

RANK RIGIDITY FOR CAT(0) CUBE COMPLEXES

PIERRE-EMMANUEL CAPRACE* AND MICHAH SAGEEV†

ABSTRACT. We prove that any group acting essentially without a fixed point at infinity on an irreducible finite-dimensional CAT(0) cube complex contains a rank one isometry. This implies that the Rank Rigidity Conjecture holds for CAT(0) cube complexes. We derive a number of other consequences for CAT(0) cube complexes, including a purely geometric proof of the Tits Alternative, an existence result for regular elements in (possibly non-uniform) lattices acting on cube complexes, and a characterization of products of trees in terms of bounded cohomology.

CONTENTS

1. Introduction	2
1.1. Rank one and contracting isometries	2
1.2. Proof ingredients	3
1.3. Lattices and regular elements	4
1.4. Euclidean alternative	5
1.5. Tits alternative	5
1.6. Dynamics on the boundary	6
1.7. Quasi-morphisms and products of trees	6
1.8. Asymptotic cones	7
Acknowledgement	8
2. Preliminaries	8
2.1. CAT(0) cube complexes	8
2.2. Hyperplanes	8
2.3. Pocsets	9
2.4. Application: Orbit quotients and skewers	10
2.5. Application: Products	11
3. Essential things	12
3.1. Essential hyperplanes	12
3.2. Essential cube complexes	12
3.3. Essential hyperplanes relative to a group action	13
3.4. Essential actions and pruning	14
3.5. Finitely many orbits of hyperplanes	15
4. The Flipping Lemma	18
4.1. Version I: no fixed point at infinity	18
4.2. Non-skewering behaviours	19
4.3. Endometries of proper metric spaces	20
4.4. Version II: hereditarily essential actions	21

Date: First draft: March 2010; revised: February 2011.

Key words and phrases. Rank rigidity, rank one isometry, cube complex, CAT(0) space, Tits alternative.

* F.R.S.-FNRS research associate. Supported in part by FNRS grant F.4520.11.

†Supported in part by ISF grant #580/07.

5.	An irreducibility criterion	25
5.1.	Strongly separated hyperplanes	25
5.2.	Finding hyperplanes in cubical sectors	25
5.3.	Proof of the Irreducibility Criterion	29
6.	Rank rigidity	30
6.1.	Strongly separated hyperplanes and contracting isometries	30
6.2.	Proof of Rank Rigidity	31
7.	Applications	32
7.1.	Cube complexes with invariant Euclidean flats	32
7.2.	Tits alternative, second version	34
7.3.	Regular elements	34
7.4.	Quasi-morphisms	35
7.5.	Asymptotic cones	37
	References	37

1. INTRODUCTION

1.1. Rank one and contracting isometries. Let X be a complete CAT(0) space. A **rank one isometry** is a hyperbolic $g \in \text{Is}(X)$ none of whose axes bounds a flat half-plane. A **contracting isometry** is a hyperbolic $g \in \text{Is}(X)$ having some axis λ such that the diameter of the orthogonal projection to λ of any ball of X disjoint from λ is bounded above. A contracting isometry is always a rank one isometry. The converse is false in general, but it holds if the ambient space X is proper, see [BF09, Theorem 5.4]. If X is Gromov hyperbolic, then every hyperbolic isometry is contracting. Even if X is not hyperbolic, a contracting isometry g acts on the visual boundary $\partial_\infty X$ with a *North-South dynamics*: g has exactly two fixed points in $\partial_\infty X$, respectively called **attracting** and **repelling**, and the positive powers of g contract the whole boundary minus the repelling fixed point to the attracting one. This feature justifies the choice of terminology, and can be exploited to derive a large number of consequences from the very existence of contracting isometries; some of these will be reviewed below. To put it short, one can say that, in the presence of a contracting isometry, the space X presents some kind of hyperbolic behaviour.

The notion of rank one and contracting isometries originates in the celebrated Rank Rigidity Theorem for Hadamard manifolds, due to W. Ballmann, M. Brin, K. Burns, P. Eberlein, R. Spatzier (see [Bal95] for more information and detailed references). A possible formulation of this result, but not the most general one, is the following. Recall that a Hadamard manifold is a simply connected, complete Riemannian manifold of non-positive curvature, and that such a manifold is said to be irreducible if it does not admit a nontrivial decomposition as a metric product of two manifolds.

Rank Rigidity Theorem ([Bal85] and [BS87]). *Let M be a Hadamard manifold and Γ be a discrete group acting properly and cocompactly on M . If M is irreducible, then either M is a higher rank symmetric space or Γ contains a rank one isometry.*

It is expected that a similar phenomenon holds for much wider classes of CAT(0) spaces. In particular, W. Ballmann and S. Buyalo [BB08] formulate the following.

Rank Rigidity Conjecture. *Let X be a locally compact geodesically complete CAT(0) space and Γ be an infinite discrete group acting properly and cocompactly on X . If X is*

irreducible, then X is a higher rank symmetric space or a Euclidean building of dimension ≥ 2 , or Γ contains a rank one isometry.

Recall that a CAT(0) space is called **geodesically complete** if every geodesic segment can be prolonged to some (not necessarily unique) bi-infinite geodesic line. Besides the manifold case, the Rank Rigidity Conjecture has been confirmed for piecewise Euclidean cell complexes of dimension 2 by W. Ballmann and M. Brin (see [BB95]), who also obtained relevant partial results in dimension 3 (see [BB00]). It also holds within the class of buildings and Coxeter groups (see [CF10]) as well as for right-angled Artin groups (see [BC10]).

The main result of this paper is the following.

Theorem A (Rank Rigidity for finite-dimensional CAT(0) cube complexes). *Let X be a finite-dimensional CAT(0) cube complex and $\Gamma \leq \text{Aut}(X)$ be a group acting without fixed point in $X \cup \partial_\infty X$. Then there is a convex Γ -invariant subcomplex $Y \subseteq X$ such that either Y is a product of two unbounded cube subcomplexes or Γ contains an element acting on Y as a contracting isometry.*

A slightly more precise version will be obtained in Theorem 6.3 below. In particular, we shall see that if X is locally finite and if the Γ -action is cocompact (but not necessarily proper), then the conclusions of Theorem A still hold even if there are Γ -fixed points at infinity.

Notice that there is no assumption on geodesic completeness or locally finiteness of X in Theorem A. Moreover, the Γ -action need not be proper or cocompact. However, the result does not hold if one does not allow to pass to a convex subcomplex in the conclusion. In order to address this issue, we shall say that a group $\Gamma \leq \text{Aut}(X)$ acts **essentially** on X if no Γ -orbit remains in a bounded neighbourhood of a half-space of X . Starting with any group $\Gamma \leq \text{Aut}(X)$, there is a way to construct a Γ -invariant cube subcomplex $Y \subseteq X$ on which Γ acts essentially provided any of the following conditions are satisfied (see Propositions 3.5 and 3.12 below):

- Γ has no fixed point at infinity.
- Γ has finitely many orbits of hyperplanes.
- Γ is finitely generated and acts properly discontinuously.

In particular, if one assumes that the Γ -action on X is essential in Theorem A, then the conclusion holds with $Y = X$. In the precise setting of the Rank Rigidity Conjecture, namely for geodesically complete spaces with a discrete cocompact group action, it turns out that the hypotheses of Theorem A become redundant:

Corollary B (Rank Rigidity for geodesically complete CAT(0) cube complexes). *Let X be a locally compact geodesically complete CAT(0) cube complex and Γ be an infinite discrete group acting properly and cocompactly on X . Then X is a product of two geodesically complete unbounded convex subcomplexes or Γ contains a rank one isometry.*

Without the assumption of geodesic completeness, Corollary B stills holds provided one passes to a Γ -invariant subcomplex on which Γ acts essentially; it also holds for non-uniform lattices (see Corollary 6.4 below).

1.2. Proof ingredients. An important elementary fact on which our proof of Rank Rigidity for CAT(0) cube complexes relies is the following tool allowing to **flip** any essential half-space \mathfrak{h} , *i.e.* to map \mathfrak{h} to a half-space \mathfrak{d} containing properly the complementary half-space \mathfrak{h}^* of \mathfrak{h} .

Flipping Lemma. *Let X be a finite-dimensional CAT(0) cube complex and $\Gamma \leq \text{Aut}(X)$ be a group acting essentially without a fixed point at infinity. Then for any half-space \mathfrak{h} , there is some $\gamma \in \Gamma$ such that $\mathfrak{h}^* \subsetneq \gamma \cdot \mathfrak{h}$.*

A slightly more precise version of the Flipping Lemma will be proved in Theorem 4.1 below.

An isometry $g \in \text{Is}(X)$ is said to **skewer** a hyperplane $\hat{\mathfrak{h}}$ bounding some half-space \mathfrak{h} if $g \cdot \mathfrak{h} \subsetneq \mathfrak{h}$ or $\mathfrak{h} \subsetneq g \cdot \mathfrak{h}$. If X is finite-dimensional, then, given $\Gamma \leq \text{Aut}(X)$, the Γ -action is essential if and only if every half-space is skewered by some isometry (see Proposition 3.2 below). This should be compared to the following consequence of the Flipping Lemma, which is an essential and flexible tool to construct ‘many’ hyperbolic elements in Γ .

Double Skewering Lemma. *Let X be a finite-dimensional CAT(0) cube complex and $\Gamma \leq \text{Aut}(X)$ be a group acting essentially without fixed point at infinity.*

Then for any two half-spaces $\mathfrak{d} \subset \mathfrak{h}$, there exists $\gamma \in \Gamma$ such that $\gamma \mathfrak{h} \subsetneq \mathfrak{d} \subset \mathfrak{h}$.

Proof. By the Flipping Lemma, we can find an element g flipping \mathfrak{d} to obtain $g\mathfrak{h}^* \subset g\mathfrak{d}^* \subset \mathfrak{d} \subset \mathfrak{h}$. Invoking the Flipping Lemma again, we then find an element a flipping $g\mathfrak{h}^*$ and, setting $\gamma = ag$, we obtain $\gamma \mathfrak{h} \subset g\mathfrak{h}^* \subset \mathfrak{d}$, as required. \square

We remark that the statement of the Double Skewering fails if one allows fixed points at infinity. A good example illustrating this is the fixator of a point ξ at infinity in the full automorphism group of a regular tree T . This group acts edge-transitively, hence essentially. Every hyperplane is skewered by a hyperbolic isometry. However, a pair $\mathfrak{d} \subset \mathfrak{h}$ of half-spaces such that \mathfrak{h} contains ξ but \mathfrak{d} does not, is not skewered by a common hyperbolic isometry.

A final ingredient in the proof of Rank Rigidity for CAT(0) cube complexes is a criterion allowing us to recognize when a CAT(0) cube complex X is **irreducible**, *i.e.* when it does not split as a product of non-trivial convex subcomplexes. A pair of hyperplanes $\hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_2$ in X is called **strongly separated** if no hyperplane crosses both $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$; in particular $\hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_2$ must be disjoint. This notion is due to J. Behrstock and R. Charney [BC10].

Irreducibility Criterion. *Let X be a finite-dimensional unbounded CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially without a fixed point at infinity.*

Then X is irreducible if and only if there is a pair of strongly separated hyperplanes.

A slightly more precise version will be established in Proposition 5.1 below.

With these tools at hand, Theorem A will be deduced from the simple observation that a hyperbolic isometry which is obtained by applying the Double Skewering Lemma to a strongly separated pair of hyperplanes must necessarily be a contracting isometry (see Lemma 6.2 below).

We now proceed to describe various applications of Theorem A.

1.3. Lattices and regular elements. When it exists, a contracting isometry of a CAT(0) space X can be viewed as an analogue of a **regular semi-simple element** in the classical case of symmetric spaces, *i.e.* a semi-simple element whose centralizer has minimal possible dimension in the ambient Lie group $\text{Is}(X)$. Indeed, if X is a proper CAT(0) space and $g \in \text{Is}(X)$ has rank one, then the centralizer $\mathcal{L}_{\text{Is}(X)}(g)$ is of dimension ≤ 1 in the sense that it is isomorphic to a closed subgroup of the isometry group of the real line \mathbf{R} .

An important property of regular semi-simple elements in the classical setting is that they form a Zariski open subset of the full isometry group. In particular, combining this with the Borel Density Theorem, it follows that any lattice in a symmetric space contains a regular semi-simple element. The following result is an analogue of that fact.

Theorem C (Existence of regular elements). *Let $X = X_1 \times \cdots \times X_n$ be a product of n irreducible unbounded locally compact CAT(0) cube complexes such that $\text{Aut}(X_i)$ acts cocompactly and essentially on X_i for all i .*

Then any (possibly non-uniform) lattice $\Gamma \leq \text{Aut}(X)$ contains an element $\gamma \in \Gamma$ which acts as a rank one isometry on each irreducible factor X_i .

As a consequence, one deduces the following result related to the Flat Closing Conjecture.

Corollary D. *Let X be a locally compact CAT(0) cube complex and Γ be a discrete group acting cocompactly on X . If X is a product of n unbounded cube subcomplexes, then Γ contains a subgroup isomorphic to \mathbf{Z}^n .*

1.4. Euclidean alternative. A simple-minded application of the Rank Rigidity of CAT(0) cube complexes is the following characterisation of those CAT(0) cube complexes whose full automorphism group stabilises some isometrically embedded Euclidean flat. By a **facing triple** of hyperplanes, we mean a triple of hyperplanes associated to a triple of pairwise disjoint half-spaces.

Theorem E. *Let X be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially without fixed point at infinity. Then either $\text{Aut}(X)$ stabilises a Euclidean flat (possibly reduced to a single point) or there is a facing triple of hyperplanes.*

In the special case when X is locally compact and $\text{Aut}(X)$ acts cocompactly, the same holds even if $\text{Aut}(X)$ has some fixed points at infinity, see Theorem 7.2 below.

1.5. Tits alternative. One of the most natural uses of the strong dynamical properties of contracting isometries is to produce Schottky pairs generating (discrete) free subgroups via the Ping Pong Lemma. In particular, there is a relation between the Rank Rigidity Conjecture for proper CAT(0) spaces and the Tits Alternative, which constitutes another major open problem in this area.

In the case of finite-dimensional CAT(0) cube complexes, a version of the Tits Alternative was already obtained by M. Sageev and D. Wise [SW05], not using contracting elements but relying instead on the Algebraic Torus Theorem.

Here we present two other versions of the Tits Alternative relying on Rank Rigidity. We emphasize that our approach is here purely geometric; most arguments use only the combinatorics of hyperplanes in CAT(0) cube complexes.

Theorem F (Tits Alternative, first version). *Let X be a finite-dimensional CAT(0) cube complex and let $\Gamma \leq \text{Aut}(X)$. Then Γ has a finite index subgroup fixing a point in $X \cup \partial_\infty X$ or Γ contains a non-Abelian free subgroup.*

Proof. Assume that Γ does not virtually fix any point in $X \cup \partial_\infty X$. It follows that there is a Γ -invariant convex subcomplex $Y \subset X$ on which Γ acts essentially and without a fixed point at infinity (see Proposition 3.5 below).

Every finite-dimensional CAT(0) cube complex Y has a canonical decomposition $Y = Y_1 \times \cdots \times Y_n$ as a finite product of irreducible subcomplexes, which is preserved by the full automorphism group $\text{Aut}(Y)$ up to permutations of possibly isomorphic factors (see Proposition 2.6 below). After replacing Γ by a finite index subgroup, we obtain in particular an essential action of Γ on each irreducible factor Y_i which does not fix any point at infinity.

Assume that for some i , the cube complex Y_i does not contain any Euclidean flat which is invariant under $\text{Aut}(Y_i)$. Then Theorem E ensures the existence of a facing triple of hyperplanes in Y_i . We then invoke the Flipping Lemma, which shows that there actually exists a facing *quadruple* of hyperplanes. We group these four hyperplanes into two pairs

and apply the Double Skewering Lemma to each of them. This provides a Schottky pair of hyperbolic isometries, from which the existence of a non-Abelian free subgroup in Γ follows via the Ping Pong Lemma.

Assume now that for all i , the factor Y_i contains some $\text{Aut}(Y_i)$ -invariant flat $F_i \subset Y_i$. Since Y_i is irreducible, it follows that F_i is one-dimensional (see Lemma 7.1 below). Thus $\text{Aut}(Y_i)$ has an index two subgroup which fixes a point at infinity, contradicting our assumption on Γ . \square

The first alternative in Theorem F might look unsatisfactory in the sense that it does not provide any algebraic information on Γ . Notice however that no assumption on local compactness of X is made. Therefore, the above statement is optimal because *any* group admits an action on a tree X with a global fixed point (take X to be an infinite star on which Γ acts by fixing the root and permuting the branches). It is nevertheless possible to combine Theorem F with some results from [CL10] to obtain a more precise description of Γ in the case of a proper action.

Corollary G (Tits Alternative, second version). *Let X be a finite-dimensional CAT(0) cube complex and Γ be a discrete group acting properly on X . Then either Γ is {locally finite}-by-{virtually Abelian} or Γ contains a non-Abelian free subgroup.*

This is a slight improvement of the statement from [SW05].

1.6. Dynamics on the boundary. It is a well known fact that if a group Γ acts on a proper CAT(0) space X and contains a rank one isometry, then the Γ -action on its limit set is **topologically minimal** (*i.e.* every orbit is dense) and the limit set is the unique Γ -minimal subset of the boundary $\partial_\infty X$ (provided the Γ -action is **non-elementary**, *i.e.* there is no global fixed point or fixed pair at infinity). We refer to [BB08] and [Ham09] for more information.

In fact, according to recent results by Gabriele Link [Lin10], even if X is a product space, it is possible to obtain very precise dynamical information on the Γ -action on its limit set if one knows that Γ contains an element which acts as a rank one isometry on each irreducible factor of X . In particular, when X is a CAT(0) cube complex, all the results recently obtained in the aforementioned paper by Gabriele Link apply to CAT(0) cube complexes under the hypotheses of Theorem C. We shall not repeat these statements here; the interested reader should consult [Lin10] and references therein.

1.7. Quasi-morphisms and products of trees. Another way to exploit the peculiar dynamical properties of contracting isometries has recently been elaborated by M. Bestvina and K. Fujiwara [BF09] in order to construct quasi-morphisms. In particular, combining the Bestvina–Fujiwara construction with Theorem A, it follows that *a discrete group acting properly, essentially and without fixed point at infinity on a locally compact finite-dimensional CAT(0) cube complex is either virtually Abelian or has an infinite-dimensional space of non-trivial quasi-morphisms*. Finer results applying also to non-discrete actions on non-proper spaces can be further obtained; we shall not give more details in this direction here. Instead, we present the following “*bounded cohomological characterisation*” of product of trees, which is in the same spirit as the main result from [BF09].

Theorem H. *Let $X = X_1 \times \cdots \times X_n$ be a product of n irreducible locally compact geodesically complete CAT(0) cube complexes such that $\text{Aut}(X_i)$ acts cocompactly on X_i for all i . For any (possibly non-uniform) lattice $\Gamma \leq \text{Aut}(X_1) \times \cdots \times \text{Aut}(X_n)$, the following conditions are equivalent.*

- (i) $\widetilde{\text{QH}}(\Gamma)$ is finite-dimensional.
- (ii) $\widetilde{\text{QH}}(\Gamma) = 0$.
- (iii) For all i , the space X_i is a semi-regular tree, and if X_i is not isometric to the real line, then the closure G_i of the projection of Γ to $\text{Aut}(X_i)$ is doubly transitive on $\partial_\infty X_i$.

Recall that $\widetilde{\text{QH}}(\Gamma)$ denotes the real vector space of quasi-morphisms of Γ modulo the subspace of trivial quasi-morphisms, which is the direct sum of the space of bounded functions on Γ and the space of genuine morphisms $\Gamma \rightarrow \mathbf{R}$. The space $\widetilde{\text{QH}}(\Gamma)$ coincides with the kernel of the canonical map from the bounded cohomology to the usual cohomology of Γ in degree two with trivial coefficients.

We point out that there is a version of Theorem H which does not require the assumption of geodesic completeness. In that case the assertion (iii) must be replaced by the fact that X_i is equivariantly quasi-isometric to some tree on which G_i acts naturally by automorphisms (see Remark 7.8 below).

As we shall see below, the details of the proof of Theorem H have little to do with CAT(0) cube complexes, but rather depend on a combination of Rank Rigidity with general arguments on quasi-morphisms and CAT(0) spaces with rank one isometries (see notably Theorem 7.4 below).

1.8. Asymptotic cones. Another typical consequence of the existence of rank one isometries concerns the aqf the corresponding groups and spaces. This phenomenon was studied in detail by C. Drutu, Sh. Mozes and M. Sapir [DMS10]. Relying on their work, it is straightforward to deduce the following consequence of Rank Rigidity.

Corollary I. *Let X be a locally compact CAT(0) cube complex such that $\text{Aut}(X)$ acts cocompact, essentially and without a fixed point at infinity. Then*

- (i) X is irreducible if and only if every asymptotic cone of X has a cut-point.
- (ii) If X is irreducible and $\Gamma \leq \text{Aut}(X)$ is a finitely generated group acting essentially without a fixed point at infinity, then every asymptotic cone of Γ has a cut-point.

The paper is organized as follows. Sections 2 and 3 are of preliminary nature. They collect a number of useful general facts on groups acting on CAT(0) cube complexes, most of which are well-known to the experts. The new material is exposed in the following sections. The proof of the Flipping Lemma is given in Section 4 and we recommend to readers who have some familiarity with CAT(0) cube complexes to start reading this paper from that point on. In fact Section 4 contains two distinct (and conceptually different) proofs of the Flipping Lemma, the first one shorter and more conceptual, the second more pedestrian but entirely self-contained. Section 5 is devoted to a criterion allowing one to recognise irreducible CAT(0) cube complexes. With this criterion at hand, we complete the proof of the Rank Rigidity theorem in Section 6 and present the applications in the final section.

Let us finally point out that part of the technical difficulties in this paper come from the fact that we have tried to let our arguments work in the most general setting possible, rather than focusing on the special case of a locally compact CAT(0) cube complex endowed with a proper cocompact action of some discrete group. The reader who is primarily interested in the latter situation will realise that many of our discussions can be simplified to a large extent in that specific case; we nevertheless decided to include discussions of finite-dimensional spaces that are possibly non-proper, or of lattices that are possibly non-uniform. This level

of generality is useful even in the study of proper actions. For example, proper actions on products may descend to improper actions on their factors.

Acknowledgement. The first-named author would like to thank the Moshav Shorashim Research Institute for its hospitality while a large part of the present work was accomplished. We express our gratitude to Frédéric Haglund and the referee for long lists of detailed comments on an earlier version of this manuscript, which helped much in improving the presentation.

2. PRELIMINARIES

We first recall some basic facts about cube complexes and their connection to the hyperplanes and half-plane systems. For more details see [Rol98], [Nic04], [CN05], [Gur05] or [Sag95].

2.1. CAT(0) cube complexes. Let X be a CAT(0) cube complex, *i.e.* a simply connected cell complex all of whose cells are Euclidean cubes with edge length one, and such that the link of each vertex is a flag complex. A theorem of M. Gromov (see *e.g.* Theorem II.5.20 from [BH99] for the finite-dimensional case and Theorem 40 from [Lea] for the general case) ensures that X , endowed with the induced length metric, is a CAT(0) space. Moreover X is complete if and only if X does not contain an ascending chain of cells (see [Lea, Theorem 31]); in particular, this is the case if X is **locally finite-dimensional**, in the sense that the supremum of the dimensions of cubes containing any given vertex is finite. We shall later focus on CAT(0) cube complexes which are **finite-dimensional**, which means that the supremum of the dimensions of all cubes in X is finite. The space X is locally compact if and only if it is locally finite, in the sense that every vertex has finitely many neighbours.

Unless specified otherwise, we shall view X as a metric space with respect to its natural CAT(0) metric (rather than focusing on the 1-skeleton of X endowed with the combinatorial metric). In some cases, the convexity properties of the CAT(0) metric constitute an advantage compared to the combinatorial metric (this is for example the case in the proof of Theorem 4.1 below), while in other cases, the combinatorial metric is easier to deal with by its very combinatorial nature (as in Lemma 6.1 below). However, both viewpoints are very often equivalent; we refer to Lemma 2.2 below for a more precise statement.

A subcomplex Y of X is called **convex** if it is convex as a CAT(0) subspace. In that case Y is itself a CAT(0) cube complex.

If Y and Z are two CAT(0) cube complexes, then the Cartesian product $X = Y \times Z$ is naturally endowed with the structure of a CAT(0) cube complex: the 1-skeleton of X is the graph theoretical Cartesian product of the 1-skeletons of Y and Z , and the metric on X coincides with $\sqrt{d_Y^2 + d_Z^2}$, where d_Y and d_Z are the respective metrics on Y and Z . For all vertices $y \in Y$ and $z \in Z$, the subcomplexes $Y \times \{z\}$ and $\{y\} \times Z$ are convex in X .

By a **cubical map** between cube complexes, we mean a cellular map so that the restriction $\sigma \rightarrow \tau$ between cubes factors as $\sigma \rightarrow \nu \rightarrow \tau$, where the first map $\sigma \rightarrow \nu$ is a natural projection onto a face of σ and the second map $\nu \rightarrow \tau$ is an isometry.

2.2. Hyperplanes. A key feature of CAT(0) cube complexes is the existence of **hyperplanes**. In fact, the midpoint of each edge in the 1-skeleton of X belongs to a unique hyperplane. A hyperplane \hat{h} is a closed convex subspace which has the property that its complement $X \setminus \hat{h}$ has exactly two connected components, both of which are convex. The closure of each of these two components is called a **half-space**. The set of hyperplanes of

X is denoted by $\hat{\mathcal{H}}(X)$ and the set of half-spaces by $\mathcal{H}(X)$. The hyperplane associated to a half-space \mathfrak{h} is denoted by $\hat{\mathfrak{h}}$ and the half-space complementary to \mathfrak{h} by \mathfrak{h}^* . By convention, we shall always use gothic characters to denote half-spaces.

We will mostly make the assumption that X is finite-dimensional. One way this assumption can be used is through the fact that any collection of pairwise crossing hyperplanes has a non-empty intersection (see [Sag95, Theorem 4.14]). Therefore, the cardinality of this set is bounded above by $\dim(X)$. This fact shall often be used through the following.

Lemma 2.1. *Let X be a finite-dimensional CAT(0) cube complex. For each $k > 0$, there is some $R(k) > 0$ such that if a geodesic path α in X crosses at least $R(k)$ hyperplanes, then α crosses a **pencil of hyperplanes** of cardinality at least k , i.e. a collection of at least k disjoint hyperplanes which are associated to a collection of k nested half-spaces.*

Proof. In view of Ramsey's theorem, it suffices to define $R(k)$ as the Ramsey number $R(k, \dim(X) + 1)$. \square

This implies the following basic observation.

Lemma 2.2. *Let X be a finite-dimensional CAT(0) cube complex. Then X is quasi-isometric to its 1-skeleton endowed with the combinatorial metric.*

Proof. Let d denote the CAT(0) metric on X and d_{ℓ^1} denote the combinatorial metric on the vertex set $X^{(0)}$. Let also $n = \dim(X)$. Then every points of X is at distance at most $\sqrt{n}/2$ from a vertex of X . Thus it suffices to show that $(X^{(0)}, d)$ is quasi-isometric to $(X^{(0)}, d_{\ell^1})$.

The definition of d implies that $d(x, y) \leq d_{\ell^1}(x, y)$ for all $x, y \in X^{(0)}$. Moreover $d_{\ell^1}(x, y)$ coincides with the number of hyperplanes separating x from y . By Lemma 2.1, there is a constant $R = R(2)$ such that any collection of R hyperplanes contains a pair of non-crossing hyperplanes. The minimal distance separating a pair of non-crossing hyperplanes in a CAT(0) cube complex is 1. This implies that for all vertices $x, y \in X^{(0)}$, we have $d(x, y) \geq d_{\ell^1}(x, y)/R$. Thus $d(x, y) \leq d_{\ell^1}(x, y) \leq Rd(x, y)$, as desired. \square

2.3. Pocsets. As before, we assume for this discussion that X is a finite-dimensional CAT(0) cube complex, although some of what is discussed below can be generalized to a somewhat broader setting. The set of half-spaces is a poset under inclusion and comes equipped with a natural order reversing involution $\mathfrak{h} \rightarrow \mathfrak{h}^*$ (hence is called a **pocset**, see [Rol98]). This pocset also satisfies the **finite interval condition**, meaning that if $\mathfrak{h}_1 \subset \mathfrak{h}_2$ are half-spaces, there are only finitely many half-spaces \mathfrak{h} satisfying $\mathfrak{h}_1 \subset \mathfrak{h} \subset \mathfrak{h}_2$. The finite dimensionality of X is reflected in the fact that the pocset has **finite width**, meaning that there is a bound on the size of a collection of **transverse** elements: halfspaces associated to intersecting hyperplanes.

Now there is also a dual construction which shows that X can be completely reconstructed from the pocset $\mathcal{H}(X)$. Given a pocset Σ satisfying the finite interval condition and having finite width, one can construct a CAT(0) cube complex $X = X(\Sigma)$ whose half-space system is naturally isomorphic to Σ . It is then an easy exercise to check that the vertices of a X are in 1-1 correspondence with ultrafilters on Σ which satisfy the descending chain condition. Namely a vertex v in X^0 may be viewed as a subset of Σ which satisfies

- (1) For every involutory invariant pair of elements $\{A, A^*\}$, exactly one of them is in v .
- (2) If $A < B$ and $A \in v$ then $B \in v$.
- (3) Every descending chain of elements of v terminates.

One joins two such vertices by an edge when they differ on a single involutory pair. That is, v and w are joined by an edge if there exists a \mathfrak{h} such that $w = (v - \mathfrak{h}) \cup \mathfrak{h}^*$. One

then attaches a cube whenever the 1-skeleton of one appears in X^1 . It is then a theorem that the resulting cube complex X is CAT(0). (By convention, if the pocset is empty then X is reduced to a single vertex.) In fact, a theorem of Roller [Rol98] then tells us that CAT(0) cube complexes and pocsets are actually dual to one another. If one starts with a finite dimensional CAT(0) cube complex X , then $X(\mathcal{H}(X)) = X$. That is the cube complex constructed from the pocset of half-spaces associated to the hyperplanes of X is X . Conversely, for any pocset Σ , we have $\mathcal{H}(X(\Sigma)) = \Sigma$. That is, if one starts with a pocset and builds a cube complex, the pocset of half-spaces associated to the cube complex is again the original pocset.

Now let X be a CAT(0) cube complex, let $\hat{\mathcal{H}} = \hat{\mathcal{H}}(X)$ be the collection of hyperplanes of X , and let $\hat{\mathcal{K}} \subset \hat{\mathcal{H}}$ be some subcollection of hyperplanes. One can carry out the above construction to build a new cube complex $X(\hat{\mathcal{K}})$ whose half-space system is isomorphic to $\hat{\mathcal{K}}$, the collection of half-spaces bounded by some $\hat{\mathfrak{h}} \in \hat{\mathcal{K}}$. One then obtains a natural quotient map $X \rightarrow X(\hat{\mathcal{K}})$. Namely a vertex in X corresponds to a choice of half-spaces in \mathcal{H} and thus gives rise to a choice of half-spaces on $\hat{\mathcal{K}}$, and it is immediate to see that the resulting ultrafilter $\hat{\mathcal{K}}$ is a vertex of $X(\hat{\mathcal{K}})$. This map on vertices is easily seen to extend to a cubical map $X \rightarrow X(\hat{\mathcal{K}})$. We call this map the **restriction quotient** arising from the subset $\hat{\mathcal{K}} \subset \hat{\mathcal{H}}$. If X furthermore comes equipped with a group action Γ and $\hat{\mathcal{K}}$ is a Γ -invariant subset of $\hat{\mathcal{H}}$, then the restriction quotient is naturally equipped with a Γ -action and the quotient map is equivariant. For example, if $\hat{\mathcal{K}}$ is the set of all hyperplanes that cross some given hyperplane $\hat{\mathfrak{h}}$, then $X(\hat{\mathcal{K}})$ is isometric to the hyperplane $\hat{\mathfrak{h}}$ endowed with its canonical structure of a cube complex inherited from X .

It is important to remark that the dimension of the cube complex $X(\hat{\mathcal{K}})$ is bounded above by $\dim(X)$, but that $X(\hat{\mathcal{K}})$ need not be proper even if X is so.

We now consider two other applications of this construction.

2.4. Application: Orbit quotients and skewers. Let $\Gamma \leq \text{Aut}(X)$ be any group of automorphisms of a CAT(0) cube complex X . We consider the orbit $\Gamma\hat{\mathfrak{h}}$ of a single hyperplane $\hat{\mathfrak{h}} \in \hat{\mathcal{H}}$. We may then form the restriction quotient $X(\Gamma\hat{\mathfrak{h}})$, which we call the **orbit quotient** of $\hat{\mathfrak{h}}$.

In order to present a useful application of this construction, we first introduce an important definition. Given a hyperplane $\hat{\mathfrak{h}} \in \hat{\mathcal{H}}(X)$ and an isometry $\gamma \in \text{Aut}(X)$, we say that γ **skewers** $\hat{\mathfrak{h}}$ (or that $\hat{\mathfrak{h}}$ is **skewered** by γ) if there exist a non-zero integer n and a half-space \mathfrak{h} bounded by $\hat{\mathfrak{h}}$ such that $\gamma^n\mathfrak{h} \subsetneq \mathfrak{h}$ (the proper inclusion is essential).

Notice that an element γ which skewers some hyperplane necessarily has positive translation length. In particular, if X is finite-dimensional, then γ acts as a hyperbolic isometry (in the usual sense of CAT(0) geometry). In that case, we have the following criterion.

Lemma 2.3. *Let X be a finite-dimensional CAT(0) cube complex, let $\gamma \in \text{Aut}(X)$ be a hyperbolic isometry and $\hat{\mathfrak{h}} \in \hat{\mathcal{H}}(X)$. The following conditions are equivalent.*

- (i) γ skewers $\hat{\mathfrak{h}}$.
- (ii) Every γ -axis meets the hyperplane $\hat{\mathfrak{h}}$ in a single point.
- (iii) Some γ -axis meets the hyperplane $\hat{\mathfrak{h}}$ in a single point.

Proof. First, note that an axis for γ and a hyperplane can intersect in at most one point, unless the axis is contained in the hyperplane.

The equivalence of (ii) and (iii) is clear as any two axes for a given hyperbolic isometry are a bounded distance apart. To see (i) \Rightarrow (iii) suppose that \mathfrak{h} is such that $\gamma^n\mathfrak{h} \subsetneq \mathfrak{h}$. Let ℓ

be an axis for γ and let p be a point on ℓ which is not in $\hat{\mathfrak{h}}$. Without loss of generality let us assume that $p \in \mathfrak{h}$. Let A be some geodesic arc joining p and $\hat{\mathfrak{h}}$. Note that A crosses only finitely many hyperplanes so that for some $m > 0$, we have that $\gamma^{nm}\hat{\mathfrak{h}} \cap A = \emptyset$. It follows that $A \subset \gamma^{mn}(h^*)$. Thus $p \in \gamma^{nm}\mathfrak{h}^*$ which implies that $\gamma^{-nm}(p) \in \mathfrak{h}^*$. Since p was chosen on ℓ and ℓ is γ -invariant, it follows that ℓ meets both \mathfrak{h} and \mathfrak{h}^* , as required.

Conversely, assume that some axis ℓ of γ meets $\hat{\mathfrak{h}}$ in a single point p . Consider the collection of positive powers $\{\gamma^n\hat{\mathfrak{h}}|n > 0\}$. This is an infinite collection of hyperplanes since the collection $\{\gamma^n p|n > 0\}$ is infinite. By finite dimensionality, there exists some $n, m > 0$ such that $\gamma^n\hat{\mathfrak{h}} \cap \gamma^m\hat{\mathfrak{h}} = \emptyset$. Thus, there exists some positive power l such that $\gamma^l\hat{\mathfrak{h}} \cap \hat{\mathfrak{h}} = \emptyset$. If \mathfrak{h} is the halfspace bounded by $\hat{\mathfrak{h}}$ containing $\gamma^l(p)$, then we have that $\gamma^l\mathfrak{h} \subset \mathfrak{h}$, as required. \square

We move on with the following useful description of orbit quotients.

Lemma 2.4. *Let X be a finite-dimensional CAT(0) cube complex, let $\Gamma \leq \text{Aut}(X)$ and $\hat{\mathfrak{h}} \in \hat{\mathcal{H}}(X)$. We have the following.*

- (i) $X(\Gamma\hat{\mathfrak{h}})$ is finite if and only if the orbit $\Gamma\hat{\mathfrak{h}}$ is finite.
- (ii) $X(\Gamma\hat{\mathfrak{h}})$ is bounded if and only if Γ has a fixed point in $X(\Gamma\hat{\mathfrak{h}})$.
- (iii) $X(\Gamma\hat{\mathfrak{h}})$ is unbounded if and only if $\hat{\mathfrak{h}}$ is skewed by some hyperbolic isometry $\gamma \in \Gamma$.

Proof. The point (i) is clear, and so is the ‘only if’ of (ii). Suppose that Γ fixes a point v in $X(\Gamma\hat{\mathfrak{h}})$. If v belong to some hyperplane of $X(\Gamma\hat{\mathfrak{h}})$, then every hyperplane of $X(\Gamma\hat{\mathfrak{h}})$ contains v . This implies that $X(\Gamma\hat{\mathfrak{h}})$ consists of a single cube, which is bounded since $\dim(X(\Gamma\hat{\mathfrak{h}})) \leq \dim(X) < \infty$. Otherwise we may assume that v is a vertex of $X(\Gamma\hat{\mathfrak{h}})$. Then every half-space of $X(\Gamma\hat{\mathfrak{h}})$ separates v from a neighbouring vertex. This implies that any minimal path in the 1-skeleton of $X(\Gamma\hat{\mathfrak{h}})$ emanating from v remains in a single cube containing v . Since the dimension of such a cube is bounded, we deduce that the 1-skeleton of $X(\Gamma\hat{\mathfrak{h}})$ is contained in a finite ball centred at v . Hence $X(\Gamma\hat{\mathfrak{h}})$ is bounded as desired.

It remains to prove (iii). If γ skews $\hat{\mathfrak{h}}$, then $\langle\gamma\rangle.\hat{\mathfrak{h}}$ contains an infinite chain of half-spaces by Lemma 2.1 and, hence the cube complex $X(\langle\gamma\rangle.\hat{\mathfrak{h}})$ is unbounded. Consequently, so is $X(\Gamma\hat{\mathfrak{h}})$. Assume conversely that $X(\Gamma\hat{\mathfrak{h}})$ is unbounded. Let d denote its dimension. By Lemma 2.1, a geodesic edge path in $X(\Gamma\hat{\mathfrak{h}})$ whose length is greater than the constant $R(3)$ necessarily crosses three pairwise disjoint hyperplanes $\hat{\mathfrak{h}}_1$, $\hat{\mathfrak{h}}_2$ and $\hat{\mathfrak{h}}_3$. Since there is a single orbit of hyperplanes in $X(\Gamma\hat{\mathfrak{h}})$, we infer that two of the half-spaces bounding the $\hat{\mathfrak{h}}_i$'s are nested, giving rise to a hyperbolic element in Γ skewering $\hat{\mathfrak{h}}$, as required. \square

2.5. Application: Products. Suppose that X is a CAT(0) cube complex which factors as a product of two CAT(0) cube complexes $X = X_1 \times X_2$. Then we have two natural projection maps $p_i : X \rightarrow X_i$ and the hyperplanes of X is partitioned as a disjoint union $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \hat{\mathcal{H}}_2$ where $\hat{\mathcal{H}}_i$ is composed of the set of hyperplanes of X_i pulled by back to X by the projection map p_i . Notice that $\hat{\mathcal{H}}_i$ is empty if and only if X_i is reduced to a single vertex. We now observe that $\hat{\mathcal{H}}_1$ and $\hat{\mathcal{H}}_2$ are transverse, meaning that every hyperplane in $\hat{\mathcal{H}}_1$ crosses every hyperplane in $\hat{\mathcal{H}}_2$. Thus, the pocset $\mathcal{H} = \mathcal{H}(X)$ admits an involution invariant decomposition $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ such that the each half-space in \mathcal{H}_1 is incomparable to any half-space in \mathcal{H}_2 . The projection map $p_i : X \rightarrow X_i$ is nothing but the restriction quotient arising from the subset $\hat{\mathcal{H}}_i \subset \hat{\mathcal{H}}$. In other words we have $X_i \cong X(\hat{\mathcal{H}}_i)$.

Conversely, we could start with two pocsets Σ_1 and Σ_2 satisfying the finite interval condition and take their disjoint union to obtain a new such pocset $\Sigma = \Sigma_1 \cup \Sigma_2$, where each element of Σ_1 is incomparable to an each element in Σ_2 . Let X denote the cube

complex associated to Σ and X_i denote the cube complex associated to Σ_i . Then we have the pocset of half-spaces associated to X is isomorphic to the pocset of half-spaces associated to $X_1 \times X_2$. Since the construction of the CAT(0) cube complex from the pocset is canonical, it follows that $X = X_1 \times X_2$. To summarize, we have the following.

Lemma 2.5. *A decomposition of a CAT(0) cube complex as a product of cube complexes corresponds to a partition of the collection of hyperplanes of X , $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \hat{\mathcal{H}}_2$ such that every hyperplane in $\hat{\mathcal{H}}_1$ meets every hyperplane in $\hat{\mathcal{H}}_2$.*

A cube complex which cannot be decomposed as above is called **irreducible**. An easy consequence is the following basic analogue of the de Rham decomposition theorem.

Proposition 2.6. *A finite-dimensional CAT(0) cube complex X admits a canonical decomposition*

$$X = X_1 \times \cdots \times X_p$$

into a product of irreducible cube complexes X_i . Every automorphism of X preserves that decomposition, up to a permutation of possibly isomorphic factors. In particular, the image of the canonical embedding

$$\text{Aut}(X_1) \times \cdots \times \text{Aut}(X_p) \hookrightarrow \text{Aut}(X)$$

has finite index in $\text{Aut}(X)$.

Proof. Since X is finite-dimensional, any product decomposition can be refined into a *finite* product of *irreducible* factors. Therefore it suffices to show that if X admits two such decompositions $X = X_1 \times \cdots \times X_p$ and $X = X'_1 \times \cdots \times X'_q$ then $p = q$ and $X_i = X'_{\sigma(i)}$ for some permutation σ of $\{1, \dots, p\}$. Consider the partitions of $\hat{\mathcal{H}}$ corresponding respectively to these product decompositions: $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \cdots \cup \hat{\mathcal{H}}_p$ and $\hat{\mathcal{H}} = \hat{\mathcal{H}}'_1 \cup \cdots \cup \hat{\mathcal{H}}'_q$, where $\hat{\mathcal{H}}_i = \hat{\mathcal{H}}(X_i)$ and $\hat{\mathcal{H}}'_j = \hat{\mathcal{H}}(X'_j)$. The second partition of $\hat{\mathcal{H}}$ induces a partition of each individual subset $\hat{\mathcal{H}}_i$. Lemma 2.5 ensures that this must be the trivial partition since X_i is irreducible. In particular we infer that $p \leq q$. By symmetry we have $p \geq q$ and the desired result follows easily. \square

3. ESSENTIAL THINGS

3.1. Essential hyperplanes. Let X be a CAT(0) cube complex. A half-space \mathfrak{h} is called **deep** if it properly contains arbitrarily large balls of X . Otherwise we say it is **shallow**. This allows us to break up the set $\hat{\mathcal{H}}$ into three types of hyperplanes: **essential**, **half-essential** and **trivial**, as follows. A hyperplane \hat{h} is called **essential** if the half-spaces \mathfrak{h} and \mathfrak{h}^* are both deep. If \mathfrak{h} and \mathfrak{h}^* are both shallow, then \hat{h} is called **trivial**. If \hat{h} is neither essential nor trivial, we call it **half-essential**. We let $\text{Ess}(X)$ (resp. $\text{Hess}(X)$, $\text{Triv}(X)$) denote the collection of essential (resp. half-essential, trivial) hyperplanes.

3.2. Essential cube complexes. A CAT(0) cube complex X is called **essential** if all its hyperplanes are essential or, in other words, if $\hat{\mathcal{H}}(X) = \text{Ess}(X)$. We define the **core** of X as the restriction quotient $X(\text{Triv}(X) \cup \text{Ess}(X))$ and the **essential core** of X as the restriction quotient $X(\text{Ess}(X))$. Clearly the essential core is always essential; it is endowed with a canonical $\text{Aut}(X)$ -action. Notice however that the core of X might be reduced to a single point in general, even if X is unbounded. For example, consider the standard squaring of the Euclidean quarter-plane, in which all hyperplanes are half-essential. In order to deal with that issue, we shall analyze the notions introduced thus far relatively to the action of a subgroup $\Gamma \leq \text{Aut}(X)$.

3.3. Essential hyperplanes relative to a group action. Let $\Gamma \leq \text{Aut}(X)$. Choose a vertex $v \in X$. A half-space \mathfrak{h} is called Γ -**deep** if it contains orbit points of v arbitrarily far from $\hat{\mathfrak{h}}$ (note that this definition is independent of the choice of v). Otherwise $\hat{\mathfrak{h}}$ is called Γ -**shallow**. Mimicking the above definitions, we again break up the set $\hat{\mathcal{H}}$ into three types of hyperplanes: Γ -**essential**, Γ -**half-essential** and Γ -**trivial**. We also define the symbols $\text{Ess}(X, \Gamma)$, $\text{Hess}(X, \Gamma)$ and $\text{Triv}(X, \Gamma)$ accordingly, as well as the Γ -**core** of X and the Γ -**essential core**.

The following statement, which we shall use repeatedly and often implicitly in the sequel, clarifies how these notions behave with respect to invariant convex subcomplexes. The proof is straightforward and will be omitted.

Lemma 3.1. *Let X be a CAT(0) cube complex, let $\Gamma \leq \text{Aut}(X)$ and $Y \subseteq X$ be a Γ -invariant convex subcomplex.*

- (i) *Each hyperplane of Y extends to a unique hyperplane of X , so that there is a natural inclusion $\hat{\mathcal{H}}(Y) \subseteq \hat{\mathcal{H}}(X)$.*
- (ii) *Each half-space \mathfrak{h} disjoint from Y is Γ -shallow. In particular we have $\text{Ess}(Y, \Gamma) = \text{Ess}(X, \Gamma)$, and the Γ -essential core of Y identifies with the Γ -essential core of X .*

The property of being Γ -essential, half-essential or trivial, can be recognized in the orbit quotient, as follows.

Proposition 3.2. *Let X be a finite-dimensional CAT(0) cube complex, let $\Gamma \leq \text{Aut}(X)$ and $\hat{\mathfrak{h}} \in \hat{\mathcal{H}}(X)$. Then the following conditions are equivalent:*

- (i) $\hat{\mathfrak{h}} \in \text{Ess}(X, \Gamma)$,
- (ii) $\hat{\mathfrak{h}}$ is skewered by some element of Γ ,
- (iii) $X(\Gamma\hat{\mathfrak{h}})$ is unbounded.

If in addition X is locally compact, then:

- (iv) $\hat{\mathfrak{h}} \in \text{Triv}(X, \Gamma)$ if and only if $X(\Gamma\hat{\mathfrak{h}})$ is finite.
- (v) $\hat{\mathfrak{h}} \in \text{Hess}(X, \Gamma)$ if and only if $X(\Gamma\hat{\mathfrak{h}})$ is infinite and bounded.

We shall make use of the following basic fact.

Lemma 3.3. *Let X be a finite-dimensional CAT(0) cube complex, let $\Gamma \leq \text{Aut}(X)$. Let $\hat{\mathfrak{h}}$ be a hyperplane such that $X(\Gamma\hat{\mathfrak{h}})$ is bounded.*

Then one of the following assertions holds true.

- (i) $\bigcap_{\gamma \in \Gamma} \gamma \cdot \hat{\mathfrak{h}}$ is non-empty and $\hat{\mathfrak{h}}$ is Γ -trivial.
- (ii) $\bigcap_{\gamma \in \Gamma} \gamma \cdot \mathfrak{h}$ contains some vertex of X and \mathfrak{h}^* is Γ -shallow.
- (iii) $\bigcap_{\gamma \in \Gamma} \gamma \cdot \mathfrak{h}^*$ contains some vertex of X and \mathfrak{h} is Γ -shallow.

Proof. Since $X(\Gamma\hat{\mathfrak{h}})$ is bounded, the group Γ has a fixed point v in $X(\Gamma\hat{\mathfrak{h}})$. If this belongs to $\hat{\mathfrak{h}}$, then it is contained in every element of $\Gamma\hat{\mathfrak{h}}$. Therefore $X(\Gamma\hat{\mathfrak{h}})$ consists of a single cube, and the orbit $\Gamma\hat{\mathfrak{h}}$ must therefore be finite. Since the hyperplanes in the orbit $\Gamma\hat{\mathfrak{h}}$ cross pairwise in $X(\Gamma\hat{\mathfrak{h}})$, they also cross pairwise in X . Therefore, Theorem 4.14 from [Sag95] ensures that $\bigcap_{\gamma \in \Gamma} \gamma \cdot \hat{\mathfrak{h}}$ is a non-empty Γ -invariant convex subset. In particular some Γ -orbit lies entirely in the hyperplane $\hat{\mathfrak{h}}$, which implies that $\hat{\mathfrak{h}}$ is Γ -trivial. We are thus in case (i).

Assume now that v lies strictly in one side of $\hat{\mathfrak{h}}$, say in the half-space \mathfrak{h} determined by $\hat{\mathfrak{h}}$. In that case, the cube supporting v must be pointwise fixed by Γ , and we may therefore assume that v is a vertex. Lifting v up to X via the quotient map $X \rightarrow X(\Gamma\hat{\mathfrak{h}})$ (see the construction of the orbit quotient in §2), we infer that the intersection $Z = \bigcap_{\gamma \in \Gamma} \gamma \cdot \mathfrak{h}$ is

non-empty in X . Thus Z is a Γ -invariant convex subset of X . In particular, some Γ -orbit is entirely contained in \mathfrak{h} . Therefore Lemma 3.1 ensures that \mathfrak{h}^* is Γ -shallow and we are in case (ii).

Finally if v lies in \mathfrak{h}^* , then the same argument shows that the situation (iii) occurs. \square

Proof of Proposition 3.2. If $X(\Gamma\hat{\mathfrak{h}})$ is unbounded, then some hyperbolic element of Γ skewers $\hat{\mathfrak{h}}$ by Lemma 2.4 and hence $\hat{\mathfrak{h}}$ is Γ -essential. Conversely, if $X(\Gamma\hat{\mathfrak{h}})$ is bounded, then \mathfrak{h} or \mathfrak{h}^* is Γ -shallow by Lemma 3.3 and, thus, $\hat{\mathfrak{h}}$ is Γ -inessential.

In view of Lemma 2.4, this proves the equivalence between (i), (ii) and (iii). Assume now in addition that X is locally compact.

By Lemma 2.4, the cube complex $X(\Gamma\hat{\mathfrak{h}})$ is finite if and only if the orbit $\Gamma\hat{\mathfrak{h}}$ is so. In that case Γ has a finite index which stabilises $\hat{\mathfrak{h}}$. Thus $\hat{\mathfrak{h}}$ is Γ -trivial.

Conversely if $\hat{\mathfrak{h}}$ is Γ -trivial, so is every element in $\Gamma\hat{\mathfrak{h}}$ and there is some $r > 0$ such that $X \subset \mathcal{N}_r(\gamma.\hat{\mathfrak{h}})$ for each $\gamma \in \Gamma$. In particular each wall of $\Gamma\hat{\mathfrak{h}}$ meets the r -ball around every vertex of X . Thus $\Gamma\hat{\mathfrak{h}}$ is finite since X is locally compact. This shows Assertion (iv). The remaining assertion is an immediate consequence of the others. \square

It is easy to deduce from the definition that any Γ -essential hyperplane intersects every Γ -trivial hyperplane. In view of Lemma 2.5, this yields the following.

Remark 3.4. Assume that $\text{Hess}(X, \Gamma)$ is empty. Then X splits as a product $Z \times C$ where $\hat{\mathcal{H}}(Z) = \text{Ess}(X, \Gamma)$ and $\hat{\mathcal{H}}(C) = \text{Triv}(X, \Gamma)$. It is good to keep in mind that if $\text{Ess}(X, \Gamma) = \emptyset$, then Z is reduced to a single vertex.

3.4. Essential actions and pruning. The Γ -action on X is called **essential** if every hyperplane is Γ -essential or, equivalently, if $\hat{\mathcal{H}}(X) = \text{Ess}(X) = \text{Ess}(X, \Gamma)$. In particular, if X admits some essential group action, then it is essential. Lemma 3.1 shows that if X contains some Γ -invariant convex subcomplex $Y \subsetneq X$, then the Γ -action on X cannot be essential. It is natural to address the question whether X contains some non-empty Γ -invariant convex subcomplex on which the Γ -action is essential. The Γ -action on the Γ -essential core of X is always essential, but this core might be reduced to a singleton. In case it is not, it is not clear *a priori* that it embeds as a convex subcomplex of X . The following result ensures that these possible pathologies do not occur under some natural conditions on the Γ -action.

Proposition 3.5. *Let X be a finite-dimensional CAT(0) cube complex and let $\Gamma \leq \text{Aut}(X)$. Assume that at least one of the following two conditions is satisfied:*

- (a) Γ has finitely many orbits of hyperplanes.
- (b) Γ has no fixed point at infinity.

Then the Γ -essential core of X is unbounded if and only if Γ has no fixed point. In that case the Γ -essential core embeds as a Γ -invariant convex subcomplex Y of X .

A key criterion allowing us to detect the existence of global fixed points at infinity is provided by the following result, which holds in arbitrary finite-dimensional CAT(0) spaces. One difficulty for finite-dimensional CAT(0) cube complexes which are not proper is that the visual boundary $\partial_\infty X$ has no reason to be non-empty *a priori*. In particular an unbounded sequence from X does not accumulate in $\partial_\infty X$ in general. This difficulty will be handled by referring to the main result of [CL10]; the latter paper is however not necessary in the case of locally compact CAT(0) cube complexes.

Proposition 3.6. *Let X be a finite-dimensional complete CAT(0) space (or more generally, a complete CAT(0) space of finite telescopic dimension, see [CL10]) and let $\Gamma \leq \text{Is}(X)$. Let also $\{Y_\alpha\}_{\alpha \in A}$ be a Γ -invariant collection of closed convex subsets of X .*

If for any finite subset $B \subseteq A$, the intersection $\bigcap_{\alpha \in B} Y_\alpha$ is non-empty, then either $\bigcap_{\alpha \in A} Y_\alpha$ is a non-empty Γ -invariant subspace or $\bigcap_{\alpha \in A} \partial_\infty Y_\alpha \subset \partial_\infty X$ is non-empty and contains a canonical circumcentre which is fixed by Γ .

Proof. Assume that the intersection $\bigcap_{\alpha \in B} Y_\alpha$ is empty. Then Theorem 1.1 from [CL10] ensures that $\bigcap_{\alpha \in A} \partial_\infty Y_\alpha \subset \partial_\infty X$ is a non-empty Γ -invariant convex subset of $\partial_\infty X$ of radius $\leq \pi/2$. Notice that if X is proper, then this can be established by an elementary and direct argument without referring to [CL10], but simply using the fact that the cone topology makes $\partial_\infty X$ into a compact space (see [CM09b, Proposition 3.2]). The existence of a fixed point in $\bigcap_{\alpha \in A} \partial_\infty Y_\alpha \subset \partial_\infty X$ is thus a consequence of [BL05, Proposition 1.4], which can be applied since $\partial_\infty X$ is finite-dimensional (see [CL10, Proposition 2.1]). \square

With this criterion at hand, we are able to provide the following.

Proof of Proposition 3.5. Let us consider the half-essential hyperplanes. We want to describe a pruning process by which we retract the complex down to a complex with no half-essential hyperplanes (in analogy with the standard pruning of a tree, which consists in removing valence-1 vertices).

Assume thus that $\text{Hess}(X, \Gamma)$ is non-empty. Then there is some Γ -deep half-space \mathfrak{h} such that $\hat{\mathfrak{h}}$ is Γ -half-essential. By Proposition 3.2 and Lemma 3.3, the set $\bigcap_{\gamma \in \Gamma} \gamma \cdot \mathfrak{h}$ is a non-empty Γ -invariant convex subset of X which contains some vertex. Therefore, the collection of all vertices contained in $\bigcap_{\gamma \in \Gamma} \gamma \cdot \mathfrak{h}$ spans a Γ -invariant convex subcomplex, which we shall denote by Y_1 .

Since Y_1 is a Γ -invariant convex subcomplex, it follows that there is a canonical embedding

$$\text{Hess}(Y_1, \Gamma) \subsetneq \text{Hess}(X, \Gamma).$$

We now repeat the above arguments, this time applied to the cube complex Y_1 . Hence if there is some Γ -deep half-space \mathfrak{h}_1 such that $\hat{\mathfrak{h}}_1 \in \text{Hess}(Y_1, \Gamma)$, then there is a non-empty Γ -invariant convex subcomplex $Y_2 \subsetneq Y_1$, and we have $\text{Hess}(Y_2, \Gamma) \subsetneq \text{Hess}(Y_1, \Gamma)$. We can then proceed inductively: provided the set $\text{Hess}(Y_n, \Gamma)$ is non-empty, we find a Γ -invariant convex subcomplex $Y_{n+1} \subsetneq Y_n$. There are two cases to consider.

If the process terminates after finitely many steps, say at step n , it means that $\text{Hess}(Y_n, \Gamma)$ is empty. By Lemma 3.1, we deduce that $\text{Ess}(Y_n, \Gamma) = \text{Ess}(X, \Gamma)$. By Remark 3.4, the subcomplex Y_n splits as a product of a bounded subcomplex and a subcomplex which is isomorphic to that Γ -essential core Z of X . Either Z is reduced to a single vertex, in which case Γ has a fixed point in Y_n , or Z contains a hyperplane, hence an essential one, in which case all Γ -orbits are unbounded. Thus we are done in this case.

Assume now that the process never terminates. Since at each step, we remove at least one Γ -orbit of half-essential hyperplanes, we deduce that the hypothesis (a) cannot be satisfied. Thus (b) holds and Γ has no fixed point at infinity. In particular, it follows from Proposition 3.6 that the intersection $Y_\infty = \bigcap_{n>0} Y_n$ is non-empty. As an intersection of convex subcomplexes, Y_∞ is itself a Γ -invariant convex subcomplex. At this point, a transfinite induction argument finishes the proof. \square

3.5. Finitely many orbits of hyperplanes. A natural question is how to compare the absolute qualifications of a hyperplane to the corresponding qualifications relative to the Γ -action. For example, it is clear that if the Γ -action on X is cocompact, then $\text{Ess}(X, \Gamma) =$

$\text{Ess}(X)$, $\text{Hess}(X, \Gamma) = \text{Hess}(X)$ and $\text{Triv}(X, \Gamma) = \text{Triv}(X)$. This fact can be generalised as follows; for another related statement, see Corollary 6.5 below.

Proposition 3.7. *Let X be a finite-dimensional locally compact CAT(0) cube complex and $\Gamma \leq \text{Aut}(X)$ be a group acting with finitely many orbits of hyperplanes.*

Then we have $\text{Triv}(X) = \text{Triv}(X, \Gamma)$, $\text{Hess}(X) = \text{Hess}(X, \Gamma)$ and $\text{Ess}(X) = \text{Ess}(X, \Gamma)$, and the set $\text{Triv}(X)$ is finite.

Proof. First, we remark that the inclusion $\text{Triv}(X) \subseteq \text{Triv}(X, \Gamma)$ is obvious. Conversely, let $\hat{\mathfrak{h}} \in \text{Triv}(X, \Gamma)$. Thus $X(\Gamma\hat{\mathfrak{h}})$ is finite by Proposition 3.2 and hence Γ has a finite index subgroup which stabilises $\hat{\mathfrak{h}}$. Since Γ has finitely many orbits on $\hat{\mathcal{H}}(X)$, so does any finite index subgroup. Upon replacing Γ by a finite index subgroup, we may thus assume that $\hat{\mathfrak{h}}$ is Γ -fixed. The distance to $\hat{\mathfrak{h}}$ is thus a Γ -invariant function defined on the collection $\hat{\mathcal{H}}(X)$ of all hyperplanes of X . Since Γ has finitely many orbits on $\hat{\mathcal{H}}(X)$, it follows that this function is uniformly bounded from above. In other words, every hyperplane of X is close to $\hat{\mathfrak{h}}$. By Lemma 2.1, this implies that the entire space X is contained in some bounded neighbourhood of $\hat{\mathfrak{h}}$. In other words $\hat{\mathfrak{h}}$ is trivial. Thus $\text{Triv}(X) = \text{Triv}(X, \Gamma)$.

We just showed that every large ball of X intersects every trivial hyperplane. Since X is locally compact, this implies that $\text{Triv}(X)$ is indeed finite.

We next point out that the inclusion $\text{Ess}(X) \supseteq \text{Ess}(X, \Gamma)$ is obvious.

In order to finish the proof, it now suffices to show that a hyperplane $\hat{\mathfrak{h}}$ such that $X(\Gamma\hat{\mathfrak{h}})$ is infinite and bounded, is necessarily half-essential. Since $X(\Gamma\hat{\mathfrak{h}})$ is infinite, the hyperplane $\hat{\mathfrak{h}}$ cannot be Γ -trivial. Therefore Lemma 3.3 implies that one of the two half-spaces associated with $\hat{\mathfrak{h}}$, say \mathfrak{h} , is Γ -shallow. We conclude that $\hat{\mathfrak{h}}$ is half-essential, as desired. \square

The following observation is sometimes useful.

Remark 3.8. Let X be a CAT(0) cube complex such that $\text{Aut}(X)$ has finitely many orbits of hyperplanes. Then there exists $N > 0$ such that a hyperplane $\hat{\mathfrak{h}}$ of X is essential if and only if both \mathfrak{h} and \mathfrak{h}^* properly contain pencils of disjoint hyperplanes of size N .

Propositions 3.5 and 3.7 show that a few things become simpler in the presence of a group action which is cofinite on the hyperplanes. The most obvious example of such an action is when $\Gamma \leq \text{Aut}(X)$ is cocompact. More generally, we have the following.


Lemma 3.9. *Let Γ be a finitely generated group acting properly discontinuously on a locally compact CAT(0) cube complex X . Assume there exists a vertex v such that, for each half-space \mathfrak{h} , the orbit $\Gamma.v$ meets both \mathfrak{h} and \mathfrak{h}^* .*

Then Γ has finitely many orbits of hyperplanes.

Proof. The group Γ is finitely generated if and only if every orbit has some tubular neighbourhood which is pathwise connected. The result follows by applying this to $\Gamma.v$. \square

Remark 3.10. The condition on the vertex v appearing in the statement of Lemma 3.9 is a **minimality condition** on the Γ -action. The Lemma could be generalised in the following way: *Let Γ be a finitely generated group acting properly discontinuously on a locally compact CAT(0) cube complex X . Then for any given vertex v , the group Γ has finitely many orbits of hyperplanes $\hat{\mathfrak{h}}$ such that the half-spaces \mathfrak{h} and \mathfrak{h}^* both contain orbit points from $\Gamma.v$.*

Remark 3.11. An example, which was communicated to us by Dani Wise, shows that the converse of Lemma 3.9 does not hold. Indeed, a countably based free group F_∞ can act properly discontinuously on a CAT(0) cube complex with only three orbits of hyperplanes.

FIGURE 1. A proper action of F_∞ with three orbits of hyperplanes

Here is a description of this example (see Figure 1). One starts with a graph consisting of a chain of bigons — take two copies $\mathbf{R}_1, \mathbf{R}_2$ of the real line, glued together along the integers. One now attaches two infinite rectangular strips. The first is attached at top and bottom to $\mathbf{R}_1, \mathbf{R}_2$ in a straightforward manner. The second is attached with a unit translation on one side. In this way one obtains a locally CAT(0) cube complex with exactly three hyperplanes and whose fundamental group is F_∞ .

The following summarizes some of the results obtained thus far in the case of a finitely generated discrete group acting properly.

Proposition 3.12. *Let Γ be a finitely generated group acting properly discontinuously on a finite-dimensional locally finite CAT(0) cube complex X . Let Y denote the Γ -essential core of X . Then*

- (i) Γ has finitely many orbits on $\hat{\mathcal{H}}(Y) = \text{Ess}(X, \Gamma)$.
- (ii) Y is unbounded if and only if Γ has no global fixed point in X .
- (iii) Y embeds as a convex Γ -invariant subcomplex of X .
- (iv) Every hyperplane of Y is skewered by some element of Γ .

Proof. Pick any vertex v . Let Z denote the convex subcomplex of X whose 0-skeleton is the set of all vertices contained in every half-space containing entirely the orbit $\Gamma.v$. Thus Z is non-empty, Γ -invariant, and for each hyperplane $\hat{h} \in \hat{\mathcal{H}}(Z)$, the orbit $\Gamma.v$ meets both \mathfrak{h} and \mathfrak{h}^* . By Lemma 3.9, this implies that Γ has finitely many orbits on $\hat{\mathcal{H}}(Z)$. By Lemma 3.1, we have $\text{Ess}(X, \Gamma) \subseteq \hat{\mathcal{H}}(Z)$. The assertion (i) follows.

We can now run the pruning process described in Proposition 3.5. This proves the assertions (ii) and (iii). Assertion (iv) follows from Lemma 2.4 and Proposition 3.2. \square

4. THE FLIPPING LEMMA

The goal of this section is to provide a proof of the Flipping Lemma, stated in the introduction. In fact, we shall give two distinct (and conceptually different) versions of the Flipping Lemma, respectively in Theorems 4.1 and 4.7, with different proofs and slightly different statements. The first one is shorter and easier; moreover it applies to CAT(0) cube complex that are possibly not locally compact. But this first proof is not self-contained as it relies on the main theorem from [CL10]. The second proof is self-contained and uses only the combinatorics of hyperplanes. However it is technically much more involved; its hypotheses do not require the absence of fixed point at infinity but do require the ambient CAT(0) space to be locally compact with a cocompact automorphism group; its conclusions are also more precise, as it yields a product decomposition of the ambient cube complex with a linelike factor. All in all, this second proof is probably less transparent, and a reader who is willing to rely on [CL10] might want to skip it.

Let X be a CAT(0) cube complex. Given a pair of half-spaces \mathfrak{h} and \mathfrak{d} associated to disjoint hyperplanes, we say that they are **nested** if $\mathfrak{h} \subset \mathfrak{d}$ or $\mathfrak{d} \subset \mathfrak{h}$. If \mathfrak{h} and \mathfrak{d} are not nested and $\mathfrak{h} \cap \mathfrak{d} \neq \emptyset$, we say that \mathfrak{h} and \mathfrak{d} are **facing**. An isometry $g \in \text{Aut}(X)$ is said to **flip** \mathfrak{h} if \mathfrak{h} and $g\mathfrak{h}$ are facing. Equivalently $\mathfrak{h}^* \subsetneq g\mathfrak{h}$.

Given a group $\Gamma \leq \text{Aut}(X)$, we say that \mathfrak{h} is **Γ -flippable** if there is an element of Γ which flips \mathfrak{h} .

4.1. Version I: no fixed point at infinity. This section is devoted to proving the following theorem.

Theorem 4.1 (Flipping Lemma, version I). *Assume that X is finite-dimensional and let $\Gamma \leq \text{Aut}(X)$ be any subgroup. Let \mathfrak{h} be a half-space which is not Γ -flippable. Then Γ has a fixed point in the visual boundary $\partial_\infty \mathfrak{h}^*$ or \mathfrak{h} is Γ -shallow.*

Notice that no assumption on Γ or on its action is made. Moreover, we do not require X to be locally compact.

It is a well known basic fact that any finite collection of pairwise crossing hyperplanes in a CAT(0) cube complex has a non-empty intersection, see [Sag95, Theorem 4.14]. A more general type Helly property exists for convex subcomplexes of CAT(0) cube complexes and was established by Gerasimov [Ger98], stating that the intersection of convex pairwise intersecting subcomplexes is non-empty. For completeness, we include a proof of a special case of this theorem, suited to our needs.

Lemma 4.2. *Let X be a CAT(0) cube complex. Let $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ be half-spaces which have pairwise a non-empty intersection. Then the intersection $\bigcap_{i=1}^n \mathfrak{h}_i$ is non-empty.*

Proof. The set $H = \{\mathfrak{h}_1, \dots, \mathfrak{h}_n\}$ is partially ordered by inclusion. Upon renumbering, we may assume that the first m half-spaces $\mathfrak{h}_1, \dots, \mathfrak{h}_m$ are the minimal elements of H . In other words, for each $i > m$, there is some $i' \leq m$ such that $\mathfrak{h}_{i'} \subset \mathfrak{h}_i$. In particular it suffices to show that $\bigcap_{i=1}^m \mathfrak{h}_i$ is non-empty. We shall show by induction on m that $\bigcap_{i=1}^m \mathfrak{h}_i$ contains a vertex.

If the $\hat{\mathfrak{h}}_i$'s are pairwise disjoint, then we have $\hat{\mathfrak{h}}_i \subset \mathfrak{h}_j$ for all $i \neq j$ and thus any vertex contained in \mathfrak{h}_i and adjacent to $\hat{\mathfrak{h}}_i$ is also contained in every other \mathfrak{h}_j .

Otherwise some $\hat{\mathfrak{h}}_i$ crosses some $\hat{\mathfrak{h}}_{i'}$. We view $\hat{\mathfrak{h}}_{i'}$ as a convex subcomplex of the first cubical subdivision of X . In this way, the hyperplane $\hat{\mathfrak{h}}_{i'}$ becomes a CAT(0) cube complex whose vertices are midpoints of edges of X . By induction, the intersection of all \mathfrak{h}_j 's such

that \hat{h}_i crosses $\hat{h}_{i'}$ contains some vertex of $\hat{h}_{i'}$. This vertex is the midpoint of an edge of X transverse to $\hat{h}_{i'}$. We let v denote the vertex of that edge contained in $\hat{h}_{i'}$. By construction v is contained in \hat{h}_j for all j such that \hat{h}_j crosses $\hat{h}_{i'}$. Since $\hat{h}_{i'}$ is entirely contained in every \hat{h}_j such that \hat{h}_j does not cross $\hat{h}_{i'}$, it follows that v is in fact contained in all the \hat{h}_i 's. \square

We are now ready for the following.

Proof of Theorem 4.1. Consider the Γ -orbit $\{\gamma \cdot \mathfrak{h}^*\}_{\gamma \in \Gamma}$. The fact that \mathfrak{h} is Γ -unflippable precisely means that any two half-spaces in that orbit have a non-empty intersection. In view of Lemma 4.2, we deduce that any finite set of half-spaces from that orbit has a non-empty intersection. We are thus in a position to invoke Proposition 3.6. This implies that either $\bigcap_{\gamma \in \Gamma} \gamma \cdot \partial_\infty \mathfrak{h}^*$ is a non-empty subset of the visual boundary which contains a Γ -fixed point, or the set $Y = \bigcap_{\gamma \in \Gamma} \gamma \cdot \mathfrak{h}^*$ is a non-empty Γ -invariant subspace of X , which we assume from now on. Since \mathfrak{h} is disjoint from Y , Lemma 3.1 implies that \mathfrak{h} is Γ -shallow, as desired. \square

Our next goal is to establish a slightly refined version of the Flipping Lemma in the case of cocompact actions. We first need to take a detour to some additional basic considerations.

4.2. Non-skewering behaviours. Let X be a CAT(0) cube complex, $\gamma \in \text{Aut}(X)$ be a hyperbolic isometry and \hat{h} be a hyperplane. In Section 2, we defined what it means for γ to skewer \hat{h} . There are also two types of non-skewering behaviours, which we describe as follows.

We say that γ is **parallel** to \hat{h} if an axis for γ (and hence all axes of γ) lies in a neighborhood of \hat{h} . If γ does not skewer \hat{h} and is not parallel to \hat{h} , we say that it is **peripheral** to \hat{h} .

Lemma 4.3. *A hyperbolic isometry γ is parallel to \hat{h} if and only if \hat{h} contains a geodesic line which is at bounded Hausdorff distance from some γ -axis.*

If in addition X is locally compact, then γ is parallel to \hat{h} if and only if there exists some $n > 0$ such that $\gamma^n \in \text{Stab}(\hat{h})$

Proof. If \hat{h} contains a geodesic line which is at bounded Hausdorff distance from some γ -axis, then this axis is contained in a bounded neighbourhood of \hat{h} . Thus γ is parallel to \hat{h} as required. Conversely, if γ has some axis ℓ in a bounded neighbourhood of \hat{h} , then the two endpoints of any of its axes are contained in the visual boundary $\partial_\infty \hat{h}$. Since \hat{h} is closed and convex, it easily follows that it must thus contain some geodesic line joining them (see for example [CM09b, Prop. 3.6]), and the desired assertion follows.

Assume now that X is locally compact. If $\gamma \in \text{Stab}(\hat{h})$ then there exists an axis of γ in \hat{h} and we are done. Conversely, suppose that an axis ℓ for γ is contained in the R -neighbourhood $\mathcal{N}_R(\hat{h})$ for some $R > 0$. Consider some p along ℓ . Then for any $n > 0$, we have that $\gamma^n(\hat{h}) \cap B_R(p) \neq \emptyset$. By local finiteness, only finitely many hyperplanes meet the ball $B_R(p)$, so by the pigeonhole principle, $\gamma^n(\hat{h}) = \gamma^m(\hat{h})$ for some $n > m > 0$. It follows that $\gamma^{n-m}(\hat{h}) = \hat{h}$, as required. \square

Lemma 4.4. *Assume that X is finite-dimensional. If a hyperbolic isometry γ is peripheral to \hat{h} , then there exists some $n > 0$ such that $\gamma^n \hat{h} \cap \hat{h} = \emptyset$.*

Proof. The elements of the set $\{\gamma^n \hat{h} | n > 0\}$ are pairwise distinct, since otherwise some positive power of γ would stabilise \hat{h} and, hence, every γ -axis would lie in a bounded

neighbourhood of \hat{h} . Since X is finite-dimensional, we infer that there exists $n > m > 0$ such that $\gamma^n \hat{h} \cap \gamma^m \hat{h} = \emptyset$. It follows that $\gamma^{n-m} \hat{h} \cap \hat{h} = \emptyset$, as required. \square

4.3. Endometries of proper metric spaces. Given a metric space X , an **endometry** of X is an injective map $\alpha : X \rightarrow X$ which is an isometry onto its image. The following general fact is independent of the CAT(0) inequality.

Proposition 4.5. *Let X be a proper metric space with a cocompact isometry group. Then every endometry $\alpha : X \rightarrow X$ is surjective.*

Proof. Given $\varepsilon > 0$, a subset $E \subset X$ is called **ε -separated** if any two distinct points of E are at distance at least ε .

Let $(\varepsilon_n)_{n \geq 0}$ be a sequence of positive reals tending to 0 as n tends to infinity. For each n , let C_n be the maximal cardinality of a ε_n -separated subset contained in a ball of radius n of X . Notice that C_n is finite since X is proper and has a cocompact isometry group. Let also $x_n \in X$ be such that the ball $B(x_n, n)$ of radius n centred at x_n contains a ε_n -separated subset, say E_n , of cardinality C_n .

Since X has cocompact isometry group, there is no loss of generality in assuming that the sequence (x_n) converges to some $x \in X$. Let now $y \in X$ be any point and $r = d(y, \alpha(x))$. Then for each n which is sufficiently larger than r , there is some $y_n \in E_n$ such that $d(y, \alpha(y_n)) \leq \varepsilon_n$. In particular the sequence $(\alpha(y_n))$ converges to y . Since X is proper, it is complete and so is $\alpha(X)$. In particular $\alpha(X)$ is closed in X and we deduce that $y \in \alpha(X)$. Thus $\alpha(X) = X$, as desired. \square

Notice that the conclusion of Proposition 4.5 fails if $\text{Is}(X)$ does not act cocompactly, as illustrated by the example of $X = \mathbf{R}^+$.

Given a hyperplane \hat{h} in a cube complex X , we denote by $\text{Ess}(\hat{h})$ the set of all hyperplanes crossing \hat{h} and which are essential in \hat{h} . Endometries will naturally occur through the following.

Lemma 4.6. *Let X be a locally compact CAT(0) cube complex with cocompact automorphism group, let \hat{h}, \hat{d} be hyperplanes and let $X_{\hat{h}}, X_{\hat{d}}$ denote their respective essential cores. Assume that every hyperplane in $\text{Ess}(\hat{h})$ crosses \hat{d} . Then there is an isometric embedding $X_{\hat{h}} \rightarrow X_{\hat{d}}$.*

Proof. We shall define a map $X_{\hat{h}} \rightarrow X_{\hat{d}}$ at the level of the 1-skeletons. For each evertex $v \in X_{\hat{h}}^{(0)}$, we need to associate a unique vertex $v' \in X_{\hat{d}}^{(0)}$. To this end, we associate to v an **ultrafilter** ϕ_v on $\text{Ess}(\hat{d})$, i.e. a function which chooses one side of each hyperplane in $\text{Ess}(\hat{d})$. We shall do this in such a way that the intersection of all half-spaces $\bigcap_{\hat{g} \in \text{Ess}(\hat{d})} \phi_v(\hat{g})$ is non-empty. This intersection must therefore contain a single vertex of $X_{\hat{d}}$. This will be our definition of v' .

In order to define the ultrafilter ϕ_v , we proceed as follows. Let $\hat{g} \in \text{Ess}(\hat{d})$. If \hat{g} does not cross \hat{h} , we define $\phi_v(\hat{g})$ so that it contains \hat{h} . If $\hat{g} \in \text{Ess}(\hat{h})$, we define $\phi_v(\hat{g})$ so that it contains v . If $\hat{g} \in \text{Hess}(\hat{h})$, we define $\phi_v(\hat{g})$ so that it contains the deep side of \hat{h} determined by \hat{g} . It remains to define $\phi_v(\hat{g})$ in case $\hat{g} \in \text{Triv}(\hat{h})$. To do this, we choose arbitrarily a point $p \in \hat{h}$ which does not lie on any hyperplane crossing \hat{h} , and we simply define $\phi_v(\hat{g})$ in such a way that it contains p .

It is now easy to see that ϕ_v is an ultrafilter that satisfies the Descending Chain Condition. Since X is finite dimensional, it follows that ϕ_v determines a vertex of $X_{\hat{d}}$ (see Section 2.3.)

This defines a map $X_{\hat{\mathfrak{h}}}^{(0)} \rightarrow X_{\hat{\mathfrak{d}}}^{(0)}$. This map preserves the relation of adjacency: Indeed, two distinct vertices of $X_{\hat{\mathfrak{h}}}$ are adjacent if and only if they are separated by a unique hyperplane, and the definition of ϕ_v is so that this will then be the case of the corresponding vertices of $X_{\hat{\mathfrak{d}}}$. Thus we have a simplicial isometric embedding $X_{\hat{\mathfrak{h}}}^{(1)} \rightarrow X_{\hat{\mathfrak{d}}}^{(1)}$. In view of the way higher dimension cubes of $X_{\hat{\mathfrak{h}}}$ and $X_{\hat{\mathfrak{d}}}$ are defined, the existence of an isometric embedding follows. \square

4.4. Version II: hereditarily essential actions. The goal of this section is to present a proof of the Flipping Lemma which does rely on Proposition 3.6. The argument is self-contained but technically more involved; moreover, it works only in the special case when X is locally compact and has a cocompact automorphism group. However, this approach has the advantage that it also works when the group Γ under consideration is allowed to fix points in the visual boundary of the ambient cube complex. This advantage will have some relevance for the applications of Rank Rigidity we shall present later, but the reader who is willing to exclude the existence of Γ -fixed points at infinity in X from the start should skip this section.

We need two more definitions.

We say that a CAT(0) cube complex X is **R-like** if there is an $\text{Aut}(X)$ -invariant geodesic line $\ell \subset X$; this is a special case of the condition that appeared in Theorem E. If in addition $\text{Aut}(X)$ acts cocompactly, then X is quasi-isometric to the real line \mathbf{R} .

Given a group $\Gamma \leq \text{Aut}(X)$ we say that the Γ -action on X is **hereditarily essential** if Γ acts essentially and if for any finite collection of pairwise crossing hyperplanes $\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n$, we have

$$\text{Ess}(\hat{\mathfrak{h}}_1 \cap \dots \cap \hat{\mathfrak{h}}_n) = \text{Ess}(\hat{\mathfrak{h}}_1 \cap \dots \cap \hat{\mathfrak{h}}_n, \text{Stab}_{\Gamma}(\hat{\mathfrak{h}}_1 \cap \dots \cap \hat{\mathfrak{h}}_n)).$$

In view of Proposition 3.2, this amounts to requiring that for any hyperplane $\hat{\mathfrak{h}}_0$ such that $\hat{\mathfrak{h}}_0 \cap \hat{\mathfrak{h}}_1 \cap \dots \cap \hat{\mathfrak{h}}_n$ is an essential hyperplane of the cube complex $\hat{\mathfrak{h}}_1 \cap \dots \cap \hat{\mathfrak{h}}_n$, there is an element in $\text{Stab}_{\Gamma}(\hat{\mathfrak{h}}_1 \cap \dots \cap \hat{\mathfrak{h}}_n)$ which skewers $\hat{\mathfrak{h}}_0$.

Basic examples of such actions are provided by groups acting cocompactly on locally compact CAT(0) cube complex; this is the main example that the reader should keep in mind at a first reading of Theorem 4.7 below. Other relevant examples will be provided by Corollary 6.5 below.

Theorem 4.7 (Flipping Lemma, version II). *Let X be a locally compact unbounded CAT(0) cube complex with a cocompact automorphism group and $G \leq \text{Aut}(X)$ be a (possibly non-closed) subgroup whose action is hereditarily essential. Let also \mathfrak{h} be a half-space which is unflippable by the action of G .*

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

- (i) X_1 is irreducible and all of its hyperplanes are compact.
- (ii) Some finite index subgroup $G' \leq G$ preserves the decomposition $X = X_1 \times X_2$.
- (iii) The G' -orbit of $\hat{\mathfrak{h}}$ is in $\hat{\mathcal{H}}_1$.
- (iv) G' fixes a point in the visual boundary $\partial_{\infty} X_1$.
- (v) If in addition G is closed, unimodular and acts cocompactly, then X_1 is **R-like**.

We start by proving the following basic special case of Theorem 4.7, which will be used in the proof of the general case.

Lemma 4.8. *Let X be a locally compact unbounded CAT(0) cube complex with a cocompact automorphism group and $G \leq \text{Aut}(X)$ be a (possibly non-closed) subgroup acting essentially. Let also \mathfrak{h} be a half-space which is unflippable by the action of G .*

If all hyperplanes of X are compact, then G fixes a point in the visual boundary $\partial_\infty X$. If in addition G is closed, unimodular and acts cocompactly on X , then X is \mathbf{R} -like.

Notice that the assumption that G be unimodular is necessary. Indeed, consider the regular trivalent tree T and let $G \leq \text{Aut}(T)$ be the stabiliser of an end $\xi \in \partial T$. Then G is closed, acts cocompactly, and no half-tree containing ξ is G -flippable. Clearly T is not \mathbf{R} -like. The point is of course that this group G is not unimodular.

Proof of Lemma 4.8. Since all hyperplanes are compact, the space X cannot contain any 2-dimensional flat. Therefore X must be Gromov hyperbolic (see [BH99, Theorem III.H.1.5]).

Next we consider the orbit $\{g \cdot \mathfrak{h} \mid g \in G\}$. Since \mathfrak{h} is not G -flippable, any two half-spaces in this orbit have a non-empty intersection. Therefore any finite subset of $\{g \cdot \mathfrak{h} \mid g \in G\}$ has a non-empty intersection by Lemma 4.2. On the other hand, the intersection $\bigcap \{g \cdot \mathfrak{h} \mid g \in G\}$ is a proper G -invariant convex subcomplex, and must therefore be empty. Since X is proper, we deduce that $\bigcap \{\partial_\infty g \cdot \mathfrak{h} \mid g \in G\}$ is a non-empty G -invariant subset of $\partial_\infty X$.

We claim that it is reduced to a single point. Indeed, if it contained two distinct points ξ, η , then these points would be joined by some geodesic line ℓ since X is Gromov hyperbolic and has a cocompact isometry group. In fact the set of all geodesic lines joining ξ to η is of the form $\ell \times K$ for some compact set K . The condition that ξ, η both belong to $\partial_\infty g \cdot \mathfrak{h}$ implies that $g \cdot \mathfrak{h}$ intersects $\ell \times K$ in some subset of the form $\ell \times K'$ for some closed convex subset $K' \subseteq K$. Since K is compact, it follows that the intersection $\bigcap \{\partial_\infty g \cdot \mathfrak{h} \mid g \in G\}$ is non-empty, a contradiction. The claim stands proven.

The claim readily implies that G fixes a point $\xi \in \partial_\infty X$. It remains to prove that if G is closed, unimodular and acts cocompactly on X , then X is \mathbf{R} -like. We shall content ourselves by providing a direct argument in the special case when G is discrete. In the general case, we refer the interested reader to [CM] or [CM09a, Theorem 3.14], which provide general statements on fixed points at infinity for unimodular groups acting cocompactly by isometries on CAT(0) spaces.

Assume thus that G is discrete and acts cocompactly. Since G fixes ξ , it preserves the collection of horoballs centred at ξ . The G -action being cocompact and essential, there must exist some $g \in G$ which permute these horoballs non-trivially. This element g must therefore be a hyperbolic isometry having ξ as one of its two unique fixed points in $\partial_\infty X$. Let $\eta \in \partial_\infty X$ be the other g -fixed point. We claim that G fixes η . Otherwise some $h \in G$ would move it, which is impossible because the subgroup generated by g and the conjugate hgh^{-1} would then be non-discrete. In particular G preserves the collection of geodesic lines joining ξ to η . As G acts cocompactly, this forces X to be \mathbf{R} -like, as desired. \square

Proof of Theorem 4.7. We are given a non-flippable half-space \mathfrak{h} and need to show that X decomposes as a product. We divide $\hat{\mathcal{H}} = \hat{\mathcal{H}}(X)$ into two subsets:

$$\hat{\mathcal{H}}_1 = \{\hat{\mathfrak{d}} \in \hat{\mathcal{H}} \mid \text{Ess}(\hat{\mathfrak{d}}) = \text{Ess}(\hat{\mathfrak{h}})\}$$

and

$$\hat{\mathcal{H}}_2 = \hat{\mathcal{H}} - \hat{\mathcal{H}}_1.$$

We aim to show that every hyperplane in $\hat{\mathcal{H}}_1$ intersects every hyperplane in $\hat{\mathcal{H}}_2$. This will give a product decomposition of X by Lemma 2.5. We proceed in several steps.

Claim 1. *For all $\hat{\mathfrak{d}} \in \hat{\mathcal{H}}$ such that $\hat{\mathfrak{d}} \subset \mathfrak{h}$, we have $\text{Ess}(\hat{\mathfrak{d}}) \subset \text{Ess}(\hat{\mathfrak{h}})$.*

Note that since \mathfrak{h} is unflippable, it follows that every element with an axis in \mathfrak{h} is parallel to $\hat{\mathfrak{h}}$, see Lemma 4.4. Consequently, every element of $\text{Stab}_G(\hat{\mathfrak{d}})$ is parallel to $\hat{\mathfrak{h}}$. This immediately tells us that $\text{Ess}(\hat{\mathfrak{d}}) \subset \text{Ess}(\hat{\mathfrak{h}})$. For if there were an essential hyperplane in $\hat{\mathfrak{d}}$ disjoint from $\hat{\mathfrak{h}}$, we would simply take an element $\alpha \in \text{Stab}_G(\hat{\mathfrak{d}})$ such that α skewers that hyperplane and then α would have an axis staying arbitrarily far away from $\hat{\mathfrak{h}}$.

Claim 2. *For all $g \in G$ such that $g\hat{\mathfrak{h}} \subset \mathfrak{h}$, we have $\text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$.*

We already have that $\text{Ess}(g\hat{\mathfrak{h}}) \subset \text{Ess}(\hat{\mathfrak{h}})$ by the first claim. To show equality, consider the essential cores $X_{g\hat{\mathfrak{h}}}$ of $g\hat{\mathfrak{h}}$ and $X_{\hat{\mathfrak{h}}}$ of $\hat{\mathfrak{h}}$. The fact that there is an inclusion of $\text{Ess}(g\hat{\mathfrak{h}})$ into $\text{Ess}(\hat{\mathfrak{h}})$ tells us that there is an isometric embedding $X_{g\hat{\mathfrak{h}}} \rightarrow X_{\hat{\mathfrak{h}}}$, see Lemma 4.6. But these two spaces are isometric and have a cocompact automorphism group since X is locally compact and $\text{Aut}(X)$ acts cocompactly. Therefore, by Proposition 4.5, we have that $X_{g\hat{\mathfrak{h}}}$ and $X_{\hat{\mathfrak{h}}}$ are isomorphic as cube complexes. Thus $\text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$. Note that this means that for any g_1, g_2 with $g_1\hat{\mathfrak{h}} \subset g_2\hat{\mathfrak{h}}$, we have that $\text{Ess}(g_1\hat{\mathfrak{h}}) = \text{Ess}(g_2\hat{\mathfrak{h}})$.

Claim 3. *For all $\hat{\mathfrak{d}} \in \hat{\mathcal{H}}$ such that $\hat{\mathfrak{d}} \subset \mathfrak{h}$, we have $\text{Ess}(\hat{\mathfrak{d}}) = \text{Ess}(\hat{\mathfrak{h}})$.*

To see this, consider some element g skewering $\hat{\mathfrak{d}}$. Since $\hat{\mathfrak{d}}$ is disjoint from $\hat{\mathfrak{h}}$, the element g must also skewer $\hat{\mathfrak{h}}$, otherwise g would be peripheral to $\hat{\mathfrak{h}}$ and some power of it would therefore flip \mathfrak{h} (see Lemma 4.4) contradicting our assumption. After possibly replacing g by some power of g , we may assume that $g\hat{\mathfrak{h}} \subset \mathfrak{h}$. If for every positive power of g , we have that $g^n\hat{\mathfrak{h}} \cap \hat{\mathfrak{d}} \neq \emptyset$, then by Remark 3.8 we would have that for some positive power of g , say g^m , $g^m\hat{\mathfrak{h}} \in \text{Ess}(\hat{\mathfrak{d}})$, which contradicts our previous claim that $\text{Ess}(\hat{\mathfrak{d}}) \subset \text{Ess}(\hat{\mathfrak{h}})$. Thus, by passing to an appropriate positive power of g , we may assume that $g\hat{\mathfrak{h}} \subset \mathfrak{d}$. Thus $\hat{\mathfrak{d}}$ separates $\hat{\mathfrak{h}}$ from $g\hat{\mathfrak{h}}$. But we have already established that $\text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$. In particular $\text{Ess}(\hat{\mathfrak{h}}) \subset \text{Ess}(g\hat{\mathfrak{h}})$. Thus $\text{Ess}(\hat{\mathfrak{h}}) \subset \text{Ess}(\hat{\mathfrak{d}})$. Note that this means that for any $g \in G$ with $\hat{\mathfrak{d}} \subset g\hat{\mathfrak{h}}$, we have that $\text{Ess}(\hat{\mathfrak{d}}) = \text{Ess}(g\hat{\mathfrak{h}})$.

Claim 4. *For all $\hat{\mathfrak{d}} \in \hat{\mathcal{H}}$ such that $\hat{\mathfrak{d}} \cap \hat{\mathfrak{h}} = \emptyset$, we have $\text{Ess}(\hat{\mathfrak{d}}) = \text{Ess}(\hat{\mathfrak{h}})$.*

We have already proved the claim for all hyperplanes $\hat{\mathfrak{d}}$ contained in \mathfrak{h} . We need to now prove the claim for all hyperplanes $\hat{\mathfrak{d}} \subset \mathfrak{h}^*$. Let \mathfrak{d} denote the half-space associated to $\hat{\mathfrak{d}}$ such that $\mathfrak{h} \subset \mathfrak{d}$. Consider some g skewering $\hat{\mathfrak{h}}$ so that $\mathfrak{h} \subset g\hat{\mathfrak{h}}$. After passing to some positive power of g , we may assume that $\hat{\mathfrak{d}} \not\subset g\hat{\mathfrak{h}}^*$. If $\hat{\mathfrak{d}} \subset g\hat{\mathfrak{h}}$, then by what we have shown so far, we have that $\text{Ess}(\hat{\mathfrak{d}}) = \text{Ess}(g\hat{\mathfrak{h}})$ and $\text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$, so we are done. Thus, we may assume that for all positive powers of g , $g^n\hat{\mathfrak{h}} \cap \hat{\mathfrak{d}} \neq \emptyset$. By Remark 3.8, we may pass to some positive power of g such that $g\hat{\mathfrak{h}}$ is essential in $\hat{\mathfrak{d}}$. Since $\text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$, we know that $\hat{\mathfrak{d}}$ is not essential in $g\hat{\mathfrak{h}}$. So there are two possibilities: either there exists an $R > 0$ such that $g\hat{\mathfrak{h}} \subset \mathcal{N}_R(\mathfrak{d})$, or there exists some $R > 0$ such that $g\hat{\mathfrak{h}} \subset \mathcal{N}_R(\mathfrak{d}^*)$.

First, suppose that $g\hat{\mathfrak{h}} \subset \mathcal{N}_R(\mathfrak{d})$. Choose some hyperbolic element a skewering $\hat{\mathfrak{d}}$, so that $\mathfrak{d} \subset a\mathfrak{d}$. Now by passing to a high enough power of a , we may assume that $g\hat{\mathfrak{h}} \cap a\hat{\mathfrak{d}} = \emptyset$. Since $g\hat{\mathfrak{h}}$ is essential in $\hat{\mathfrak{d}}$, we may find an element $b \in \text{Stab}_G(\hat{\mathfrak{d}})$ such that b skewers $g\hat{\mathfrak{h}}$ so that $bg\hat{\mathfrak{h}} \subset g\hat{\mathfrak{h}}$. We claim that after passing to some positive power of b , we have that $ba\hat{\mathfrak{d}} \subset g\hat{\mathfrak{h}}$. To see this consider some geodesic edge path α from $a\hat{\mathfrak{d}}$ to $\hat{\mathfrak{d}}$. Let p denote the terminal endpoint α adjacent to $\hat{\mathfrak{d}}$. We may apply some power of b so that $bp \in g\hat{\mathfrak{h}}$. Now $ba\hat{\mathfrak{d}}$ is a path from $ba\hat{\mathfrak{d}}$ to bp which does not cross $g\hat{\mathfrak{h}}$. Thus $ba\hat{\mathfrak{d}} \subset g\hat{\mathfrak{h}}$. Setting $c = ba$, we have that c skewers $\hat{\mathfrak{d}}$ with $c\mathfrak{d} \subset \mathfrak{d}$. We claim further that c skewers $g\hat{\mathfrak{h}}$. For otherwise a positive power of c would flip $g\hat{\mathfrak{h}}$ by Lemma 4.4 and hence there would be an element

flipping \mathfrak{h} as well. By passing to some power of c , we may further assume that $c\mathfrak{h} \subset \mathfrak{h}$. We thus have that $\text{Ess}(cg\hat{\mathfrak{h}}) = \text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$. We also have that $\text{Ess}(c\hat{\mathfrak{d}}) = \text{Ess}(g\hat{\mathfrak{h}})$. Thus $\text{Ess}(c\hat{\mathfrak{d}}) = \text{Ess}(cg\hat{\mathfrak{h}})$, so that $\text{Ess}(\hat{\mathfrak{d}}) = \text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$.

It remains to handle the case when $g\hat{\mathfrak{h}} \subset \mathcal{N}_R(\mathfrak{d}^*)$. We may assume that $g\hat{\mathfrak{h}} \not\subset \mathcal{N}_R(\mathfrak{d})$ for otherwise, we are in case 1. Thus there exist points of $g\hat{\mathfrak{h}} \cap \mathfrak{d}^*$ arbitrarily far away from $\hat{\mathfrak{d}}$. In particular, we can find an arbitrarily large pencil of parallel hyperplanes crossing $g\hat{\mathfrak{h}}$ and lying entirely in \mathfrak{d}^* (see Lemma 2.1). Amongst them, some must belong to $\text{Ess}(g\hat{\mathfrak{h}})$ by Remark 3.8, and none of them crosses $\hat{\mathfrak{h}}$ since they are contained in \mathfrak{d}^* . This contradicts the fact that $\text{Ess}(g\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{h}})$, thereby concluding the proof of the claim.

Claim 5. *Every hyperplane in $\hat{\mathcal{H}}_1$ crosses every hyperplane in $\hat{\mathcal{H}}_2$.*

Consider a hyperplane $\hat{\mathfrak{d}}_1 \in \hat{\mathcal{H}}_1$ and $\hat{\mathfrak{d}}_2 \in \hat{\mathcal{H}}_2$ and suppose that $\hat{\mathfrak{d}}_1 \cap \hat{\mathfrak{d}}_2 = \emptyset$. Let \mathfrak{d}_1 and \mathfrak{d}_2 be the respective halfspaces satisfying $\mathfrak{d}_1 \subset \mathfrak{d}_2$. Choose some $g \in G$ such that g skewers $\hat{\mathfrak{d}}_1$, so that $g\mathfrak{d}_1 \subset \mathfrak{d}_1$. We know that g is not parallel to $\hat{\mathfrak{h}}$, otherwise we would have that $\hat{\mathfrak{d}}_1$ is essential in $\hat{\mathfrak{h}}$, contradicting the fact that $\text{Ess}(\hat{\mathfrak{d}}_1) = \text{Ess}(\hat{\mathfrak{h}})$. Consequently, g either skewers $\hat{\mathfrak{h}}$ or g is peripheral to $\hat{\mathfrak{h}}$. In either case, we may replace g by some positive power of itself such that $g\hat{\mathfrak{h}} \cap \hat{\mathfrak{h}} = \emptyset$ and $g\hat{\mathfrak{h}} \subset \mathfrak{d}_1$. Thus $g\hat{\mathfrak{h}} \cap \hat{\mathfrak{d}}_2 = \emptyset$. Applying g^{-1} we see that $g^{-1}\hat{\mathfrak{d}}_2$ and $g^{-1}\hat{\mathfrak{h}}$ are disjoint from $\hat{\mathfrak{h}}$. Thus $\text{Ess}(g^{-1}\hat{\mathfrak{h}}) = \text{Ess}(g^{-1}\hat{\mathfrak{d}}_2) = \text{Ess}(\hat{\mathfrak{h}})$ by the preceding claim. Now we apply g to see that $\text{Ess}(\hat{\mathfrak{d}}_2) = \text{Ess}(\hat{\mathfrak{h}})$ contradicting the fact that $\hat{\mathfrak{d}}_2 \in \hat{\mathcal{H}}_2$.

At this point, Lemma 2.5 ensures that X factors as a product $X_1 \times X_2$ corresponding to the partition above as $\hat{\mathcal{H}} = \hat{\mathcal{H}}_1 \cup \hat{\mathcal{H}}_2$.

Notice that $\text{Ess}(\hat{\mathfrak{h}}) \subseteq \hat{\mathcal{H}}_2$. In particular, if $\hat{\mathcal{H}}_2 = \emptyset$, then $\text{Ess}(\hat{\mathfrak{h}}) = \emptyset$ and, hence $\text{Ess}(\hat{\mathfrak{d}}) = \emptyset$ for all $\hat{\mathfrak{d}} \in \hat{\mathcal{H}}$. In view of Proposition 3.5, which can be applied since $\text{Aut}(X)$ has finitely many orbits of hyperplanes, it follows that every hyperplane in $X = X_1$ is compact. In particular X has a unique non-compact irreducible factor. Since the G -action is essential, it follows that X has no non-trivial compact factor, and we deduce that X must thus be irreducible. Thus assertion (i) holds, and assertions (ii), (iii) are trivial in this case. Moreover assertions (iv) and (v) follow from Lemma 4.8.

Assume now that there exists a hyperplane in $\hat{\mathcal{H}}_2$. Notice that by the definition of $\hat{\mathcal{H}}_1$ and $\hat{\mathcal{H}}_2$, we have $\text{Ess}(\hat{\mathfrak{h}}) = \text{Ess}(\hat{\mathfrak{d}}) \subset \hat{\mathcal{H}}_2$ for any $\hat{\mathfrak{d}} \in \hat{\mathcal{H}}_1$. Moreover the factor X_1 is isomorphic to the restriction quotient $X(\hat{\mathcal{H}}_1)$ (see §2). It follows that for every hyperplane $\hat{\mathfrak{d}}$ of X_1 , the set of essential hyperplanes of $\hat{\mathfrak{d}}$ (in the complex X_1) is empty. As shown in the case $\hat{\mathcal{H}}_2 = \emptyset$, this implies that every hyperplane of X_1 is compact. Since X_1 is essential (because X is so), it follows that X_1 is irreducible. Proposition 2.6 thus ensures the existence of a finite index subgroup of $\text{Aut}(X)$ which preserves the decomposition $X = X_1 \times X_2$.

It only remains to prove (iv) and (v) for the subcomplex X_1 . To do this there is no loss of generality in assuming $X = X_1$. In this way, we are reduced to the case $\hat{\mathcal{H}}_2 = \emptyset$ which has already been treated. \square

Corollary 4.9. *Let X be a locally compact unbounded CAT(0) cube complex with cocompact automorphism group and $G \leq \text{Aut}(X)$ be a group whose action is hereditarily essential.*

Then X decomposes as a product $X = X_1 \times \cdots \times X_p \times Y$ of subcomplexes (possibly $p = 0$ or Y is reduced to a single vertex). This decomposition is preserved by some finite index subgroup $G' \leq G$. Moreover G' has no fixed point in $\partial_\infty Y$, every half-space of Y is G' -flippable, and for all $i \in \{1, \dots, p\}$, every hyperplane of X_i is compact.

If in addition G is closed, unimodular and acts cocompactly, then X_i is \mathbf{R} -like for all i .

Proof. If every half-space of X is G -flippable, then G cannot fix any point at infinity. We can take thus take $Y = X$ and we are done in this case.

If there is some G -unflippable half-space \mathfrak{h} , then Theorem 4.7 yields a splitting $X = X_1 \times X'$. The desired result then follows by considering the cube complex X' and using induction on dimension. \square

5. AN IRREDUCIBILITY CRITERION

5.1. Strongly seperated hyperplanes. The goal of this section is to establish the final ingredient needed for the proof of the Rank Rigidity Theorem. Recall that two hyperplanes $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ in a CAT(0) cube complex are called **strongly seperated** if no hyperplane $\hat{\mathfrak{d}}$ has a non-empty intersection with both $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$. In particular $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ must be disjoint.

Proposition 5.1. *Let X be a finite-dimensional unbounded CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially without a fixed point at infinity. Then the following conditions are equivalent.*

- (i) X is irreducible.
- (ii) There is a pair of strongly separated hyperplanes.
- (iii) For each half-space \mathfrak{h} there is a pair of half-spaces $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h}_1 \subset \mathfrak{h} \subset \mathfrak{h}_2$ and the hyperplanes $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ are strongly separated.

Notice that a general irreducibility criterion in terms of hyperplanes is already available from Lemma 2.5. The power of Proposition 5.1 is that irreducibility can be detected on a single pair of hyperplanes, rather than on a global property of the collection of all hyperplanes.

The proof of Proposition 5.1 is postponed to Section 5.3 below. We first need to assemble a few subsidiary lemmas.

5.2. Finding hyperplanes in cubical sectors. Let $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ be a pair of intersecting hyperplanes. The complementary components of $\hat{\mathfrak{h}}_1 \cup \hat{\mathfrak{h}}_2$ will be called **sectors**. Note that $\hat{\mathfrak{h}}_2$ meets $\hat{\mathfrak{h}}_1$ in a hyperplane of $\hat{\mathfrak{h}}_1$, so that $\hat{\mathfrak{h}}_1$ is divided into two **half-hyperplanes**, $\hat{\mathfrak{h}}_1 \cap \hat{\mathfrak{h}}_2$ and $\hat{\mathfrak{h}}_1 \cap \hat{\mathfrak{h}}_2^*$. Similarly, $\hat{\mathfrak{h}}_1$ subdivides $\hat{\mathfrak{h}}_2$ into two half-hyperplanes. Thus a sector of $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ is bounded by two half-hyperplanes. Two sectors determined by $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ are called **opposite** if they do not share a common half-hyperplane.

Our main technical lemma is the following criterion ensuring the existence of hyperplanes contained in opposite sectors determined by a pair of intersecting hyperplanes.

Lemma 5.2. *Let X be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially without a fixed point at infinity. Suppose that X is irreducible.*

Then for any pair of crossing hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{d}}$, there exist two disjoint hyperplanes that are respectively contained in two opposite sectors determined by $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{d}}$.

Before undertaking the proof of Lemma 5.2, we first need to establish a number of auxiliary facts.

Lemma 5.3. *Let X be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially without a fixed point at infinity.*

Let $\hat{\mathfrak{h}}, \hat{\mathfrak{d}}$ be a pair of crossing hyperplanes. If one of the four sectors determined by $\hat{\mathfrak{h}}, \hat{\mathfrak{d}}$ contains a hyperplane, then there exist two disjoint hyperplanes respectively contained in opposite sectors determined by $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{d}}$.




FIGURE 2. Proof of Lemma 5.3

Proof. Let \hat{d}' be a hyperplane contained in one of the four sectors determined by \hat{h}, \hat{d} . We may assume without loss of generality that $\hat{d}' \subset \hat{h} \cap \hat{d}$. Applying the Double Skewering Lemma (see §1.2), we find $g_1 \in \text{Aut}(X)$ such that $g_1 \cdot \hat{h} \subset \hat{d}'$. In particular g_1 skewers both \hat{h} and \hat{d}' . Since $\hat{d}' \subset \hat{d}$, it follows that no g_1 -axis is contained in a bounded neighbourhood of \hat{d} . Thus g_1 either skewers \hat{d} or is peripheral to \hat{d} . We deduce from Lemmas 2.3 and 4.4 that there exists some $n > 0$ such that $g_1^n \cdot \hat{d} \subset \hat{d}$. Upon replacing g_1 by g_1^n , we may and shall assume that $g_1 \cdot \hat{d} \subset \hat{d}$.

Invoking again the Double Skewering Lemma, we now find an element $g_2 \in \text{Aut}(X)$ such that $g_2 \cdot \hat{d} \subset g_1 \cdot \hat{d}'$. Setting $\gamma = g_2 g_1$, we notice that

$$\gamma \cdot \hat{h} = g_2 g_1 \cdot \hat{h} \subset g_2 \cdot \hat{d}' \subset g_2 \cdot \hat{d} \subset g_1 \cdot \hat{d}' \subset g_1 \cdot \hat{h} \subset \hat{d}',$$

from which it follows that γ skewers both \hat{h} and \hat{d}' . Moreover we also have

$$\gamma \cdot \hat{d} = g_2 g_1 \cdot \hat{d} \subset g_2 \cdot \hat{d} \subset g_1 \cdot \hat{d}' \subset \hat{d}',$$

from which it follows that either $\gamma \cdot \hat{d} \subset \hat{d}'$, or $\gamma \cdot \hat{d}^* \subset \hat{d}'$.

Assume first that $\gamma \cdot \hat{d} \subset \hat{d}'$. Thus we have $\hat{h} \cup \hat{d} \subset \gamma^{-1} \cdot \hat{d}'$. Passing to the complementary half-spaces, this means that $\hat{h}^* \cap \hat{d}^* \supset \gamma^{-1} \cdot (\hat{d}')^*$. In particular the hyperplane $\gamma^{-1} \cdot \hat{d}'$ is contained in the sector $\hat{h}^* \cap \hat{d}^*$, which is opposite to the sector $\hat{h} \cap \hat{d}$ containing the hyperplane \hat{d} . Thus we are done in this case.

Assume now that $\gamma \cdot \hat{d}^* \subset \hat{d}'$ (see Figure 2). In that case the hyperplane $\gamma^{-1} \cdot \hat{d}'$ is contained in the sector $\hat{h}^* \cap \hat{d}$. We then repeat the above construction with the roles of \hat{h} and \hat{d} interchanged. This yields an element $\gamma' \in \text{Aut}(X)$ such that $\gamma' \cdot \hat{d} \subset \hat{d}'$ and either $\gamma' \cdot \hat{h} \subset \hat{d}'$, or $\gamma' \cdot \hat{h}^* \subset \hat{d}'$. In the former case, we deduce as above that $(\gamma')^{-1} \cdot \hat{d}'$ is contained in the sector $\hat{h}^* \cap \hat{d}^*$ and we are done. Otherwise $(\gamma')^{-1} \cdot \hat{d}'$ is contained in the sector $\hat{h} \cap \hat{d}^*$. This means that we have found two hyperplanes, namely $\gamma^{-1} \cdot \hat{d}'$ and $(\gamma')^{-1} \cdot \hat{d}'$, which are respectively contained in the opposite sectors $\hat{h}^* \cap \hat{d}$ and $\hat{h} \cap \hat{d}^*$. \square

We are now ready for the following.

Proof of Lemma 5.2. We denote by $\hat{\mathcal{H}}$ the set of all hyperplanes which are equal to or disjoint from $\hat{\mathfrak{h}}$, and by $\hat{\mathcal{K}}$ the set of hyperplanes which are equal to or disjoint from $\hat{\mathfrak{d}}$. Thus $\hat{\mathfrak{h}} \in \hat{\mathcal{H}}$ and $\hat{\mathfrak{d}} \in \hat{\mathcal{K}}$. Moreover, if $\hat{\mathcal{H}} \cap \hat{\mathcal{K}} \neq \emptyset$, then some hyperplane is disjoint from both $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{d}}$, and Lemma 5.3 then yields the desired conclusion. We assume henceforth that $\hat{\mathcal{H}}$ is disjoint from $\hat{\mathcal{K}}$. We shall consider two cases.




FIGURE 3. Proof of Lemma 5.2

Assume first that there exist hyperplanes $\hat{\mathfrak{h}}' \in \hat{\mathcal{H}}$ and $\hat{\mathfrak{d}}' \in \hat{\mathcal{K}}$ such that $\hat{\mathfrak{h}}'$ does not cross $\hat{\mathfrak{d}}'$. There is no loss of generality in assuming $\mathfrak{h}' \subset \mathfrak{h}$ and $\mathfrak{d}' \subset \mathfrak{d}$ (see Figure 3). We then apply Lemma 5.3 to the pair of crossing hyperplanes $\hat{\mathfrak{h}}, \hat{\mathfrak{d}}'$ and the hyperplane $\hat{\mathfrak{h}}'$. This implies in particular that one of the two sectors $\mathfrak{h} \cap \mathfrak{d}'$ or $\mathfrak{h}^* \cap \mathfrak{d}'$ contains some hyperplane $\hat{\mathfrak{h}}''$. Notice that in particular we have $\hat{\mathfrak{h}} \subset \mathfrak{d}$. We then invoke Lemma 5.3 once more, this time to the crossing pair $\hat{\mathfrak{h}}, \hat{\mathfrak{d}}$ and the hyperplane $\hat{\mathfrak{h}}''$. This provides the required pair of hyperplanes.

Assume next that every hyperplane in $\hat{\mathcal{H}}$ meets every hyperplane in $\hat{\mathcal{K}}$. We shall show that this implies that X admits a non-trivial product decomposition, which contradicts the hypothesis that X is irreducible. To this end, we define $\hat{\mathcal{H}}'$ (resp. $\hat{\mathcal{K}}'$) to be the collection of all those hyperplanes $\hat{\mathfrak{a}}$ which are disjoint from some hyperplane belonging to $\hat{\mathcal{H}}$ (resp. $\hat{\mathcal{K}}$). Thus we have $\hat{\mathcal{H}} \subseteq \hat{\mathcal{H}}'$ and $\hat{\mathcal{K}} \subseteq \hat{\mathcal{K}}'$.

We claim that for each half-space \mathfrak{a} such that $\hat{\mathfrak{a}} \in \hat{\mathcal{H}}'$, there exist two half-spaces $\mathfrak{b}, \mathfrak{b}'$ such that $\mathfrak{b} \subset \mathfrak{a} \subset \mathfrak{b}'$ and $\hat{\mathfrak{b}}, \hat{\mathfrak{b}}'$ both belong to $\hat{\mathcal{H}}$.

Indeed, since $\hat{\mathfrak{a}} \in \hat{\mathcal{H}}'$, there is some $\hat{\mathfrak{c}} \in \hat{\mathcal{H}}$ which is disjoint from $\hat{\mathfrak{a}}$. Thus $\hat{\mathfrak{c}}$ is disjoint from $\hat{\mathfrak{h}}$, or $\hat{\mathfrak{c}} = \hat{\mathfrak{h}}$. If $\hat{\mathfrak{a}}$ is disjoint from $\hat{\mathfrak{h}}$, then $\hat{\mathfrak{a}}$ belongs to $\hat{\mathcal{H}}$ and the claim follows using the fact that all hyperplanes are essential. Otherwise $\hat{\mathfrak{a}}$ crosses $\hat{\mathfrak{h}}$, and we can apply Lemma 5.3 to the crossing pair $\hat{\mathfrak{a}}, \hat{\mathfrak{h}}$ and the hyperplane $\hat{\mathfrak{c}}$. This yields a pair of disjoint hyperplanes $\hat{\mathfrak{b}}, \hat{\mathfrak{b}}'$ which are separated by both $\hat{\mathfrak{a}}$ and $\hat{\mathfrak{h}}$. Therefore $\hat{\mathfrak{b}}, \hat{\mathfrak{b}}'$ both belong to $\hat{\mathcal{H}}$ and the claim follows.

Similar arguments show the corresponding claim for $\hat{\mathcal{K}}'$, namely: for each half-space \mathfrak{a} such that $\hat{\mathfrak{a}} \in \hat{\mathcal{K}}'$, there exist two half-spaces $\mathfrak{b}, \mathfrak{b}'$ such that $\mathfrak{b} \subset \mathfrak{a} \subset \mathfrak{b}'$ and $\hat{\mathfrak{b}}, \hat{\mathfrak{b}}'$ both belong to $\hat{\mathcal{K}}$.

These two claims imply that every hyperplane in $\hat{\mathcal{H}}'$ crosses every hyperplane in $\hat{\mathcal{K}}'$. Let $\mathcal{R} = \hat{\mathcal{H}}(X) \setminus (\hat{\mathcal{H}}' \cup \hat{\mathcal{K}}')$. Thus every hyperplane in \mathcal{R} crosses every hyperplane in $\hat{\mathcal{H}} \cup \hat{\mathcal{K}}$ and, hence, every hyperplane in $\hat{\mathcal{H}}' \cup \hat{\mathcal{K}}'$ by the above two claims. Finally, notice that the partition

$$\hat{\mathcal{H}}(X) = \hat{\mathcal{H}}' \cup \hat{\mathcal{K}}' \cup \mathcal{R}$$

is non-trivial since $\hat{\mathfrak{h}}'' \in \hat{\mathcal{H}}'$ and $\hat{\mathfrak{d}} \in \hat{\mathcal{K}}'$. Therefore, we can invoke Lemma 2.5, which yields the absurd conclusion that X is a product. \square

We shall also need the following variant of Lemma 5.3.

Lemma 5.4. *Let X be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially without a fixed point at infinity.*

Let $\mathfrak{h}_1 \supseteq \mathfrak{h}_2 \supseteq \mathfrak{h}_3$ be a chain of half-spaces and \mathfrak{d} be a half-space such that $\hat{\mathfrak{d}}$ crosses $\hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_2$ and $\hat{\mathfrak{h}}_3$, and the sector $\mathfrak{h}_3 \cap \mathfrak{d}$ contains some half-space \mathfrak{d}' .

Then there exists a pair of disjoint hyperplanes $\hat{\mathfrak{h}}'$ and $\hat{\mathfrak{h}}''$ that are separated by $\hat{\mathfrak{d}}$ and at least two of the three hyperplanes $\hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_2$ and $\hat{\mathfrak{h}}_3$.




FIGURE 4. Proof of Lemma 5.4

Proof. We first apply Lemma 5.3 to the crossing pair $\hat{\mathfrak{h}}_1, \hat{\mathfrak{d}}$ and the hyperplane $\hat{\mathfrak{d}}'$. This provides in particular a hyperplane $\hat{\mathfrak{h}}$ disjoint from $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{d}}$, which is either contained in $\mathfrak{h}_1^* \cap \mathfrak{d}^*$ or in $\mathfrak{h}_1^* \cap \mathfrak{d}$.

In the former case, we have two disjoint hyperplanes $\hat{\mathfrak{h}}' = \hat{\mathfrak{d}}'$ and $\hat{\mathfrak{h}}' = \hat{\mathfrak{h}}$ which are separated by each of the elements of $\{\hat{\mathfrak{d}}, \hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_2, \hat{\mathfrak{h}}_3\}$. Thus we are done in this case.

Assume now that $\hat{\mathfrak{h}} \subset \hat{\mathfrak{h}}_1^* \cap \mathfrak{d}$. We now invoke Lemma 5.3 again, this time applied to the crossing pair $\hat{\mathfrak{d}}, \hat{\mathfrak{h}}_2$. This provides in particular a hyperplane $\hat{\mathfrak{h}}''$ which is either contained in $\mathfrak{h}_2 \cap \mathfrak{d}^*$ or in $\mathfrak{h}_2^* \cap \mathfrak{d}^*$.

In the former case, we have two disjoint hyperplanes $\hat{\mathfrak{h}}' = \hat{\mathfrak{h}}$ and $\hat{\mathfrak{h}}''$ which are separated by each of the elements of $\{\hat{\mathfrak{d}}, \hat{\mathfrak{h}}_1, \hat{\mathfrak{h}}_2\}$, and we are done.

In the latter case, we have two disjoint hyperplanes $\hat{\mathfrak{h}}' = \hat{\mathfrak{d}}'$ and $\hat{\mathfrak{h}}''$ which are separated by each of the elements of $\{\hat{\mathfrak{d}}, \hat{\mathfrak{h}}_2, \hat{\mathfrak{h}}_3\}$, and we are equally done. \square

5.3. Proof of the Irreducibility Criterion.

Proof of Proposition 5.1. The implication (ii) \Rightarrow (i) follows easily from Lemma 2.5, while (iii) \Rightarrow (ii) is obvious. We need to show that (i) \Rightarrow (iii).

Assume thus that X is irreducible and let \mathfrak{h} be any half-space. Suppose for a contradiction that (iii) fails. Thus for each pair of half-spaces $\mathfrak{h}', \mathfrak{h}''$ such that $\mathfrak{h}' \subsetneq \mathfrak{h} \subsetneq \mathfrak{h}''$, there is some hyperplane $\hat{\mathfrak{d}}$ crossing both \mathfrak{h}' and \mathfrak{h}'' .

Since Γ acts essentially, there is some $\gamma \in \Gamma$ which skewers $\hat{\mathfrak{h}}$ (see Proposition 3.2). Upon replacing γ by an appropriate power, we have $\gamma^{-1}\mathfrak{h} \subsetneq \mathfrak{h} \subsetneq \gamma\mathfrak{h}$. Set $\mathfrak{h}_0 = \gamma^{-1}\mathfrak{h}$ and $\mathfrak{h}'_0 = \gamma\mathfrak{h}$.

We shall now construct inductively an infinite sequence $(\mathfrak{h}_n, \mathfrak{h}'_n, \mathfrak{d}_n)_{n>0}$ of triples of half-spaces which, together with \mathfrak{h}_0 and \mathfrak{h}'_0 , satisfies the following conditions for all $n > 0$:

- (a) $\hat{\mathfrak{d}}_n$ crosses $\hat{\mathfrak{h}}_{n-1}$ and $\hat{\mathfrak{h}}'_{n-1}$.
- (b) $\hat{\mathfrak{d}}_n$ separates $\hat{\mathfrak{h}}_n$ from $\hat{\mathfrak{h}}'_n$.
- (c) $\mathfrak{h}_n \subsetneq \mathfrak{h}_{n-1} \subsetneq \mathfrak{h} \subsetneq \mathfrak{h}'_{n-1} \subsetneq \mathfrak{h}'_n$.

For all $n > 0$, we now describe an inductive construction of the triple $(\mathfrak{h}_n, \mathfrak{h}'_n, \mathfrak{d}_n)$. We apply the Double Skewering Lemma to the pair $\mathfrak{h}_{n-1} \subsetneq \mathfrak{h}'_{n-1}$. This yields an element $\gamma \in \Gamma$ such that $\mathfrak{h}'_{n-1} \subsetneq \gamma\mathfrak{h}_{n-1} \subsetneq \gamma\mathfrak{h}'_{n-1}$. Since the pair $\mathfrak{h}_{n-1}, \gamma\mathfrak{h}'_{n-1}$ cannot be strongly separated by assumption, there is some hyperplane $\hat{\mathfrak{d}}$ crossing both \mathfrak{h}_{n-1} and $\gamma\mathfrak{h}'_{n-1}$. By Lemma 5.2, one of the sectors $\mathfrak{d} \cap \mathfrak{h}_{n-1}$ or $\mathfrak{d}^* \cap \mathfrak{h}_{n-1}$ must contain properly some half-space. Upon replacing \mathfrak{d} by \mathfrak{d}^* , we are thus in a position to invoke Lemma 5.4 to the chain $\mathfrak{h}_{n-1} \subsetneq \mathfrak{h}'_{n-1} \subsetneq \gamma\mathfrak{h}'_{n-1}$ and the half-space \mathfrak{d} . This yields two hyperplanes $\hat{\mathfrak{h}}'$ and $\hat{\mathfrak{h}}''$ which are separated by $\hat{\mathfrak{d}}$ and at least two of the three hyperplanes $\hat{\mathfrak{h}}_{n-1}, \hat{\mathfrak{h}}'_{n-1}$ and $\gamma\hat{\mathfrak{h}}'_{n-1}$.

If $\hat{\mathfrak{h}}'$ and $\hat{\mathfrak{h}}''$ are separated by $\hat{\mathfrak{h}}_{n-1}$ and $\hat{\mathfrak{h}}'_{n-1}$, then we define \mathfrak{h}_n as the half-space bounded by $\hat{\mathfrak{h}}'$ and contained in \mathfrak{h}_{n-1} , and we define \mathfrak{h}'_n as the half-space bounded by $\hat{\mathfrak{h}}''$ and containing \mathfrak{h}'_{n-1} . We also set $\mathfrak{d}_n = \mathfrak{d}$ in this case.

If $\hat{\mathfrak{h}}'$ and $\hat{\mathfrak{h}}''$ are separated by $\hat{\mathfrak{h}}'_{n-1}$ and $\gamma\hat{\mathfrak{h}}'_{n-1}$, then the hyperplanes $\gamma^{-1}\hat{\mathfrak{h}}'$ and $\gamma^{-1}\hat{\mathfrak{h}}''$ are separated by $\gamma^{-1}\hat{\mathfrak{h}}'_{n-1}, \hat{\mathfrak{h}}_{n-1}$ and $\hat{\mathfrak{h}}'_{n-1}$. We then define \mathfrak{h}_n as the half-space bounded by $\gamma^{-1}\hat{\mathfrak{h}}'$ and contained in \mathfrak{h}_{n-1} , and we define \mathfrak{h}'_n as the half-space bounded by $\gamma^{-1}\hat{\mathfrak{h}}''$ and containing \mathfrak{h}'_{n-1} . We also set $\mathfrak{d}_n = \gamma^{-1}\mathfrak{d}$ in this case.

In either case, the hyperplane $\hat{\mathfrak{d}}_n$ crosses both \mathfrak{h}_{n-1} and \mathfrak{h}'_{n-1} . Moreover $\hat{\mathfrak{d}}_n$ separates $\hat{\mathfrak{h}}_n$ from $\hat{\mathfrak{h}}'_n$ and we have $\mathfrak{h}_n \subsetneq \mathfrak{h}_{n-1} \subsetneq \mathfrak{h} \subsetneq \mathfrak{h}'_{n-1} \subsetneq \mathfrak{h}'_n$, as desired.

This inductive construction yields an infinite sequence of triples $(\mathfrak{h}_n, \mathfrak{h}'_n, \mathfrak{d}_n)$ satisfying the conditions (a), (b), (c).

Notice that $\hat{\mathfrak{d}}_n$ crosses both $\hat{\mathfrak{h}}_{n-1}$ and $\hat{\mathfrak{h}}'_{n-1}$, it must in fact cross $\hat{\mathfrak{h}}_m$ and $\hat{\mathfrak{h}}'_m$ for all $m < n$ by (c). In particular we have $\hat{\mathfrak{d}}_n \neq \hat{\mathfrak{d}}_m$ for all $m < n$ by (b). Moreover (b) also implies that $\hat{\mathfrak{d}}_n$ must cross $\hat{\mathfrak{d}}_m$ for all $m < n$. It follows that the hyperplanes $\hat{\mathfrak{d}}_1, \hat{\mathfrak{d}}_2, \dots$ are pairwise distinct and pairwise crossing. This contradicts the fact that X is finite-dimensional (see Lemma 2.1). \square

6. RANK RIGIDITY

The goal of this section is to provide a proof of Theorem A from the introduction.

6.1. Strongly separated hyperplanes and contracting isometries. As mentioned in the introduction, the last missing piece in the proof of Rank Rigidity for CAT(0) cube complexes is that a hyperbolic isometry which double skewers a pair of strongly separated hyperplanes must necessarily be contracting. This contracting behaviour will be deduced from the following key lemma, which readily implies that such an isometry must be rank one.

Lemma 6.1. *Let X be a finite-dimensional CAT(0) cube complex, let \mathfrak{h} be a half-space and $\gamma \in \text{Aut}(X)$ be a hyperbolic isometry with axis ℓ such that $\gamma\mathfrak{h} \subsetneq \mathfrak{h}$. Assume that the hyperplanes $\hat{\mathfrak{h}}$ and $\gamma\hat{\mathfrak{h}}$ are strongly separated.*

Then there is a constant C , depending only on γ , such that each geodesic segment crossing at least three walls in the orbit $\langle \gamma \rangle \hat{\mathfrak{h}}$ has a non-empty intersection with the C -neighbourhood of ℓ .

Proof. Let $p_0 \in \ell \cap \hat{\mathfrak{h}}$ and set $\mathfrak{h}_n = \gamma^n \mathfrak{h}$ and $p_n = \gamma^n p_0$ for all $n \in \mathbf{Z}$. Let also N be the number of hyperplanes crossed by $[p_0, p_1]$.

Let $a, b \in X$ be two points such that the geodesic segment $[a, b]$ crosses $\hat{\mathfrak{h}}_{i-1}$, $\hat{\mathfrak{h}}_i$ and $\hat{\mathfrak{h}}_{i+1}$ for some $i \in \mathbf{Z}$. Let $x \in [a, b] \cap \hat{\mathfrak{h}}_i$. We shall show that x is at distance at most N' away from ℓ , where N' is some constant depending only on N .

By hypothesis, the two pairs of hyperplanes $\hat{\mathfrak{h}}_{i-1}, \hat{\mathfrak{h}}_i$ and $\hat{\mathfrak{h}}_i, \hat{\mathfrak{h}}_{i+1}$ are both strongly separated. Therefore, none of the hyperplanes separating p_i from x can cross $\hat{\mathfrak{h}}_{i-1}$ or $\hat{\mathfrak{h}}_{i+1}$. Since $[a, b]$ is geodesic, it crosses each hyperplane at most once. It follows that each of the hyperplanes separating p_i from x must cross either $[p_{i-1}, p_i]$ or $[p_i, p_{i+1}]$. We conclude that the number of hyperplanes separating p_i from x is at most $2N$. The desired conclusion follows since the CAT(0) metric is quasi-isometric to the hyperplane distance (see Lemma 2.2). \square

Lemma 6.2. *Let X be a finite-dimensional CAT(0) cube complex, let $\mathfrak{h}' \subsetneq \mathfrak{h}''$ be a nested pair of half-spaces and $\gamma \in \text{Aut}(X)$ be such that $\mathfrak{h}'' \subsetneq \gamma\mathfrak{h}'$.*

If the hyperplanes $\hat{\mathfrak{h}}', \hat{\mathfrak{h}}''$ are strongly separated, then γ is a contracting isometry.

Proof. Let ℓ be some γ -axis (notice that the hypotheses imply that γ is hyperbolic). Let also $p_0 \in \ell \cap \hat{\mathfrak{h}}'$ and for all $n \in \mathbf{Z}$, set $p_n = \gamma^n p_0$ and $\mathfrak{h}_n = \gamma^n \mathfrak{h}'$. The hypotheses imply that for all $i \neq j \in \mathbf{Z}$, the hyperplanes $\hat{\mathfrak{h}}_i, \hat{\mathfrak{h}}_j$ are strongly separated.

Suppose for a contradiction that γ is not contracting. Then there are two sequences (x_n) and (y_n) in X such that $d(x_n, y_n) \leq d(x_n, \ell)$ and that $\lim_n d(x'_n, y'_n) = \infty$, where x'_n and y'_n respectively denote the orthogonal projection of x_n and y_n on ℓ .

Since $[p_0, p_1]$ is a fundamental domain for the $\langle \gamma \rangle$ -action on ℓ , there is no loss of generality in assume that $x'_n \in [p_0, p_1]$ for all n . Upon extracting and reversing the orientation on ℓ , we may further assume that the n hyperplanes $\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n$ separate x'_n from y'_n for all $n > 0$.

Let C be the constant from Lemma 6.1. We claim that the number of hyperplanes in $\{\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n\}$ that can be crossed by $[x_n, x'_n]$ or $[y_n, y'_n]$ is at most $C + 3$. Indeed, since x'_n is the orthogonal projection of x_n on ℓ , it follows that $d(z, \ell) = d(z, x'_n)$ for all $z \in [x_n, x'_n]$. Therefore, for $z \in [x_n, x'_n]$ with $d(z, x'_n) > C$, we deduce from Lemma 6.1 that $[z, x_n]$ cannot cross more than two walls in $\{\hat{\mathfrak{h}}_1, \dots, \hat{\mathfrak{h}}_n\}$. This proves the claim for $[x_n, x'_n]$; the argument for $[y_n, y'_n]$ is similar.

We next consider the geodesic quadrilateral with vertices x_n, x'_n, y'_n, y_n . Each hyperplane separating x'_n from y'_n must cross one of the three geodesic segments $[x_n, y_n]$, $[x_n, x'_n]$ or

$[y_n, y'_n]$. From the previous claim, we deduce that at least $n - 2C - 6$ of the hyperplanes $\hat{h}_1, \dots, \hat{h}_n$ must cross $[x_n, y_n]$. Let i be the minimal index such that $[x_n, y_n]$ crosses \hat{h}_i . Let $q_i \in [x_n, y_n] \cap \hat{h}_i$ and $q_{i+2} \in [x_n, y_n] \cap \hat{h}_{i+2}$. By Lemma 6.1, there is some $q \in [q_i, q_{i+2}]$ which is C -close to ℓ . Thus we have

$$\begin{aligned} d(x_n, y_n) &\leq d(x_n, x'_n) \\ &= d(x_n, \ell) \\ &\leq d(x_n, q) + d(q, \ell) \\ &\leq d(x_n, q) + C, \end{aligned}$$

whence $d(q, y_n) = d(x_n, y_n) - d(x_n, q) \leq C$. In particular, the number of hyperplanes crosses by $[q, y_n]$ is at most C .

Recalling that the number of hyperplanes in $\hat{h}_1, \dots, \hat{h}_n$ crossed by $[x_n, y_n]$ is at least $n - 2C - 6$, we deduce that the number of hyperplanes crossed by $[q, y_n]$ must be at least $n - 2C - 8$. Thus we get a contradiction as soon as $n > 3C + 8$. \square

6.2. Proof of Rank Rigidity. In view of Proposition 3.5, Theorem A from the introduction is an immediate consequence of the following.

Theorem 6.3. *Let X be a finite-dimensional CAT(0) cube complex and $\Gamma \leq \text{Aut}(X)$ be a group acting essentially without fixed point at infinity.*

Then X is a product of two cube subcomplexes or every hyperplane of X is skewered by a contracting isometry in Γ .

If in addition X is locally compact and Γ acts cocompactly, then the same conclusion holds even if Γ fixes a point at infinity.

Proof. Suppose that X is irreducible and let \mathfrak{h} be a half-space. By Proposition 5.1, there is a pair of half-spaces $\mathfrak{h}', \mathfrak{h}''$ such that $\mathfrak{h}' \subsetneq \mathfrak{h} \subsetneq \mathfrak{h}''$ and the hyperplanes \hat{h}' and \hat{h}'' are strongly separated. From the Double Skewering Lemma, we deduce that there is some $\gamma \in \Gamma$ such that $\mathfrak{h}'' \subsetneq \gamma \mathfrak{h}'$. In particular γ skewers \hat{h} . We conclude by invoking Lemma 6.2, which ensures that γ is a contracting isometry.

Assume now that X is locally compact and that Γ acts cocompactly (but not necessarily properly), and assume that X is irreducible. We then invoke Corollary 4.9. Two situations can occur. The first is that all hyperplanes are compact; in that case the Flat Plane Theorem (see [BH99, Theorem III.H.1.5]) implies that X is Gromov hyperbolic, and it is then clear that all hyperbolic isometries are contracting. The second is that Γ has no fixed point at infinity, and we are then reduced to a situation which has already been dealt with. \square

Corollary 6.4. *Let X be an unbounded locally compact CAT(0) cube complex such that $\text{Aut}(X)$ acts cocompactly and let $\Gamma \leq \text{Aut}(X)$ be a possibly non-uniform lattice. We have the following.*

- (i) Γ acts essentially on the essential core Y of X .
- (ii) Y embeds as an $\text{Aut}(X)$ -invariant convex subcomplex in X .
- (iii) Y decomposes as a product $Y \cong Y_1 \times \dots \times Y_m$ of $m \geq 1$ irreducible unbounded convex subcomplexes, and for each i , there is some $\gamma_i \in \Gamma$ preserving that decomposition and acting on Y_i as a rank one isometry.

Proof. We remark that, since Γ is possibly non-uniform, its action on X is not necessarily cocompact. The goal is thus to reduce to a situation where Γ acts essentially without a fixed point at infinity, so that the first part of Theorem 6.3 will give the desired conclusion.

We start by applying Proposition 3.5 to the whole group $\text{Aut}(X)$ (see also Proposition 3.7). This shows that there is thus no loss of generality in assuming that $\text{Aut}(X)$

acts essentially on X . By Proposition 2.6 we have a canonical product decomposition $X \cong X_1 \times \cdots \times X_m$, where X_i is an irreducible subcomplex. Moreover, upon replacing Γ by a finite index subgroup, we can assume that Γ preserves this decomposition.

For all $i = 1, \dots, m$, we need to show that Γ acts essentially on X_i and that it contains an element γ_i acting as a rank one isometry on X_i . We already know from Theorem 6.3 that $\text{Aut}(X_i)$ contains rank one isometries.

Since $\text{Aut}(X_i)$ acts cocompactly, there is a non-empty minimal closed convex $\text{CAT}(0)$ subspace $Z_i \subseteq X_i$, which need not be a subcomplex. Since $\text{Aut}(X_i)$ contains rank one elements, it follows that Z_i is an irreducible $\text{CAT}(0)$ space. If Z_i is flat, then it is isometric to the real line and X is \mathbf{R} -like. The desired result is then clear. We assume henceforth that Z_i is not flat.

We claim that $\text{Aut}(X_i)$ has no fixed point in $\partial_\infty X_i$. This follows from [CM] (for a special case, see also Theorem 3.14 from [CM09a]); a direct argument in the current specific setting could also be obtained using Corollary 4.9.

Now, from the ‘geometric Borel density theorem’ proved in [CM09a, Theorem 2.4], we now infer that Γ has no fixed point at infinity of Z_i and that it acts minimally on Z_i , *i.e.* Z_i contains no non-empty proper Γ -invariant closed $\text{CAT}(0)$ subspace. Since every hyperplane of X_i separates Z_i , it follows in particular that for any $z \in Z_i$, the orbit $\Gamma.z$ contains points on both sides of every hyperplane. Recall that each hyperplane is $\text{Aut}(X_i)$ -essential. In particular, given a half-space \mathfrak{h} , there is a half-space $\mathfrak{d} \subset \mathfrak{h}$ such that the distance from any point of \mathfrak{d} to $\hat{\mathfrak{h}}$ is arbitrarily large. Since \mathfrak{d} intersects non-trivially the orbit $\Gamma.z$, we deduce that \mathfrak{h} contains points of the orbit $\Gamma.z$ which are arbitrarily far away from the hyperplane $\hat{\mathfrak{h}}$. In other words, this means that \mathfrak{h} is Γ -deep. Since \mathfrak{h} was arbitrary, this proves that every hyperplane of X is Γ -essential or, equivalently, that the Γ -action on X_i is essential.

We have seen that Γ has no fixed point at infinity of Z_i . Since $Z_i \subseteq X_i$ is $\text{Aut}(X_i)$ -invariant and since $\text{Aut}(X_i)$ acts cocompactly on X_i , it follows that Γ has no fixed point at infinity of X_i . At this point, we are able to invoke Theorem 6.3, which provides an element γ_i in Γ acting as a rank one isometry on X_i . \square

Corollary 6.5. *Let X be a locally compact $\text{CAT}(0)$ cube complex such that $\text{Aut}(X)$ acts cocompactly and essentially. Then the action of any (possibly non-uniform) lattice $\Gamma \leq \text{Aut}(X)$ is hereditarily essential.*

Proof. For each hyperplane $\hat{\mathfrak{d}}$, the stabilizer $\text{Stab}_{\text{Aut}(X)}(\hat{\mathfrak{d}})$ is an open subgroup of $\text{Aut}(X)$ which acts cocompactly on $\hat{\mathfrak{d}}$. It follows that $\text{Stab}_\Gamma(\hat{\mathfrak{d}}) = \Gamma \cap \text{Stab}_{\text{Aut}(X)}(\hat{\mathfrak{d}})$ is a lattice in $\text{Stab}_{\text{Aut}(X)}(\hat{\mathfrak{d}})$ to which Corollary 6.4 applies. This shows in particular that $\text{Stab}_\Gamma(\hat{\mathfrak{d}})$ acts essentially on the essential core of $\hat{\mathfrak{d}}$, and the desired result follows by induction on dimension. \square

We can now complete the proof of Corollary B from the introduction, which is concerned with the geodesically complete case.

Proof of Corollary B. A group acting cocompactly on a geodesically complete $\text{CAT}(0)$ space necessarily acts minimally, see [CM09b, Lemma 3.13]. In view of Lemma 3.3 and Remark 3.4, this implies that every hyperplane is $\text{Aut}(X)$ -essential. The result then follows from Corollary 6.4. \square

7. APPLICATIONS

7.1. Cube complexes with invariant Euclidean flats. We start with an elementary observation (see §4.4 for the definition of an \mathbf{R} -like complex).

Lemma 7.1. *Let X be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially. Then $\text{Aut}(X)$ stabilizes some n -dimensional flat $\mathbf{R}^n \subseteq X$ if and only if X decomposes as a product $X = X_1 \times \cdots \times X_n$ of subcomplexes, each of which is essential and \mathbf{R} -like.*

Proof. Clearly, an essential CAT(0) cube complex X_i that is \mathbf{R} -like is irreducible. Therefore, a product $X = X_1 \times \cdots \times X_n$ of n complexes of this form has the property that $\text{Aut}(X)$ preserves some Euclidean flat, since this product decomposition is preserved by some finite index subgroup of $\text{Aut}(X)$ (see Proposition 2.6).

Assume conversely that $\text{Aut}(X)$ stabilizes some n -dimensional flat $F \subseteq X$. Since X is essential, every hyperplane must separate F into two non-empty disjoint pieces. More precisely, for each $\hat{h} \in \hat{\mathcal{H}}(X)$, the intersection $\hat{h} \cap F$ is a Euclidean hyperplane of F . Let $\hat{\mathcal{H}}_1$ be the collection of all hyperplanes $\hat{d} \in \hat{\mathcal{H}}(X)$ such that $\hat{d} \cap F$ is parallel to $\hat{h} \cap F$ in the sense of Euclidean geometry. Let also $\hat{\mathcal{H}}_2 = \hat{\mathcal{H}}(X) - \hat{\mathcal{H}}_1$. If X is not \mathbf{R} -like, that is to say if $n > 1$, then $\hat{h} \cap F$ is non-compact and it follows that $\hat{\mathcal{H}}_2$ is not empty. Lemma 2.5 then yields a product decomposition of X into two subcomplexes and the result follows by induction on dimension. \square

Theorem 7.2. *Let X be a finite-dimensional CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially and satisfies at least one of the following conditions:*

- (a) $\text{Aut}(X)$ has no fixed point at infinity.
- (b) $\text{Aut}(X)$ acts cocompactly and X is locally compact.

Then $\text{Aut}(X)$ stabilises some Euclidean flat if and only if there is no facing triple of hyperplanes.

Proof. Assume first that X contains an $\text{Aut}(X)$ -invariant flat F (possibly reduced to a singleton). Since every hyperplane is $\text{Aut}(X)$ -essential, every hyperplane separates F into two non-empty disjoint convex pieces. Therefore, it follows from the definition of a hyperplane that for each $\hat{h} \in \hat{\mathcal{H}}(X)$, the intersection $\hat{h} \cap F$ is a Euclidean hyperplane of F . The desired assertion follows since there is no facing triple of hyperplanes in Euclidean geometry.

Assume now that X does not contain any $\text{Aut}(X)$ -invariant flat F ; in particular it is unbounded. Let $X = X_1 \times \cdots \times X_n$ be the canonical decomposition provided by Proposition 2.6. It suffices to prove that one of the factors X_i contains a facing triple of hyperplanes. By Lemma 7.1, one of the irreducible factors, say X_i , is not \mathbf{R} -like. Thus there is no loss of generality in assuming that $X = X_i$ or, equivalently, that X is irreducible but not \mathbf{R} -like.

If condition (b) holds, then Corollary 4.9 (applied with $G = \text{Aut}(X)$) implies that $\text{Aut}(X)$ has no fixed point at infinity since $X = X_i$ is irreducible. Thus it suffices to prove the existence of a facing triple under the assumption that condition (a) holds.

Assume first that each hyperplane of X is compact. In that case, the existence of a facing triple of hyperplanes is easy to establish; details are left to the reader. We assume henceforth that some hyperplane, say \hat{h} , is non-compact.

By Proposition 5.1, there is a pair of strongly separated hyperplanes \hat{h}', \hat{h}'' such that $\hat{h}' \subsetneq \hat{h} \subsetneq \hat{h}''$. By the Double Skewering Lemma, there is some $g \in \text{Aut}(X)$ such that $\hat{h}'' \subsetneq g\hat{h}'$. In particular, each pair of hyperplanes in the $\langle g \rangle$ -orbit of \hat{h} is strongly separated.

Since \hat{h} is not compact, infinitely many hyperplanes cross it. Amongst them, none can cross $g^{-1}\hat{h}$ or $g\hat{h}$, and only finitely many separate $g^{-1}\hat{h}$ from $g\hat{h}$. It follows that there is some halfspace \mathfrak{d} containing both $g^{-1}\hat{h}$ and $g\hat{h}$, and such that \hat{d} crosses \hat{h} . It follows that \hat{d} , $g^{-1}\hat{h}$ and $g\hat{h}$ forms a facing triple of hyperplanes. \square

7.2. Tits alternative, second version.

Proof of Corollary G. It follows from [CL10, Theorem 1.7] that Γ is amenable if and only if it is {locally finite}-by-{virtually Abelian}. Thus it suffices to show that if Γ is not amenable, then it contains a non-Abelian free subgroup. We assume henceforth that Γ is non-amenable. From [CL, Theorem A.5], we deduce that Γ does not fix any point in the *ultrafilter bordification* of X (see the Appendix of [CL] for details). In particular, this means that there is some Γ -invariant restriction quotient Y of X such that the Γ -action on Y has no fixed point in $Y \cup \partial_\infty Y$. Theorem F then ensures that Γ contains a free subgroup. \square

7.3. Regular elements. We first need to recapitulate the information we have obtained so far on lattices of locally compact CAT(0) cube complexes.

Proposition 7.3. *Let X be a locally compact CAT(0) cube complex such that $\text{Aut}(X)$ acts essentially and cocompactly and let $\Gamma \leq \text{Aut}(X)$ be a (possibly non-uniform) lattice. Then:*

- (i) X decomposes as a product $X = X_1 \times \cdots \times X_p \times Y$ of subcomplexes such that X_i is \mathbf{R} -like for each i and Y has no \mathbf{R} -like factor.
- (ii) Every automorphism preserves the decomposition upon permuting some possibly isomorphic factors X_i .
- (iii) The Γ -action on X (and hence on Y) is hereditarily essential.
- (iv) Γ does not fix any point in $\partial_\infty Y$.
- (v) Every pair of disjoint hyperplanes is double-skewed by some element of Γ .

Proof. First notice that, by Proposition 2.6, there is a canonical decomposition $X = X_1 \times \cdots \times X_k$ into irreducible subcomplexes, which is virtually invariant under $\text{Aut}(X)$. Upon reordering, we can assume that the \mathbf{R} -like factors are precisely the first p ones. In particular every automorphism permutes the first p factors among themselves.

Next consider any point $x = (x_1, \dots, x_p) \in X_1 \times \cdots \times X_p$. Since the stabiliser $\text{Aut}(X)_x$ of x in $\text{Aut}(X)$ is open, it follows that Γ_x is a lattice in $\text{Aut}(X)_x$ and therefore acts as a lattice on $Y = X_{p+1} \times \cdots \times X_k$. We then observe that the Γ -action on X , as well as the Γ_x -action on Y , is hereditarily essential by Corollary 6.5. We can thus invoke Corollary 4.9. Since Y does not have any \mathbf{R} -like factor, this implies in particular that Γ_x does not fix any point in $\partial_\infty Y$. Thus (iv) holds.

It only remains to prove (v). To this end, remark that the splitting $X = X_1 \times \cdots \times X_p \times Y$ comes with a decomposition of the set of hyperplanes $\hat{\mathcal{H}}(X) = \hat{\mathcal{H}}(X_1) \cup \cdots \cup \hat{\mathcal{H}}(X_p) \cup \hat{\mathcal{H}}(Y)$. Clearly any pair of disjoint hyperplanes must belong to the same component. If they belong to $\hat{\mathcal{H}}(Y)$, then we conclude by using any of the two versions of the Double Skewing Lemma. If they belong to $\hat{\mathcal{H}}(X_i)$, we just remark that any hyperbolic isometry of X_i will double-skewer the given pair of hyperplanes as desired, since X_i is \mathbf{R} -like and all its hyperplanes are compact. \square

Proof of Theorem C. Let $\hat{\mathcal{H}}_i = \hat{\mathcal{H}}(X_i)$ for all i , so that the given splitting $X = X_1 \times \cdots \times X_n$ yields a decomposition $\hat{\mathcal{H}}(X) = \hat{\mathcal{H}}_1 \cup \cdots \cup \hat{\mathcal{H}}_n$.

For each i , pick two hyperplanes $\hat{\mathfrak{h}}_i, \hat{\mathfrak{d}}_i \in \hat{\mathcal{H}}_i$ such that $\mathfrak{h}_i \subset \mathfrak{d}_i$ and that any element g of $\text{Aut}(X_i)$ with $g \cdot \mathfrak{d}_i \subset \mathfrak{h}_i$ is rank one. The existence of such a pair of hyperplanes is guaranteed by Proposition 5.1 and Lemma 6.2. In particular, it suffices to exhibit an element $g \in \Gamma$ such that $g \cdot \mathfrak{d}_i \subset \mathfrak{h}_i$ simultaneously for all i . We shall do this by induction on n , the case $n = 1$ being covered by Proposition 7.3.

Let x be any point of X_n contained in the projection of $\hat{\mathfrak{d}}_n$ and consider its stabilisers Γ_x and $\text{Aut}(X)_x$. Since X_n is locally compact, there are finitely many hyperplanes in X_n

containing the point x . Therefore Γ_x has a finite index subgroup Γ'_x which stabilizes $\hat{\mathfrak{d}}_n$. Upon replacing Γ'_x by a subgroup of index two, we may further assume that Γ'_x stabilizes both \mathfrak{d}_n and \mathfrak{d}_n^* . Moreover, since the stabiliser $\text{Aut}(X)_x$ of x in $\text{Aut}(X)$ is an open subgroup of $\text{Aut}(X)$, it follows that $\Gamma_x = \Gamma \cap \text{Aut}(X)_x$ is a lattice in $\text{Aut}(X)_x$, and so is thus Γ'_x . The group $\text{Aut}(X_n)_x$ being compact, it follows that the image of the projection of Γ'_x to $\text{Aut}(X_1) \times \cdots \times \text{Aut}(X_{n-1})$ is a lattice. For all i , since $\hat{\mathfrak{d}}_i$ and $\hat{\mathfrak{h}}_i$ are essential hyperplanes of X , it follows that they are also essential as hyperplanes of X_i . In particular, they are Γ'_v -essential for all $i < n$ (see Corollary 6.5), and the induction hypothesis then yields an element $a \in \Gamma'_x$ such that $a.\mathfrak{d}_i \subset \mathfrak{h}_i$ for all $i < n$ and $a.\mathfrak{d}_n = \mathfrak{d}_n$.

We now pick a point $y \in X_1 \times \cdots \times X_{n-1}$ contained in the projection of $a.(\hat{\mathfrak{d}}_1 \cap \cdots \cap \hat{\mathfrak{d}}_{n-1})$. For the same reason as before, we can use the induction hypothesis to find an element $b \in \Gamma_y$ which stabilizes $a.\mathfrak{d}_i$ for all $i < n$ and such that $b.\mathfrak{d}_n \subset \mathfrak{h}_n$.

It remains to set $g = ba$ and observe that $g.\mathfrak{d}_i \subset \mathfrak{h}_i$ for all i , as required. \square

Proof of Corollary D. There is no loss of generality in assuming that Γ acts essentially (see Corollary 6.4) and that X decomposes as a product $X = X_1 \times \cdots \times X_n$ of n irreducible subcomplexes. In particular $\text{Aut}(X)$ is virtually isomorphic to $\text{Aut}(X_1) \times \cdots \times \text{Aut}(X_n)$ by Proposition 2.6.

Let $\gamma \in \Gamma$ be a regular element as provided by Theorem C. For each irreducible factor X_i of X , the centraliser in $\text{Aut}(X_i)$ of the projection of γ stabilises the pair consisting of its attracting and repelling fixed points at infinity. In particular this centraliser stabilises some geodesic line in X_i . It follows that the centraliser of γ in $\text{Aut}(X)$ stabilises some n -flat embedded in X on which it acts cocompactly.

By a Lemma of Selberg [Sel60], the centraliser $\mathcal{Z}_\Gamma(\gamma)$ is a cocompact lattice in the centraliser $\mathcal{Z}_{\text{Aut}(X)}(\gamma)$. We have just seen that, up to some compact normal subgroup, the latter is isomorphic to closed subgroup of $\text{Is}(\mathbf{R}^n)$. Thus $\mathcal{Z}_\Gamma(\gamma)$ is a Bieberbach group of rank n , and thus contains \mathbf{Z}^n . \square

7.4. Quasi-morphisms. Theorem H from the introduction will be deduced from the following.

Theorem 7.4. *Let $X = X_1 \times \cdots \times X_n$ be a product of n geodesically complete locally compact CAT(0) spaces. Assume that $\text{Is}(X_i)$ acts cocompactly on X_i and contains some rank one isometry. For any lattice $\Gamma \leq \text{Is}(X_1) \times \cdots \times \text{Is}(X_n)$, the following conditions are equivalent.*

- (i) $\widetilde{\text{QH}}(\Gamma)$ is finite-dimensional.
- (ii) $\widetilde{\text{QH}}(\Gamma) = 0$.
- (iii) For all i , the space X_i is either a tree or a rank one symmetric space, and if X_i is not isometric to the real line \mathbf{R} , then the closure G_i of the projection of Γ on $\text{Is}(X_i)$ acts doubly transitively on $\partial_\infty X_i$.

The following basic facts shows that the set of rank one elements is ‘large’ as soon as it is non-empty.

Lemma 7.5. *Let X be a proper CAT(0) space such that $\text{Is}(X)$ contains a rank one isometry. Then any group $\Gamma \leq \text{Is}(X)$ acting without fixed point at infinity and with full limit set also contains a rank one isometry.*

Proof. Let $g \in \text{Is}(X)$ be a rank one isometry. Let also $\xi_+, \xi_- \in \partial_\infty X$ be the attracting and repelling fixed points of g respectively. Since Γ has full limit set, there is a sequence $\gamma_n \in \Gamma$ such that for some (and hence for all) point $p \in X$, we have $\lim_{n \rightarrow \infty} \gamma_n.p = \xi_+$. Upon extracting a subsequence, we can assume moreover that the sequence $\gamma_n^{-1}.p$ also converges

as n tends to infinity, say, to some point $\eta \in \partial_\infty X$. (Clearly $(\gamma_n^{-1}.p)$ is unbounded, otherwise $(\gamma_n.p)$ would be bounded.)

If $\eta = \xi_+$, then we pick any element $\lambda \in \Gamma$ such that $\lambda.\xi_+ \neq \xi_+$ and we set $\gamma'_n = \gamma_n\lambda$ for all n . In that case we still have $\lim_{n \rightarrow \infty} \gamma'_n.p = \xi_+$, and we have also

$$\lim_{n \rightarrow \infty} (\gamma'_n)^{-1}.p = \lambda^{-1}.(\lim_{n \rightarrow \infty} \gamma_n^{-1}.p) = \lambda^{-1}.\xi_+ \neq \xi_+.$$

It then follows from Lemmas III.3.1 and III.3.2 from [Bal95] that γ'_n is rank one for all sufficiently large $n > 0$. \square

We shall also use the following trichotomy, which is obtained in [CF10] as a consequence of the construction of quasi-morphisms by M. Bestvina and K. Fujiwara in [BF09].

Proposition 7.6 (Theorem 1.8 from [CF10]). *Let X be a proper CAT(0) space with cocompact isometry group, and let $G \leq \text{Is}(X)$ be a group containing some rank one isometry. Then one of the following assertions hold.*

- (i) G fixes a point in $\partial_\infty X$ or stabilises some geodesic line, and the closure $\overline{G} \leq \text{Is}(X)$ is amenable.
- (ii) \overline{G} acts doubly transitively on $\partial_\infty X$ and the space $\widetilde{\text{QH}}_c(\overline{G})$ of continuous non-trivial quasi-morphisms vanishes.
- (iii) \overline{G} is not doubly transitive on $\partial_\infty X$, and the spaces $\widetilde{\text{QH}}(G)$ and $\widetilde{\text{QH}}_c(\overline{G})$ are both infinite-dimensional.

A final ingredient need for the proof of Theorem 7.4 is the following result, which is a consequence of [CM09b, Theorem 1.3].

Proposition 7.7. *Let X be a locally compact geodesically complete CAT(0) space with cocompact isometry group and assume that X is not reduced to a single point.*

Then $\text{Is}(X)$ is doubly transitive on $\partial_\infty X$ if and only if X is isometric either to a rank one symmetric space or to a locally finite semi-regular tree.

Proof. The ‘if’ part is clear. We assume henceforth that $\text{Is}(X)$ is doubly transitive on $\partial_\infty X$. This hypothesis implies in particular that any two points of the boundary can be joined by a geodesic line. In particular any two points of the boundary are at distance π in the angular metric. Hence X is a visibility space, and it follows from Theorems II.9.33 and III.1.5 in [BH99] that X is Gromov hyperbolic.

Fix a point $\xi \in \partial_\infty X$ and pick any base point $p \in X$. The geodesic ray $[p, \xi)$ can be extended to some geodesic line ℓ . Notice that the set $P(\ell)$ consisting the union of all geodesic lines at bounded Hausdorff distance from ℓ is contained in bounded neighbourhood of ℓ , since otherwise $\partial_\infty X$ would contain some point at distance $\pi/2$ from ξ .

For any other point $x \in X$, there is a geodesic line ℓ' containing x and having ξ as one of its endpoints. Therefore there is an element of $\text{Is}(X)_\xi$ which maps ℓ' to a line parallel to ℓ . This shows that $X = \text{Is}(X)_\xi.P(\ell)$. Notice that $\text{Is}(X)$ contains some hyperbolic isometry; this follows *e.g.* from the fact that X is Gromov hyperbolic and unbounded. Since $\text{Is}(X)$ is doubly transitive on the boundary, it follows in particular that $\text{Is}(X)_\xi$ contains some element acting cocompactly on $P(\ell)$. Combining this with the fact that $X = \text{Is}(X)_\xi.P(\ell)$, we infer that $\text{Is}(X)_\xi$ acts cocompactly on X . The desired result then follows from [CM09b, Theorem 1.3]. \square

Proof of Theorem 7.4. Since $\text{Is}(X_i)$ acts cocompactly on X_i , it has full limit set. From Proposition 2.9 in [CM09a], we deduce that the projection Γ_i of Γ to $\text{Is}(X_i)$ also has full limit set. Lemma 7.5 thus ensures that Γ_i contains some rank one isometry. We are now ready to show the desired equivalences.

- (ii) \Rightarrow (i) Obvious.
- (iii) \Rightarrow (ii) In case each X_i is a tree, the vanishing of $\widetilde{\text{QH}}(\Gamma)$ is explicitly stated as Corollary 26 in [BM02], modulo the fact that $\widetilde{\text{QH}}(\Gamma)$ coincides with the kernel of the canonical map from the bounded cohomology of Γ with trivial coefficients to the usual cohomology in degree 2 (see *e.g.* [Mon01, Cor. 13.3.2]). In case some X_i is a rank one symmetric space, the same proof applies *verbatim*. Thus $\widetilde{\text{QH}}(\Gamma) = 0$ as desired.
- (i) \Rightarrow (iii) Assume that (iii) fails. In view of Proposition 7.7, this amounts to saying that for some i , the space X_i is not isometric to the real line and that the closure G_i of Γ_i does not act doubly transitively on $\partial_\infty X_i$. By the trichotomy of Proposition 7.6, it follows that either G_i stabilises some geodesic line in X_i or that $\widetilde{\text{QH}}(\Gamma)$ is infinite-dimensional. But the former possibility implies that X_i is isometric to the real line, since any group acting cocompactly on a geodesically complete CAT(0) space does not preserve any non-empty proper subspace (see [CM09b, Lemma 3.13]), and is thus precluded. Hence (i) fails as well, as was to be shown. \square

Remark 7.8. A similar result holds without the hypothesis that X_i is geodesically complete. In that case, the condition (iii) must be weakened and replaced by the fact that X_i is quasi-isometric to a rank one symmetric space or a semi-regular tree, and the Γ -action on X_i is **quasi-distance-transitive** in the sense that there is some $\delta > 0$ such that for all $x, y, x', y' \in X_i$, there is some $\gamma \in \Gamma$ such that $d_{X_i}(x, \gamma.x') < \delta$ and $d_{X_i}(y, \gamma.y') < \delta$. This can be established by following the same proof; the main change is that Proposition 7.7 must be replaced by the ‘quasi-isometric’ version we have just described.

Proof of Theorem H. As in the proof of C, we remark that $\text{Aut}(X_i)$ contains a lattice for each i . By Corollary B, it follows that $\text{Aut}(X_i)$ contains some rank one isometry. The desired result then follows from Theorem 7.4. \square

7.5. Asymptotic cones. The existence of cut-points in asymptotic cones of a space or a group has been studied notably in [DMS10]. A function $\text{Div} : \mathbf{N} \rightarrow \mathbf{R} \cup \{\infty\}$, called the **divergence**, is associated to any path-connected metric X . In case X is proper and has a cocompactly isometry group, it is shown in [DMS10, Prop. 1.1] that Div grows linearly with n if and only if no asymptotic cone of X has a cut-point.

Proof of Corollary I. We do not provide all the details; the reader is assumed to have some familiarity with [DMS10].

Since X is essential (see Proposition 3.7), it follows that any irreducible factor of X is non-compact. In particular, if X is not irreducible, then it follows that it has linear divergence and, hence, no asymptotic cone has a cut-point. On the other hand, if X is irreducible, then Theorem A ensures that $\text{Aut}(X)$ contains a rank one element. In particular X contains some rank one geodesic. By Proposition 3.24 from [DMS10], this implies that every asymptotic cone of X has a cut-point.

We now turn to the subgroup Γ . Theorem A again implies that Γ contains some rank one element. By Proposition 3.26 from [DMS10], this implies that Γ contains Morse elements (as defined in [DMS10]) relatively to its own word metric (even if Γ is distorted in X), and this in turn implies the existence of asymptotic cut-points by Proposition 3.24 from *loc. cit.* \square

REFERENCES

- [Bal85] Werner Ballmann, *Nonpositively curved manifolds of higher rank*, Ann. of Math. (2) **122** (1985), no. 3, 597–609.

- [Bal95] ———, *Lectures on spaces of nonpositive curvature*, DMV Seminar, vol. 25, Birkhäuser Verlag, Basel, 1995, With an appendix by Misha Brin.
- [BB95] Werner Ballmann and Michael Brin, *Orbihedra of nonpositive curvature*, Inst. Hautes Études Sci. Publ. Math. (1995), no. 82, 169–209 (1996).
- [BB00] ———, *Rank rigidity of Euclidean polyhedra*, Amer. J. Math. **122** (2000), no. 5, 873–885.
- [BB08] Werner Ballmann and Sergei Buyalo, *Periodic rank one geodesics in Hadamard spaces*, Geometric and probabilistic structures in dynamics, Contemp. Math., vol. 469, Amer. Math. Soc., Providence, RI, 2008, pp. 19–27.
- [BC10] Jason Behrstock and Ruth Charney, *Divergence and quasimorphisms of right-angled Artin groups*, preprint, 2010.
- [BF09] Mladen Bestvina and Koji Fujiwara, *A characterization of higher rank symmetric spaces via bounded cohomology*, Geom. Funct. Anal. **19** (2009), no. 1, 11–40.
- [BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften 319, Springer, Berlin, 1999.
- [BL05] Andreas Balseer and Alexander Lytchak, *Centers of convex subsets of buildings*, Ann. Global Anal. Geom. **28** (2005), no. 2, 201–209.
- [BM02] M. Burger and N. Monod, *Continuous bounded cohomology and applications to rigidity theory*, Geom. Funct. Anal. **12** (2002), no. 2, 219–280.
- [BS87] Keith Burns and Ralf Spatzier, *Manifolds of nonpositive curvature and their buildings*, Inst. Hautes Études Sci. Publ. Math. (1987), no. 65, 35–59.
- [CF10] Pierre-Emmanuel Caprace and Koji Fujiwara, *Rank one isometries of buildings and quasimorphisms of Kac–Moody groups*, Geom. Funct. Anal. **19** (2010), no. 5, 1296–1319.
- [CL] Pierre-Emmanuel Caprace and Jean Lécureux, *Combinatorial and group-theoretic compactifications of buildings*, Ann. Inst. Fourier, to appear.
- [CL10] Pierre-Emmanuel Caprace and Alexander Lytchak, *At infinity of finite-dimensional CAT(0) spaces*, Math. Ann. **346** (2010), no. 1, 1–21.
- [CM] Pierre-Emmanuel Caprace and Nicolas Monod, *Fixed points and amenability in non-positive curvature*, In preparation.
- [CM09a] ———, *Isometry groups of non-positively curved spaces: discrete subgroups*, J. Topol. **2** (2009), no. 4, 701–746.
- [CM09b] ———, *Isometry groups of non-positively curved spaces: Structure theory*, J. Topol. **2** (2009), no. 4, 661–700.
- [CN05] Indira Chatterji and Graham Niblo, *From wall spaces to CAT(0) cube complexes*, Internat. J. Algebra Comput. **15** (2005), no. 5-6, 875–885.
- [DMS10] Cornelia Drutu, Shahar Mozes, and Mark Sapir, *Divergence in lattices in semisimple Lie groups and graphs of groups*, Trans. Amer. Math. Soc. **362** (2010), 2451–2505.
- [Ger98] V. Gerasimov, *Fixed-point-free actions on cubings*, Siberian Adv. Math. **8** (1998), 36–58.
- [Gur05] Dan Guralnik, *Coarse decompositions of boundaries for CAT(0) groups*, PhD thesis, Technion, Haifa, 2005.
- [Ham09] Ursula Hamenstädt, *Rank one isometries of proper CAT(0)-spaces*, Discrete groups and geometric structures, Contemp. Math., vol. 501, Amer. Math. Soc., Providence, RI, 2009, pp. 43–59.
- [Lea] Ian Leary, *A metric kan–thurston theorem*, Preprint (2010).
- [Lin10] Gabriele Link, *Asymptotic geometry in the product of Hadamard spaces with rank one isometries*, Geom. Topol. **14** (2010), 1063–1094.
- [Mon01] Nicolas Monod, *Continuous bounded cohomology of locally compact groups*, Lecture Notes in Mathematics, vol. 1758, Springer-Verlag, Berlin, 2001.
- [Nic04] Bogdan Nica, *Cubulating spaces with walls*, Algebr. Geom. Topol. **4** (2004), 297–309 (electronic).
- [Rol98] Martin A. Roller, *Poc sets, median algebras and group actions. an extended study of dunwoody’s construction and sageev’s theorem*, Habilitationsschrift, Regensburg, 1998.
- [Sag95] Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) **71** (1995), no. 3, 585–617. MR MR1347406 (97a:20062)
- [Sel60] Atle Selberg, *On discontinuous groups in higher-dimensional symmetric spaces.*, Contrib. Function Theory, Int. Colloqu. Bombay, Jan. 1960, 147–164 (1960)., 1960.
- [SW05] Michah Sageev and Daniel T. Wise, *The Tits alternative for CAT(0) cubical complexes*, Bull. London Math. Soc. **37** (2005), no. 5, 706–710.

UCLouvain – IRMP, CHEMIN DU CYCLOTRON 2, 1348 LOUVAIN-LA-NEUVE, BELGIUM
E-mail address: `pe.caprace@uclouvain.be`

DEPT. OF MATH., TECHNION, HAIFA 32000, ISRAEL
E-mail address: `sageevm@techunix.technion.ac.il`