N_1 \sqcup N_2$, but the symmetry of this case is not useful.

4.D.2. Proofs needed in step III.c.

LEMMA 4.16. *Let X be a CAT(0) cube complex and A be any cubical subset of X (that is, A is a union of cubes, not necessarily connected). If Y is the smallest strongly convex subcomplex of X containing A , then*

$$\hat{\mathfrak{H}}(Y) = \hat{\mathfrak{H}}(A) \sqcup \{\hat{h} \in \hat{\mathfrak{H}}(Y) : \hat{h} \text{ separates } A \text{ in at least two non-trivial subsets}\}.$$

Proof. We only need to check that every $\hat{h} \in \hat{\mathfrak{H}}(Y) \setminus \hat{\mathfrak{H}}(A)$ separates A in non-trivial subsets. Take $\hat{h} = (h, h^*) \in \hat{\mathfrak{H}}(Y)$ and assume by contradiction that $A \subseteq h$. Then any geodesic between two points of A is also contained in h (otherwise this geodesic would cross h twice). Hence $Y \subseteq h$, contradicting that $h \in \hat{\mathfrak{H}}(Y)$. \square

PROPOSITION 4.17. *Let X be a CAT(0) cube complex and $\Gamma \rightarrow \text{Aut}(X)$ be an essential action. Let $\mathfrak{H}' \subset \mathfrak{H}(X)$ be a Γ -invariant subset of halfspaces and $X_\alpha \subset X$ be a Γ -invariant family of subcomplexes such that $\hat{\mathfrak{H}}(X_\alpha) = \hat{\mathfrak{H}}'$. Let Y be the smallest strongly convex subcomplex containing $A := \bigcup_\alpha X_\alpha$. Then $Y = X$ and $\overline{X} = \overline{X(\hat{\mathfrak{H}}')} \times Z$.*

Proof. Since Y is Γ -invariant and the action is essential, then $Y = X$.

Because of Lemma 4.16, the hyperplanes of Y are of two types: either they are in $\hat{\mathfrak{H}}' = \hat{\mathfrak{H}}(A)$ and they separate one (equivalently, any) of the X_α or they separate a X_α from a $X_{\alpha'}$. Any hyperplane \hat{h} of this second type will cross any hyperplane $\hat{k} \in \hat{\mathfrak{H}}'$. Indeed, if $\hat{h} = (h, h^*)$ and $\hat{k} = (k, k^*)$, it is easy to see that the four intersections in (2.2) are non-empty. Hence X is a product. \square

LEMMA 4.18. *If $|H_\mu| = \infty$ and H_μ contains strongly separated halfspaces, then H_μ^+ satisfies the DCC.*

Proof. Let $h, k \in \mathfrak{H}(X)$ be a pair of strongly separated halfspaces in H_μ with $h \subset k$. There is the following decomposition:

$$H_\mu^+ = P(h) \cup P(k), \tag{4.6}$$

where $P(h)$ and $P(k)$ are the μ -heavy halfspaces that are parallel, respectively, to h and k . Note that while $P(h)$ and $P(k)$ are not necessarily disjoint, their union is the whole of H_μ^+ since h and k are strongly separated.

Let $h_n \in H_\mu^+$ be a descending chain, that is, $h_{n+1} \subset h_n$. We must show that the chain terminates. By passing to a subsequence, we may assume that h_n belong to the same set for all $n \in \mathbb{N}$ and it is hence enough to consider, for example, the case $h_n \in P(h)$ for all $n \in \mathbb{N}$.

Since $h_n \in H_\mu^+$ and $h \in H_\mu$, we cannot have that $h_n \subset h$ or $h \subset h_n^*$. Let us suppose that $h \subset h_n$. Since between h and h_n there are only finitely many halfspaces, and since no μ -heavy halfspace can be contained in a balanced one, the chain must terminate. Likewise the same argument applied to $h^* \subset h_n$ shows that the chain must terminate. \square

LEMMA 4.19. *For every measure μ either \hat{H}_μ contains a pair of strongly separated hyperplanes or there exists a pair $h \in H_\mu^-, k \in H_\mu^+$ of halfspaces, such that the hyperplanes \hat{h} and \hat{k} are strongly separated and, for every $x \in H_\mu$, $\hat{x} \subset h^* \cap k$.*

Proof. Suppose that H_μ does not contain strongly separated halfspaces. We first show that, for every $x \in H_\mu$, there exist $k_0(x), k_3(x) \in H_\mu^\pm$ such that $\hat{k}_0(x)$ and $\hat{k}_3(x)$ are strongly separated and $\hat{x} \subset k_0^*(x) \cap k_3(x)$. For ease of notation, we drop the dependence on x .

In fact, since X is irreducible, given $x \in H_\mu$, there exist halfspaces k_1, k_2 such that \hat{k}_1 and \hat{k}_2 are strongly separated hyperplanes and $k_1 \subset x \subset k_2$. Then at least one between the k_1 and k_2 must be in H_μ^\pm , but perhaps not both of them. Then double skewer k_2 into k_1 and k_1^* into k_2^* to obtain

$$\gamma k_2 \subset k_1 \subset x \subset k_2 \subset \gamma^{-1} k_1,$$

where the pairs $\gamma k_2, k_2$ and $k_1, \gamma^{-1} k_1$ are strongly separated. Since all hyperplanes corresponding to pairs of halfspaces in the sequence $k_0 \subset k_1 \subset k_2 \subset k_3$ are strongly separated, there can be at most one halfspace that belongs to H_μ . By measure considerations, this halfspace can only

be either k_1 or k_2 , so that $k_0, k_3 \in H_\mu^\pm$, and the assertion is proved. In particular, $\gamma k_1 \in H_\mu^-$ and $\gamma^{-1}k_2 \in H_\mu^+$.

Double skewer once again to get $h \in H_\mu^-$ and $k \in H_\mu^+$ with \hat{h}, \hat{k} strongly separated, such that

$$h \subset k_0 \subset x \subset k_3 \subset k.$$

We show now that, given any $y \in H_\mu$, we have $\hat{y} \subset h^* \cap k$. In fact, we cannot have $k \subset y$ or $k \subset y^*$, since $y, y^* \in H_\mu$ and $k \in H_\mu^+$. Analogously, we cannot have that $\hat{y} \cap \hat{k}$, because otherwise \hat{y} could not intersect \hat{k}_3 and hence it would have to contain it, which is impossible again by measure considerations. Hence $y \subset k$. An analogous argument shows that $h \subset y$, thus completing the proof. \square

The above argument can be extended to show the following lemma.

LEMMA 4.20. *Let $\mu_i \in \mathcal{P}(\overline{X})$ be measures such that \hat{H}_{μ_i} does not contain strongly separated hyperplanes for all i and $H_{\mu_i} \cap H_{\mu_j} \neq \emptyset$ for all i, j . Then there exists a pair of halfspaces $h \subset k$ such that \hat{h}, \hat{k} are strongly separated and, for every $x \in H_{\mu_j}$, $\hat{x} \subset h^* \cap k$.*

Proof. Fix μ_0 and apply Lemma 4.19 to find halfspaces $h_2 \subset h_3$ such that \hat{h}_2, \hat{h}_3 are strongly separated and

$$\hat{x} \subset h_2^* \cap h_3. \tag{4.7}$$

Use the double skewering lemma several times to find a chain $h_0 \subset \dots \subset h_5$ of halfspaces with corresponding pairwise strongly separated hyperplanes. We will use that (4.7) holds in particular for every $x_j \in H_{\mu_0} \cap H_{\mu_i}$ to show that $\hat{y} \subset h_0 \cap h_5$ for every $y \in H_{\mu_j}$ and every j .

Consider in fact $y \in H_{\mu_i}$. Observe that \hat{y} can be transverse to at most one $\hat{h}_i, 0 \leq i \leq 5$, since these are pairwise strongly separated. If it is transverse to any \hat{h}_i for $1 \leq i \leq 4$, we are done, since then $\hat{y} \subset h_0^* \cap h_5$. Suppose instead that \hat{y} is transverse to \hat{h}_0 . Then \hat{h}_1 and \hat{h}_2 are nested in between \hat{y} and \hat{x}_j , which is impossible by Lemma 4.6 part(3) and because \hat{H}_{μ_j} does not contain strongly separated hyperplanes. A similar argument shows that \hat{y} cannot be transverse to \hat{h}_5 .

If instead \hat{y} is parallel to all \hat{h}_i , for $0 \leq i \leq 5$, then we have to check that $\hat{y} \subset h_0$ and $\hat{y} \subset h_5^*$ cannot happen. In fact, if $\hat{y} \subset h_0$, as before, this would force \hat{h}_1, \hat{h}_2 to be in \hat{H}_{μ_j} , which is not possible because they are a strongly separated pair. The case in which $\hat{y} \subset h_5^*$ can be excluded analogously. \square

COROLLARY 4.21. *Assume that, for almost every $\mu \in \mathcal{P}(\overline{X})$, there are no strongly separated pairs in H_μ . If $H_\mu \cap H_\nu \neq \emptyset$ for almost every pair (μ, ν) , then the Γ -action is non-essential.*

Proof. Fix a generic measure μ_0 with a generic Γ -invariant set B_0 such that, for every $\nu \in B_0$, we have that $H_{\mu_0} \cap H_\nu \neq \emptyset$.

Lemma 4.20 implies the existence of a pair of halfspaces $h \subset k$ such that $h^* \cap k$ contains all the hyperplanes in \hat{H}_μ for $\nu \in B_0$, in particular those in $\gamma \hat{H}_{\mu_0} = \hat{H}_{\gamma_* \mu_0}$ for all $\gamma \in \Gamma$. This shows that the two halfspaces h and k are not Γ -flippable, which contradicts either that the action is essential or that it is without fixed points on the CAT(0) boundary [20, Theorem 4.1]. \square

5. Proof of Theorems 1.1 and 1.2

Let X be a finite-dimensional CAT(0) cube complex, $\Gamma \rightarrow \text{Aut}(X)$ be a non-elementary action and $Y \subset X$ be the essential core of X . Let (B, ϑ) be any strong Γ -boundary. In order to prove our main result, we constructed in Section 4 a measurable Γ -equivariant boundary map $\varphi : B \rightarrow \partial X$ to the Roller boundary ∂X . The precomposition of the median cocycle c with $\varphi : B \rightarrow \partial X$ yields a Γ -equivariant cocycle defined on B^3 , which we will show is non-zero on a set of positive measure (Proposition 5.1 and Lemma 5.3). According to (3.8), this ensures the existence of a non-trivial cohomology class on Γ . Then [9] ensures that the median class of the Γ -action $\rho^*(\mathfrak{m}_{(n,R)}) \in H_b^2(\Gamma, \ell^p(\mathfrak{H}(X)^n))$, $n \geq 2$, corresponds to the cohomology class $c \circ \varphi^3$ on B^3 and hence does not vanish.

5.A. Passing from a cocycle on ∂X to a cocycle on B

We give a condition on a Γ -equivariant cocycle $d : (\partial X)^3 \rightarrow E$ to guarantee that $d \circ \varphi^3 : B^3 \rightarrow E$ is non-zero on a set of positive measure, where $\varphi : B \rightarrow \partial X$ is a measurable Γ -equivariant map.

Let K be compact metrizable Γ -space. A measure $\lambda \in \mathcal{P}(K)$ is quasi-invariant if λ and $\gamma_*\lambda$ have the same null sets for all $\gamma \in \Gamma$.

If $h \in \mathfrak{H}(X)$ is a halfspace, then we set

$$\bar{h} := \{x \in \bar{X} : x \in h\}$$

and

$$\partial h := \bar{h} \cap \partial X.$$

PROPOSITION 5.1. *Let Γ be a group with a non-elementary and essential action $\Gamma \rightarrow \text{Aut}(Y)$ on a finite-dimensional CAT(0) cube complex Y . If (B, ϑ) is a strong Γ -boundary, let $\varphi : B \rightarrow \partial Y$ be a Γ -equivariant measurable map. Let $d : (\partial Y)^3 \rightarrow E$ be an everywhere defined alternating bounded Γ -equivariant Borel cocycle with values in a coefficient Γ -module E . If there exist halfspaces $h_i \in \mathfrak{H}(Y)$ such that $d(\xi_1, \xi_2, \xi_3) \neq 0$ for every $(\xi_1, \xi_2, \xi_3) \in \partial h_1 \times \partial h_2 \times \partial h_3$, then $d \circ \varphi^3$ is a non-trivial element of $\mathcal{ZL}_{\text{alt},*}^\infty(B^3, E)^\Gamma$.*

The proof is almost an immediate consequence of the following lemma.

LEMMA 5.2. *Let Y be a finite-dimensional CAT(0) cube complex and $\Gamma \rightarrow \text{Aut}(Y)$ be a non-elementary essential action. If $\lambda \in \mathcal{P}(\partial Y)$ is any quasi-invariant probability measure, then $\lambda(\partial h) > 0$ for any halfspace $h \in \mathfrak{H}(Y)$.*

Proof. If $\lambda(\partial h) = 0$, then $\lambda(\partial h^*) = 1$. By the Flipping Lemma [20, Theorem 4.1], there exists $\gamma \in \Gamma$ such that $h^* \subset \gamma h$. But this is a contradiction because $\partial h^* \subset \partial(\gamma h) = \gamma \partial h$, while, by quasi-invariance, $\lambda(\gamma \partial h) = 0$. \square

Proof of Proposition 5.1. If there exist halfspaces $h_i \in \mathfrak{H}(Y)$ such that $d(\xi_1, \xi_2, \xi_3) \neq 0$ for every $(\xi_1, \xi_2, \xi_3) \in \partial h_1 \times \partial h_2 \times \partial h_3$, then $d \circ \varphi^3(x_1, x_2, x_3) \neq 0$ for almost every $(x_1, x_2, x_3) \in \varphi^{-1}(\partial h_1) \times \varphi^{-1}(\partial h_2) \times \varphi^{-1}(\partial h_3) =: S \subset B^3$.

By Lemma 5.2 applied to the quasi-invariant probability measure $\varphi_*\vartheta \in \mathcal{P}(\partial Y)$, the set S has positive ϑ^3 -measure and hence $d \circ \varphi^3$ is a non-trivial element of $\mathcal{ZL}_{\text{alt},*}^\infty(B^3, E)^\Gamma$. \square

5.B. Proof of Theorems 1.1 and 1.2

We start with a lemma ensuring that the cocycle is non-vanishing on a set of positive measure.

LEMMA 5.3. *Let X be a finite-dimensional CAT(0) cube complex, with a non-elementary action $\Gamma \rightarrow \text{Aut}(X)$. Then, for every essential $h \in \mathfrak{H}(X)$ and for every $n \geq 2$, there is a positive measure set $A_{(h,n)} \subset \partial X^3$ and $R_h > 0$ so that, for every $R > R_h$, the restriction $c_{(n,R)}|_{A_{(h,n)}}$ does not vanish.*

Proof. Let γ and $\gamma' \in \Gamma$ be such that the triple $h, \gamma h, \gamma' h$ is an über-parallel facing triple in an orbit as in Lemma 2.37. By the Flipping Lemma, there exists $\mu \in \Gamma$ such that $\mu h^* \subset h$ and let $\eta \in \Gamma$ be an element that skewers μh into γh , so as to obtain

$$\mu h \supset h^* \supset \gamma h \supset \eta \mu h.$$

The pair μh and $\mu \eta h$ is über-parallel and hence so is any consecutive pair in the sequence

$$\mu h \supset \eta \mu h \supset \eta^2 \mu h \supset \dots \supset \eta^{n-1} \mu h.$$

Because of Lemma 5.2, the set

$$A_{(h,n)} := \partial(\mu h^*) \times \partial(\eta^{(n-1)} \mu h) \times \partial(\gamma' h)$$

has positive measure. Since, for every $(\xi_1, \xi_2, \xi_3) \in A_{(h,n)}$, the set $[[\xi_3, \xi_2]]_{(n,R)} \setminus ([[\xi_3, \xi_1]]_{(n,R)} \cup [[\xi_1, \xi_2]]_{(n,R)})$ is not empty, provided R is larger than the translation length R_h of η , Lemma 3.11 ensures that $c_{(n,R)}|_{A_{(h,n)}}$ does not vanish. \square

Proof of Theorem 1.1. Let Y be the essential core of the Γ -action on X . According to [9], the class of $c_{(n,R)} \circ \varphi^3$ is the isometric image of the median class $m_{(n,R)}$ under the isomorphism (3.8), where $\varphi : B \rightarrow \partial Y$ is the boundary map constructed in Theorem 4.1. Let $R_\Gamma = \min\{R_h \mid h \text{ essential}\}$ and where the R_h s are as defined in Lemma 5.3. Then, Proposition 5.1 and Lemma 5.3 ensure that $c_{(n,R)} \circ \varphi^3$ is non-trivial if $n \geq 2$ and $R \geq R_\Gamma$. \square

Proof of Theorem 1.2. If the Γ -action is elementary, by definition there exists a finite orbit in $X \cup \partial_{\triangleleft} X$. If the finite Γ -orbit is in X , then there is a subgroup of finite index $\Gamma_0 < \Gamma$ that fixes a point $x \in X$. Hence the median class of the Γ_0 -action on X vanishes. Since the map $H_b^2(\Gamma, \ell^p(\mathfrak{H}(X)^n)) \rightarrow H_b^2(\Gamma_0, \ell^p(\mathfrak{H}(X)^n))$ is injective [38], the median class of the Γ -action vanishes.

If, on the other hand, there exists a finite Γ -orbit in $\partial_{\triangleleft} X$, then we can apply Proposition 2.26 and deduce that either there is a finite orbit in ∂X , in which case we conclude as in the first part and the median class of the Γ -action vanishes, or there exists a finite index subgroup $\Gamma' < \Gamma$ and a Γ' -invariant subcomplex $X' \subset \partial X$ in which the Γ' -action is non-elementary. By Theorem 1.1, the median class of the Γ' -action on X' does not vanish, and, since X' corresponds to a lifting decomposition of halfspaces, by Proposition 3.3 it is the restriction of the median class of the Γ -action on X . \square

6. Applications

6.A. Rigidity of actions

Proof of Theorem 1.5. We will need our action to satisfy the property that the barycenter of every face has trivial stabilizer. This is a natural generalization to CAT(0) cube complexes of the notion of no edge inversions in the context of actions on trees. For this reason, we start with an arbitrary action and then pass to its cubical subdivision.

Let Y' be the cubical subdivision of Y . Observe that $\text{Aut}(Y) \hookrightarrow \text{Aut}(Y')$ and the image acts with the property that the barycenter of every face in Y' has trivial stabilizer. Moreover, Γ acts essentially on Y' and Y' is irreducible.

The proof follows very closely the strategy of the proof in [52]. Namely, if we denote by e_i the identity in G_i , then we aim to show that there is an $i \in \{1, \dots, \ell\}$ for which the set

$$Y_i := \{x \in Y' : \text{if } \gamma_m \in \Gamma \text{ such that } \text{pr}_i(\gamma_m) \longrightarrow e_i, \\ \text{then there exists } N > 0 \text{ such that } \gamma_m x = x \text{ for all } m \geq N\}$$

is not empty, where $\text{pr}_i : G \rightarrow G_i$ is the i th projection. It is easy to see that the set Y_i is Γ -invariant. Indeed, if $\gamma_m \in \Gamma$ is a sequence such that $\text{pr}_i(\gamma_m) \rightarrow e_i$, then, for every $\gamma \in \Gamma$, we have that $\text{pr}_i(\gamma^{-1}\gamma_m\gamma) \rightarrow e_i$. Moreover, Y_i is convex with respect to the CAT(0) metric: in fact, let $x_1, x_2 \in Y_i$ and let $\gamma_m \in \Gamma$ be a sequence such that $\text{pr}_i(\gamma_m) \rightarrow e_i$. Then, by definition of Y_i there exists N sufficiently large such that $\gamma_m x_j = x_j$ for all $m \geq N$ and $j = 1, 2$. Since Γ acts by isometries, if $m \geq N$, then γ_m will also fix pointwise the unique CAT(0) geodesic between x_1 and x_2 .

We claim now that if Y_i is not empty, then it is in fact a subcomplex of Y' . To see this, let us write Y' as the disjoint union of k -dimensional faces, where a k -dimensional face is the interior of a k -dimensional cube if $1 \leq k \leq \dim(Y)$ and is the boundary of a one-dimensional cube if $k = 0$. Let F_k be a k -dimensional face that has non-empty intersection with Y_i . Then $F_k \subset Y_i$. Since we are acting on the cubical subdivision Y' , we have also that if γ_m eventually fixes a face F_k , then it fixes all lower-dimensional faces that are contained in its closure $\overline{F_k}$, thus showing that $\overline{F_k} \subset Y_i$. Thus Y_i is a CAT(0) cube subcomplex of Y' .

We are then left to show that there exists $i \in \{1, \dots, \ell\}$ such that $Y_i \neq \emptyset$.

Let $i \in \{1, \dots, \ell\}$ and let $\mathcal{H}_i \subset \ell^2(\mathfrak{H}(Y)^n)$ be the (possibly trivial) subspace on which the isometric action of Γ extends continuously to G via the projection $\text{pr}_i : G \rightarrow G_i$. By [17, Theorem 16], $H_b^2(\Gamma, \ell^2(\mathfrak{H}(Y)^n)) = \bigoplus_{i=1}^{\ell} H_b^2(G_i, \mathcal{H}_i)$, so that by Theorem 1.1 there exists an $i \in \{1, \dots, \ell\}$ (not necessarily unique) such that $\mathcal{H}_i \neq \{0\}$.

The space Y_i will be constructed from these data as follows.

Define on $\mathfrak{H}(Y)^n$ an equivalence relation, namely if $s, s' \in \mathfrak{H}(Y)^n$, then we say that $s \sim s'$ if $f(s) = f(s')$ for all $f \in \mathcal{H}_i$. Since these functions are square summable, all of the equivalence classes are finite, with the possible exception of the class where all functions in \mathcal{H}_i vanish. Moreover, Γ permutes all the finite equivalence classes and leaves invariant the only infinite one (if it exists). Therefore, the complement $\mathfrak{H}(Y)_0^n$ of the infinite class in $\mathfrak{H}(Y)^n$ is Γ -invariant.

CLAIM 6.1. *Let $[s] \in \mathfrak{H}(Y)_0^n / \sim$ be a finite equivalence class and $\text{Stab}_\Gamma([s])$ be its stabilizer. If $\gamma_m \in \Gamma$ is a sequence such that $\text{pr}_i(\gamma_m) \rightarrow e_i$, then there exists $N > 0$ such that, for all $m \geq N$, $\gamma_m \in \text{Stab}_\Gamma([s])$.*

We assume the claim for the moment and show that Y_i is not empty. Fix $[s] \in \mathfrak{H}(Y)_0^n$ and let \hat{h} be a hyperplane corresponding to one of the halfspaces appearing in an element of $[s]$. By Lemma 2.14, there exists $\gamma \in \Gamma$ such that \hat{h} and $\gamma\hat{h}$ are strongly separated. By [3, Lemma 2.2], the bridge $b(\hat{h}, \gamma\hat{h})$ consists of a single geodesic (of finite length). Observe that since $\mathfrak{H}(X)_0^n$ is invariant, the class $[\gamma s]$ is finite. Hence there are finitely many halfspaces in the set $\{h : h \in s', \text{ for } s' \in [s']\}$, both if $s' = s$ and if $s' = \gamma s$. It follows that if we define $L := \text{Stab}_\Gamma([s]) \cap \text{Stab}_\Gamma([\gamma s])$, then the L -orbit of $b(\hat{h}, \gamma\hat{h})$ is finite, therefore bounded, and its circumcenter is an L -fixed point.

If $\gamma_m \in \Gamma$ is a sequence such that $\text{pr}_i(\gamma_m) \rightarrow e_i$, Claim 6.1 implies that, for m large enough, the sequence γ_m belongs to L and hence fixes the circumcenter, thus showing that $Y_i \neq \emptyset$. Then Proposition 4.3 in [52] shows that the action of Γ on Y_i extends to G by factoring through G_i .

Since the action of Γ on Y' is essential, it follows that $Y_i = Y'$. Observe however that, since $\text{Aut}(Y)$ is closed in $\text{Aut}(Y')$ in the topology of the pointwise convergence, the extension of the action to G is in $\text{Aut}(Y)$. □

Proof of Claim 6.1. Let $s \in \mathfrak{H}(Y)_0^n$ and let $f \in \mathcal{H}_i$ so that $f(s) \neq 0$. Since $\lim_{m \rightarrow \infty} \|\gamma_m f - f\|_2 = 0$, it follows that $\lim_{m \rightarrow \infty} f(\gamma_m s) = f(s)$. Because f is square summable, it takes finitely many values in a $|f(s)|/2$ -neighborhood of $f(s)$, so that we conclude there exists $N(f, s)$ such that $f(\gamma_m s) = f(s)$ for all $m \geq N(f, s)$. In particular, $\{\gamma_m s : m \geq 1\}$ is finite.

If $\gamma_{m_k} s \not\sim s$ for some subsequence m_k , then, by passing to a further subsequence, we may assume that $s_0 := \gamma_{m_k} s \not\sim s$. But then there is $g \in \mathcal{H}_i$ such that $g(s) \neq 0$ and $g(s_0) \neq g(s)$, which, together with

$$g(s_0) = \lim_k g(\gamma_{m_k} s) = g(s),$$

is a contradiction. □

The proof of the above theorem does rely on the assumption that Y is irreducible and essential. In general, we can pass to the essential core Y of the Γ -action on X and to its cubical subdivision Y' . Let $\Gamma' < \Gamma$ be the finite index subgroup that acts on each of the irreducible factors in Y' and let $G'_i := \overline{\text{pr}_i(\Gamma')}$. By applying Theorem 1.5 to each of the irreducible factors, we obtain that the action of Γ' on Y extends continuously to an action of G' , where $G' = G'_1 \times \cdots \times G'_\ell$, by factoring via one of the factors.

We have hence proved the following.

COROLLARY 6.2. *Let X be a finite-dimensional CAT(0) cube complex and Γ be an irreducible lattice in the product of locally compact groups $G_1 \times \cdots \times G_\ell =: G$. Let $\Gamma \rightarrow \text{Aut}(X)$ be a non-elementary action on X . Then the action of Γ on the essential core of X virtually extends to a continuous action of an open finite index subgroup in G , by factoring via one of the factors.*

6.B. The class \mathcal{C}_{reg}

We now prove Corollary 1.8 concerning the class of groups \mathcal{C}_{reg} . The idea of the proof is as follows. If the action is proper, then in particular the vertex stabilizers are finite. We can then find, for each $n > 1$, an n -tuple $s \in \mathfrak{H}(X)^n$ with finite stabilizers such that $\ell^p(\Gamma \cdot s) \hookrightarrow \ell^p(\Gamma)$. We then prove that s can be chosen in such a way that this map does not vanish on the image of the cocycle.

Proof of Corollary 1.8. Let $h \subset k$ be two strongly separated halfspaces in \mathfrak{H} . Since the action of Γ is also by CAT(0)-isometries, the stabilizer of $\{h, k\}$ must also stabilize their CAT(0) bridge. Since the action is proper, the stabilizer of the CAT(0) bridge $b(\hat{h}, \hat{k})$ is finite. It follows that if $s = (h_1, \dots, h_n)$ is a über-separated sequence of halfspaces of consecutive distance less than or equal to R , then $\text{Stab}_\Gamma(s)$, the stabilizer of s in Γ , is finite.

Fix $s \in \mathfrak{H}(Y)^n$, and consider the Γ -equivariant map

$$\sigma_s : \ell^p(\Gamma \cdot s) \longrightarrow \ell^p(\Gamma) \tag{6.1}$$

defined by $\sigma_s f(\gamma) := f(\gamma s)$. Since $\|\sigma_s f\|_p = |\text{Stab}_\Gamma(s)| \|f\|_p$, the map is injective.

Now, observe that $\ell^p(\mathfrak{H}(Y)^n) = \bigoplus_{s \in S} \ell^p(\Gamma \cdot s)$, where S is a choice of Γ -orbit representatives.

Since the cocycle $c : \bar{Y} \times \bar{Y} \times \bar{Y} \rightarrow \ell^p(\mathfrak{H}(Y)^n)$ is Γ -equivariant, for every s , we may postcompose c with σ_s and precompose with the boundary map $\varphi : B \rightarrow \partial Y$ and obtain

$$c_s = \sigma_s \circ c \circ \varphi^3 : B \times B \times B \longrightarrow \ell^p(\Gamma \cdot s).$$

We now choose s so as to guarantee that c_s is non-vanishing on a set of positive measure:

As in the proof of Lemma 5.3, we fix $h \in \mathfrak{H}(Y)$ and find $\gamma, \gamma' \in \Gamma$ so that $h^*, \gamma h, \gamma' h$ form a facing über-parallel triple. Let $s = (h, \gamma h, \dots, \gamma^{n-1} h)$. Then, c_s restricted to the set

$$A_{(h,n)} = \partial h \times \partial(\gamma^{(n-1)} h^*) \times \partial(\gamma' h^*)$$

is non-zero, as is demonstrated in the proof of Lemma 5.3. \square

Appendix A. Some more proofs

A.A. The measurability of certain key maps

The notation used in this section refers to Section 4.

LEMMA A.1. *Let $I \subseteq [0, 1]$ be a subinterval, possibly open, half open or closed. Let $H_\mu^I = \{h \in \mathfrak{H}(X) : \mu(h) \in I\}$. The map $\mathcal{P}(\overline{X}) \rightarrow 2^{\mathfrak{H}(X)}$, defined by $\mu \mapsto H_\mu^I$, is measurable with respect to the weak-* topology on $\mathcal{P}(\overline{X})$.*

As a consequence, the map $N : \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$ defined by $N(\mu) = |H_\mu^I|$ is measurable.

Proof. Recall the definition of cylinder sets: Let $F_1, F_2 \in 2^{\mathfrak{H}(X)}$ be two finite sets. The cylinder set associated to them is

$$C(F_1, F_2) = \{H \in 2^{\mathfrak{H}} : F_1 \subseteq H \text{ and } F_2 \subseteq H^c\}.$$

Cylinder sets form a basis for the topology on $2^{\mathfrak{H}(X)}$. Therefore, it is sufficient to show that

$$K(F_1, F_2) = \{\mu : H_\mu^I \in C(F_1, F_2)\}$$

is measurable.

To this end, observe that h is open and closed as a subset of \overline{X} so that its characteristic function $\mathbf{1}_h$ is continuous. Therefore, the set $E_I(h) = \{\mu : \mu(h) \in I\}$, for $h \in \mathfrak{H}(X)$, is weak-* open, half open or closed according to I , and therefore is measurable. Then the following completes the proof:

$$K(F_1, F_2) = \bigcap_{h \in F_1} E_I(h) \cap \bigcap_{h \in F_2} (E_I(h))^c. \quad \square$$

COROLLARY A.2. *The following maps are measurable:*

- (1) $C_1 : \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by $C_1(\mu) := |H_\mu|$;
- (2) $C_2 : \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by $C_2(\mu, \nu) := |H_\mu \cap H_\nu|$;
- (3) $T : \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by

$$T(\mu, \nu) := |\tau((H_\mu \cap H_\nu^+) \cup (H_\nu \cap H_\mu^+))|;$$

- (4) $N_\nu : \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by $N_\nu(\mu) := |(H_\mu \times H_\nu) \cap \mathcal{S}|$, where

$$\mathcal{S} := \{(h_1, h_2) \in \mathfrak{H}(X) \times \mathfrak{H}(X) : h_1, h_2 \text{ are strongly separated}\};$$

- (5) $C_3 : \mathcal{P}(\overline{X}) \times \mathcal{P}(\overline{X}) \rightarrow \mathbb{N} \cup \{\infty\}$, defined by

$$C_3(\mu, \nu) := |[H_\mu^+ \cap (H_\nu \setminus H_\mu)] \cup [H_\nu^+ \cap (H_\mu \setminus H_\nu)]|.$$

Proof. The proofs of statements (1), (2) and (5) are consequences of Lemma A.1 and the observation that the product and composition of Γ -equivariant measurable maps are again Γ -equivariant and measurable. Statement (4) follows by considering the fact that \mathcal{S} is a Γ -invariant set. Statement (3) follows from the following result. \square

LEMMA A.3. *The map $p : 2^{\mathfrak{H}(X)} \rightarrow 2^X$ defined by $\mathbf{1}_S \mapsto \mathbf{1}_{\bigcap_{h \in S} h}$ is measurable.*

Proof. Choose an enumeration $\mathfrak{H}(X) = \{h_n : n \in \mathbb{N}\}$. Recall that the standard projection $\pi_N : 2^{\mathfrak{H}(X)} \rightarrow 2^{\{h_1, \dots, h_N\}}$ is continuous. Next define $p_N : 2^{\{h_1, \dots, h_N\}} \rightarrow 2^X$ as $\mathbb{1}_S \mapsto \mathbb{1}_{(\bigcap_{h \in S} h)}$, which is also continuous as $2^{\{h_1, \dots, h_N\}}$ is endowed with the discrete topology.

Observe that $p(\mathbb{1}_S) = \sup\{p_N \circ \pi_N(\mathbb{1}_S) : n \in \mathbb{N}\}$ and is hence measurable as the supremum of continuous functions. \square

Recall that, in (3.2), we set the notation $\mathcal{E}_p := \ell^q(\mathfrak{H}(X)^n)$ if $1/p + 1/q = 1$ and $1 < p < \infty$, and \mathcal{E}_1 to be the Banach space of functions on $\mathfrak{H}(X)^n$ that vanish at infinity.

LEMMA A.4. *For all $1 \leq p < \infty$, the cocycle $c : \overline{X}^3 \rightarrow \ell^p(\mathfrak{H}^n)$ is a Borel map, where $\overline{X}^3 \subset 2^{\mathfrak{H}(X)}$ has the induced product topology and $\ell^p(\mathfrak{H}(X)^n)$ has the weak-* topology as the dual of \mathcal{E}_p .*

Proof. Choose an enumeration of \mathfrak{H} and let $\mathfrak{H}_N := \{h_1, \dots, h_N\}$. Let us define the finite space

$$\mathfrak{H}_N^n := \{s \in \mathfrak{H}^n : s \subset \mathfrak{H}_N\},$$

and, for any subsets $E, F \subset \mathfrak{H}$, the set

$$(E \setminus F)_N^n := \{s \in \mathfrak{H}_N^n : s \subset E \text{ and } s \not\subset F\}.$$

The map $c_N^+ : 2^{\mathfrak{H}} \times 2^{\mathfrak{H}} \times 2^{\mathfrak{H}} \rightarrow C_0(\mathfrak{H}^n)$ defined as

$$c_N^+(F_1, F_2, F_3) := \mathbb{1}_{(F_3 \setminus F_2)_N^n} + \mathbb{1}_{(F_1 \setminus F_3)_N^n} + \mathbb{1}_{(F_2 \setminus F_1)_N^n}$$

factors through the canonical projection $2^{\mathfrak{H}} \rightarrow 2^{\mathfrak{H}_N}$ on triples and hence is continuous. Then the map $c_N(F_1, F_2, F_3) := c_N^+(F_1, F_2, F_3) - c_N^+(F_1, F_3, F_2)$ is also continuous and in particular is continuous when restricted to the subset $\overline{X}^3 \subset (2^{\mathfrak{H}})^3$. (Here we use the identification of a vertex $v \in X$ with the principal ultrafilter containing v .)

For any $f \in \mathcal{E}_p$, the function on \overline{X}^3 defined by

$$(x, y, z) \mapsto \langle c_N(x, y, z), f \rangle$$

is continuous. Its pointwise limit

$$(x, y, z) \mapsto \lim_{N \rightarrow \infty} \langle c_N(x, y, z), f \rangle$$

is measurable and, in fact,

$$\langle c(x, y, z), f \rangle = \lim_{N \rightarrow \infty} \langle c_N(x, y, z), f \rangle,$$

thus showing that the cocycle c restricted to \overline{X}^3 is Borel. \square

A.B. A lemma in graph theory

Let $\mathcal{G}(V, E)$ be a *tournament*, that is a complete directed finite graph with vertices V and edges E . We denote by $s, t : E \rightarrow V$, respectively, the *source* and the *target* of an edge. We allow the possibility that there are two edges between two vertices, one in each direction. Given a vertex $v \in V$, we denote by $o(v)$ (respectively, $i(v)$) the number of outgoing (respectively, incoming) edges at v . Since the graph is complete,

$$o(v) + i(v) \geq |V| - 1, \tag{A.1}$$

for every $v \in V$.

The next lemma shows that if the graph is complete, there is at least one vertex that has ‘many’ outgoing edges.

LEMMA A.5. If $\mathcal{G} := \mathcal{G}(V, E)$ is a complete directed finite graph and $|V| = D$, then there exists $v \in V$ such that $o(v) \geq (D - 1)/2$.

Proof. From (A.1), we have that

$$\sum_{v \in V} o(v) + i(v) \geq D(D - 1).$$

We have also that

$$\sum_{v \in V} o(v) = \sum_{v \in V} i(v),$$

so that

$$\sum_{v \in V} o(v) \geq \frac{D(D - 1)}{2}.$$

Since $|V| = D$, the assertion follows readily. \square

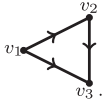
DEFINITION A.6. We say that a complete directed finite graph $\mathcal{G}(V, E)$ with $|V| = D$ is *strictly upper triangular* (or *transitive tournament*) if there exists a numbering v_1, \dots, v_D of its vertices, such that, for all $j = 1, \dots, D$,

$$\begin{aligned} o(v_j) &= D - j, \\ i(v_j) &= j - 1. \end{aligned}$$

The terminology is inspired from the fact that the corresponding $D \times D$ adjacency matrix M with coefficients

$$M_{ij} := \begin{cases} 1 & \text{if there exists } e \in E \text{ with } s(e) = v_i \text{ and } t(e) = v_j, \\ 0 & \text{otherwise,} \end{cases}$$

is strictly upper triangular, namely v_1 is connected by an outgoing vertex to all of the remaining v_2, \dots, v_D , v_2 is connected to v_3, \dots, v_D and so on.

EXAMPLE A.7. A strictly upper triangular graph with $d = 2$ corresponds to the matrix $M = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$ and is of the form .

LEMMA A.8. Let $\mathcal{G} = \mathcal{G}(V, E)$ be a complete directed graph (not necessarily finite) and $D \in \mathbb{N}$. If $|V| \geq 5^D$, then there exist D vertices v_1, \dots, v_D such that the induced complete directed subgraph on v_1, \dots, v_D is strictly upper triangular.

Proof. The idea is to construct the strictly upper triangular graph inductively. By Lemma A.5, there exists $v_1 \in V$ with $o(v_1) \geq (|V| - 1)/2 \geq |V|/5$ outgoing edges. Denote by

$$O(v_1) := \{e \in E : s(e) = v_1\}$$

the set of outgoing edges (so that $|O(v_1)| = o(v_1)$). We consider now the induced complete directed subgraph $\mathcal{G}(v_1)$ on v_1 and on the vertices at the end of the edges in $O(v_1)$, namely $\mathcal{G}_1 := \mathcal{G}(V(v_1), E(v_1)) \subset \mathcal{G}$, where

$$V(v_1) := \{v_1\} \sqcup \{t(e) : e \in O(v_1)\} \quad \text{and} \quad O(v_1) \subsetneq E(v_1) \subseteq E,$$

and the edges $E(v_1)$ are exactly the edges in E needed to complete the graph on the vertices $V(v_1)$. Remark that, by construction, $|V(v_1)| = o(v_1) + 1 \geq |V|/5 \geq 5^{D-1}$ and the induced

complete directed subgraph on any ordered pair of vertices v_1, v , with $v \in V(v_1)$, is (trivially) a strictly upper triangular graph.

We could now proceed to formulate a rigorous proof by induction, but we prefer showing how to move to the next step, as we believe that the very simple idea of the proof will be more transparent.

We repeat now exactly the same construction as before, applied to the graph $\mathcal{G}_1(V(v_1), E(v_1))$ instead of $\mathcal{G}(V, E)$. Namely, let $v_2 \in V(v_1)$ be the vertex, whose existence is asserted by Lemma A.5, such that if

$$O(v_2) := \{e \in E(v_1) : s(e) = v_2\},$$

then $o(v_2) = |O(v_2)| \geq (|V(v_1)| - 1)/2 \geq |V(v_1)|/5$. By construction there is an outgoing edge from v_1 to v_2 and from v_1 to any other vertex in \mathcal{G}_1 . From the graph \mathcal{G}_1 , we retain now only those vertices $w \in V(v_1)$ that are at the end of an outgoing edge from v_2 , so that v_1, v_2, w is strictly upper triangular, and eliminate all of the other vertices. Namely, let \mathcal{G}_2 be the induced complete directed graph on v_2 and on the vertices that are the targets of the $o(v_2)$ edges in $E(v_1)$ outgoing from v_2 ; that is, $\mathcal{G}_2 := \mathcal{G}(V(v_2), E(v_2)) \subset \mathcal{G}_1$, where

$$V(v_2) := \{v_2\} \sqcup \{t(e) : e \in O(v_2)\} \quad \text{and} \quad O(v_2) \subsetneq E(v_2) \subseteq E(v_1) \subseteq E.$$

Now we have a complete directed graph on $|V(v_2)| \geq o(v_2) + 1 \geq |V(v_1)|/5 \geq 5^{D-2}$ vertices from which we can continue choosing vertices v_3, \dots, v_D such that, at each step, we increase by one our strictly upper triangular graph and we reduce by a factor of 5 the cardinality of the vertex set. \square

Appendix B. *Boundary stabilizers in CAT(0) cube complexes*
by Pierre-Emmanuel Caprace

Let X be a (not necessarily proper) CAT(0) cube complex. A group $\Gamma \leq \text{Aut}(X)$ is called *locally X-elliptic* if every finitely generated subgroup of Γ fixes a point of X . The goal of this appendix is to establish the following.

THEOREM B.1. *Let X be a finite-dimensional CAT(0) cube complex and α be a point in the Roller compactification.*

Then the stabilizer $\text{Stab}_{\text{Aut}(X)}(\alpha)$ has a locally X-elliptic normal subgroup N such that the quotient $\text{Stab}_{\text{Aut}(X)}(\alpha)/N$ is finitely generated and virtually abelian of rank $\leq \dim(X)$.

The following consequence of Theorem B.1 is immediate.

COROLLARY B.2. *Let Γ be a finitely generated group satisfying the following two conditions.*

- (1) *Every finite index subgroup of Γ has finite abelianization.*
- (2) *For every Γ -action on a finite-dimensional CAT(0) cube complex, there is a finite Γ -orbit in the Roller compactification.*

Then every Γ -action on a finite-dimensional CAT(0) cube complex has a fixed point.

REMARK B.3. As pointed out to me by Talia Fernos, the converse statement to Corollary B.2 holds as well, namely: A finitely generated group Γ all of whose actions on finite-dimensional CAT(0) cube complexes have fixed points, automatically satisfies (1) and (2). Indeed, Property (2) is straightforward, while the existence of a finite index subgroup with infinite abelianization can be used to produce an unbounded action on the standard cubulation of the Euclidean n -space, for n large enough.

The proof of Theorem B.1 uses a relation between the Roller boundary and the *simplicial boundary* of X . The latter boundary was constructed by Hagen [29], under the hypothesis that X is finite-dimensional. Before reviewing its construction, we start with an abstract tool that will be used to produce the abelian quotient appearing in Theorem B.1.

Following Yves de Cornulier [22], we say that two subsets M, N of a set X are *commensurate* if their symmetric difference $M\Delta N$ is finite. We say that M is *commensurated* by the action of a group G acting on X if, for all $g \in G$, the sets M and gM are commensurate.

PROPOSITION B.4 ([22, Proposition 4.H.1]). *Let G be a group, X be a discrete G -set and $M \subset X$ be a commensurated subset.*

Then the map

$$tr_M: G \longrightarrow \mathbf{Z}: g \longmapsto \#(M \setminus g^{-1}M) - \#(g^{-1}M \setminus M)$$

is a homomorphism, called the transfer character. Moreover, for any $N \subset X$ commensurate to M and any $g \in G$, we have $tr_M(g) = tr_N(g)$.

We now turn back to our geometric setting: in the rest of this note we let X be a CAT(0) cube complex. Its set of halfspaces (respectively, hyperplanes) is denoted by $\mathfrak{H}(X)$ (respectively, $\mathfrak{W}(X)$). The map $\mathfrak{H}(X) \rightarrow \mathfrak{W}(X): h \mapsto \hat{h}$ associates to each halfspace its boundary hyperplane. In order to define the simplicial boundary, we recall some of Mark Hagen's terminology from [29]. A set of hyperplanes $\mathcal{U} \subset \mathfrak{W}(X)$ is called:

- (i) *inseparable* if each hyperplane separating two elements of \mathcal{U} belongs to \mathcal{U} ;
- (ii) the *inseparable closure* of a set of hyperplanes \mathcal{V} if \mathcal{U} is the smallest inseparable set containing \mathcal{V} ;
- (iii) *unidirectional* if for each $\hat{h} \in \mathcal{U}$, at least one halfspace bounded by \hat{h} contains only finitely many elements of \mathcal{U} ;
- (iv) a *facing triple* if \mathcal{U} consists of the three boundary hyperplanes of three pairwise disjoint halfspaces;
- (v) a *UBS* if \mathcal{U} is infinite, inseparable, unidirectional and contains no facing triple (UBS stands for *unidirectional boundary set*);
- (vi) a *minimal UBS* if every UBS \mathcal{U}' contained in \mathcal{U} is commensurate to \mathcal{U} ;
- (vii) *almost transverse* to a set of hyperplanes \mathcal{V} if each $\hat{h} \in \mathcal{U}$ crosses all but finitely many hyperplanes in \mathcal{V} , and each $\hat{k} \in \mathcal{V}$ crosses all but finitely many hyperplanes in \mathcal{U} .

The following result is due to Mark Hagen.

PROPOSITION B.5 ([29, Theorem 3.10]). *Assume that X is finite-dimensional.*

Given a UBS \mathcal{V} , there exists a UBS \mathcal{V}' commensurate to \mathcal{V} such that \mathcal{V}' is partitioned into a finite union of minimal UBS, say $\mathcal{U}_1 \cup \dots \cup \mathcal{U}_k$, where $k \leq \dim(X)$, such that, for $i \neq j \in \{1, \dots, k\}$, the set \mathcal{U}_i is almost transverse to \mathcal{U}_j .

Furthermore, if \mathcal{V}'' is a UBS which is commensurate to \mathcal{V} and is partitioned into a finite union of minimal UBS, say $\mathcal{U}'_1 \cup \dots \cup \mathcal{U}'_{k'}$, which are pairwise almost transverse, then $k = k'$ and, up to reordering, the set \mathcal{U}'_i is commensurate to \mathcal{U}_i for all i .

Following Hagen [29], the *simplicial boundary* of X , denoted by $\partial_\Delta X$, is defined (when X is finite-dimensional) as the abstract simplicial complex whose underlying poset is the set of commensuration classes of UBS, with the natural order relation induced by inclusion. Its vertices thus correspond to the commensuration classes of minimal UBS, and two vertices are adjacent if they are represented by two UBS that are almost transverse to one another. The set of simplices of $\partial_\Delta X$ is denoted by $\mathfrak{S}X$.

The following observation, which is implicit in [29] (see the proof of Lemma 3.7 in [29]), characterizes the minimal UBS.

LEMMA B.6. *Assume that X is finite-dimensional. Let $h_0 \supsetneq h_1 \supsetneq \dots$ be an infinite chain of halfspaces. Then the inseparable closure of $\{\hat{h}_i \mid i \geq 0\}$ is a minimal UBS. Moreover, every minimal UBS contains a (necessarily cofinite) UBS arising in this way.*

Proof. Let $\mathcal{V} = \{\hat{h}_i \mid i \geq 0\}$ and \mathcal{U} be the inseparable closure of \mathcal{V} . It is clear that \mathcal{U} is infinite and inseparable. Observe moreover that each hyperplane in \mathcal{U} separates \hat{h}_0 from \hat{h}_i for some sufficiently large i . This implies that \mathcal{U} is unidirectional and does not contain any facing triple. Thus \mathcal{U} is a UBS. Proposition B.5 implies that \mathcal{U} must be minimal.

That every minimal UBS arises in this way follows since any UBS contains an infinite set of pairwise disjoint hyperplanes by the finite dimensionality of X . □

We now briefly recall that definition of the Roller compactification, following Martin Roller [45]. A section $\alpha: \mathfrak{W}(X) \rightarrow \mathfrak{H}(X)$ of the map $h \mapsto \hat{h}$ is called an *ultrafilter* if, for every finite set $\mathcal{F} \subset \mathfrak{W}(X)$, the intersection $\bigcap_{\hat{h} \in \mathcal{F}} \alpha(\hat{h})$ is non-empty. An ultrafilter is *principal* if $\bigcap_{\hat{h} \in \mathfrak{W}(X)} \alpha(\hat{h})$ is non-empty, in which case that intersection contains a unique vertex. Conversely, every vertex v of X gives rise to a unique ultrafilter, also denoted by v , which maps each hyperplane \hat{h} to the unique halfspace bounded by \hat{h} and containing v . We identify henceforth each vertex with the corresponding principal ultrafilter. The collection of all ultrafilters is denoted by \overline{X} . The subset of non-principal ultrafilters is denoted by ∂X . With respect to the topology of pointwise convergence, the set $\overline{X} = X^{(0)} \cup \partial X$ is compact; it is a compactification of the (discrete) set of principal ultrafilters $X^{(0)}$. The set \overline{X} (respectively, ∂X) is called the *Roller compactification* (respectively, *Roller boundary*) of X . The following observation provides a link between the Roller boundary and the simplicial boundary.

LEMMA B.7. *Let $\alpha \in \partial X$. The following hold for all $x, y \in X^{(0)}$.*

- (1) *The set $\mathcal{U}(x, \alpha) = \{\hat{h} \in \mathfrak{W}(X) \mid x(\hat{h}) \neq \alpha(\hat{h})\}$ is a UBS.*
- (2) *The sets $\mathcal{U}(x, \alpha)$ and $\mathcal{U}(y, \alpha)$ are commensurate.*

In particular, the map $\Sigma: \partial X \rightarrow \mathfrak{S}X$, associating to α the commensuration class of the UBS $\mathcal{U}(x, \alpha)$, where x is a fixed vertex, is well-defined and $\text{Aut}(X)$ -equivariant.

Proof. (1) The set $\mathcal{U}(x, \alpha)$ is infinite since otherwise α would be principal, because x is so. That $\mathcal{U}(x, \alpha)$ is inseparable is clear. That $\mathcal{U}(x, \alpha)$ is unidirectional follows from the fact that x is principal. Finally, given any facing triple of hyperplanes, the maps x and α must coincide on at least one of them. Hence $\mathcal{U}(x, \alpha)$ is a UBS.

(2) Any hyperplane in the symmetric difference $\mathcal{U}(x, \alpha) \Delta \mathcal{U}(y, \alpha)$ separates x from y . Therefore, $\mathcal{U}(x, \alpha) \Delta \mathcal{U}(y, \alpha)$ is finite. □

The final ingredient needed for the proof of Theorem B.1 is the following result, due to Sageev.

PROPOSITION B.8. *Assume that X is finite-dimensional. Let $\Gamma \leq \text{Aut}(X)$ be a finitely generated group acting without any fixed point on X . Then there exists $\gamma \in \Gamma$ and $h \in \mathfrak{H}(X)$ such that $\gamma h \subsetneq h$. In particular, a (possibly infinitely generated) group Λ is locally X -elliptic if and only if every element of Λ has a fixed point in X .*

Proof. This follows from the proof of Theorem 5.1 in [47]. □

Proof of Theorem B.1. If $\alpha \in X^{(0)}$, then the desired conclusion is clear. We suppose henceforth that α belongs to ∂X . Lemma B.7 then provides a k -simplex $\sigma = \Sigma(\alpha)$ fixed by $\Gamma = \text{Stab}_{\text{Aut}(X)}(\alpha)$. Note that $k + 1 \leq \dim(X)$ by Proposition B.5. We denote the vertices of σ by v_0, \dots, v_k . For each j , the stabilizer $\text{Stab}_\Gamma(v_j)$ is of finite index in Γ , and commensurates any UBS representing v_j . We denote by χ_j the transfer character associated to this commensurating action by means of Proposition B.4. Thus the sum $\bigoplus_{j=0}^k \chi_j$ is a homomorphism to \mathbf{Z}^{k+1} which is defined on a finite index subgroup of Γ . We denote its kernel by Γ^0 , and claim that Γ^0 is locally X -elliptic. The desired conclusion follows from that claim.

By Proposition B.8, it suffices to show that every element of Γ^0 has a fixed point. Suppose for a contradiction that an element $g \in \Gamma^0$ has none. Applying Proposition B.8 to the cyclic group generated by g , we find a halfspace h and a positive integer n such that $g^n h \subsetneq h$. Set $\hat{h}_i = g^{ni} \hat{h}$ for all $i \in \mathbf{Z}$. Since g fixes α , the collection $\{\alpha(\hat{h}_i) \mid i \in \mathbf{Z}\}$ is $\langle g^n \rangle$ -invariant, and must therefore be a chain. Upon replacing g by g^{-1} , we may assume that $\alpha(\hat{h}_0) \supsetneq \alpha(\hat{h}_1)$. Let $x \in \alpha(\hat{h}_0) \setminus \alpha(\hat{h}_1)$ be a vertex. In particular $x(\hat{h}_i) \neq \alpha(\hat{h}_i)$ for all $i > 0$. It follows that the UBS $\mathcal{U}(x, \alpha)$, which represents $\sigma = \Sigma(\alpha)$, contains \hat{h}_i for all $i > 0$. We now apply Proposition B.5 to $\mathcal{U}(x, \alpha)$. This yields a finite set of pairwise almost transverse minimal UBS $\mathcal{U}_0, \dots, \mathcal{U}_k$ contained in $\mathcal{U}(x, \alpha)$, each representing a vertex v_j of σ , and such that the union $\bigcup_{j=0}^k \mathcal{U}_j$ is cofinite in $\mathcal{U}(x, \alpha)$. Since the hyperplanes \hat{h}_i have pairwise empty intersection, we deduce moreover from Proposition B.5 that there is some $j \in \{0, \dots, k\}$ such that $\hat{h}_i \in \mathcal{U}_j$ for all i larger than some fixed I . By Lemma B.6, we may assume that \mathcal{U}_j is the inseparable closure of $\{\hat{h}_i \mid i > I\}$. Since $g^n(\hat{h}_i) = \hat{h}_{i+1}$ for all i , we infer that g^n maps \mathcal{U}_j properly inside itself, thereby implying that $0 \neq \chi_j(g^n) = n\chi_j(g)$. This contradicts the fact that $g \in \Gamma^0 \leq \text{Ker}(\chi_j)$. \square

Acknowledgements. We thank the Forschungsinstitute für Mathematik at ETH, Zürich, the Institut Mittag-Leffler in Djursholm, Sweden, the Institut Henri Poincaré in Paris and the Mathematical Sciences Research Institute in Berkeley, CA for their hospitality. Part of this research was conducted while Indira Chatterji was at Jawaharlal Nehru University as a visiting NBHM professor, and she thanks them for their hospitality. Our thanks go also to Marc Burger, Ruth Charney and Nicolas Monod for useful discussions during various phases of the preparation of this paper. We are very thankful to Michah Sageev for having suggested a particular case of Lemma 2.6 that led to the proof of Theorem 1.9 and to Pierre-Emmanuel Caprace for pointing out Corollary 1.7 to us. We are grateful to Max Forester and Jing Tao for pointing out a mistake in an earlier version. Precisely, in an earlier version we used the concept of *tightly nested n -tuples of halfspaces*, but the cocycle obtained that way failed to be bounded in some cases, as, for instance, in Example 3.10. Pierre-Emmanuel Caprace enthusiastically thanks Yves de Cornulier for pleasant discussions and useful comments on Appendix B.

References

1. U. BADER, B. DUCHESNE and J. LÉCUREAUX, ‘Furstenberg maps for CAT(0) targets of finite telescopic dimension’, *Ergodic Theory Dynam. Systems* 1 (2016) 1–20.
2. U. BADER, A. FURMAN and A. SHAKER, ‘Superrigidity, Weyl groups, and actions on the circle’, Preprint, 2006, arXiv:math/0605276v4.
3. J. BEHRSTOCK and R. CHARNEY, ‘Divergence and quasimorphisms of right-angled Artin groups’, *Math. Ann.* 352 (2012) 339–356.
4. M. BESTVINA, K. BROMBERG and K. FUJIWARA, ‘Bounded cohomology with coefficients in uniformly convex Banach spaces’, *Comment. Math. Helv.*, to appear.
5. M. R. BRIDSON and A. HAEFLIGER, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences] 319 (Springer, Berlin, 1999).
6. 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.

7. M. BUCHER, M. BURGER and A. IOZZI, ‘A dual interpretation of the Gromov–Thurston proof of Mostow rigidity and volume rigidity for representations of hyperbolic lattices’, *Trends in harmonic analysis*, Springer INdAM Series 3 (Springer, Milan, 2013) 47–76.
8. M. BUCHER, M. BURGER and A. IOZZI, ‘Rigidity of representations of hyperbolic lattices $\Gamma < \mathrm{PSL}(2, \mathbf{C})$ into $\mathrm{PSL}(n, \mathbf{C})$ ’, Preprint, 2014, arXiv:1412.3428.
9. M. BURGER and A. IOZZI, ‘Boundary maps in bounded cohomology’. Appendix to: ‘Continuous bounded cohomology and applications to rigidity theory’ [*Geom. Funct. Anal.* 12 (2002) 219–280; by Burger and N. Monod], *Geom. Funct. Anal.* 12 (2002) 281–292.
10. M. BURGER and A. IOZZI, ‘Bounded Kähler class rigidity of actions on Hermitian symmetric spaces’, *Ann. Sci. École Norm. Sup. (4)* 37 (2004) 77–103.
11. M. BURGER and A. IOZZI, ‘Bounded cohomology and totally real subspaces in complex hyperbolic geometry’, *Ergodic Theory Dynam. Systems* 32 (2012) 467–478.
12. M. BURGER, A. IOZZI, F. LABOURIE and A. WIENHARD, ‘Maximal representations of surface groups: symplectic Anosov structures’, *Pure Appl. Math. Q.* 1 (2005) 543–590.
13. M. BURGER, A. IOZZI and A. WIENHARD, ‘Hermitian symmetric spaces and Kähler rigidity’, *Transform. Groups* 12 (2007) 5–32.
14. M. BURGER, A. IOZZI and A. WIENHARD, ‘Tight homomorphisms and Hermitian symmetric spaces’, *Geom. Funct. Anal.* 19 (2009) 678–721.
15. M. BURGER, A. IOZZI and A. WIENHARD, ‘Surface group representations with maximal Toledo invariant’, *Ann. of Math. (2)* 172 (2010) 517–566.
16. M. BURGER and N. MONOD, ‘Bounded cohomology of lattices in higher rank Lie groups’, *J. Eur. Math. Soc. (JEMS)* 1 (1999) 199–235.
17. M. BURGER and N. MONOD, ‘Continuous bounded cohomology and applications to rigidity theory’, *Geom. Funct. Anal.* 12 (2002) 219–280.
18. P.-E. CAPRACE and A. LYTCHAK, ‘At infinity of finite-dimensional CAT(0) spaces’, *Math. Ann.* 346 (2010) 1–21.
19. P.-E. CAPRACE and N. MONOD, ‘Fixed points and amenability in non-positive curvature’, *Math. Ann.* 356 (2013) 1303–1337.
20. P.-E. CAPRACE and M. SAGEEV, ‘Rank rigidity for CAT(0) cube complexes’, *Geom. Funct. Anal.* 21 (2011) 851–891.
21. W. W. COMFORT and S. NEGREPONTIS, *The theory of ultrafilters* (Springer, New York, 1974). Die Grundlehren der mathematischen Wissenschaften, Band 211.
22. Y. CORNULIER, ‘Group actions with commensurated subsets, wallings and cubings’, Preprint, 2013, arXiv:1302.5982.
23. Y. CORNULIER, ‘Irreducible lattices, invariant means, and commensurating actions’, *Math. Z.* 279 (2015) 1–26.
24. W. M. GOLDMAN, ‘Discontinuous groups and the Euler class’, Thesis, University of California at Berkeley, 1980.
25. W. M. GOLDMAN, ‘Characteristic classes and representations of discrete subgroups of Lie groups’, *Bull. Amer. Math. Soc. (N.S.)* 6 (1982) 91–94.
26. M. GROMOV, ‘Hyperbolic groups’, *Essays in group theory*, Mathematical Sciences Research Institute Publications 8 (Springer, New York, 1987) 75–263.
27. M. GROMOV, ‘Asymptotic invariants of infinite groups’, *Geometric group theory, vol. 2* (Sussex, 1991), London Mathematical Society Lecture Note Series 182 (Cambridge University Press, Cambridge, 1993) 1–295.
28. D. GURALNIK, ‘Coarse decomposition of boundaries for CAT(0) groups’, Preprint, 2006, arXiv:math/0611006v2.
29. M. F. HAGEN, ‘The simplicial boundary of a CAT(0) cube complex’, *Algebr. Geom. Topol.* 13 (2013) 1299–1367.
30. U. HAMENSTÄDT, ‘Isometry groups of proper CAT(0)-spaces of rank one’, *Groups Geom. Dyn.* 6 (2012) 579–618.
31. L. HERNÁNDEZ LAMONEDA, ‘Maximal representations of surface groups in bounded symmetric domains’, *Trans. Amer. Math. Soc.* 324 (1991) 405–420.
32. M. HULL and D. OSIN, ‘Induced quasicocycles on groups with hyperbolically embedded subgroups’, *Algebr. Geom. Topol.* 13 (2013) 2635–2665.
33. A. IOZZI, ‘Bounded cohomology, boundary maps, and rigidity of representations into $\mathrm{Homeo}_+(\mathbf{S}^1)$ and $\mathrm{SU}(1, n)$ ’, *Rigidity in dynamics and geometry* (Cambridge, 2000) (Springer, Berlin, 2002) 237–260.
34. V. A. KAIMANOVICH, ‘Double ergodicity of the Poisson boundary and applications to bounded cohomology’, *Geom. Funct. Anal.* 13 (2003) 852–861.
35. S. MATSUMOTO, ‘Some remarks on foliated S^1 bundles’, *Invent. Math.* 90 (1987) 343–358.
36. J. MILNOR, ‘On the existence of a connection with curvature zero’, *Comment. Math. Helv.* 32 (1958) 215–223.
37. I. MINEYEV, N. MONOD and Y. SHALOM, ‘Ideal bicombings for hyperbolic groups and applications’, *Topology* 43 (2004) 1319–1344.
38. N. MONOD, *Continuous bounded cohomology of locally compact groups*, Lecture Notes in Mathematics 1758 (Springer, Berlin, 2001).

39. N. MONOD, ‘Superrigidity for irreducible lattices and geometric splitting’, *J. Amer. Math. Soc.* 19 (2006) 781–814.
40. N. MONOD and Y. SHALOM, ‘Cocycle superrigidity and bounded cohomology for negatively curved spaces’, *J. Differential Geom.* 67 (2004) 395–455.
41. N. MONOD and Y. SHALOM, ‘Orbit equivalence rigidity and bounded cohomology’, *Ann. of Math.* (2) 164 (2006) 825–878.
42. G. MOUSSONG, *Hyperbolic Coxeter groups* (ProQuest LLC, Ann Arbor, MI, 1988). PhD Thesis, The Ohio State University.
43. A. NEVO and M. SAGEEV, ‘The Poisson boundary of CAT(0) cube complex groups’, *Groups Geom. Dyn.* 7 (2013) 653–695.
44. G. A. NIBLO and L. D. REEVES, ‘The geometry of cube complexes and the complexity of their fundamental groups’, *Topology* 37 (1998) 621–633.
45. M. ROLLER, ‘Poc sets, median algebras and group actions. An extended study of Dunwoody’s construction and Sageev’s theorem’, Preprint, 1998, <http://www.personal.soton.ac.uk/gan/Roller.pdf>.
46. P. ROLLI, ‘Notes on CAT(0) cube complexes’, Preprint, 2010.
47. M. SAGEEV, ‘Ends of group pairs and non-positively curved cube complexes’, *Proc. London Math. Soc.* (3) 71 (1995) 585–617.
48. P. SCOTT, ‘Ends of pairs of groups’, *J. Pure Appl. Algebra* 11 (1977/78) 179–198.
49. Z. SELA, ‘Uniform embeddings of hyperbolic groups in Hilbert spaces’, *Israel J. Math.* 80 (1992) 171–181.
50. J.-P. SERRE, ‘Amalgames et points fixes’, *Proceedings of the Second International Conference on the Theory of Groups (Australian National University, Canberra, 1973)*, Lecture Notes in Mathematics 372 (Springer, Berlin, 1974) 633–640.
51. J.-P. SERRE, *Trees* (Springer, Berlin, 1980). Translated from the French by John Stillwell.
52. Y. SHALOM, ‘Rigidity of commensurators and irreducible lattices’, *Invent. Math.* 141 (2000) 1–54.
53. D. TOLEDO, ‘Representations of surface groups in complex hyperbolic space’, *J. Differential Geom.* 29 (1989) 125–133.
54. R. J. ZIMMER, *Ergodic theory and semisimple groups*, Monographs in Mathematics 81 (Birkhäuser, Basel, 1984).

Indira Chatterji
 Université de Nice
 Laboratoire de Mathématiques J.A.
 Dieudonné
 06108 Nice cedex 02
 France

indira.chatterji@math.cnrs.fr

Alessandra Iozzi
 Department Mathematik
 ETHZ
 Rämistrasse 101
 CH-8092 Zürich
 Switzerland

iozzi@math.ethz.ch

Talia Fernós
 Department of Mathematics and Statistics
 University of North Carolina at Greensboro
 317 College Avenue
 Greensboro, NC 27412
 USA

t_fernos@uncg.edu

Pierre-Emmanuel Caprace
 UCLouvain
 1348 Louvain-la-Neuve
 Belgium

pierre-emmanuel.caprace@uclouvain.be