

DEFORMING CUBULATIONS OF HYPERBOLIC GROUPS

ELIA FIORAVANTI AND MARK HAGEN

ABSTRACT. We describe a procedure to deform cubulations of hyperbolic groups by “bending hyperplanes”. Our construction is inspired by related constructions like Thurston’s Mickey Mouse example [Thu79, Example 8.7.3], walls in fibred hyperbolic 3–manifolds [Duf12] and free-by- \mathbb{Z} groups [HW16], and Hsu–Wise *turns* [HW15b, Definition 4.1].

As an application, we show that every cocompactly cubulated Gromov–hyperbolic group admits a proper, cocompact, essential action on a CAT(0) cube complex with a single orbit of hyperplanes. This answers (in the negative) a question of Wise [Wis12b, Problem 5.2], who proved the result in the case of free groups.

We also study those cubulations of a general group G that are not susceptible to trivial deformations. We name these *bald cubulations* and observe that every cocompactly cubulated group admits at least one bald cubulation. We then apply the hyperplane-bending construction to prove that every cocompactly cubulated hyperbolic group G admits infinitely many bald cubulations, provided G is not a virtually free group with $\text{Out}(G)$ finite. By contrast, we show that the Burger–Mozes examples each admit a unique bald cubulation.

CONTENTS

1. Introduction.	1
2. Preliminaries.	10
3. Bending hyperplanes.	24
4. The number of bald cubulations.	37
References	41

1. INTRODUCTION.

The theory of group actions on CAT(0) cube complexes, and in particular its applications to 3–manifolds [Ago13, Wis11], has recently exerted a large influence in group theory and topology. “Cubulating” a group — constructing a proper action on a CAT(0) cube complex, usually via the method introduced by Sageev in [Sag95] — reveals a great deal about the group’s structure. This is particularly true when the group G is hyperbolic: in this case, when the *codimension–1 subgroups* used to cubulate the group are quasiconvex, the action is also cocompact [Sag97, NR03, HW14]; we say

Hagen was supported by EPSRC grant EP/R042187/1.

G is *cocompactly cubulated* if it acts properly and cocompactly on a CAT(0) cube complex. In this case, work of Agol [Ago13] and Haglund–Wise [HW08] shows that such a hyperbolic group G has many useful properties, e.g. \mathbb{Z} –linearity, separability of quasiconvex subgroups, etc.

Many of the procedures for cubulating hyperbolic groups G arising in nature make it clear that cubulations of G are non-canonical and proving their existence is often non-constructive. Proofs that a given G is cubulated often proceed as follows. First, one describes a general procedure for finding quasiconvex codimension–1 subgroups in G . Then, one shows that any two points in the Gromov boundary of G can be separated by the limit set of some coset of a codimension–1 subgroup of the given type: one constructs a particular subgroup “targeted” at the given pair of boundary points. Then one applies a theorem of Bergeron–Wise [BW12], relying on a compactness argument, to extract a finite collection of codimension–1 subgroups that suffice to ensure a proper action on a CAT(0) cube complex.

For example, when G is the fundamental group of a hyperbolic 3–manifold M , the codimension–1 subgroups can be taken to be fundamental groups of quasi-Fuchsian surfaces immersed in M . The work of Kahn and Markovic [KM12] shows that these are enough to separate any two points in the boundary of G .

While a lot of information about G can be gleaned from the mere fact of it being cubulated, one usually does not know much about the specific cube complex. It is therefore natural to want some sort of “space of all cocompact cubulations” of a given hyperbolic group G .

One way to proceed is by analogy to deformation spaces of actions on trees, introduced by Forester in [For02]. There, one considers all *minimal* actions of G on trees in which the set of elliptic subgroups is held fixed. In our setting, one might wish to consider all of the proper, cocompact actions of G on CAT(0) cube complexes. The right notion of a “minimal” action on a CAT(0) cube complex X should at least include the requirement that there is no G –invariant convex subcomplex, so one should restrict to actions that are *essential* in the sense of Caprace–Sageev [CS11]; this avoids distractions like attaching a leaf edge to each vertex, or taking the product of X with a finite cube complex. A result in [CS11] makes this a safe restriction, since any cocompact cubulation can be replaced with an essential one without really changing much.

However, one should also impose some additional restrictions that are best illustrated by considering the simple case where $G = \mathbb{Z}$. The most obvious cubulation, the action by translations on the tiling of \mathbb{R} by 1–cubes, seems intuitively better than the action on the cube complex obtained by, say, stringing together countably many squares, with each intersecting the next in a single vertex. Both of these cubulations are essential, and in both cases the hyperplane–stabilisers are trivial, but only in the first case is the action of $\{1\}$ on each hyperplane the “right” cubulation of the trivial group.

So, we ask that X is *hyperplane-essential*: every hyperplane-stabiliser acts essentially on its hyperplane. This, too, turns out to be reasonable, in the sense that, if G admits a proper, cocompact action on a CAT(0) cube complex, then it admits a proper, cocompact, hyperplane-essential action [HT19]. Passing to a hyperplane-essential action is somewhat more violent than making an action essential, since it seriously changes which subgroups are hyperplane-stabilisers.

Remark (Hyperplane-essentiality). Essentiality and hyperplane-essentiality are defined precisely in Section 2. When G is the fundamental group of a hyperbolic 3-manifold, the cubulations provided by using Kahn–Markovic surfaces are automatically essential and hyperplane-essential (see Remark 2.24).

There are also stronger conditions that one might want to impose on a cubulation $G \curvearrowright X$. For example, X has *codimension- k hyperplanes*, each of which is the intersection of k pairwise-transverse hyperplanes; one could ask that the stabiliser of each hyperplane of each codimension acts essentially on it. This condition is strictly stronger than hyperplane-essentiality (which imposes restrictions only on codimension-1 hyperplanes) and is equivalent to X having no free faces; equivalently, the CAT(0) metric has the *geodesic extension property* [BH99, Proposition II.5.10]. However, it is not at all clear whether one can pass from an arbitrary cocompact cubulation to a cubulation with no free faces, so we work with hyperplane-essential actions instead of the stronger version.

Restricting to proper, cocompact, essential, hyperplane-essential actions seems to be reasonable for the purpose of considering the “space of cubulations” of G because of the following theorem of Beyrer and the first author [BF19a]. For hyperbolic groups G acting on cube complexes X with the above properties, the action $G \curvearrowright X$ is completely determined, up to G -equivariant cubical isomorphism, by the *length function* $\ell_X: G \rightarrow \mathbb{N}$. This is the function $\ell_X(g) = \inf_{x \in X} d(x, gx)$, where d is the ℓ_1 metric on X . Throughout this paper, we will say that cubulations $G \curvearrowright X, G \curvearrowright Y$ are *equivalent* if there exists a G -equivariant cubical isomorphism $X \rightarrow Y$.

This suggests a natural topology for the space of such cubulations [BF19b]. First, we allow ourselves to continuously vary the ℓ_1 metric on X by replacing the cubes by *cubeoids* — the side-lengths need not be 1, and we can continuously vary length functions by rescaling edges (always assigning the same length to parallel edges). From this point of view, passing from X to a cubical subdivision — equivariantly replacing each hyperplane with several parallel copies but keeping the metric fixed — has no effect on the length function, which now takes values in $\mathbb{R}_{\geq 0}$. Regarding each geometric, essential, hyperplane-essential action on a CAT(0) cuboid complex as a length function gives a map from the set of such cubulations of G to $\mathbb{R}^G - \{0\}$, which we equip with the product topology. One can also projectivise, regarding as equivalent any two cubulations inducing homothetic length functions.




FIGURE 1. Bending hyperplanes into a crooked hyperplane.

This suggests that we should not consider cubulations $G \curvearrowright X, G \curvearrowright Y$ essentially different if they admit a common subdivision. This motivates us to consider only cubulations in which no two halfspaces are at finite Hausdorff distance (see the notion of “bald” cubulation in Definition 1.1 below).

In this paper, we concern ourselves with deformations of a cubulation $G \curvearrowright X$, i.e. with moving around in the space of cubulations. We leave discussion of the subject from the point of view of the above topology for later work. Instead, we are concerned with a much more basic question:

Question. *For which (hyperbolic) cocompactly cubulated groups G is the space of essential, hyperplane-essential cubulations infinite, even up to subdivision?*

One way to move in the space of cubulations is using the action of $\text{Out}(G)$; this varies the G -action, but not the underlying cube complex. It is not hard to show that, if $\text{Out}(G)$ is infinite, then the existence of a cocompact cubulation of G ensures that there are infinitely many with the above properties. However, a common situation is where G is a one-ended hyperbolic group that does not admit any two-ended splittings, so $\text{Out}(G)$ is finite.

In general, one needs to vary the cube complex as well as the action. In this paper, we *bend hyperplanes* to transform one cubulation into another and answer the above question.

Bending hyperplanes. Let G be a one-ended hyperbolic group acting properly and cocompactly on an essential, hyperplane-essential $\text{CAT}(0)$ cube complex X . We now describe the *bending* procedure for deforming $G \curvearrowright X$ into a new cubulation.

The idea is to produce, given the hyperplanes of X , a new *crooked hyperplane* \mathcal{C} built from pieces of old hyperplanes. Each piece is obtained from some hyperplane $u \subseteq X$ by cutting u along a family $\{w_i\}$ of pairwise-disjoint hyperplanes that intersect it. Each w_i is itself cut along a family of hyperplanes containing u , and also contributes a piece to the crooked hyperplane. Finally, the various pieces are glued along their boundaries, which are codimension-2 hyperplanes in X . This is depicted in Figure 1.

Since G is hyperbolic, results of Agol [Ago13] and Haglund–Wise [HW08] imply that hyperplane-stabilisers are separable. This allows one to choose the pieces so that their bounding codimension-2 hyperplanes are all far apart, which in turn allows one to produce a crooked hyperplane that is 2-sided and quasiconvex in X . If the pieces are constructed equivariantly with respect to

a finite-index subgroup of G , then the crooked hyperplane is also acted upon cocompactly by its stabiliser. Hence each crooked hyperplane corresponds to a quasiconvex codimension-1 subgroup of G , along with a specified wall; applying Sageev's construction [Sag95, Sag97] yields a cocompact, essential G -action on a new cube complex Y with a single orbit of hyperplanes.

With a bit more care, we can do the bending in such a way that every infinite-order $g \in G$ has its axis cut by some translate of the crooked hyperplane, ensuring, by [BW12], that the G -action on Y is proper.

As an application of the bending procedure, we can therefore answer a question asked by Wise in [Wis12b, Problem 5.2]:

Theorem A. *Let G be a Gromov-hyperbolic group that admits a proper, cocompact action on a CAT(0) cube complex. Then there exists a CAT(0) cube complex X on which G acts properly, cocompactly, and essentially with a single orbit of hyperplanes.*

The above theorem was proved by Wise in the case where G is a free group. He used an ingenious *antenna* construction to produce a codimension-1 subgroup H of G , and an associated H -wall, so that the action on the resulting dual cube complex has the claimed properties. In fact, Wise goes considerably further in the free group case: his construction shows that one can choose H to be an *arbitrary* finitely generated infinite-index subgroup.

As described above, our proof proceeds along completely different lines in the one-ended case, relying on bending hyperplanes. When G is a surface group, the resulting cubulation essentially originates from a single (necessarily non-simple) filling closed curve on the surface.

In the general case, we split G as a finite graph of groups with finite edge groups and use a hybrid technique: we apply a version of Wise's antenna construction to the Bass-Serre tree, apply the bending construction to the various one-ended vertex groups, and glue up the pieces to get the required wall.

Remark. In the one-ended case, our proof of Theorem A actually shows more. Let $K \leq G$ be a hyperplane-stabiliser in a proper, cocompact action of G on an essential, hyperplane-essential CAT(0) cube complex. Then, for every open neighbourhood $U \subseteq \partial_\infty G$ of the limit set of K , the cubulation in Theorem A can be chosen to have a hyperplane-stabiliser H with limit set contained in U . The two sides of the H -wall can similarly be picked arbitrarily close to the two sides of the K -wall.

In other words, an arbitrarily small deformation of one of the original walls always suffices to obtain a proper action with a single orbit of hyperplanes.

Bald cubulations. Let G be a (not necessarily hyperbolic) group acting properly and cocompactly on a CAT(0) cube complex X . It is not difficult to produce infinitely many cocompact cubulations of G , no two of which are equivalent. This is because of various relatively uninteresting procedures:

- we can cubically subdivide X indefinitely (or just subdivide a G -orbit of hyperplanes);
- given a vertex $v \in X$ and a finite CAT(0) cube complex F , we can G -equivariantly attach a copy of F to each vertex in $G \cdot v$;
- for every $n \geq 2$, we can *breed* a G -orbit of hyperplanes of X with an n -cube. This procedure is described in the case $n = 2$ in [BF19a, Example 5.5], but the extension to a general n is straightforward.

Such procedures for creating new cubulations from old ones are not very interesting because they always result in actions $G \curvearrowright X$ with some of the following properties:

- $G \curvearrowright X$ *inessentially*: there is some hyperplane \mathfrak{w} of X and some component \mathfrak{h} of $X - \mathfrak{w}$ such that each G -orbit intersects \mathfrak{h} in a set at bounded distance from \mathfrak{w} ;
- G acts *hyperplane-inessentially*: there is some hyperplane \mathfrak{w} such that the action of $\text{Stab}_G(\mathfrak{w})$ on \mathfrak{w} is inessential;
- X contains two hyperplanes $\mathfrak{w}_1, \mathfrak{w}_2$ that have associated halfspaces $\mathfrak{h}_1, \mathfrak{h}_2$ lying at finite Hausdorff distance from each other.

If G is hyperbolic and acts properly and cocompactly on several essential CAT(0) cube complexes X_1, \dots, X_k , there is an additional “cheap” procedure to create new cubulations. We can cubulate G using all of the codimension-1 subgroups arising as hyperplane-stabilisers in the various $G \curvearrowright X_i$. Again, the resulting action $G \curvearrowright X$ can fail to be hyperplane-essential, even if all original actions $G \curvearrowright X_i$ had this property.

As discussed above, we wish to restrict to essential actions, in light of [CS11, Proposition 3.5]. A procedure called *panel collapse* reduces an arbitrary hyperplane-inessential cocompact cubulation to a hyperplane-essential one [HT19], so it makes sense to consider only hyperplane-essential actions.

This motivates the following definition:

Definition 1.1 (Bald). A *bald cubulation* of a group G is a proper, cocompact, essential, hyperplane-essential action by cubical automorphisms on a CAT(0) cube complex X with the following additional property. Suppose that $\mathfrak{w}_1, \mathfrak{w}_2$ are hyperplanes of X that bound halfspaces at finite Hausdorff distance. Then there is a cubical splitting $X = \mathbb{R} \times Y$ such that \mathfrak{w}_1 and \mathfrak{w}_2 are hyperplanes of the \mathbb{R} -factor.

As an example, it follows from the Cubical Flat Torus theorem of Woodhouse–Wise [WW17] that, for every bald cubulation of a free abelian group, the underlying CAT(0) cube complex is isomorphic to the standard tiling of \mathbb{R}^n with vertex set the integer lattice (Proposition 4.5).

To make an arbitrary cocompact cubulation bald, we *shave* it. This is Proposition 2.29, the core of which lies in [HT19].

Proposition B. *Every cocompactly cubulated group admits a bald cubulation.*

Because of the uninteresting procedures described above, every group G has infinitely many different cocompact cubulations, as soon as it has one. It is a much less trivial question whether G admits infinitely many different *bald* cubulations. The reason is that shaving — in particular, panel collapse — can radically shrink the hyperplane-stabilisers. We are not aware of any general technique to ensure that different essential cubulations will not get shaved into the same bald cubulation.

In fact, there can be no general procedure to produce distinct bald cubulations of a cocompactly cubulated group. In Proposition 4.8 we show:

Proposition C. *Let T_1, T_2 be locally finite trees with all vertices of degree at least 3. Let $U_1, U_2 \leq \text{Aut}(T_1), \text{Aut}(T_2)$ be closed, locally primitive subgroups generated by edge-stabilisers and satisfying Tits' independence property. Let $\Gamma \leq U_1 \times U_2$ be a uniform lattice with dense projections to U_1 and U_2 . Then the standard action of Γ on $T_1 \times T_2$ is the only bald cubulation of Γ .*

Proposition C implies that the irreducible lattices in products of trees constructed by Burger–Mozes in [BM00a, BM00b] have a unique bald cubulation (see in particular [BM00b, Theorem 6.3]).

An interesting question is to what extent the hypotheses of Proposition C can be relaxed. Specifically, a BMW group (for Burger–Mozes–Wise) is a group Γ admitting a free, vertex-transitive action $\Gamma \curvearrowright \text{Aut}(T_1) \times \text{Aut}(T_2)$, where T_1, T_2 are finite valence regular trees [Cap19]. Many examples of BMW groups have been studied, beginning with the aforementioned work of Burger–Mozes and contemporaneous work of Wise [Wis96, Wis07]. Which irreducible BMW groups have a unique bald cubulation?

The proof of Proposition C proceeds by studying the de Rham decomposition of X , where $\Gamma \curvearrowright X$ is a bald cubulation provided by Proposition B. Combining results from [Tit70, Sha00] enables us to apply the superrigidity theorem of Chatterji–Fernós–Iozzi [CFI16], which, in conjunction with a result of Caprace–de Medts [CDM11], implies that each factor of the de Rham decomposition is a CAT(0) cube complex with compact hyperplanes. Baldness — in particular, hyperplane-essentiality — then implies that each factor is a tree. A result in [BMZ09] finally shows that X is equivariantly isomorphic to the product of two trees we started with.

Proposition C stands in sharp contrast to the situation for (most) hyperbolic groups:

Theorem D (Infinitely many bald cubulations). *Let G be a nonelementary Gromov-hyperbolic group acting properly and cocompactly on a CAT(0) cube complex. Suppose that **at least one** of the following holds:*

- G is not virtually free;
- $\text{Out}(G)$ is infinite.

Then G admits infinitely many pairwise-inequivalent bald cubulations.

The case where $\text{Out}(G)$ is infinite is straightforward and is dealt with in Lemma 4.3. The main content of Theorem D is the case where G splits as a

(possibly trivial) finite graph of groups with finite edge groups and at least one vertex group one-ended.

Groups to which the theorem applies include fundamental groups of hyperbolic surfaces, fundamental groups of hyperbolic 3-manifolds [KM12, BW12], non-virtually free groups with finite $C'(\frac{1}{6})$ presentations [Wis04] (and hence random groups at sufficiently low density in Gromov’s model [OW11]), non-virtually free hyperbolic Coxeter groups [NR03], hyperbolic free-by-cyclic groups [HW15a], non-virtually free one-relator groups with torsion [Wis11], Bourdon groups [Bou97], and others.

The theorem does not apply in the case where G is virtually free and $\text{Out}(G)$ is finite; such groups were characterised by Pettet [Pet97] and include certain right-angled Coxeter groups (see [Hea21, Proposition 5.5] and, more generally, [LS07, Theorem 1.1] and [GPR12, Theorem 1.4]). We discuss the virtually free case more below.

Remark (Strategy of the proof of Theorem D). Lemmas 4.3 and Lemma 3.19 reduce the claim to the case where G is one-ended.

The strategy in the one-ended case is as follows. First, we assume for a contradiction that G admits only finitely many bald cubulations. From this, in Lemma 4.1, we deduce that some such cubulation $G \curvearrowright X$ has a hyperplane \mathfrak{w} whose limit set in $\partial_\infty G$ is “minimal”, in the sense that it does not properly contain the limit set of any hyperplane of any bald cubulation.

We choose an infinite-order element $g \in G$ whose axis is cut by \mathfrak{w} . We then apply hyperplane-bending along \mathfrak{w} to produce crooked hyperplanes \mathfrak{u}_n with three key properties:

- (i) the fixed points of g in $\partial_\infty G$ lie in different components of the complement in $\partial_\infty G$ of the limit set of \mathfrak{u}_n ;
- (ii) the limit set of \mathfrak{w} is not contained in that of \mathfrak{u}_n ;
- (iii) every neighbourhood in $\partial_\infty G$ of the limit set of \mathfrak{w} contains the limit set of \mathfrak{u}_n for all sufficiently large n .

For each n , we add to the hyperplanes of X the G -orbit of \mathfrak{u}_n and cubulate, to get a new cubulation $G \curvearrowright X_n$. Shaving these cubulations, we obtain bald cubulations $G \curvearrowright Y_n$ with the property that there exist hyperplanes $\mathfrak{v}_n \subseteq Y_n$ cutting the axis of g , whose limit set is contained in that of \mathfrak{u}_n . This is our only form of control on how shaving affects hyperplane-stabilisers.

By property (iii), the limit sets of the \mathfrak{v}_n must Hausdorff-converge in $\partial_\infty G$ to a subset of the limit set of \mathfrak{w} . From the assumption that there are only finitely many bald cubulations, we see that there are only $\langle g \rangle$ -finitely many limit sets \mathfrak{v}_n . From this it is straightforward to conclude that one of the \mathfrak{v}_n has its limit set contained in that of \mathfrak{w} and, from property (ii), this is a proper inclusion. This contradicts the “minimality” of \mathfrak{w} .

In this argument, it is crucial that \mathfrak{v}_n have nonempty limit set, which is where one-endedness of G comes in: no cubulation of a one-ended group can have a bounded hyperplane.

Further questions. Our results and techniques raise various questions.

The application of hyperplane-bending used to prove Theorem A in the one-ended case, and its combination with the ideas from [Wis12b] in the general case, does not in general yield a hyperplane-essential cubulation. This raises the following question:

Question 1. Does there exist a hyperbolic group G such that no hyperplane-essential cubulation of G has a single orbit of hyperplanes? Does there exist such a G that is one-ended?

When G is a free group, the “exotic” cubulations from [Wis12b] are not hyperplane-essential. They are thus susceptible to the panel collapse procedure from [HT19] (summarised in Proposition 2.29 below), which shrinks the hyperplane-stabilisers. It is unknown whether every cubulation of a free group with a single orbit of hyperplanes panel collapses to a tree.

In some cases, the proof of Theorem D relies on twisting a fixed cubulation by the action of $\text{Out}(G)$, and in other cases, it does not. This motivates the following question:

Question 2. Let G be a cocompactly cubulated one-ended hyperbolic group. Are there infinitely many bald cubulations up to the $\text{Out}(G)$ -action?

When G has no two-ended splitting, $\text{Out}(G)$ is finite [BF95, Corollary 1.3], and one gets a positive answer from Theorem D. More generally, by Levitt’s characterisation of hyperbolic groups with infinite outer automorphism group [Lev05, Theorem 1.4], $\text{Out}(G)$ is finite provided G does not split over a two-ended subgroup with infinite centre. So, for example, there are examples of hyperbolic right-angled Coxeter groups G that split over D_∞ but have $\text{Out}(G)$ finite, and so Theorem D applies to give a positive answer to the question in those cases.

One can ask more refined versions of the question by measuring the complexity of bald cubulations $G \curvearrowright X$ in some way, and then asking if there are infinitely many bald cubulations, up to the action of $\text{Out}(G)$, with at most a given complexity. Examples of complexity include the dimension of X , the number of G -orbits of hyperplanes in X , etc.

When $\text{Out}(G)$ is infinite, one can often make fixed elements $g \in G$ have arbitrarily large translation length in bald cubulations of G . This motivates:

Question 3. Given a cocompactly cubulated hyperbolic group that is not virtually free, can each infinite-order element $g \in G$ become arbitrarily long in the bald cubulations of G ?

Finally, the groups we have shown to admit unique bald cubulations are not virtually special (irreducible BMW groups). This motivates:

Question 4. Which (non-hyperbolic) virtually special groups admit infinitely many bald cubulations?

Among hyperbolic cocompactly cubulated groups, in view of Theorem D, the only remaining question is about virtually free groups with finite outer automorphism groups.

We observe that if G is a virtually free group and $\text{Out}(G)$ is finite, then the existence of infinitely many bald cubulations of G will require the construction of bald cubulations $G \curvearrowright X$ where X has some infinite hyperplane-stabilisers. Indeed, if all hyperplane-stabilisers are finite, then baldness implies that X is a tree. All of the proper, cocompact actions of G on trees belong to the same deformation space \mathcal{D} in the sense of [GL07], and \mathcal{D} is $\text{Out}(G)$ -finite (up to projectivising), by [GL07, Proposition 8.6]. So it appears some other idea is needed, possibly along the lines of the proof of Theorem D.

Outline of the paper. In Section 2, we first discuss background on $\text{CAT}(0)$ cube complexes and cubulating groups. We then prove various technical lemmas which will be used later. We also discuss the notion of an *abstract hyperplane*. The procedure for “shaving” a $\text{CAT}(0)$ cube complex into a bald one is also discussed in this section, proving Proposition B. In Section 3, we describe hyperplane-bending, and also generalise Wise’s *antenna* construction, to prove Theorem A. In Section 4.1, we prove Theorem D, and in Section 4.2 we prove Proposition C.

Acknowledgements. We are grateful to Pierre-Emmanuel Caprace for directing us to Lemma 1.4.7 in [BMZ09], to Anthony Genevois for pointing out [Lev05], and to Daniel Groves, Jason Manning and Henry Wilton for discussions. We also thank the anonymous referee for their many helpful comments.

Fioravanti thanks Max Planck Institute for Mathematics in Bonn for its hospitality and financial support.

2. PRELIMINARIES.

For basic notions related to $\text{CAT}(0)$ cube complexes, we direct the reader to e.g. [Che00, Hag07, Hag08, Sag14, Sag95, Wis12a]. We recall some of these presently.

Throughout this section, X denotes a $\text{CAT}(0)$ cube complex.

2.1. $\text{CAT}(0)$ cube complexes.

2.1.1. *Hyperplanes, halfspaces, separation, transversality.* We denote by $\mathcal{W}(X)$ the set of hyperplanes of X and by $\mathcal{H}(X)$ the set of halfspaces. For each $\mathfrak{w} \in \mathcal{W}(X)$, the two components of $X - \mathfrak{w}$ are the halfspaces $\mathfrak{h}, \mathfrak{h}^*$ associated to \mathfrak{w} . Each $\mathfrak{h} \in \mathcal{H}(X)$ is associated to (bounded by) a unique hyperplane \mathfrak{w} , and \mathfrak{h}^* always denotes the other halfspace associated to \mathfrak{w} .

Given $\mathfrak{w} \in \mathcal{W}(X)$ and $A, B \subseteq X$, we say that \mathfrak{w} *separates* A and B if there is a halfspace \mathfrak{h} associated to \mathfrak{w} such that $A \subseteq \mathfrak{h}$ and $B \subseteq \mathfrak{h}^*$. Let $\mathcal{W}(A|B)$ denote the set of hyperplanes \mathfrak{w} separating A from B . For ease of notation, we will write $\mathcal{W}(x|y, z)$, rather than $\mathcal{W}(x|\{y, z\})$.

Hyperplanes $\mathfrak{u}, \mathfrak{w}$ are *transverse* if they are distinct and satisfy $\mathfrak{u} \cap \mathfrak{w} \neq \emptyset$. Equivalently, letting $\mathfrak{a}, \mathfrak{b}$ be halfspaces associated to $\mathfrak{u}, \mathfrak{w}$ respectively, each

of the four intersections $\mathfrak{a} \cap \mathfrak{b}$, $\mathfrak{a}^* \cap \mathfrak{b}$, $\mathfrak{a}^* \cap \mathfrak{b}^*$, $\mathfrak{a} \cap \mathfrak{b}^*$ is nonempty. We also say that the halfspaces \mathfrak{a} and \mathfrak{b} are *transverse*.

Definition 2.1 (Facing triple, chain). The pairwise disjoint hyperplanes $\mathfrak{u}, \mathfrak{v}, \mathfrak{w}$ of X form a *facing triple* in X if no two of $\mathfrak{u}, \mathfrak{v}, \mathfrak{w}$ are separated by the third; equivalently, there exist disjoint halfspaces $\mathfrak{a}, \mathfrak{b}, \mathfrak{c}$ respectively associated to $\mathfrak{u}, \mathfrak{v}, \mathfrak{w}$.

The distinct hyperplanes $\mathfrak{w}_1, \mathfrak{w}_2, \dots$ form a *chain* if, for each i , there exists a halfspace \mathfrak{h}_i associated to \mathfrak{w}_i such that (up to relabelling), we have $\mathfrak{h}_i \subsetneq \mathfrak{h}_{i+1}$ for all i .

For each $\mathfrak{w} \in \mathscr{W}(X)$, recall that \mathfrak{w} is a CAT(0) cube complex whose cubes are *midcubes* of cubes c of X with $c \cap \mathfrak{w} \neq \emptyset$. Accordingly, we will sometimes abuse language and refer to a “vertex of \mathfrak{w} ” — by this we mean a 0-cube of \mathfrak{w} when the latter is regarded as a cube complex; equivalently, vertices of \mathfrak{w} are midpoints of edges of X dual to \mathfrak{w} . The hyperplanes of \mathfrak{w} are exactly the subspaces $\mathfrak{w} \cap \mathfrak{u}$, as \mathfrak{u} varies over the hyperplanes of X transverse to \mathfrak{w} .

2.1.2. *The ℓ_1 metric.* In this paper, we always work with the ℓ_1 metric on X , which we denote d . We will only ever be interested in distances between vertices of X , or between vertices of the cubical subdivision of X . Accordingly, we just need the following facts about d :

- If $x, y \in X^{(0)}$, then $d(x, y) = \#\mathscr{W}(x|y)$.
- The metric d restricts on $X^{(0)}$ to the metric induced by the usual graph metric on $X^{(1)}$.
- In particular, combinatorial geodesics in $X^{(1)}$ are exactly combinatorial paths containing at most one edge intersecting each hyperplane.

2.1.3. *The median.* Recall from e.g. [Che00] that a graph Γ is *median* if there exists a ternary operator $\mu: (\Gamma^{(0)})^3 \rightarrow \Gamma^{(0)}$ such that (letting d denote the usual graph metric), we have $d(x_i, x_j) = d(x_i, \mu(x_1, x_2, x_3)) + d(x_j, \mu(x_1, x_2, x_3))$ for $i \neq j$ and all vertices x_1, x_2, x_3 . A *discrete median algebra* is the vertex-set of a median graph, equipped with the median operator. (This is not the standard definition, but it is equivalent by [Rol98, Proposition 2.17].)

By [Che00, Theorem 6.1], $X^{(1)}$, with the graph-metric d , is a median graph, and conversely each median graph is the 1-skeleton of a uniquely determined CAT(0) cube complex. Letting μ denote the median on $X^{(0)}$, we have for all $x, y, z \in X^{(0)}$ that $\mathscr{W}(x|\mu(x, y, z)) = \mathscr{W}(x|y, z)$.

Fixing $p \in X^{(0)}$, the Gromov product at p therefore satisfies $(x \cdot y)_p = \#\mathscr{W}(p|x, y)$. Indeed:

$$(x \cdot y)_p = d(p, \mu(p, x, y)) = \#\mathscr{W}(p|\mu(p, x, y)) = \#\mathscr{W}(p|x, y).$$

2.1.4. *Convexity, gate-projection, and bridges.* A subset $S \subseteq X^{(0)}$ is *convex* if $\mu(x, y, z) \in S$ for all $x, y \in S$, $z \in X$. A subcomplex Y of X is *convex* if $Y^{(0)}$ is convex and Y is *full*, in the sense that Y contains every cube c of X

for which $c^{(0)} \subseteq Y^{(0)}$. Equivalently, Y is the largest subcomplex contained in the intersection of all halfspaces containing Y . (For subcomplexes, this notion of convexity agrees with CAT(0) metric convexity [Hag07].)

If $Y \subseteq X$ is a convex subcomplex, then any combinatorial geodesic with endpoints on Y lies in Y . Moreover, Y is itself a CAT(0) cube complex. We identify $\mathscr{W}(Y)$ with the subset of $\mathscr{W}(X)$ of hyperplanes that intersect Y .

Given $A \subseteq X$, its *cubical convex hull* is the intersection of all convex subcomplexes containing A . It is common to use the term *interval* to refer to the set $I(x, y)$ of vertices z such that $\mu(x, y, z) = z$, where $x, y \in X^{(0)}$. The interval $I(x, y)$ is just the convex hull of $\{x, y\}$.

If $Y \subseteq X$ is a convex subcomplex, there is a *gate-projection* $\pi_Y: X^{(0)} \rightarrow Y^{(0)}$ characterised by the property that any hyperplane \mathfrak{w} separates $x \in X^{(0)}$ from $\pi_Y(x)$ if and only if \mathfrak{w} separates x from Y . If $x, y \in X^{(0)}$, then $\mathscr{W}(\pi_Y(x)|\pi_Y(y)) = \mathscr{W}(x|y) \cap \mathscr{W}(Y)$, so π_Y is 1-Lipschitz.

The vertex $\pi_Y(x)$ is the unique closest point of Y to x . In fact, one can extend π_Y to a cubical map $\pi_Y: X \rightarrow Y$; see [BHS17, Section 2.1].

Let Y, Z be convex subcomplexes of X . Then $\pi_Y(Z)$ is a convex subcomplex of X , and the hyperplanes intersecting $\pi_Y(Z)$ are exactly the hyperplanes intersecting both Y and Z . In particular, if there is no such hyperplane, $\pi_Y(Z)$ is a single vertex.

The convex hull $B(Y, Z)$ of $\pi_Y(Z) \cup \pi_Z(Y)$ is therefore a CAT(0) cube complex whose set of hyperplanes has the form $(\mathscr{W}(Y) \cap \mathscr{W}(Z)) \sqcup \mathscr{W}(Y|Z)$. By e.g. [CS11, Proposition 2.5], we get $B(Y, Z) \cong \pi_Y(Z) \times H \cong \pi_Z(Y) \times H$, where H is isomorphic to the interval $I(\pi_Z(y), \pi_Y(\pi_Z(y)))$ for any vertex $y \in Y$. In particular, $\pi_Y(Z)$ and $\pi_Z(Y)$ are isomorphic CAT(0) cube complexes; the maps π_Y, π_Z restrict to cubical isomorphisms on these sets. The subcomplex $B(Y, Z)$ is the disjoint union of the intervals $I(y, z)$ as (y, z) varies over the pairs in $Y \times Z$ with $d(y, z) = d(Y, Z)$. We refer to $B(Y, Z)$ as the *bridge* between Y and Z .

2.1.5. *Walls and median subalgebras.* This will only be used in Section 4.2.

A *subalgebra* of $X^{(0)}$ is a subset A such that $\mu(a, b, c) \in A$ whenever $a, b, c \in A$. Given a subalgebra A , a subset $B \subseteq A$ is *median-convex* in A if $\mu(a, b, c) \in B$ whenever $a, b \in B$ and $c \in A$. Note that this coincides with our usual notion of convexity when $A = X^{(0)}$.

A *wall* in A is a partition $A = \mathfrak{a} \sqcup \mathfrak{a}^*$, where $\mathfrak{a}, \mathfrak{a}^*$ are nonempty and median-convex in A . When $A = X^{(0)}$, such partitions always originate from hyperplanes of X , by intersecting $X^{(0)}$ with the two associated halfspaces. For a general subalgebra, we still refer to the sets $\mathfrak{a}, \mathfrak{a}^*$ as *halfspaces*. Let $\mathscr{W}(A)$ and $\mathscr{H}(A)$ be the set of walls and the set of halfspaces of A .

If $S \subseteq X^{(0)}$ is a convex subset, and A is a subalgebra of $X^{(0)}$, then $S \cap A$ is median-convex in A . It follows that if $\mathfrak{w} \in \mathscr{W}(X)$ is a hyperplane such that A intersects both associated halfspaces $\mathfrak{h}, \mathfrak{h}^*$, then the partition $(\mathfrak{h} \cap A) \sqcup (\mathfrak{h}^* \cap A)$ is a wall in A . By [Bow13, Lemma 6.5], all walls of A actually arise this way.

2.1.6. *Cubical subdivision.* Recall that X admits a *cubical subdivision* X' — see Definition 2.4 in [Hag07] — which is the CAT(0) cube complex constructed as follows.

Given a cube $c \cong [-\frac{1}{2}, \frac{1}{2}]^n$, let c' be the cube complex obtained by subdividing each factor $[-\frac{1}{2}, \frac{1}{2}]$ so that it is a graph isomorphic to $K_{1,2}$ (but with edges of length $\frac{1}{2}$), and taking the product cell structure.

The subdivision X' is formed from X by replacing each cube c by c' . Then X' is a CAT(0) cube complex.

The obvious identity maps $c \rightarrow c'$ induce a map $X \rightarrow X'$; the preimage of the vertex set of X' is the set of barycentres of cubes of X . Letting d' be the ℓ_1 metric on X' (regarded as an abstract CAT(0) cube complex whose cubes have side length 1), we have $d'(x, y) = 2d(x, y)$.

Each hyperplane of X is a convex subcomplex of X' , and $X' \rightarrow X$ induces a two-to-one map $\mathscr{W}(X') \rightarrow \mathscr{W}(X)$ in an obvious way.

Letting μ' denote the median on $(X')^{(0)}$, we have that $\mu'(x, y, z) = \mu(x, y, z)$ for $x, y, z \in X^{(0)} \subset (X')^{(0)}$. By working in X' , we can thus extend the notion of convexity to subspaces of X that become subcomplexes upon subdivision, and this is in particular true for hyperplanes and halfspaces. In particular, it makes sense to talk about the gate projection $\pi_{\mathfrak{h}}: X \rightarrow \mathfrak{h}$ where \mathfrak{h} is a hyperplane or halfspace, the bridge between two hyperplanes or two halfspaces, etc.

2.1.7. *Facts about group actions.* We denote by $\text{Aut}(X)$ the group of cubical automorphisms of X . The action of $\text{Aut}(X)$ is an isometric action on (X, d) and an action by median isomorphisms on (X, μ) . It induces natural actions on the sets $\mathscr{W}(X)$ and $\mathscr{H}(X)$.

We implicitly assume all group actions $G \curvearrowright X$ to be by cubical automorphisms. We say that a group G is *cubulated* if there exists a CAT(0) cube complex X and a *proper* action $\rho: G \rightarrow \text{Aut}(X)$. If, in addition, ρ can be chosen to be cocompact, then we say G is *cocompactly cubulated*.

We will discuss further properties of actions later; for the moment, we recall some facts from [Hag07]. Let $g \in \text{Aut}(X)$ and let $\mathfrak{w} \in \mathscr{W}(X)$. We say that g *has an inversion along* \mathfrak{w} if $g\mathfrak{w} = \mathfrak{w}$ and $g\mathfrak{h} = \mathfrak{h}^*$, where \mathfrak{h} is one of the halfspaces associated to \mathfrak{w} . We say that g *acts without inversions* if g does not have an inversion along any hyperplane, and g *acts stably without inversions* if g^n acts without inversions for all $n \in \mathbb{Z}$. We have:

- If $g \in \text{Aut}(X)$ acts stably without inversions and does not fix a vertex, then g is *combinatorially hyperbolic*, i.e. there is a combinatorial geodesic γ preserved by g , on which g acts as a nontrivial translation.
- $\text{Aut}(X)$ acts naturally on the cubical subdivision X' , and every $g \in \text{Aut}(X)$ acts stably without inversions on X' .

If X is finite-dimensional, then any $g \in \text{Aut}(X)$ has a power acting stably without inversions. Indeed, the hyperplanes along which the powers of g have inversions are pairwise transverse.

2.2. Pocsets. A *pocset* is a triple $(\mathcal{P}, \preceq, *)$, where the pair (\mathcal{P}, \preceq) is a poset and $*$ is an order-reversing involution. Two distinct elements $a, b \in \mathcal{P}$ are *incomparable* if $a \not\preceq b$, $b \not\preceq a$ and $a \neq b^*$. We say that a and b are *transverse* if a and a^* are incomparable with b and b^* . The *dimension* of \mathcal{P} is the maximal cardinality of a subset of pairwise-transverse elements.

An *ultrafilter* is a subset $\sigma \subseteq \mathcal{P}$ such that:

- (1) there do not exist $a, b \in \sigma$ with $a \preceq b^*$;
- (2) for every $a \in \mathcal{P}$, we have $\#(\sigma \cap \{a, a^*\}) = 1$.

Equivalently, σ is a maximal subset satisfying (1). We say that σ is a *DCC ultrafilter* if, in addition, σ does not contain any infinite descending chains. We denote by $\min \sigma \subseteq \sigma$ the subset of \preceq -minimal elements. Two ultrafilters are *almost equal* if their symmetric difference is finite.

For every CAT(0) cube complex X , the triple $(\mathcal{H}(X), \subseteq, *)$ is a pocset, and the notions of transversality and dimension coincide with the usual ones. For every $v \in X^{(0)}$, the set $\{\mathfrak{h} \in \mathcal{H}(X) \mid v \in \mathfrak{h}\}$ is a DCC ultrafilter, and, if X is finite-dimensional, all DCC ultrafilters arise this way (in particular, any two of them are almost equal).

Conversely, to each pair (\mathcal{P}, σ) , where \mathcal{P} is a pocset and $\sigma \subseteq \mathcal{P}$ is an ultrafilter, we can associate a CAT(0) cube complex $X = X(\mathcal{P}, \sigma)$. Vertices of X are exactly ultrafilters on \mathcal{P} that are almost equal to σ . Two vertices are joined by an edge exactly when the symmetric difference of the corresponding ultrafilters has only two elements (the minimal possible size). See [Sag95, Rol98, Gur07] for details of the construction.

Lemma 2.2. *Let \mathcal{P} be a finite-dimensional pocset.*

- (1) *For every maximal subset $\tau \subseteq \mathcal{P}$ of pairwise-transverse elements, there exists a unique DCC ultrafilter $\sigma \subseteq \mathcal{P}$ with $\tau \subseteq \min \sigma$.*
- (2) *For every DCC ultrafilter $\sigma \subseteq \mathcal{P}$, there exists a subset $\tau \subseteq \min \sigma$ such that τ is a maximal pairwise-transverse subset of \mathcal{P} .*
- (3) *Any two DCC ultrafilters on \mathcal{P} are almost equal.*

Let a group Δ act on \mathcal{P} preserving the pocset structure.

- (4) *If there are only finitely many Δ -orbits of maximal pairwise-transverse subsets of \mathcal{P} , then the induced action $\Delta \curvearrowright X(\mathcal{P}, \sigma)$ is cocompact.*

Proof. Parts (1)–(3) follow from Proposition 3.1 and Corollary 3.3 in [Gur07]. In particular, Δ leaves invariant the almost-equality class of any DCC ultrafilter σ , thus inducing an action $\Delta \curvearrowright X(\mathcal{P}, \sigma)$. Finally, part (4) is immediate from part (1) and (2). \square

2.3. Actions, essentiality, hyperplane-essentiality, and skewering. Given a group action $G \curvearrowright X$ and a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$, we denote by $G_{\mathfrak{w}} \leq G$ the stabiliser of \mathfrak{w} . The following is e.g. [Sag14, Exercise 1.6].

Lemma 2.3. *Let a group G act cocompactly on X . For every hyperplane $\mathfrak{w} \in \mathcal{W}(X)$, the action $G_{\mathfrak{w}} \curvearrowright \mathfrak{w}$ is cocompact.*

Proof. Let K be a compact subcomplex of X such that $G \cdot K = X$. Since K is compact, there are only finitely many translates $g_1\mathfrak{w}, \dots, g_k\mathfrak{w}$ of \mathfrak{w} that are dual to edges of K . Let $L = \bigcup_{j=1}^k g_j^{-1}K$. Then L is compact.

Let e be an edge dual to \mathfrak{w} . Choose $g \in G$ such that $e \subseteq gK$. Then $g^{-1}e$ is dual to $g_i\mathfrak{w}$ for some $i \leq k$. Let $h = g_i^{-1}g^{-1}$. Then e and he are dual to \mathfrak{w} , and he is also dual to $h\mathfrak{w}$. Hence $h \in G_{\mathfrak{w}}$. Now, he is an edge of $g_i^{-1}K$ and is thus an edge of L . So $G_{\mathfrak{w}} \cdot (\mathfrak{w} \cap L)$ contains every vertex of \mathfrak{w} , as required. \square

Definition 2.4 (Skewering). Let $g \in \text{Aut}(X)$ and let $\mathfrak{w} \in \mathscr{W}(X)$. We say that g *skewers* \mathfrak{w} if there is a halfspace \mathfrak{h} associated to \mathfrak{w} and an integer $n \neq 0$ such that $g^n\mathfrak{h} \subsetneq \mathfrak{h}$. In this case, we also say g *skewers* the halfspace \mathfrak{h} .

Definition 2.5 (Essential stuff). The CAT(0) cube complex X is *essential* if, for each hyperplane \mathfrak{w} , each of the associated halfspaces contains points in X arbitrarily far from \mathfrak{w} . If $G \curvearrowright X$, we say that the action is *essential* if, for some (hence any) $x_0 \in X^{(0)}$, and each hyperplane \mathfrak{w} , each of the associated halfspaces contains points in $G \cdot x_0$ arbitrarily far from \mathfrak{w} . In the latter case, we also say that X is *G-essential*.

The cube complex X is *hyperplane-essential* if each hyperplane \mathfrak{w} , regarded itself as a CAT(0) cube complex, is essential. The action of G is *hyperplane-essential* if each hyperplane \mathfrak{w} has the property that $G_{\mathfrak{w}}$ acts essentially on \mathfrak{w} .

Remark 2.6.

- (1) Suppose X is finite-dimensional. By [CS11, Proposition 3.2], the action $G \curvearrowright X$ is essential if and only if every hyperplane of X is skewered by an element of G . Similarly, $G \curvearrowright X$ is hyperplane-essential if and only if, whenever $\mathfrak{u}, \mathfrak{w} \in \mathscr{W}(X)$ are transverse, there exist $g \in G_{\mathfrak{w}}$ skewering \mathfrak{u} and $h \in G_{\mathfrak{u}}$ skewering \mathfrak{w} .
- (2) If G acts cocompactly, then X is essential if and only if it is G -essential. Similarly, by Lemma 2.3, X is a hyperplane-essential cube complex if and only if $G \curvearrowright X$ is a hyperplane-essential action.

The following is [BF19b, Proposition 2.11] and a proof is given in [Hag20, Proposition 1].

Proposition 2.7. *Let X be cocompact, locally finite, essential, hyperplane-essential and irreducible. For any two transverse halfspaces \mathfrak{h}_1 and \mathfrak{h}_2 , there exists a halfspace $\mathfrak{k} \subseteq \mathfrak{h}_1 \cap \mathfrak{h}_2$.*

Remark 2.8 (Boundaries). Given a CAT(0) cube complex X , we denote by $\partial_{\infty}X$ its visual boundary (with the cone topology). Since ℓ_1 -convex subcomplexes, hyperplanes, and halfspaces are convex in the CAT(0) metric, we have the following. If A is a convex subcomplex, hyperplane, or halfspace, the inclusion $A \hookrightarrow X$ extends to a continuous injection $\partial_{\infty}A \hookrightarrow \partial_{\infty}X$.

Throughout this paper, we will often be in the situation where X admits a proper, cocompact action by a hyperbolic group G . In this case, $\partial_{\infty}X$ is G -equivariantly homeomorphic to the Gromov boundary $\partial_{\infty}G$ of G .

We say that X is *reducible* if there exist nontrivial CAT(0) cube complexes A, B with $X \cong A \times B$. In this case, every hyperplane of A, B determines a hyperplane of X , thus giving rise to a partition $\mathcal{W}(X) = \mathcal{W}(A) \sqcup \mathcal{W}(B)$. Every hyperplane in the set $\mathcal{W}(A)$ is transverse to every hyperplane in $\mathcal{W}(B)$. If X is not reducible, we say that X is *irreducible*.

We will often require the following fact about actions on CAT(0) cube complexes, which is Proposition 2.6 in [CS11].

Proposition 2.9 (De Rham decomposition). *Let X be finite-dimensional. Then there is a canonical decomposition $X = \prod_{i=1}^m X_i$, for irreducible CAT(0) cube complexes X_1, \dots, X_m , which is preserved by $\text{Aut}(X)$ (possibly permuting the factors). Hence the canonical embedding $\text{Aut}(X_1) \times \dots \times \text{Aut}(X_m) \hookrightarrow \text{Aut}(X)$ has finite-index image.*

Later, when working with a geometric action $G \curvearrowright X$, it will often be useful to assume that G is one-ended, enabling use of the following lemma:

Lemma 2.10 (One-ended cube complexes). *Let X be one-ended and essential. Then there does not exist a partition $\mathcal{W}(X) = \mathcal{A} \sqcup \mathcal{B}$ such that \mathcal{A}, \mathcal{B} are nonempty and no element of \mathcal{A} is transverse to an element of \mathcal{B} .*

Proof. Suppose for the sake of contradiction that such a partition of $\mathcal{W}(X)$ exists. By Lemma 2 in [Nib02], there exists a vertex $v \in X$ such that $X - \{v\}$ is disconnected. Since X is essential, each connected component of $X - \{v\}$ is unbounded. This shows that X has at least two ends. \square

Definition 2.11 (Halfspace-stabiliser). Let G be a group acting on X . Let $\mathfrak{w} \in \mathcal{W}(X)$ and let $\mathfrak{h}, \mathfrak{h}^*$ be the associated halfspaces. The *halfspace-stabiliser* $G_{\mathfrak{w}}^0$ is the kernel of the natural action of $G_{\mathfrak{w}}$ on $\{\mathfrak{h}, \mathfrak{h}^*\}$ (which has index at most 2 in $G_{\mathfrak{w}}$).

Recall that a subgroup H of a group G is *separable* if for all $g \in G - H$, there is a finite-index subgroup $G' \leq G$ such that $H \leq G'$ and $g \notin G'$.

Lemma 2.12 (Large-girth covers). *Let a group G act properly, cocompactly and with separable halfspace-stabilisers on X . Then, for every $n \geq 1$, there exists a finite-index subgroup $H \triangleleft G$ such that:*

- $H \curvearrowright X$ has no hyperplane inversions;
- for every $\mathfrak{w} \in \mathcal{W}(X)$, any two distinct elements of $H \cdot \mathfrak{w}$ are disjoint and at distance $\geq n$.

Proof. Let $\mathfrak{w}_1, \dots, \mathfrak{w}_k \in \mathcal{W}(X)$ be such that $G \cdot \{\mathfrak{w}_1, \dots, \mathfrak{w}_k\} = \mathcal{W}(X)$.

Since halfspace-stabilisers are separable, there exist subgroups $H_i \leq G$ such that any two elements of $H_i \cdot \mathfrak{w}_i$ are at distance $\geq n$ and no element of H_i swaps the sides of \mathfrak{w}_i .

Up to passing to further finite-index subgroups, we can assume that $H_i \triangleleft G$. Set $H = H_1 \cap \dots \cap H_k$.

Given any $\mathfrak{w} \in \mathcal{W}(X)$, there exist $1 \leq i \leq k$ and $g \in G$ with $\mathfrak{w} = g\mathfrak{w}_i$. If $h \in H$ and $\mathfrak{w} \neq h\mathfrak{w}$, we have $\mathfrak{w}_i \neq g^{-1}h\mathfrak{w}_i$. Since H is normal in G , the

element $g^{-1}hg$ lies in H , hence $d(\mathfrak{w}, h\mathfrak{w}) = d(\mathfrak{w}_i, g^{-1}hg\mathfrak{w}_i) \geq n$. A similar argument shows that $H \curvearrowright X$ has no hyperplane inversions. \square

Lemma 2.13. *Let G be a Gromov-hyperbolic group, with a proper cocompact action $G \curvearrowright X$. Given essential hyperplanes $\mathfrak{w}_1, \mathfrak{w}_2$ and $n \geq 1$, there exists $\mathfrak{w}'_2 \in G \cdot \mathfrak{w}_2$ such that $d(\mathfrak{w}_1, \mathfrak{w}'_2) \geq n$.*

Proof. By Proposition 3.2 in [CS11], the orbits $G \cdot \mathfrak{w}_1$ and $G \cdot \mathfrak{w}_2$ contain infinite chains of hyperplanes; moreover, we can assume that $G \cdot \mathfrak{w}_1 \neq G \cdot \mathfrak{w}_2$. If every element of $G \cdot \mathfrak{w}_1$ were transverse to an element of $G \cdot \mathfrak{w}_2$, this would violate Theorem 3.3 of [Gen16]. It follows that some $\mathfrak{w}''_2 \in G \cdot \mathfrak{w}_2$ is disjoint from \mathfrak{w}_1 and we can achieve the required distance from \mathfrak{w}_1 by considering a hyperplane $\mathfrak{w}'_2 = g^N \mathfrak{w}''_2$, where g skewers \mathfrak{w}''_2 and N is large. \square

Given a geodesic $\gamma \subseteq X$, we denote by $\mathscr{W}(\gamma)$ the set of hyperplanes crossed by γ .

Lemma 2.14. *There exists a constant $D = D(\delta)$ such that the following holds. For every δ -hyperbolic CAT(0) cube complex X and every geodesic $\gamma \subseteq X$, there exists a hyperplane $\mathfrak{w} \in \mathscr{W}(\gamma)$ with $\text{diam } \pi_{\mathfrak{w}}(\gamma) \leq D$.*

Proof. There exists a constant $C = C(\delta)$ such that, given any two transverse chains of halfspaces, one of them must have cardinality $< C$ (see e.g. [Gen16, Theorem 3.3]).

Set $d = \dim X$, $\Delta = Cd(2d + 1)$, $D = 2\Delta$ and observe that d (hence D) is bounded above in terms of δ . Since gate-projections are 1-Lipschitz, we can assume that the length of γ is at least D .

Among any $2Cd + 1$ halfspaces entered consecutively by γ , there exists a chain $\mathfrak{h}_{-C} \supseteq \cdots \supseteq \mathfrak{h}_C$. Let $x_- \in \mathfrak{h}_0^* \cap \gamma$ and $x_+ \in \mathfrak{h}_0 \cap \gamma$ be adjacent vertices of X and let \mathfrak{w} be the hyperplane bounding \mathfrak{h}_0 . Note that $\mathscr{W}(x_- | \mathfrak{h}_C)$ and $\mathscr{W}(x_+ | \mathfrak{h}_{-C}^*)$ each contain at most $2Cd$ hyperplanes (not the optimal bound).

Suppose for the sake of contradiction that there exists $y \in \gamma \cap \mathfrak{h}_C$ with $d(\pi_{\mathfrak{w}}(x_-), \pi_{\mathfrak{w}}(y)) > \Delta$. Then there exists a chain $\mathfrak{k}_1 \supseteq \cdots \supseteq \mathfrak{k}_{C(2d+1)}$ of halfspaces with $\pi_{\mathfrak{w}}(x_-) \in \mathfrak{k}_1^*$ and $\pi_{\mathfrak{w}}(y) \in \mathfrak{k}_{C(2d+1)}$. For all $i, j \geq 1$, we have $y \in \mathfrak{h}_i \cap \mathfrak{k}_j$, $\pi_{\mathfrak{w}}(y) \in \mathfrak{h}_i^* \cap \mathfrak{k}_j$ and $x_- \in \mathfrak{h}_i^* \cap \mathfrak{k}_j^*$. Thus, either $\mathfrak{h}_i \subsetneq \mathfrak{k}_j$ or \mathfrak{h}_i and \mathfrak{k}_j are transverse. If $j > 2Cd$, the halfspaces \mathfrak{h}_C and \mathfrak{k}_j must be transverse, as $\#\mathscr{W}(x_- | \mathfrak{h}_C) \leq 2Cd$. It follows that \mathfrak{h}_i and \mathfrak{k}_j are transverse for all $1 \leq i \leq C$ and $2Cd + 1 \leq j \leq 2Cd + C$, violating our choice of C .

This proves that $d(\pi_{\mathfrak{w}}(x_-), \pi_{\mathfrak{w}}(y)) \leq \Delta$ for every $y \in \gamma \cap \mathfrak{h}_C$ and a similar argument shows that $d(\pi_{\mathfrak{w}}(x_+), \pi_{\mathfrak{w}}(z)) \leq \Delta$ for all $z \in \gamma \cap \mathfrak{h}_{-C}^*$. We conclude that the projection $\pi_{\mathfrak{w}}(\gamma \cap (\mathfrak{h}_C \sqcup \mathfrak{h}_{-C}^*))$ is contained in the Δ -neighbourhood of $\pi_{\mathfrak{w}}(x_-) = \pi_{\mathfrak{w}}(x_+)$. Since $\pi_{\mathfrak{w}}(\gamma)$ is an (unparametrised) geodesic, it all lies in the Δ -neighbourhood of $\pi_{\mathfrak{w}}(x_+)$, hence $\text{diam } \pi_{\mathfrak{w}}(\gamma) \leq 2\Delta = D$. \square

2.4. Hyperbolic groups and abstract hyperplanes. Let G be a Gromov-hyperbolic group. Every infinite-order element $g \in G$ has exactly two fixed points in the Gromov boundary $\partial_{\infty}G$. We denote by g^+ the stable fixed point and by g^- the unstable one. The following is classical:

Lemma 2.15. *The set $\{(g^+, g^-) \mid g \text{ infinite-order}\}$ is dense in $\partial_\infty G \times \partial_\infty G$.*

If $H \leq G$ is a quasiconvex subgroup, we denote by $\Lambda H \subseteq \partial_\infty G$ its limit set. The following is Lemma 2.6 in [GMRS98]:

Lemma 2.16. *We have $\Lambda(H \cap K) = \Lambda H \cap \Lambda K$ for any two quasiconvex subgroups $H, K \leq G$.*

If G acts properly and cocompactly on a geodesic metric space X , there exists a unique G -equivariant homeomorphism $\phi: \partial_\infty G \rightarrow \partial_\infty X$. Given a subset $\Omega \subseteq X$, we denote by $\partial_\infty \Omega \subseteq \partial_\infty X$ its limit set in the visual boundary of X and by $\Lambda \Omega = \phi^{-1}(\partial_\infty \Omega) \subseteq \partial_\infty G$ the pull-back to G .

Let us now fix a Cayley graph $\Gamma(G)$ of G with respect to a finite generating set. Given a subgroup $H \leq G$, we denote by $N_r(H)$ its closed r -neighbourhood in $\Gamma(G)$.

The following is Lemma 7.3 in [HW14], although it originally appeared implicitly in [Sag97, GMRS98]. See also Lemma 7 in [NR03] and Theorem 1.1 in [HW09] for additional details on its proof.

Lemma 2.17. *Given $D, \kappa \geq 0$ there exists a constant C such that the following holds for all κ -quasiconvex subgroups $H_1, \dots, H_k \leq G$. If*

$$N_D(g_1 H_{i_1}), \dots, N_D(g_n H_{i_n})$$

pairwise intersect, there exists $g \in G$ that is C -close to all $g_1 H_{i_1}, \dots, g_n H_{i_n}$.

We will need the following result in Section 3.3.

Lemma 2.18. *Let X be a locally connected, proper, geodesic, δ -hyperbolic space and, for some $R \geq 0$, let $U \subseteq X$ be a closed R -quasiconvex subset. For every subset $A \subseteq X - U$, let us set $\tilde{A} := A \sqcup U$. Then:*

- (1) *if A is a union of connected components of $X - U$, the set \tilde{A} is R -quasiconvex;*
- (2) *if A and B are unions of connected components of $X - U$ such that $A \cap B = \emptyset$, then $\partial_\infty \tilde{A} \cap \partial_\infty \tilde{B} = \partial_\infty U$;*
- (3) *$\partial_\infty X - \partial_\infty U$ is a disjoint union of the open subsets $\partial_\infty \tilde{C} - \partial_\infty U$, where C is a connected component of $X - U$.*

Proof. In order to prove part (1), consider a geodesic $\gamma \subseteq X$ joining two points $x, y \in \tilde{A}$. Let $x', y' \in \gamma$ be the points furthest from x and y , respectively, such that the sub-segments $xx', yy' \subseteq \gamma$ are entirely contained in the closure $\bar{A} \subseteq X$. If x or y do not lie in \bar{A} , we set $x' = x$ or $y' = y$. The points x' and y' always lie in U , so the sub-segment of γ joining them is contained in the R -neighbourhood of $U \subseteq \tilde{A}$. Since the rest of γ is contained in \bar{A} , the entire γ lies inside the R -neighbourhood of \tilde{A} , showing part (1).

Given A, B as in part (2) and a point $\xi \in \partial_\infty \tilde{A} \cap \partial_\infty \tilde{B}$, we consider quasi-geodesic rays $r_A \subseteq \tilde{A}$ and $r_B \subseteq \tilde{B}$ representing ξ . Let $x_n \in r_A$ and $y_n \in r_B$ be diverging sequences with $\sup d(x_n, y_n) < +\infty$. Observing that any geodesic joining x_n to y_n must intersect U , we deduce that r_A and r_B stay at bounded distance from U . Hence $\xi \in \partial_\infty U$, which proves part (2).

Now, let \mathcal{C} be the set of all connected components $C \subseteq X - U$. Given a point $\xi \in \partial_\infty X - \partial_\infty U$ and a ray $r \subseteq X$ representing it, a sub-ray $r' \subseteq r$ must be disjoint from U . It follows that r' is contained in some $C \in \mathcal{C}$, hence $\xi \in \partial_\infty \tilde{C} - \partial_\infty U$. This shows that $\partial_\infty X$ is the union of the sets $\partial_\infty \tilde{C} - \partial_\infty U$ with $C \in \mathcal{C}$; by part (2), this is a disjoint union. Finally, observe that, for each $C \in \mathcal{C}$, the boundary $\partial_\infty X$ is union of the two closed subsets $\partial_\infty \tilde{C}$ and $\partial_\infty(X - C)$. Again by part (2), this shows that $\partial_\infty \tilde{C} - \partial_\infty U$ has closed complement in $\partial_\infty X - \partial_\infty U$ and is therefore open. \square

The following is the key notion in this subsection. It allows us to cubulate hyperbolic groups with as few non-canonical choices as possible.

Definition 2.19. An *abstract hyperplane* for a hyperbolic group G is a pair (H, \mathfrak{H}) , where $H \leq G$ is quasiconvex and \mathfrak{H} is an H -invariant partition $\partial_\infty G - \Lambda H = \mathfrak{H}^+ \sqcup \mathfrak{H}^-$, where \mathfrak{H}^\pm are nonempty open subsets.

We say that \mathfrak{H}^\pm are the *sides* of \mathfrak{H} and that \mathfrak{H} is *subordinate* to H . Two points $\xi, \eta \in \partial_\infty G$ are *separated* by \mathfrak{H} if they lie on opposite sides.

Observe that, by H -invariance, the closures $\overline{\mathfrak{H}^+}, \overline{\mathfrak{H}^-} \subseteq \partial_\infty G$ are exactly $\mathfrak{H}^+ \sqcup \Lambda H$ and $\mathfrak{H}^- \sqcup \Lambda H$.

Lemma 2.20. *Given any abstract hyperplane (H, \mathfrak{H}) , there exist two H -invariant open subsets $H^\pm \subseteq \Gamma(G)$ and a constant $D > 0$ such that:*

- (1) $\Gamma(G) - N_D(H) = H^+ \sqcup H^-$;
- (2) $H^+ \sqcup N_D(H)$, $H^- \sqcup N_D(H)$ and $N_D(H)$ are connected;
- (3) $\Lambda H^+ = \overline{\mathfrak{H}^+}$ and $\Lambda H^- = \overline{\mathfrak{H}^-}$.

Proof. Given $L > 0$, we denote by $A_L^+ \subseteq \Gamma(G)$ the closed L -neighbourhood of the weak hull of $\overline{\mathfrak{H}^+} = \mathfrak{H}^+ \sqcup \Lambda H$. (Recall that the weak hull of $\overline{\mathfrak{H}^+}$ is the union of all $\Gamma(G)$ -geodesics joining distinct points in $\overline{\mathfrak{H}^+}$.)

The set A_L^- is defined similarly. Observe that, for every sufficiently large value of L , there exists $D > 0$ such that:

- $\Lambda A_L^\pm = \overline{\mathfrak{H}^\pm}$ and $\Gamma(G) = A_L^+ \cup A_L^-$;
- $A_L^+ \cap A_L^- \subseteq N_D(H)$;
- the sets $A_L^+ \cup N_D(H)$, $A_L^- \cup N_D(H)$ and $N_D(H)$ are connected.

Thus, the sets $H^+ = A_L^+ - N_D(H)$ and $H^- = A_L^- - N_D(H)$ are open, H -invariant, and satisfy (1) and (2). It is clear from the construction that $\mathfrak{H}^+ \subseteq \Lambda H^+ \subseteq \Lambda A_L^+ = \mathfrak{H}^+ \sqcup \Lambda H$. Since H^+ is nonempty and H -invariant, we also have $\Lambda H \subseteq \Lambda H^+$, which shows (3). \square

Remark 2.21. Lemma 2.20 shows that, if there exists an abstract hyperplane subordinate to H , then H must be a codimension-one subgroup of G . Conversely, it is not hard to see that, for every quasiconvex codimension-one subgroup $H \leq G$, there exist abstract hyperplanes subordinate to H .

Remark 2.22. Let $(H, \mathfrak{H}), (K, \mathfrak{K})$ be abstract hyperplanes and let H^\pm, K^\pm be the sets constructed in Lemma 2.20. The constant D can always be

enlarged, so, without loss of generality, it is the same for both. If $N_D(H)$ and $N_D(K)$ are disjoint, then a side of \mathfrak{H} is disjoint from a side of \mathfrak{K} .

Indeed, since $N_D(H)$ is connected, we have either $N_D(H) \subseteq K^+$ or $N_D(H) \subseteq K^-$; without loss of generality, let us assume that the former holds. Similarly, we have $N_D(K) \subseteq H^-$ without loss of generality. It follows that the connected set $H^+ \sqcup N_D(H)$ is disjoint from $N_D(K)$ and thus contained in a single connected component of its complement. Since $N_D(H) \subseteq K^+$, we have $H^+ \subseteq K^+$ and $\Lambda H^+ \subseteq \Lambda K^+$. Hence $\mathfrak{H}^+ \cap \mathfrak{K}^- = \emptyset$.

The following is little more than a rephrasing in terms of $\partial_\infty G$ of well-known results from [Sag97, BW12].

Proposition 2.23. *Let \mathcal{H} be a G -invariant set of abstract hyperplanes.*

- (1) *If \mathcal{H} contains only finitely many G -orbits, \mathcal{H} gives rise to a cocompact G -action on an essential CAT(0) cube complex $X(\mathcal{H})$.*
- (2) *In this case, the action $G \curvearrowright X(\mathcal{H})$ is proper if and only if $g^+, g^- \in \partial_\infty G$ are separated by an element of \mathcal{H} , for every infinite-order $g \in G$.*

Proof. The collection $\mathcal{P} = \{\overline{\mathfrak{H}^\epsilon} \mid \mathfrak{H} \in \mathcal{H}, \epsilon \in \{\pm\}\}$ has a natural structure of poset coming from inclusions. We promote this to a structure of pocset by setting $(\overline{\mathfrak{H}^+})^* = \overline{\mathfrak{H}^-}$.

Observe that $\overline{\mathfrak{H}^+} \subseteq \overline{\mathfrak{K}^+}$ if and only if $\mathfrak{H}^+ \cap \mathfrak{K}^- = \emptyset$. Thus, $\overline{\mathfrak{H}^+}$ and $\overline{\mathfrak{K}^+}$ are transverse if and only if both \mathfrak{H}^+ and \mathfrak{H}^- intersect both \mathfrak{K}^+ and \mathfrak{K}^- .

Since there are finitely many G -orbits in \mathcal{H} , Lemma 2.20 provides a constant D that works for every element of \mathcal{H} . By Remark 2.22, every set of k pairwise-transverse elements of \mathcal{P} corresponds to a collection of k cosets of uniformly quasiconvex subgroups of G whose D -neighbourhoods pairwise intersect. Lemma 2.17 shows that \mathcal{P} is finite-dimensional and contains only finitely many G -orbits of maximal pairwise-transverse subsets. Lemma 2.2 thus yields a natural cocompact action on a CAT(0) cube complex $X(\mathcal{H})$.

For every $\mathfrak{H} \in \mathcal{H}$, Lemma 2.15 shows that there exists an infinite-order element $g \in G$ with $g^+ \in \mathfrak{H}^+$ and $g^- \in \mathfrak{H}^-$. A power of g must then skewer the hyperplane of $X(\mathcal{H})$ determined by \mathfrak{H} . We conclude that $X(\mathcal{H})$ is essential. Finally, part (2) follows from Proposition 1.3 in [BW12]. \square

Remark 2.24. Let H_1, \dots, H_k be quasiconvex subgroups of G with the property that, for each $i \leq k$, the difference $\partial_\infty G - \Lambda H_i$ has exactly two connected components, and these are left invariant by the H_i -action. Each H_i determines a unique abstract hyperplane \mathfrak{H}_i and we can consider the collection $\mathcal{H} = G \cdot \mathfrak{H}_1 \cup \dots \cup G \cdot \mathfrak{H}_k$. In this case, the cube complex $X(\mathcal{H})$ provided by part (1) of Proposition 2.23 is automatically hyperplane-essential.

In order to see this, consider abstract hyperplanes $\mathfrak{H}, \mathfrak{K} \in \mathcal{H}$, with stabilisers H, K respectively. If $\Lambda H \cap \mathfrak{K}^+ = \emptyset$, the connected set \mathfrak{K}^+ is partitioned into the two open sets $\mathfrak{K}^+ \cap \mathfrak{H}^\pm$. It follows that one of these two sets is empty and, in particular, \mathfrak{H} and \mathfrak{K} are not transverse.

Thus, if \mathfrak{H} and \mathfrak{K} are transverse, the four intersections $\Lambda H \cap \mathfrak{K}^\pm$ and $\Lambda K \cap \mathfrak{H}^\pm$ must all be nonempty and open in the respective limit sets. By

Lemma 2.15, there exists an infinite-order element $h \in H$ with $h^+ \in \mathfrak{K}^+$ and $h^- \in \mathfrak{K}^-$; in particular, a sufficiently large power of h skewers the hyperplane of $X(\mathcal{H})$ determined by \mathfrak{K} . Similarly, there exists $k \in K$ skewering the hyperplane determined by \mathfrak{H} . This shows that the action $G \curvearrowright X(\mathcal{H})$ is hyperplane-essential.

2.5. Shaving cocompact cubulations. Let X be a CAT(0) cube complex.

Definition 2.25. Two distinct hyperplanes of X are said to be *effectively parallel* if they are disjoint and bound halfspaces at finite Hausdorff distance.

As an example, the cubical subdivision X' contains a pair of effectively parallel hyperplanes for every hyperplane of X .

Note that, in general, two disjoint hyperplanes can be at finite Hausdorff distance without being effectively parallel. For instance, in a tree with all vertices of degree ≥ 3 , any two distinct hyperplanes are at finite Hausdorff distance, but no two of them are effectively parallel.

We also observe that, if X is hyperplane essential and $\mathfrak{h}_1, \mathfrak{h}_2$ are distinct halfspaces at finite Hausdorff distance, then the hyperplanes $\mathfrak{w}_1, \mathfrak{w}_2$ bounding them are necessarily disjoint. Thus, \mathfrak{w}_1 and \mathfrak{w}_2 are effectively parallel.

Given a subset $A \subseteq \mathcal{H}(X)$, we employ the notation $A^* = \{\mathfrak{h}^* \mid \mathfrak{h} \in A\}$.

Lemma 2.26. *Let X be locally finite, cocompact, essential and hyperplane-essential. Let $\mathcal{P} \subseteq \mathcal{H}(X)$ be a maximal set of halfspaces pairwise at finite Hausdorff distance. Then:*

- (1) *the subset $\mathcal{P} \subseteq \mathcal{H}(X)$ is totally ordered by inclusion;*
- (2) *\mathcal{P}^* is also a maximal set of halfspaces at finite Hausdorff distance;*
- (3) *if $\mathfrak{w} \in \mathcal{W}(X)$ doesn't bound an element of \mathcal{P} , then \mathfrak{w} is either transverse to all, contained in all, or disjoint from all elements of \mathcal{P} .*

Proof. Since any two elements of \mathcal{P} are at finite Hausdorff distance and X is hyperplane-essential, no two elements of \mathcal{P} are transverse. Since X is essential, no two elements of \mathcal{P} can be disjoint or have disjoint complements. This shows part (1), while part (2) is clear. We now prove part (3).

Let $\mathfrak{u}, \mathfrak{v}$ be hyperplanes bounding elements of \mathcal{P} . By essentiality of X , no $\mathfrak{w} \in \mathcal{W}(X)$ can form a facing triple with \mathfrak{u} and \mathfrak{v} . Since \mathcal{P} is maximal, if \mathfrak{w} does not bound an element of \mathcal{P} , then \mathfrak{w} cannot separate \mathfrak{u} and \mathfrak{v} either. Thus, if \mathfrak{w} is not transverse to any element of \mathcal{P} , then \mathfrak{w} must be either contained in all elements of \mathcal{P} , or contained in all of their complements.

Finally, suppose that \mathfrak{w} is transverse to an element of \mathcal{P} , but not to all of them. In this case, \mathfrak{w} and the elements of \mathcal{P} all originate from a single de Rham factor of X and Proposition 2.7 provides a hyperplane forming a facing triple with two hyperplanes bounding elements of \mathcal{P} . This is a contradiction and it concludes the proof of part (3). \square

Given X as in Lemma 2.26 and $\mathfrak{h} \in \mathcal{H}(X)$, we denote by $\mathcal{P}(\mathfrak{h}) \subseteq \mathcal{H}(X)$ the subset of all halfspaces at finite Hausdorff distance from \mathfrak{h} . We define:

$$\text{Para}(X) = \{\mathcal{P}(\mathfrak{h}) \mid \mathfrak{h} \in \mathcal{H}(X)\}.$$

Given distinct elements $\mathcal{Q}_1, \mathcal{Q}_2 \in \text{Para}(X)$, Lemma 2.26 shows that whether $\mathfrak{h}_1 \in \mathcal{Q}_1$ is contained in $\mathfrak{h}_2 \in \mathcal{Q}_2$ is independent of the choice of \mathfrak{h}_1 and \mathfrak{h}_2 . When this happens, we write $\mathcal{Q}_1 \preceq \mathcal{Q}_2$. We obtain a pocset $(\text{Para}(X), \preceq, *)$ and a surjective pocset homomorphism $\mathcal{P}: \mathcal{H}(X) \rightarrow \text{Para}(X)$.

The dimension of the pocset $\text{Para}(X)$ coincides with $\dim X < +\infty$. Lemma 2.2 thus guarantees the existence of a unique class of DCC ultrafilters on $\text{Para}(X)$, which gives rise to a CAT(0) cube complex $\text{Cmp}(X)$. We refer to $\text{Cmp}(X)$ as the *compression* of X . Observe that $\text{Para}(X)$ and $\text{Cmp}(X)$ are naturally equipped with an $\text{Aut}(X)$ -action.

The preimage under \mathcal{P} of any ultrafilter on $\text{Para}(X)$ is an ultrafilter on $\mathcal{H}(X)$. Assuming for a moment that all fibres of \mathcal{P} are finite, preimages of DCC ultrafilters are again DCC. In this case, we obtain an $\text{Aut}(X)$ -equivariant injection $\iota: \text{Cmp}(X)^{(0)} \hookrightarrow X^{(0)}$, which does not shrink distances.

As the next result shows, the condition on the fibres of \mathcal{P} corresponds to X having no \mathbb{R} -factors in its de Rham decomposition.

Lemma 2.27. *Let X be locally finite, cocompact, essential and hyperplane-essential. Assume that X has no factors isomorphic to \mathbb{R} in its de Rham decomposition. Then:*

- (1) *the fibres of the map \mathcal{P} are uniformly finite and ι is bi-Lipschitz;*
- (2) *$g \in \text{Aut}(X)$ skewers $\mathfrak{h} \in \mathcal{H}(X)$ if and only if it skewers the halfspace of $\text{Cmp}(X)$ determined by $\mathcal{P}(\mathfrak{h})$;*
- (3) *$\text{Cmp}(X)$ is locally finite, cocompact, essential, hyperplane-essential and it has no halfspaces at finite Hausdorff distance.*

Proof. Let n be the number of orbits of the action $\text{Aut}(X) \curvearrowright \mathcal{H}(X)$. If an element $\mathcal{Q} \in \text{Para}(X)$ contains $> n$ halfspaces, there exist $g \in \text{Aut}(X)$ and $\mathfrak{h} \in \mathcal{Q}$ such that $\mathfrak{h} \subsetneq g\mathfrak{h} \in \mathcal{Q}$. In this case, the halfspaces $g^n\mathfrak{h}$ are pairwise at finite Hausdorff distance, hence $\{g^n\mathfrak{h}\}_{n \in \mathbb{Z}} \subseteq \mathcal{Q}$. Part (1) of Lemma 2.26 implies that \mathcal{Q} is a bi-infinite chain and, by part (3) of Lemma 2.26, every halfspace in $\mathcal{Q} \sqcup \mathcal{Q}^*$ is transverse to all the halfspaces in $\mathcal{H}(X) - (\mathcal{Q} \sqcup \mathcal{Q}^*)$. This contradicts the assumption that X has no de Rham factors isomorphic to \mathbb{R} . We conclude that all fibres of \mathcal{P} have cardinality $\leq n$, hence ι is n -Lipschitz. This yields part (1). Parts (2) and (3) follow immediately. \square

We now make Definition 1.1 from the introduction a bit more precise (recalling that, in a hyperplane-essential cube complex, halfspaces at finite Hausdorff distance are always bounded by effectively parallel hyperplanes).

Definition 2.28. A CAT(0) cube complex X is *bald* if it is essential, hyperplane-essential and, moreover, the following holds. If $\mathfrak{w}_1, \mathfrak{w}_2 \in \mathcal{W}(X)$ are effectively parallel, there exists a factor L in the de Rham decomposition of X such that $L \cong \mathbb{R}$ and $\mathfrak{w}_1, \mathfrak{w}_2 \in \mathcal{W}(L)$.

A *bald cubulation* is a proper, cocompact action on a bald cube complex.

Recall that, as defined in Section 2.4, if a hyperbolic group G acts properly and cocompactly on a CAT(0) cube complex X , then every subset $\Omega \subseteq X$ determines limit sets $\Lambda\Omega \subseteq \partial_\infty G$ and $\partial_\infty \Omega \subseteq \partial_\infty X$.

Proposition 2.29. *Let a group G act properly and cocompactly on a CAT(0) cube complex X . Then there exists another CAT(0) cube complex X_\bullet and a proper cocompact action $G \curvearrowright X_\bullet$ such that:*

- (1) X_\bullet is bald;
- (2) if G is hyperbolic and $g \in G$ skewers $\mathfrak{w} \in \mathscr{W}(X)$, there exists a hyperplane $\mathfrak{u} \in \mathscr{W}(X_\bullet)$ such that \mathfrak{u} is skewered by g and $\Lambda \mathfrak{u} \subseteq \Lambda \mathfrak{w}$.

Proof. By replacing X with the cubical subdivision, we can assume that G acts on X without inversions. Let $\#(X) = (n_0, \dots, n_{\dim X - 2})$, where for each i , n_i is the number of G -orbits of i -cubes.

If X is hyperplane-essential, then, by passing to the G -essential core, we can assume that the cube complex is essential and hyperplane-essential (cf. Proposition 3.5 in [CS11]).

If not, then, by Theorem A of [HT19], X contains a G -invariant subspace Y that has the structure of a CAT(0) cube complex, with $Y^{(0)} = X^{(0)}$. Moreover, the set of hyperplanes of Y is exactly the set of components of subspaces of the form $\mathfrak{u} \cap Y$, where \mathfrak{u} is a hyperplane of X . Finally, the action of G on Y is without inversions, and $\#(Y) < \#(X)$ (in lexicographic order).

Iterating finitely many times, we find a hyperplane-essential cocompact action $G \curvearrowright Z$, where Z is a CAT(0) cube complex G -equivariantly embedded in X ; by replacing Z with its G -essential core, we have that Z is essential and hyperplane-essential, and $Z^{(0)} \subseteq X^{(0)}$.

Since G acts on X properly and $Z \hookrightarrow X$ is G -equivariant, each 0-cube of X , and hence each 0-cube of Z , has finite stabiliser in G . Since G acts on Z cocompactly, the action of G on Z is therefore proper. We are left to deal with effectively parallel hyperplanes, in order to ensure that Z is bald.

Isolating the \mathbb{R} -factors in the de Rham decomposition of Z , we obtain a splitting $Z = \mathbb{R}^m \times W$, where $m \geq 0$ and W is a CAT(0) cube complex with no \mathbb{R} -factors. Observe that G leaves invariant this decomposition of Z , and the induced action $G \curvearrowright W$ is cocompact.

We set $X_\bullet = \mathbb{R}^m \times \text{Cmp}(W)$. Observe that the G -action descends to X_\bullet . By part (1) of Lemma 2.27, X_\bullet is G -equivariantly quasi-isometric to Z ; hence $G \curvearrowright X_\bullet$ is proper and cocompact. By part (3) of Lemma 2.27, X_\bullet is bald. This completes the proof of part (1).

We now assume that G is hyperbolic and prove part (2). Hyperbolicity implies that $m = 0$, so $Z = W$ and $X_\bullet = \text{Cmp}(Z)$.

Suppose that $g \in G$ skewers a hyperplane \mathfrak{w} of X . By Theorem A of [HT19], $\mathfrak{w} \cap Z = \bigsqcup_{i \in \mathcal{I}} \mathfrak{w}_i$, where each \mathfrak{w}_i is a hyperplane of Z . In particular, $\Lambda \mathfrak{w}_i \subseteq \Lambda \mathfrak{w}$ for each $i \in \mathcal{I}$.

Let γ be an axis for g in Z , so that the endpoints of γ are $g^\pm \in \partial_\infty Z$. Suppose that no \mathfrak{w}_i is skewered by γ . Then for each i , we can choose a component \mathfrak{w}_i^+ of $Z - \mathfrak{w}_i$ so that $\gamma \subseteq \bigcap_i \mathfrak{w}_i^+$. Hence γ is a $\langle g \rangle$ -invariant embedded path in X which is disjoint from \mathfrak{w} , so g^-, g^+ lie on the same side of $\Lambda \mathfrak{w}$ in $\partial_\infty Z$, contradicting that g skewers \mathfrak{w} .

In conclusion, g skewers a hyperplane $\mathfrak{w}_i \in \mathscr{W}(Z)$ and $\Lambda \mathfrak{w}_i \subseteq \Lambda \mathfrak{w}$. Part (2) of the proposition now follows from part (2) of Lemma 2.27. \square

3. BENDING HYPERPLANES.

3.1. Controlling families of hyperplanes. Let a group G act properly and cocompactly on a CAT(0) cube complex X . As usual, we endow X with its ℓ_1 metric and set $d = \dim X$.

We are interested in the case when the quotient $G \backslash X$ is a *special* cube complex, in the sense of [HW08, Definition 3.2]. Recall that two distinct hyperplanes $\mathfrak{a}_1, \mathfrak{a}_2$ in $G \backslash X$ *inter-oscuate* if they both intersect and osculate: there is a square whose barycentre is in $\mathfrak{a}_1 \cap \mathfrak{a}_2$, and there are also 1-cubes dual to $\mathfrak{a}_1, \mathfrak{a}_2$ that share a vertex but do not lie in a common square. A special cube complex never has inter-osculating hyperplanes.

Lemma 3.1. *Suppose that the quotient $G \backslash X$ is a special cube complex. Given disjoint hyperplanes \mathfrak{w}_1 and \mathfrak{w}_2 , there exists a finite-index subgroup $H \leq G$ such that any two elements of $H \cdot \mathfrak{w}_1 \cup H \cdot \mathfrak{w}_2$ are disjoint and at distance $\geq \frac{1}{d} \cdot d(\mathfrak{w}_1, \mathfrak{w}_2)$.*

Proof. By Corollary 7.9 in [HW08], halfspace-stabilisers are separable in G . Thus, Lemma 2.12 allows us to assume that any two hyperplanes in the same G -orbit are disjoint and at distance $\geq \frac{1}{d} \cdot d(\mathfrak{w}_1, \mathfrak{w}_2)$. Let $n \geq 0$ be maximal such that $\mathscr{W}(\mathfrak{w}_1 | \mathfrak{w}_2)$ contains n pairwise-disjoint hyperplanes. By Dilworth's lemma, we have $n \geq \frac{1}{d} \cdot (d(\mathfrak{w}_1, \mathfrak{w}_2) - 1)$.

Claim: *Given disjoint $\mathfrak{u}_1, \mathfrak{u}_2 \in \mathscr{W}(X)$ and $m \geq 0$ maximal such that $\mathscr{W}(\mathfrak{u}_1 | \mathfrak{u}_2)$ contains m pairwise-disjoint hyperplanes, there exists a finite-index subgroup $L \leq G$ such that every element of $L \cdot \mathfrak{u}_2$ is separated from \mathfrak{u}_1 by at least m pairwise-disjoint hyperplanes (just disjoint from \mathfrak{u}_1 if $m = 0$).*

Applying the claim to $\mathfrak{u}_i = \mathfrak{w}_i$ clearly concludes the proof. We will prove the claim by induction on $m \geq 0$.

The base step $m = 0$ is immediate taking $L = G$, as the quotient $G \backslash X$ has no inter-osculating hyperplanes.

When $m \geq 1$, we can pick $\mathfrak{u} \in \mathscr{W}(\mathfrak{u}_1 | \mathfrak{u}_2)$ such that $\mathscr{W}(\mathfrak{u} | \mathfrak{u}_1)$ contains $m - 1$ pairwise-disjoint hyperplanes. The inductive hypothesis yields a finite-index subgroup $K \leq G$ such that every element of $K \cdot \mathfrak{u}$ is separated from \mathfrak{u}_1 by at least $m - 1$ pairwise-disjoint hyperplanes and such that no element of $K \cdot \mathfrak{u}$ is transverse to \mathfrak{u}_2 or \mathfrak{u}_1 .

Since no two elements of $K \cdot \mathfrak{u}$ are transverse, the corresponding restriction quotient (see [CS11, Section 2.3]) of X is a tree \mathcal{T} . Recall that the preimage in \mathcal{T} of the midpoint of any edge is a hyperplane in $K \cdot \mathfrak{u}$, every hyperplane in $K \cdot \mathfrak{u}$ is sent to the midpoint of an edge, and all other hyperplanes are collapsed to vertices.

Since the hyperplanes \mathfrak{u}_1 and \mathfrak{u}_2 are contained in distinct connected components of $X - \bigcup K \cdot \mathfrak{u}$, they get collapsed to distinct vertices $v_1, v_2 \in \mathcal{T}$.

Let $k \in K$ be such that $kv_2 \neq v_1$. Then kv_2 and v_1 are separated by an edge of \mathcal{T} ; the preimage of the midpoint of this edge is a hyperplane $\mathfrak{v} \in K \cdot \mathfrak{u}$ that separates the preimages of kv_2 and v_1 . In particular, \mathfrak{v} separates \mathfrak{u}_1 from $k\mathfrak{u}_2$. Hence, the set

$$\mathcal{W}(k\mathfrak{u}_2|\mathfrak{u}_1) \supseteq \{\mathfrak{v}\} \sqcup \mathcal{W}(\mathfrak{v}|\mathfrak{u}_1)$$

contains at least m pairwise-disjoint hyperplanes. The K -stabiliser of v_2 is separable in K by Corollary 7.9 in [HW08]. It follows that there exists a finite-index subgroup $L \leq K$ such that $v_1 \notin L \cdot v_2$. Every element of $L \cdot \mathfrak{u}_2$ is then separated from \mathfrak{u}_1 by at least m pairwise-disjoint hyperplanes. \square

Proposition 3.2. *Suppose that G is one-ended and that X is essential, hyperplane-essential and δ -hyperbolic. For every $n > 0$, there exist $m \geq 4$, hyperplanes $\mathfrak{u}_1, \dots, \mathfrak{u}_m$ and a finite-index subgroup $H \leq G$ such that:*

- $G \cdot \{\mathfrak{u}_1, \dots, \mathfrak{u}_m\} = \mathcal{W}(X)$;
- \mathfrak{u}_i is transverse to \mathfrak{u}_{i+1} for every $1 \leq i < m$;
- any two elements of $H \cdot \mathfrak{u}_{i-1} \cup H \cdot \mathfrak{u}_{i+1}$ are at distance $\geq n$, for every $1 < i < m$;
- $H \cdot \mathfrak{u}_i \neq H \cdot \mathfrak{u}_j$ whenever $i \neq j$.

Proof. By Theorem 1.1 in [Ago13], we can assume that the quotient $G \backslash X$ is a special cube complex. By Lemma 2.12, we can further assume that any two hyperplanes in the same G -orbit are disjoint and at distance $\geq n$. By Lemma 2.10, there exists a sequence $\mathfrak{u}_1, \dots, \mathfrak{u}_m$ of (not necessarily distinct) hyperplanes satisfying the first two conditions.

Since hyperplane-stabilisers are separable [HW08, Theorem 1.3], the fourth condition can always be ensured by passing to a further finite-index subgroup, as long as the other three conditions are satisfied and the \mathfrak{u}_i are pairwise distinct. We will progressively modify $\mathfrak{u}_1, \dots, \mathfrak{u}_m$ in order to ensure that we are in this situation.

Consider $1 < j < m$ and a finite index-subgroup $H \leq G$ such that $\mathfrak{u}_1, \dots, \mathfrak{u}_j$ are pairwise distinct and the third condition holds for all $i < j$. Since the action $G_{\mathfrak{u}_j} \curvearrowright \mathfrak{u}_j$ is proper and cocompact by Lemma 2.3, Lemma 2.13 yields $g \in G_{\mathfrak{u}_j}$ such that $d(g\mathfrak{u}_{j+1}, \mathfrak{u}_{j-1}) \geq dn$. We can moreover ensure that $g\mathfrak{u}_{j+1} \notin \{\mathfrak{u}_1, \dots, \mathfrak{u}_j\}$. For $l \geq j+1$, we replace each \mathfrak{u}_l with $g\mathfrak{u}_l$. Note that the new hyperplanes still satisfy the first two conditions and, for $i < j$, also the third. Since now $d(\mathfrak{u}_{j-1}, \mathfrak{u}_{j+1}) \geq dn$, Lemma 3.1 yields a finite-index subgroup $K \leq H$ such that any two elements of $K \cdot \mathfrak{u}_{j-1} \cup K \cdot \mathfrak{u}_{j+1}$ are at distance $\geq n$. Replacing H with K , the third condition is now satisfied for $i \leq j$ and $\mathfrak{u}_1, \dots, \mathfrak{u}_{j+1}$ are pairwise distinct. \square

Now let X be δ -hyperbolic and consider hyperplanes $\mathfrak{v}_0, \dots, \mathfrak{v}_m$ such that:

- \mathfrak{v}_i is transverse to \mathfrak{v}_{i+1} for every $0 \leq i < m$;
- $d(\mathfrak{v}_{i-1}, \mathfrak{v}_{i+1}) \geq n$ for every $0 < i < m$.

If $x \in \mathfrak{v}_0$ and $y \in \mathfrak{v}_m$ are vertices of the respective hyperplanes, we set $x_0 = x$ and, for each $0 \leq i \leq m-1$, we define inductively x_{i+1} as the gate-projection

of x_i to \mathbf{v}_{i+1} . Note that $x_{i+1} \in \mathbf{v}_i \cap \mathbf{v}_{i+1}$. Finally, set $x_{m+1} = y$. Joining each x_i to x_{i+1} by an ℓ_1 geodesic, we obtain a path $\eta \subseteq \mathbf{v}_0 \cup \dots \cup \mathbf{v}_m \subseteq X$, which we will refer to as a *standard path* from x to y .

By construction, the segment of η that joins x_i to x_{i+2} is a geodesic for all $0 \leq i \leq m-1$. It follows that η is a k -local geodesic, with:

$$k \geq \min_{1 \leq i \leq m-1} d(x_i, x_{i+1}) \geq \min_{1 \leq i \leq m-1} d(\mathbf{v}_{i-1}, \mathbf{v}_{i+1}) \geq n.$$

When X is δ -hyperbolic and $n > 8\delta$, Theorem III.H.1.13 in [BH99] guarantees that η is a $(3, 2\delta)$ -quasi-geodesic. By the Morse Lemma, there exists a constant $K = K(\delta)$ such that every geodesic in X connecting x and y is at Hausdorff distance $\leq K$ from any standard path η .

Lemma 3.3. *Let X and $\mathbf{v}_0, \dots, \mathbf{v}_m$ be as above, with $m \geq 2$ and $n > 8\delta$.*

- (1) *Any geodesic from a point $x \in \mathbf{v}_0$ to a point $y \in \mathbf{v}_m$ has Hausdorff distance $\leq K$ from any standard path $\eta \subseteq \mathbf{v}_0 \cup \dots \cup \mathbf{v}_m$ from x to y . In particular, the union $\mathbf{v}_0 \cup \dots \cup \mathbf{v}_m$ is K -quasi-convex in X .*
- (2) *We have $d(\mathbf{v}_0, \mathbf{v}_m) \geq \frac{n(m-1)}{3} - 2\delta$. Thus, \mathbf{v}_0 and \mathbf{v}_m are disjoint.*

Proof. We have already shown part (1) in the above discussion. Regarding part (2), pick vertices $x \in \mathbf{v}_0$ and $y \in \mathbf{v}_m$ with $d(x, y) = d(\mathbf{v}_0, \mathbf{v}_m)$.

Let η be a standard path from x to y . As shown above, η is an n -local geodesic and it contains points x_1, \dots, x_m satisfying $d(x_i, x_{i+1}) \geq n$. It follows that the domain of η has length $\geq n(m-1)$ and we know that η is a $(3, 2\delta)$ -quasi-geodesic. Thus:

$$d(\mathbf{v}_0, \mathbf{v}_m) = d(x, y) \geq \frac{n(m-1)}{3} - 2\delta > 0,$$

as required. \square

3.2. Systems of switches. Let G be a one-ended hyperbolic group. We consider a proper cocompact G -action on an essential, hyperplane-essential, δ -hyperbolic CAT(0) cube complex X . We fix $M \geq 1$ such that, for every $\mathfrak{w} \in \mathscr{W}(X)$ and every vertex $x \in \mathfrak{w}$, the orbit $G_{\mathfrak{w}} \cdot x$ is M -dense in \mathfrak{w} , using Lemma 2.3.

Let us denote by $\text{Trans}(X)$ the collection of subsets $\{\mathbf{u}, \mathbf{v}\} \subseteq \mathscr{W}(X)$ such that \mathbf{u} is transverse to \mathbf{v} . Given a subset $\mathcal{S} \subseteq \text{Trans}(X)$ and $\mathbf{u} \in \mathscr{W}(X)$, let $\mathcal{S}_{\mathbf{u}} \subseteq \mathscr{W}(X)$ be the collection of all those \mathbf{v} with $\{\mathbf{u}, \mathbf{v}\} \in \mathcal{S}$.

Definition 3.4. An n -system of switches is a pair $\mathcal{S} = (\mathcal{S}, H)$, where:

- $H \triangleleft G$ is a finite-index subgroup;
- $\mathcal{S} \subseteq \text{Trans}(X)$ is an H -invariant subset;
- for every $\mathbf{u} \in \mathscr{W}(X)$, any two elements of $\mathcal{S}_{\mathbf{u}}$ are at distance $\geq n$.

The *support* of \mathcal{S} is the set $\text{supp } \mathcal{S} := \{\mathbf{u} \in \mathscr{W}(X) \mid \mathcal{S}_{\mathbf{u}} \neq \emptyset\}$. We say that \mathcal{S} is *full* if $G \cdot \text{supp } \mathcal{S} = \mathscr{W}(X)$.

Lemma 3.5. *For every $n > 0$, there exists a full n -system of switches.*

Proof. Let u_1, \dots, u_m and H be as provided by Proposition 3.2; passing to a further finite-index subgroup, we can assume that H is normal in G . Let \mathcal{S} be the union of the sets $H \cdot \{u_i, u_{i+1}\}$ for $1 \leq i < m$. Then the pair $\mathcal{S} = (\mathcal{S}, H)$ is a full n -system of switches.

Indeed, we have $\mathfrak{w} \in \mathcal{S}_{u_j}$ if and only if there exist $h \in H$ and i satisfying $h \cdot \{\mathfrak{w}, u_j\} = \{u_i, u_{i+1}\}$. Since $H \cdot u_l \neq H \cdot u_{l'}$ whenever $l \neq l'$, we must have $j \in \{i, i+1\}$ and $hu_j = u_j$. It follows that $\mathcal{S}_{u_j} = H_{u_j} \cdot u_{j-1} \sqcup H_{u_j} \cdot u_{j+1}$, any two elements of which are at distance $\geq n$. \square

We write $U(\mathcal{S}) \subseteq X$ for the union of all hyperplanes in $\text{supp } \mathcal{S}$, and $U_{\mathfrak{h}}(\mathcal{S}) \subseteq X$ for the union of all intersections $\mathfrak{u} \cap \mathfrak{v}$ with $\{\mathfrak{u}, \mathfrak{v}\} \in \mathcal{S}$. We have:

$$U(\mathcal{S}) = U_{\mathfrak{h}}(\mathcal{S}) \sqcup (U(\mathcal{S}) - U_{\mathfrak{h}}(\mathcal{S})),$$

where every point of $U(\mathcal{S}) - U_{\mathfrak{h}}(\mathcal{S})$ lies in a unique hyperplane of $\text{supp } \mathcal{S}$, while points of $U_{\mathfrak{h}}(\mathcal{S})$ each lie in two distinct elements of $\text{supp } \mathcal{S}$.

We will refer to $U_{\mathfrak{h}}(\mathcal{S})$ and $U(\mathcal{S}) - U_{\mathfrak{h}}(\mathcal{S})$, respectively, as the *singular* and *regular* part of $U(\mathcal{S})$ (and to their points as *singular* and *regular points*). Note that every vertex of each hyperplane in $\text{supp } \mathcal{S}$ is a regular point.

We denote by $\text{comp } \mathcal{S}$ the set of all connected components of the regular part of $U(\mathcal{S})$. Each element of $\text{comp } \mathcal{S}$ is a connected component of a set of the form $\mathfrak{w} - \bigcup_{\mathfrak{v} \in \mathcal{S}_{\mathfrak{w}}} \mathfrak{v}$ with $\mathfrak{w} \in \text{supp } \mathcal{S}$. Instead, note that connected components of the singular part $U_{\mathfrak{h}}(\mathcal{S})$ are in one-to-one correspondence with elements of \mathcal{S} . For every regular point $x \in U(\mathcal{S})$, we write $[x]$ for the only element of $\text{comp } \mathcal{S}$ that contains x .

To each n -system of switches $\mathcal{S} = (\mathcal{S}, H)$ we can associate a bipartite graph $\mathcal{G}(\mathcal{S})$ equipped with a cocompact H -action. The vertex set is the disjoint union $\mathcal{S} \sqcup \text{comp } \mathcal{S}$. We join each vertex $\{\mathfrak{u}, \mathfrak{v}\} \in \mathcal{S}$ to the four elements of $\text{comp } \mathcal{S}$ that contain $\mathfrak{u} \cap \mathfrak{v}$ in their closure; two are contained in \mathfrak{u} and two in \mathfrak{v} . We call vertices of $\mathcal{G}(\mathcal{S})$ *regular* if they originate from elements of $\text{comp } \mathcal{S}$ and *singular* if they originate from elements of \mathcal{S} .

All singular vertices have degree 4 in $\mathcal{G}(\mathcal{S})$, while regular vertices can have infinite degree in general.

Every hyperplane $\mathfrak{w} \in \text{supp } \mathcal{S}$ determines a subgraph $\mathcal{G}(\mathfrak{w}) \subseteq \mathcal{G}(\mathcal{S})$ spanned by the elements of $\text{comp } \mathcal{S}$ that are contained in \mathfrak{w} . Note that $\mathcal{G}(\mathfrak{w})$ is naturally isomorphic to the barycentric subdivision of the restriction quotient of X corresponding to the set $\mathcal{S}_{\mathfrak{w}} \subseteq \mathcal{W}(X)$. In particular, since the definition of a system of switches means that $\mathcal{S}_{\mathfrak{w}}$ consists of pairwise disjoint hyperplanes, the graph $\mathcal{G}(\mathfrak{w})$ is a tree.

Example 3.6. It can be helpful to have the following special case in mind in this subsection. Let G be the fundamental group of a closed, orientable surface Σ . Let X be a 2-dimensional cubulation constructed by applying Sageev's construction to a suitable finite filling collection of curves \mathcal{F} on Σ .

A choice of a system of switches \mathcal{S} corresponds to a choice of a finite cover $\Sigma' \rightarrow \Sigma$ and a set of intersections between the lifts of \mathcal{F} to Σ' . We will perform surgery on these intersections and obtain a new family of curves in

Σ' (sometimes, one of the new curves will fill Σ'). Here, the main difference from the general case is that regular vertices of $\mathcal{G}(\mathcal{S})$ always have degree 2.

Definition 3.7. We say that a connected subgraph $A \subseteq \mathcal{G}(\mathcal{S})$ is:

- (1) *two-sided* if every singular vertex that lies in A has degree 2 in A ;
- (2) *star-complete* if, whenever a regular vertex lies in A , its star in $\mathcal{G}(\mathcal{S})$ is also contained in A .

Given a connected subgraph $A \subseteq \mathcal{G}(\mathcal{S})$, let $U(A) \subseteq U(\mathcal{S})$ be obtained by taking the union of the closures of the elements of $\text{comp } \mathcal{S}$ that lie in A .

We say that a subspace $Y \subset X$ is *locally codimension-1* if for all $y \in Y$, there is a neighbourhood N of y in X equipped with a homeomorphism $(Y \cap N) \times [-\frac{1}{2}, \frac{1}{2}] \rightarrow N$ restricting to the inclusion $Y \cap N \rightarrow N$ on $(Y \cap N) \times \{0\}$.

Proposition 3.8. *There exists $K = K(\delta)$ such that all the following hold for $n > 8\delta$ and any n -system of switches $\mathcal{S} = (\mathcal{S}, H)$.*

- (1) *The graph $\mathcal{G}(\mathcal{S})$ is a forest.*
- (2) *Consider $\mathbf{u}, \mathbf{v} \in \text{supp } \mathcal{S}$ such that $\mathcal{G}(\mathbf{u})$ and $\mathcal{G}(\mathbf{v})$ are contained in the same connected component of $\mathcal{G}(\mathcal{S})$. Then \mathbf{u} and \mathbf{v} are transverse if and only if $\{\mathbf{u}, \mathbf{v}\} \in \mathcal{S}$.*
- (3) *For every connected subgraph $A \subseteq \mathcal{G}(\mathcal{S})$, the subset $U(A) \subseteq U(\mathcal{S})$ is connected and K -quasiconvex.*
- (4) *A connected subgraph $A \subseteq \mathcal{G}(\mathcal{S})$ is star-complete and two-sided if and only if the subset $U(A) \subseteq X$ is locally codimension-1.*
- (5) *The orbit $G \cdot U(A)$ is locally finite in X if and only if no regular vertex of $\mathcal{G}(\mathcal{S})$ lies in infinitely many pairwise-distinct H -translates of A .*

Proof. Consider an immersed path $\gamma \subseteq \mathcal{G}(\mathcal{S})$ between two regular vertices $[x], [y] \in \text{comp } \mathcal{S}$. Let $\mathbf{v}_0, \dots, \mathbf{v}_k \in \text{supp } \mathcal{S}$ be the hyperplanes containing the elements of $\gamma \cap \text{comp } \mathcal{S}$, in the same order as they appear moving from $[x]$ to $[y]$ along γ . In particular, $x \in [x] \subseteq \mathbf{v}_0$, $y \in [y] \subseteq \mathbf{v}_k$, the hyperplane \mathbf{v}_i is transverse to \mathbf{v}_{i+1} for each $0 \leq i \leq k-1$ and we have $d(\mathbf{v}_{i-1}, \mathbf{v}_{i+1}) \geq n$ for every $0 < i < k$.

It follows that $\mathbf{v}_0, \dots, \mathbf{v}_k$ satisfy the hypotheses of Lemma 3.3. Denoting by $\tilde{\gamma} \subseteq U(\mathcal{S}) \subseteq X$ any standard path joining x and y , any geodesic joining x and y in X is at Hausdorff distance $\leq K$ from $\tilde{\gamma}$. This shows part (3), while parts (1) and (2) follow from part (2) of Lemma 3.3. Part (4) is obvious.

Finally, we prove part (5). Observe that $G \cdot U(A)$ is locally finite if and only if $H \cdot U(A)$ is. This fails if and only if a point $z \in U(\mathcal{S})$ lies in infinitely many pairwise-distinct H -translates of $U(A)$, say $h_i U(A) = U(h_i A)$. Without loss of generality, z is a vertex of a hyperplane of X . The above then happens if and only if the vertex $[z] \in \mathcal{G}(\mathcal{S})$ lies in the pairwise-distinct subgraphs $h_i A$. \square

Definition 3.9. A *crooked hyperplane* is a connected, two-sided, star-complete subtree $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ such that no regular vertex of $\mathcal{G}(\mathcal{S})$ lies in infinitely many pairwise-distinct H -translates of Γ .

Given a crooked hyperplane $\Gamma \subseteq \mathcal{G}(\mathcal{S})$, we denote by G_Γ the G -stabiliser of the subset¹ $U(\Gamma) \subseteq X$. Part (5) of Proposition 3.8 and Lemma 2.3 in [HS18] show that the action $G_\Gamma \curvearrowright U(\Gamma)$ is proper and cocompact. By part (3), the subgroup $G_\Gamma \leq G$ is quasiconvex. Moreover, from part (4) and Mayer–Vietoris, we see that $X - U(\Gamma)$ has exactly two connected components. We write $G_\Gamma^0 \leq G_\Gamma$ for the subgroup (of index at most two) that leaves invariant both connected components of $X - U(\Gamma)$.

Remark 3.10. For every crooked hyperplane $\Gamma \subseteq \mathcal{G}(\mathcal{S})$, there is a natural map $\iota: \partial_\infty \Gamma \rightarrow \partial_\infty U(\Gamma)$ taking the endpoint of a ray γ in the tree Γ to the endpoint at infinity of any standard path $\tilde{\gamma} \subseteq U(\Gamma)$. The map ι is a homeomorphism onto its image, and it satisfies:

$$\partial_\infty U(\Gamma) = \iota(\partial_\infty \Gamma) \sqcup \bigcup_{c \in \Gamma \cap \text{comp } \mathcal{S}} \partial_\infty c.$$

Thus, $\partial_\infty U(\Gamma)$ is nonempty even when each element of $\text{comp } \mathcal{S}$ is bounded.

Proposition 3.11. *Let \mathcal{S} be an n -system of switches. Every compact, two-sided subtree $A \subseteq \mathcal{G}(\mathcal{S})$ is contained in a crooked hyperplane.*

Proof. Let us write $\mathcal{G} = \mathcal{G}(\mathcal{S})$ for simplicity and let $\mathcal{G}_2 \subseteq \mathcal{G}$ be the subset of singular vertices; recall that the action $H \curvearrowright \mathcal{G}$ is cocompact. Since the action $H \curvearrowright \mathcal{G}$ has separable vertex- and edge-stabilisers by Corollary 7.9 in [HW08], there exists a finite-index subgroup $L \triangleleft H$ such that A projects injectively to the quotient $\bar{\mathcal{G}} = L \backslash \mathcal{G}$ and such that every element of \mathcal{G}_2 projects to a degree-4 vertex of $\bar{\mathcal{G}}$. Let $\pi: \mathcal{G} \rightarrow \bar{\mathcal{G}}$ denote the quotient projection.

For each $v \in \pi(\mathcal{G}_2)$, we choose a partition of the four edges incident to v into two pairs. We do so ensuring that, if $v \in \pi(A)$, one element of the partition consists precisely of the two edges lying in $\pi(A)$. We now lift these partitions to \mathcal{G} . Let \mathcal{G}' be the graph obtained from \mathcal{G} by replacing every vertex in \mathcal{G}_2 with two vertices of degree 2, according to the chosen partitions. The graph \mathcal{G}' naturally comes equipped with an immersion $f: \mathcal{G}' \rightarrow \mathcal{G}$.

By construction, there exists a connected component $\Gamma \subseteq \mathcal{G}'$ such that $A \subseteq f(\Gamma)$. It is clear that $f(\Gamma)$ is a connected, two-sided, star-complete subtree of \mathcal{G} . Each regular vertex of \mathcal{G} lies in at most one L -translate of $f(\Gamma)$. It follows that every regular vertex of \mathcal{G} lies in finitely many H -translates of $f(\Gamma)$. This shows that $f(\Gamma)$ is a crooked hyperplane, concluding the proof. \square

If $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ is a crooked hyperplane, we denote by \mathcal{T}_Γ the connected component of $\mathcal{G}(\mathcal{S})$ that contains Γ .

Let D be as in Lemma 2.14, let K be as in Proposition 3.8 and let M be the constant chosen at the beginning of this subsection.

¹The G -stabiliser of Γ itself would not make sense, as, although H acts on $\mathcal{G}(\mathcal{S})$, the group G does not.

Proposition 3.12. *Consider $n > 2(M + D + 2K)$ and a full n -system of switches $\mathcal{S} = (\mathcal{S}, H)$. For every non-constant geodesic $\gamma \subseteq X$, there exists a regular vertex $v \in \mathcal{G}(\mathcal{S})$ with the following property. For every crooked hyperplane $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ intersecting the orbit $H \cdot v$, there exists $g \in G$ such that $g\gamma$ intersects $U(\mathcal{T}_\Gamma)$ in a single point, which lies in $U(\Gamma)$.*

Before proving Proposition 3.12, we need to obtain a couple of lemmas.

Lemma 3.13. *Let \mathcal{T} be a connected component of $\mathcal{G}(\mathcal{S})$. Consider a regular point $p \in U(\mathcal{T})$ and let \mathfrak{w} be the hyperplane containing $[p] \in \text{comp } \mathcal{S}$. Then, for every $x \in U(\mathcal{T}) - [p]$, we have $d(p, \pi_{\mathfrak{w}}(x)) \geq d(p, \mathfrak{w} - [p]) - 2K$.*

Proof. Let $\tilde{\gamma}$ be a standard path from p to x and let α be an ℓ_1 geodesic joining p and x . Let $y \in \tilde{\gamma}$ be the first point that does not lie in $[p]$ and let $y' \in \alpha$ be a point that is closest to y ; in particular, we have $d(y, y') \leq K$. It follows that:

$$(x \cdot y)_p \geq (x \cdot y')_p - K = d(p, y') - K \geq d(p, y) - 2K \geq d(p, \mathfrak{w} - [p]) - 2K.$$

Since $p, y \in \mathfrak{w}$, every element of $\mathcal{W}(p|x, y)$ is transverse to \mathfrak{w} ; hence we have $\mathcal{W}(p|x, y) \subseteq \mathcal{W}(p|\pi_{\mathfrak{w}}(x))$. As $(x \cdot y)_p = \#\mathcal{W}(p|x, y)$, this concludes the proof. \square

Lemma 3.14. *Let a hyperplane $\mathfrak{w} \in \text{supp } \mathcal{S}$ and a vertex $x \in \mathfrak{w}$ be given. If $n > 2M$, there exists $g \in G_{\mathfrak{w}}$ such that $d(gx, \mathfrak{w} - [gx]) \geq \frac{n}{2} - M$.*

Proof. Since \mathcal{S} is $H_{\mathfrak{w}}$ -invariant and the action $H_{\mathfrak{w}} \curvearrowright \mathfrak{w}$ is essential, the component $[x] \subseteq \mathfrak{w}$ must have at least two boundary components. Given that any two boundary components of $[x]$ are at distance $\geq n$ from each other and $[x]$ is connected, there exists a vertex $q \in \mathfrak{w}$ such that $[q] = [x]$ and $d(q, \mathfrak{w} - [q]) \geq \frac{n}{2}$. By the definition of M , there exists an element $g \in G_{\mathfrak{w}}$ such that $d(gx, q) \leq M < \frac{n}{2}$. Hence $[gx] = [q]$ and we obtain $d(gx, \mathfrak{w} - [gx]) \geq d(q, \mathfrak{w} - [q]) - M \geq \frac{n}{2} - M$. \square

Proof of Proposition 3.12. By Lemma 2.14 there exists a hyperplane $\mathfrak{w} \in \mathcal{W}(\gamma)$ for which $\text{diam } \pi_{\mathfrak{w}}(\gamma) \leq D$. Since \mathcal{S} is full, we can replace γ with a G -translate and assume that $\mathfrak{w} \in \text{supp } \mathcal{S}$. Let x be the point of intersection between γ and \mathfrak{w} . Let $\mathcal{T} \subseteq \mathcal{G}(\mathcal{S})$ be the connected component that contains the vertex $[x] \in \mathcal{G}(\mathcal{S})$. Up to replacing γ with a $G_{\mathfrak{w}}$ -translate, Lemma 3.14 allows us to assume that $d(x, \mathfrak{w} - [x]) > 2K + D$. Then, Lemma 3.13 gives:

$$d(x, \pi_{\mathfrak{w}}(U(\mathcal{T}) - [x])) > D \geq \text{diam } \pi_{\mathfrak{w}}(\gamma).$$

We conclude that γ and $U(\mathcal{T}) - [x]$ are disjoint.

If $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ is a crooked hyperplane containing $[x]$, we have $\mathcal{T}_\Gamma = \mathcal{T}$ and:

$$\gamma \cap U(\mathcal{T}_\Gamma) = \gamma \cap [x] = \{x\}.$$

Finally, set $v = [x]$. Let Γ' be a crooked hyperplane containing hv for some $h \in H$. Then $h\gamma$ intersects $U(\mathcal{T}_{\Gamma'})$ in the single point $hx \in U(\Gamma')$. \square

We say that $g \in G$ *skewers* a crooked hyperplane Γ if we have $g\overline{C} \subseteq C$ for one of the two connected components $C \subseteq X - U(\Gamma)$, where \overline{C} denotes the closure. In this case, we have $d(gC, U(\Gamma)) > 0$, by cocompactness of $G_\Gamma^0 \curvearrowright U(\Gamma)$. We remark that, if Γ is skewered by an element of G , the subgroup $G_\Gamma^0 \leq G$ is codimension-one.

Corollary 3.15. *Consider $n > 2(M + D + 2K)$, a full n -system of switches $\mathcal{S} = (\mathcal{S}, H)$ and a crooked hyperplane $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ projecting surjectively to $H \setminus \mathcal{G}(\mathcal{S})$. Then:*

- (1) *every non-constant geodesic $\gamma \subseteq X$ has a G -translate intersecting $U(\Gamma)$ in a single point;*
- (2) *for every infinite-order element $g \in G$, a G -conjugate of a power of g skewers $U(\Gamma)$.*

Proof. Part (1) is immediate from Proposition 3.12 and the fact that $U(\Gamma)$ is contained in $U(\mathcal{T}_\Gamma)$.

Regarding part (2), we can replace g with a power and assume that g admits an axis $\gamma \subseteq X$ [Hag07]. Again by Proposition 3.12, we can replace g with a conjugate and assume that γ intersects $U(\mathcal{T}_\Gamma)$ in a single point $x \in U(\Gamma)$. It follows that one connected component $C_+ \subseteq X - U(\Gamma)$ contains the positive half of $\gamma - \{x\}$, while the other component $C_- \subseteq X - U(\Gamma)$ contains the negative half.

Let us pick $m > 0$ such that $g^m \in H$. Note that $g^m U(\Gamma)$ and $U(\Gamma)$ are disjoint. Otherwise, we would have:

$$\emptyset \neq g^m U(\Gamma) \cap U(\Gamma) = U(g^m \Gamma \cap \Gamma).$$

Hence $g^m \mathcal{T}_\Gamma = \mathcal{T}_\Gamma$ and

$$g^m x \in g^m U(\mathcal{T}_\Gamma) = U(\mathcal{T}_\Gamma),$$

contradicting the fact that $\gamma \cap U(\mathcal{T}_\Gamma) = \{x\}$.

Now, since $g^m U(\Gamma)$ and $U(\Gamma)$ are disjoint, we have $g^m U(\Gamma) \subseteq C_+$. Note moreover that $g^m C_- \cap C_- \neq \emptyset$, as both sets contain a sub-ray of γ . We conclude that $g^m \overline{C}_+ \subseteq C_+$, i.e. g^m skewers Γ . \square

Part (2) of Corollary 3.15 and Proposition 2.23 now yield:

Corollary 3.16. *Consider $n > 2(M + D + 2K)$, a full n -system of switches $\mathcal{S} = (\mathcal{S}, H)$ and a crooked hyperplane $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ projecting surjectively to $H \setminus \mathcal{G}(\mathcal{S})$. There exists an essential CAT(0) cube complex X_Γ and a proper cocompact action $G \curvearrowright X_\Gamma$ with a single orbit of hyperplanes. All hyperplane-stabilisers of $G \curvearrowright X_\Gamma$ are conjugate to $G_\Gamma \leq G$.*

The following proves Theorem A in the one-ended case.

Corollary 3.17. *Every cocompactly cubulated one-ended hyperbolic group admits an essential, cocompact cubulation with a single orbit of hyperplanes.*

Proof. Let G be a one-ended hyperbolic group with a proper cocompact action on a CAT(0) cube complex X . By Proposition 2.29, we can assume

that X is essential and hyperplane-essential. Lemma 3.5 provides a full n -system of switches \mathcal{S} with $n > 2(M + D + 2K)$, where D, K, M are as above. Every connected component $\mathcal{T} \subseteq \mathcal{G}(\mathcal{S})$ is a tree that projects surjectively to $H \backslash \mathcal{G}(\mathcal{S})$. Since X is hyperplane-essential, \mathcal{T} has no leaves. The stabiliser $H_{\mathcal{T}}$ acts cocompactly on \mathcal{T} , hence minimally. It follows that there exists a compact, two-sided subtree $A \subseteq \mathcal{T} \subseteq \mathcal{G}(\mathcal{S})$ that projects surjectively to the finite graph $H \backslash \mathcal{G}(\mathcal{S})$. By Proposition 3.11, there exists a crooked hyperplane $\Gamma \subseteq \mathcal{G}(\mathcal{S})$ containing A and we can apply Corollary 3.16. \square

Remark 3.18. The cubulation provided by Corollary 3.16 is in general not hyperplane-essential. This is due to the following configuration, in which we may find two crooked hyperplanes $\mathcal{C}_1 = U(\Gamma)$ and $\mathcal{C}_2 = gU(\Gamma)$. Denoting by \mathcal{C}_i^{\pm} the two connected components of $X - \mathcal{C}_i$, all four intersections $\mathcal{C}_1^{\pm} \cap \mathcal{C}_2^{\pm}$ might contain points arbitrarily far from $\mathcal{C}_1 \cup \mathcal{C}_2$, even if, say, the intersection $\mathcal{C}_1 \cap \mathcal{C}_2^+$ is bounded. In this case, the cubulation of G arising from $G \cdot U(\Gamma)$ has transverse hyperplanes $\mathfrak{w}_1, \mathfrak{w}_2$ arising from $\mathcal{C}_1, \mathcal{C}_2$, but it is impossible to skewer $\mathfrak{w}_1 \cap \mathfrak{w}_2$ with a hyperbolic element stabilising \mathfrak{w}_1 .

3.3. The infinitely-ended case. In this subsection, we complete the proof of Theorem A by addressing the case where G is not one-ended.

The idea is to construct ‘‘antennae’’ (in the sense of [Wis12b]) in the maximal Bass-Serre tree and to attach crooked hyperplanes constructed for the one-ended vertex groups. We now describe the construction in detail.

Let G be a cocompactly cubulated hyperbolic group. By [Dun79, Dun85], G is the fundamental group of a finite graph of groups \mathcal{G} where edge groups are finite and vertex groups are either finite or one-ended. Let $G \curvearrowright \mathcal{T}$ be the action on the corresponding Bass-Serre tree. Let G_1, \dots, G_k be the one-ended vertex groups and let v_1, \dots, v_k be their fixed vertices in \mathcal{T} .

3.3.1. The orbicomplex \overline{X} . By Proposition 1.2 in [Bow98], each G_i is quasi-convex in G and, by Theorem H in [Hag08], each G_i is cocompactly cubulated. By Proposition 2.29, we can pick a proper cocompact action of each G_i on an essential, hyperplane-essential CAT(0) cube complex X_i . Cubically subdividing if necessary, we can assume that each $G_i \curvearrowright X_i$ has no hyperplane inversions. Each finite subgroup $F \leq G_i$ has a global fixed point in X_i ; hence F preserves a cube of X_i and, since there are no hyperplane inversions, F must fix a vertex.

We now construct a specific ‘‘orbicomplex’’ \overline{X} with $G = \pi_1 \overline{X}$. We start with the disjoint union of the quotient orbicomplexes $\overline{X}_i := G_i \backslash X_i$ for $1 \leq i \leq k$, plus a singleton for every finite vertex group of \mathcal{G} . For each edge of \mathcal{G} , we add an edge connecting the corresponding orbicomplexes \overline{X}_i or singletons. If F is the associated edge group, we ensure that the attaching vertex in \overline{X}_i is the projection of a vertex of X_i that is fixed by the image of the homomorphism $F \rightarrow G_i$.

Let $G \curvearrowright \overline{X}$ be the action on the universal cover of \overline{X} . This is a proper cocompact action on an essential, hyperplane-essential CAT(0) cube complex.

We do not want to identify effectively parallel hyperplanes, as this can alter the action $G \curvearrowright \mathcal{T}$. However, the construction of X will also be required later on in the proof of Theorem D. For that purpose, we observe that Proposition 2.29 yields:

Lemma 3.19. *Let G be a cocompactly cubulated hyperbolic group and let G_1, \dots, G_k be the one-ended factors of the maximal splitting of G over finite subgroups. Given bald cubulations $G_i \curvearrowright X_i$, there exists a bald cubulation $G \curvearrowright X$ such that, for each i , the G_i -essential core of the restriction $G_i \curvearrowright X$ is G_i -equivariantly isomorphic to X_i .*

In fact, the case when G has torsion will only be needed in Section 4.1 when we prove Theorem D.

In the remainder of this section, we can and shall assume that $\overline{X} = G \backslash X$ is a genuine cube complex by passing to a torsion-free finite-index subgroup of G (whose existence is guaranteed by specialness [Ago13, HW08]) and applying the following trick:

Lemma 3.20. *Let G be a hyperbolic group with a finite-index subgroup $H \leq G$. Then, if H admits a cocompact cubulation with a single orbit of hyperplanes, so does G .*

Proof. Let $H \curvearrowright X$ be a cocompact cubulation with a single H -orbit of hyperplanes. We pick a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$, with associated halfspaces $\mathfrak{h}^+, \mathfrak{h}^-$ and halfspace-stabiliser $H_{\mathfrak{w}}^0 \leq H$. Consider the two sets $\mathfrak{H}^{\pm} := \Lambda \mathfrak{h}^{\pm} - \Lambda H_{\mathfrak{w}}^0$, where limit sets are taken in $\partial_{\infty} H = \partial_{\infty} G$. Since the action $H \curvearrowright X$ is necessarily essential, the sets \mathfrak{H}^{\pm} are nonempty and we obtain an abstract hyperplane $(H_{\mathfrak{w}}^0, \mathfrak{H})$ for both H and G .

We now apply Proposition 2.23 to the collection $\mathcal{H} = G \cdot \mathfrak{H}$. We obtain a cocompact action $G \curvearrowright X(\mathcal{H})$ with a single G -orbit of hyperplanes. If $g \in G$ has infinite-order, a power of g lies in H , where it skewers a hyperplane of X . It follows that the points $g^{\pm} \in \partial_{\infty} G = \partial_{\infty} H$ are separated by an abstract hyperplane in $H \cdot \mathfrak{H} \subseteq \mathcal{H}$. This shows that the action $G \curvearrowright X(\mathcal{H})$ is proper and, thus, the desired cubulation of G . \square

We thus assume, in the rest of the discussion, that G is **torsion-free**.

The CAT(0) cube complex X constructed right before Lemma 3.19 comes equipped with a natural G -equivariant projection $\pi: X \rightarrow \mathcal{T}$. For every open edge $e \subseteq \mathcal{T}$, the preimage $\pi^{-1}(e) \subseteq X$ consists of a single separating (open) edge of X . For every vertex $v \in \mathcal{T}$, the preimage $X_v := \pi^{-1}(v) \subseteq X$ is a convex subcomplex of X with a proper, cocompact, essential, hyperplane-essential action $G_v \curvearrowright X_v$; here $G_v \leq G$ denotes the stabiliser of the vertex $v \in \mathcal{T}$. We identify $X_{v_i} = \pi^{-1}(v_i)$ with X_i and the action $G_{v_i} \curvearrowright X_{v_i}$ with $G_i \curvearrowright X_i$.

3.3.2. Antennae. Recall that $G \curvearrowright X$ without inversions. Thus, by [Hag07], every $g \in G - \{1\}$ admits an axis $\gamma \subseteq X$. The projection $\pi(\gamma) \subseteq \mathcal{T}$ is either a single vertex (if g is elliptic in \mathcal{T}), or an axis for g in \mathcal{T} .

Let $G_{BS} \subseteq G$ be the subset of elements that admit an axis $\gamma \subseteq X$ that projects *injectively* to \mathcal{T} . In other words, this is an axis that does not contain any edges lying in one of the nontrivial fibres of the map $\pi: X \rightarrow \mathcal{T}$.

Since X is locally finite and the action $G \curvearrowright X$ is cocompact, we can find a finite collection \mathcal{P} of length-two paths in \mathcal{T} with the following property. Every element of G_{BS} has a conjugate whose axis (in \mathcal{T}) contains an element of \mathcal{P} as a sub-path.

Possibly replacing each element of \mathcal{P} by a G -translate, there exists a geodesic segment $\alpha_1 \subseteq \mathcal{T}$ that intersects each element of \mathcal{P} in its middle vertex. Replacing each vertex $v_i \in \mathcal{T}$ with a G -translate if necessary, there exists another geodesic segment $\alpha_2 \subseteq \mathcal{T}$ containing v_1, \dots, v_k (recall that the v_i are representatives of the G -orbits of vertices of \mathcal{T} with infinite G -stabiliser, as chosen at the beginning of Section 3.3). We can moreover assume that α_1 and α_2 intersect at an endpoint and only at that endpoint.

Let $\mathcal{A} \subseteq \mathcal{T}$ be the union of α_1 , α_2 and all elements of \mathcal{P} , shown in Figure 2. This is an antenna with some missing arms (cf. Section 2.1 in [Wis12b]). We also choose a finite tree $A \subseteq X$ with $\pi(A) = \mathcal{A}$.




FIGURE 2. The subtree $\mathcal{A} \subseteq \mathcal{T}$.

3.3.3. *The cube complex \overline{U} .* As in the proof of Corollary 3.17, there exist systems of switches \mathcal{S}_i in X_i and crooked hyperplanes $\Gamma_i \subseteq \mathcal{G}(\mathcal{S}_i)$ that satisfy the hypotheses of Corollary 3.15. Thus, every $g \in G_i - \{1\}$ has a conjugate of a power skewering $U(\Gamma_i) \subseteq X_i$, and every geodesic in X_i has a G_i -translate intersecting $U(\Gamma_i)$ in a single point.

Remark 3.21. Replacing Γ_i with a G_i -translate, we can assume that there exists an element $a_i \in G_i$ such that $a_i U(\Gamma_i)$ separates $U(\Gamma_i)$ from $A \cap X_i$. This is a purely technical assumption to avoid an issue in the proof of Lemma 3.23 below.

Let us write $U_i = U(\Gamma_i)$ for short and fix a shortest path $\beta_i \subseteq X_i$ from U_i to $A \cap X_i$ (which is nonempty since $v_i \in A$).

Let $L_i \leq G_i$ denote the stabiliser of U_i and let $L_i^0 \leq L_i$ be the subgroup (of index at most two) that leaves invariant both connected components of $X_i - U_i$. Set $\overline{U}_i = L_i^0 \backslash U_i$. This is a compact cube complex with a natural cubical immersion $\overline{U}_i \hookrightarrow \overline{X}'_i$ (recall that \overline{X}'_i is the cubical subdivision of \overline{X}_i).

In fact, by construction, U_i is CAT(0), because it is a tree of spaces whose vertex spaces are CAT(0) cube complexes and whose edge spaces are convex subcomplexes of the incident vertex spaces. Since $L_i^0 \curvearrowright U_i$ freely, we can identify L_i^0 with $\pi_1 \overline{U}_i$, i.e. the immersion $\overline{U}_i \hookrightarrow \overline{X}'_i$ induces the inclusion $L_i^0 \hookrightarrow G_i$ at the level of fundamental groups.




FIGURE 3. The cube complex $\bar{U} \looparrowright G \backslash X$.

We now assemble a cube complex \bar{U} as in Figure 3 by taking a copy of the tree $A \cup \beta_1 \cup \dots \cup \beta_k \subseteq X$ and attaching a copy of \bar{U}_i at the end of β_i that does not lie on A . This comes equipped with a π_1 -injective immersion $\bar{U} \looparrowright \bar{X}$.

The immersion $\bar{U} \looparrowright \bar{X}$ lifts to an embedding $U \hookrightarrow X$, where U is the universal cover of \bar{U} ; we also use the notation U for the image of this embedding, which is shown in Figure 4. (As shown in the figure, U contains the tree $A \cup \beta_1 \cup \dots \cup \beta_k \subseteq X$.) We identify the fundamental group $\pi_1 \bar{U}$ with a subgroup $L \leq G$ that stabilises U and acts cocompactly on it.

3.3.4. *Quasiconvexity of L .* For each $v_i \in \mathcal{T}$, the intersection $U \cap X_i$ is the union of U_i and all the L_i^0 -translates of the path $\beta_i \subseteq X_i$. In particular, $U \cap X_i$ is at finite Hausdorff distance from U_i , hence quasiconvex. For an arbitrary vertex $v \in \mathcal{T}$, the intersection $U \cap \pi^{-1}(v)$ is an L -translate of either some $U \cap X_i$, or some sub-path of the finite tree $A \subseteq X$. It follows that the intersections $U \cap \pi^{-1}(v)$ are uniformly quasiconvex

Observe that X decomposes as a tree of spaces with respect to the connected components of the various sets $\pi^{-1}(v)$, and U is also a tree of spaces with respect to the components of $U \cap \pi^{-1}(v)$. It follows from uniform quasiconvexity of the latter sets that U is quasiconvex in X . Since the action $L \curvearrowright U$ is proper and cocompact, $L \leq G$ is a quasiconvex subgroup.




FIGURE 4. The subcomplex $U \subseteq X$.

3.3.5. *Cutting using L .* We now proceed to analyse the connected components of $X - U$. Note that the projection $\mathcal{T}_U := \pi(U) \subseteq \mathcal{T}$ is an L -invariant subtree and that the quotient $L \backslash \mathcal{T}_U$ is naturally identified with \mathcal{A} . We define an L -invariant map $\mathfrak{p}: X \rightarrow \mathcal{A}$ by composing the projection $\pi: X \rightarrow \mathcal{T}$ with the nearest-point projection $\mathcal{T} \rightarrow \mathcal{T}_U$ and, finally, the quotient projection $\mathcal{T}_U \rightarrow L \backslash \mathcal{T}_U \cong \mathcal{A}$.

We also consider the L -invariant convex subset $\mathcal{U} = \pi^{-1}(\mathcal{T}_U) \subseteq X$, which contains U , and the L -equivariant gate-projection $p_{\mathcal{U}}: X \rightarrow \mathcal{U}$. Observe that $\mathfrak{p} \circ p_{\mathcal{U}} = \mathfrak{p}$. We denote by C_i^+ and C_i^- the two connected components of $X_i - U_i$ and by $\mathcal{C}_i^{\pm} \subseteq \mathcal{U}$ the unions of all their L -translates.

Lemma 3.22.

- (1) For every vertex $w \in \mathcal{A}$, the set $\mathfrak{p}^{-1}(w) - U$ is an L -invariant union of connected components of $X - U$. Every connected component of $X - U$ is contained in one of these sets.
- (2) If C is a connected component of $\mathfrak{p}^{-1}(v_i) - U$, the projection $p_{\mathcal{U}}(C)$ is contained in either \mathcal{C}_i^+ or \mathcal{C}_i^- (and this property is L -invariant).

Proof. Recall that the map $\mathfrak{p}: X \rightarrow \mathcal{A}$ is continuous and L -invariant. If $x \in \mathcal{A}$ is a point in the interior of an edge, we have $\mathfrak{p}^{-1}(x) \subseteq U$. Hence $X - U$ is a disjoint union of the finitely many, closed, L -invariant subsets $\mathfrak{p}^{-1}(w) - U$, where $w \in \mathcal{A}$ is a vertex. This proves part (1).

Recall that the L -stabiliser of $X_i \subseteq X$ is L_i^0 . Given that L_i^0 does not swap the two sides of $U_i \subseteq X_i$, the sets \mathcal{C}_i^+ and \mathcal{C}_i^- are disjoint. Since $\mathfrak{p} \circ p_{\mathcal{U}} = \mathfrak{p}$, a point $x \in X$ lies in $\mathfrak{p}^{-1}(v_i)$ for some $1 \leq i \leq k$ if and only if the projection $p_{\mathcal{U}}(x)$ lies in an L -translate of the subset $X_i \subseteq X$. In particular, $p_{\mathcal{U}}(\mathfrak{p}^{-1}(v_i) - U) \subseteq \mathcal{C}_i^+ \sqcup \mathcal{C}_i^-$. Observing that the \mathcal{C}_i^{\pm} are open in the union $\mathcal{C}_i^+ \sqcup \mathcal{C}_i^-$, we obtain part (2). \square

Recall that each element $P \in \mathcal{P}$ is a length-two sub-path $P \subseteq \mathcal{A}$; we denote by $z_P^{\pm} \in \mathcal{A}$ its two endpoints. Let $\mathcal{H}^- \subseteq X$ be the union of all connected components of $X - U$ that are contained either in $\mathfrak{p}^{-1}(z_P^-)$ for some $P \in \mathcal{P}$, or in $p_{\mathcal{U}}^{-1}(\mathcal{C}_i^-)$ for some $1 \leq i \leq k$. We also set $\mathcal{H}^+ := X - (\mathcal{H}^- \sqcup U)$. In particular, \mathcal{H}^+ contains all connected components of $X - U$ that are either contained in some $\mathfrak{p}^{-1}(z_P^+)$ or in some $p_{\mathcal{U}}^{-1}(\mathcal{C}_i^+)$.

We obtain an L -invariant partition $X = \mathcal{H}^- \sqcup U \sqcup \mathcal{H}^+$. By part (3) of Lemma 2.18, this gives rise to an abstract hyperplane (L, \mathfrak{H}) (cf. Definition 2.19), where $\mathfrak{H}^{\pm} = \Lambda(U \sqcup \mathcal{H}^{\pm}) - \Lambda U$. Note that the sets \mathfrak{H}^{\pm} are both nonempty as, for each $1 \leq i \leq k$, the intersection $p_{\mathcal{U}}(U) \cap X_i$ is at finite Hausdorff distance from U_i and $\partial_{\infty} \mathcal{C}_i^{\pm} - \partial_{\infty} U_i \neq \emptyset$.

Now, applying Proposition 2.23 to the collection of abstract hyperplanes $G \cdot \mathfrak{H}$, we obtain a cocompact, essential G -action on a CAT(0) cube complex with a single orbit of hyperplanes. In order to complete the proof of Theorem A, we are only left to show that this action is proper. By part (2) of Proposition 2.23, this amounts to the following:

Lemma 3.23. *Every $g \in G - \{1\}$ has a conjugate h with $h^+ \in \mathfrak{H}^+$ and $h^- \in \mathfrak{H}^-$ (or $h^- \in \mathfrak{H}^+$ and $h^+ \in \mathfrak{H}^-$).*

Proof. There are two cases to consider, depending on whether g lies in G_{BS} .

If $g \in G_{BS}$, we can replace g with a conjugate so that a path $P \in \mathcal{P}$ is contained in its axis $\gamma \subseteq \mathcal{T}$. Any axis $\gamma' \subseteq X$ will satisfy $\pi(\gamma') = \gamma$ and $\pi(\gamma' \cap U) = P$. It follows that $\gamma' - U$ contains sub-rays lying in \mathcal{H}^+ and \mathcal{H}^- . Without loss of generality, we have $g^+ \in \Lambda(U \sqcup \mathcal{H}^+)$ and $g^- \in \Lambda(U \sqcup \mathcal{H}^-)$. Since no power of g stabilises U , Lemma 2.16 guarantees that $g^\pm \notin \Lambda U$. Thus $g^+ \in \mathfrak{H}^+$ and $g^- \in \mathfrak{H}^-$, as required.

Suppose instead that $g \notin G_{BS}$. Then, replacing g with a conjugate, we can assume that g admits an axis $\gamma \subseteq X$ that intersects one of the spaces $X_i \subseteq X$ in a non-trivial geodesic. By our choice of the crooked hyperplanes Γ_i , the geodesic $\gamma \cap X_i$ intersects U_i in a single point. Thus, either $\gamma \subseteq X_i$ and C_i^\pm each contain a sub-ray of γ , or $\gamma \cap X_i$ is a finite segment with one endpoint in C^+ and one endpoint in C^- .

We conclude that the sets \mathcal{H}^\pm each contain a sub-ray of γ , unless possibly if $\gamma \cap X_i$ is a segment with one of its endpoints in $A \cap X_i$. This issue can be avoided simply by conjugating g by the element $a_i \in G_i$ mentioned in Remark 3.21.

Without loss of generality, we have $g^+ \in \Lambda(U \sqcup \mathcal{H}^+)$ and $g^- \in \Lambda(U \sqcup \mathcal{H}^-)$ once again. It is clear that no power of g stabilises U . As before, we conclude that $g^+ \in \mathfrak{H}^+$ and $g^- \in \mathfrak{H}^-$. \square

We have proved:

Corollary 3.24. *Every cocompactly cubulated hyperbolic group admits a cocompact, essential cubulation with a single orbit of hyperplanes.*

4. THE NUMBER OF BALD CUBULATIONS.

4.1. Bald cubulations of hyperbolic groups. Let G be a cocompactly cubulated, non-elementary hyperbolic group. We first assume that G is one-ended.

Lemma 4.1. *Suppose that G admits only finitely many bald cubulations up to equivariant cubical isomorphism. Then there exists a bald cubulation $G \curvearrowright X$ and a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$ with the following property. Let $G \curvearrowright Y$ be a bald cubulation. Then for each hyperplane $\mathfrak{u} \in \mathcal{W}(Y)$, $\Lambda \mathfrak{u}$ is not properly contained in $\Lambda \mathfrak{w}$.*

Proof. If the lemma did not hold, there would exist an infinite sequence of bald cubulations $G \curvearrowright X_n$ and hyperplanes $\mathfrak{w}_n \in \mathcal{W}(X_n)$ such that $\Lambda \mathfrak{w}_{n+1} \subsetneq \Lambda \mathfrak{w}_n$. Since G is virtually special, it has a finite-index torsion-free subgroup $H \leq G$. Given that G has only finitely many bald cubulations, and each has only finitely many H -orbits of hyperplanes, there exist $h \in H$ and $m < n$ with $h\Lambda \mathfrak{w}_m = \Lambda \mathfrak{w}_n \subsetneq \Lambda \mathfrak{w}_m$.

Observing that h has infinite order, it follows that $h^+ \in \Lambda \mathfrak{w}_m$. Hence, by Lemma 2.16, we have $\Lambda(\langle h \rangle \cap \text{Stab}_G(\mathfrak{w}_m)) = \{h^\pm\} \cap \Lambda \mathfrak{w}_m \neq \emptyset$. It follows

that $\langle h \rangle \cap \text{Stab}_G(\mathfrak{w}_m)$ is infinite, i.e. a positive power of h stabilises $\Lambda \mathfrak{w}_m$. This is a contradiction. \square

Suppose G has finitely many bald cubulations. Let $G \curvearrowright X$ and $\mathfrak{w} \in \mathscr{W}(X)$ be the cubulation and hyperplane provided by Lemma 4.1. By essentiality of $G \curvearrowright X$, there exists an element $g \in G$ skewering \mathfrak{w} . By [Hag07], we can replace g with a power to ensure that it admits an axis $\gamma \subseteq X$. Let p be the vertex of \mathfrak{w} that lies on γ . Let K be the constant in Proposition 3.8 and define M as at the beginning of Section 3.2.

For every $n > 2M$, let $\mathcal{S}_n = (\mathcal{S}_n, H_n)$ be a full n -system of switches (its existence is guaranteed by Lemma 3.5). Replacing \mathcal{S}_n with

$$k \cdot \mathcal{S}_n := (k\mathcal{S}_n, kH_nk^{-1}) = (k\mathcal{S}_n, H_n)$$

for some $k \in G$, we can assume that $\mathfrak{w} \in \text{supp } \mathcal{S}_n$. Let $[p]_n \in \text{comp } \mathcal{S}_n$ denote the component that contains p . Again replacing \mathcal{S}_n with $k \cdot \mathcal{S}_n$ for some $k \in G_{\mathfrak{w}}$, Lemma 3.14 enables us to assume that $d(p, \mathfrak{w} - [p]_n) \geq \frac{n}{2} - M$.

Now, Proposition 3.11 guarantees the existence of a crooked hyperplane $\Gamma_n \subseteq \mathcal{G}(\mathcal{S}_n)$ that contains the vertex $[p]_n \in \mathcal{G}(\mathcal{S}_n)$, but not the entire subtree $\mathcal{G}(\mathfrak{w}) \subseteq \mathcal{G}(\mathcal{S}_n)$. We will need the following observations.

Lemma 4.2.

- (1) We have $\partial_{\infty} \mathfrak{w} - \partial_{\infty} U(\Gamma_n) \neq \emptyset$.
- (2) If $n > 2(2K + M + \text{diam } \pi_{\mathfrak{w}}(\gamma))$, then $U(\Gamma_n)$ is skewered by a power of g .
- (3) For every open neighbourhood \mathcal{V} of $\Lambda \mathfrak{w}$ in $\partial_{\infty} G$, there exists \bar{n} such that $\Lambda U(\Gamma_n) \subseteq \mathcal{V}$ for all $n \geq \bar{n}$.

Proof. Since $\mathcal{G}(\mathfrak{w}) \not\subseteq \Gamma_n$, there exists a ray r_n contained in the subtree $\mathcal{G}(\mathfrak{w}) \subseteq \mathcal{G}(\mathcal{S}_n)$ and disjoint from Γ_n . As in Remark 3.10, the corresponding standard path $\tilde{r}_n \subseteq \mathfrak{w} \subseteq X$ is a quasi-geodesic ray defining a point of $\partial_{\infty} \mathfrak{w} - \partial_{\infty} U(\Gamma_n)$. This shows part (1). Part (2) follows from Lemma ??, since $d(p, \mathfrak{w} - [p]_n) \geq \frac{n}{2} - M > 2K + \text{diam } \pi_{\mathfrak{w}}(\gamma)$.

Finally, Lemma 3.13 yields $d(p, \pi_{\mathfrak{w}}(U(\Gamma_n) - \mathfrak{w})) \geq \frac{n}{2} - M - 2K$, which diverges for $n \rightarrow +\infty$. Recall that, by Remark 3.10, points of $\partial_{\infty} U(\Gamma_n)$ are represented by uniform quasi-geodesic rays contained in $U(\Gamma_n)$. It follows that the limit sets $\partial_{\infty} U(\Gamma_n)$ Hausdorff-converge to $\partial_{\infty} \mathfrak{w}$ with respect to the visual metric on $\partial_{\infty} X$ determined by the point p . This proves part (3). \square

We are now ready to prove Theorem D.

Proof of Theorem D. Let us assume for a moment that the theorem has been proved in the one-ended case. If G is virtually free, then the theorem follows from Lemma 4.3 below. If G is neither virtually free nor one-ended, G has at least one one-ended factor in its maximal splitting over finite subgroups and the theorem follows from Lemma 3.19. It remains to handle the one-ended case.

The one-ended case: Suppose now that G is one-ended. Let the cubulation $G \curvearrowright X$ and the crooked hyperplanes Γ_n be those constructed above.

We denote by \mathcal{H} the collection of abstract hyperplanes arising from the hyperplanes of X , and by \mathfrak{H}_n the abstract hyperplane determined by $U(\Gamma_n)$.

Let $G \curvearrowright X_n$ be the essential cubulation arising from the set $\mathcal{H} \cup G \cdot \mathfrak{H}_n$ via Proposition 2.23. Let $G \curvearrowright (X_n)_\bullet$ be the bald cubulation provided by Proposition 2.29. Recall that, by part (2) of Lemma 4.2, the hyperplane of X_n corresponding to \mathfrak{H}_n is skewered by a power of g for all large n . Thus, part (2) of Proposition 2.29 guarantees that a power of g also skewers a hyperplane $\mathbf{u}_n \in \mathscr{W}((X_n)_\bullet)$ with $\Lambda \mathbf{u}_n \subseteq \Lambda U(\Gamma_n)$.

In each bald cubulation of G , only finitely many $\langle g \rangle$ -orbits of hyperplanes are skewered by a power of g . If G admitted only finitely many bald cubulations, infinitely many limit sets $\Lambda \mathbf{u}_n$ would lie in the same $\langle g \rangle$ -orbit. There are two cases to consider. Note that $\Lambda \mathbf{u}_n \neq \emptyset$ for all n , as G is one-ended.

Case 1: *there exist two diverging sequences (a_k) and (b_k) with the property that $g^{b_k} \Lambda \mathbf{u}_{a_0} = \Lambda \mathbf{u}_{a_k}$.* Since a power of g skewers \mathbf{u}_{a_0} , the subsets $g^{b_k} \Lambda \mathbf{u}_{a_0}$ Hausdorff-converge to g^+ . This contradicts part (3) of Lemma 4.2, as $\Lambda \mathbf{u}_{a_k} \subseteq \Lambda U(\Gamma_{a_k})$ and $g^+ \notin \Lambda \mathfrak{w}$.

Case 2: *there exists a diverging sequence (a_k) such that $\Lambda \mathbf{u}_{a_k}$ is constant.* Call Δ this subset of $\partial_\infty G$. Note that $\Delta \subseteq \bigcap_k \Lambda U(\Gamma_{a_k})$, which is contained in $\Lambda \mathfrak{w}$ by part (3) of Lemma 4.2. By minimality of $\Lambda \mathfrak{w}$, we have $\Delta = \Lambda \mathfrak{w}$. This implies that $\Lambda \mathfrak{w} \subseteq \Lambda U(\Gamma_{a_k})$, contradicting part (1) of Lemma 4.2. \square

Lemma 4.3. *Let G be a group with $\text{Out}(G)$ infinite, and suppose that G admits a proper, cocompact action on a CAT(0) cube complex. Then G admits infinitely many bald cubulations, no two of which are G -equivariantly isomorphic.*

Proof. By Proposition 2.29, G admits a bald cubulation $\rho: G \rightarrow \text{Aut}(X)$. Each $\phi \in \text{Aut}(G)$ defines an action $\rho \circ \phi: G \rightarrow \text{Aut}(X)$, which is again a bald cubulation. For simplicity, given $g \in G$ and $x \in X$, we write gx to mean $\rho(g)(x)$.

Let $\{x_1, \dots, x_k\}$ contain exactly one vertex of X from each $\rho(G)$ -orbit. Let $g_1, \dots, g_m \in G$ generate G . Consider the constants $s = \max_j d(x_1, x_j)$ and $r = \max_i d(x_1, g_i x_1)$.

Let $\phi \in \text{Aut}(G)$ and suppose that $\rho \circ \phi$ and ρ are equivalent. By definition, there is $\iota \in \text{Aut}(X)$ such that $\iota(hx) = \phi(h)\iota(x)$ for all $h \in G, x \in X$.

Choose $h \in G$ so that $\iota(x_1) = hx_j$ for some $j \leq k$. Then, for each $i \leq m$, we have $\phi(g_i)hx_j = \phi(g_i)(\iota(x_1)) = \iota(g_i x_1)$. So $d(\phi(g_i)hx_j, hx_j) = d(g_i x_1, x_1) \leq r$, from which the triangle inequality gives $d(h^{-1}\phi(g_i)hx_1, x_1) \leq r + 2s$ for all $i \leq m$. Hence we can re-choose ϕ in its outer class so that each $\phi(g_i)$ displaces x_1 by at most a distance depending only on ρ and the (fixed) generating set of G . There are finitely many possible choices for each $\phi(g_i)$, and hence there are only finitely many $\Phi \in \text{Out}(G)$ such that Φ has a representative $\phi \in \text{Aut}(G)$ with ρ and $\rho \circ \phi$ equivalent.

Thus, if there were only finitely many equivalence classes of actions $\rho: G \rightarrow \text{Aut}(X)$, we would have that $\text{Out}(G)$ is finite, a contradiction. \square

Remark 4.4. In order to deal with general virtually free groups in the proof of Theorem D, one might be tempted to behave as in Section 3.3.2: work in a torsion-free finite-index subgroup, and then use the same idea as Lemma 3.20 to cubulate the original group. Unfortunately, this does not preserve hyperplane-essentiality.

4.2. Groups with few bald cubulations. In this subsection, we prove Proposition C (cf. Proposition 4.8 below) and the following result mentioned in the introduction.

Proposition 4.5. *Let X be an essential, hyperplane-essential CAT(0) cube complex endowed with a proper, cocompact action of \mathbb{Z}^n . Then X is isomorphic to the standard tiling of \mathbb{R}^n .*

Proof. By the Cubical Flat Torus theorem [WW17, Theorem 3.6], there exists an invariant convex subcomplex $Y \subseteq X$ that splits as product of quasi-lines C_1, \dots, C_n . Since the \mathbb{Z}^n -action is essential, we have $Y = X$. Since X is essential and hyperplane-essential, so is each C_j . Every hyperplane of an essential quasi-line is bounded. Thus, since each C_j is hyperplane-essential, it follows that $C_j \cong \mathbb{R}$. In conclusion, $X \cong \mathbb{R}^n$. \square

Before proving Proposition 4.8, we need to obtain a couple of lemmas.

Lemma 4.6. *Let G act cocompactly on a bald cube complex X with no \mathbb{R} -factors. If $A \subseteq X^{(0)}$ is a G -invariant, nonempty median subalgebra, then $A = X^{(0)}$.*

Proof. Since the action $G \curvearrowright X$ is essential, every halfspace of X intersects A nontrivially. We obtain a G -equivariant map $r_A: \mathcal{H}(X) \rightarrow \mathcal{H}(A)$ that takes each halfspace of X to its intersection with A . By Lemma 6.5 in [Bow13], this map is surjective. By Lemma 2.3, every halfspace of X is at finite Hausdorff distance from its intersection with A . By definition, no two halfspaces of X are at finite Hausdorff distance, so the fibres of r_A are singletons. We conclude that r_A is a bijection, hence $A = X^{(0)}$. \square

Lemma 4.7. *Let G be a group such that no finite-index subgroup of G admits a nontrivial additive homomorphism to \mathbb{R} . Let G act cocompactly on a proper, unbounded CAT(0) space \mathcal{X} . Then every G -orbit in the visual boundary $\partial_\infty \mathcal{X}$ is infinite.*

Proof. By our assumptions, the visual boundary $\partial_\infty \mathcal{X}$ is nonempty. Suppose for the sake of contradiction that G has a finite orbit in $\partial_\infty \mathcal{X}$, so a finite-index subgroup $G_0 \leq G$ fixes a point $\xi \in \partial_\infty \mathcal{X}$.

Let $b_\xi: \mathcal{X} \rightarrow \mathbb{R}$ be any Busemann function determined by ξ . Given any $x \in \mathcal{X}$, the map $\phi: G_0 \rightarrow \mathbb{R}$ defined by $\phi(g) = b_\xi(gx) - b_\xi(x)$ is easily seen to be an additive homomorphism. By our assumption on G , the map ϕ must vanish identically. Hence G_0 leaves invariant each horosphere around ξ , contradicting cocompactness of $G_0 \curvearrowright \mathcal{X}$. \square

Proposition 4.8. *For $i = 1, 2$, let T_i be locally finite trees with all vertices of degree ≥ 3 . Let $U_i \leq \text{Aut}(T_i)$ be closed, locally primitive subgroups generated by edge stabilisers and satisfying Tits' independence property. Then, for any uniform lattice $\Gamma \leq U_1 \times U_2$ with dense projections to U_1 and U_2 , the standard action $\Gamma \curvearrowright T_1 \times T_2$ is the only bald cubulation of Γ .*

Proof. Let $\Gamma \curvearrowright X$ be a bald cubulation. Let $X_1 \times \dots \times X_k$ be the de Rham decomposition² of X and let $\Gamma_0 \leq \Gamma$ be a finite-index subgroup leaving each factor invariant. Each X_j is a locally finite, bald cube complex endowed with a cocompact Γ_0 -action.

Observe that U_1 and U_2 are simple groups by the argument in [Tit70] (see e.g. Theorem 3.3 in [CDM11]). Theorem 0.8 in [Sha00] thus implies that every additive homomorphism $\Gamma_0 \rightarrow \mathbb{R}$ vanishes identically, and the same holds for any finite-index subgroup of Γ_0 . Lemma 4.7 then yields that there are no finite Γ_0 -orbits in the visual boundaries $\partial_\infty X_j$.

Again by simplicity, U_i has no finite-index open subgroups, so the projection of Γ_0 to U_i is dense. By Theorem 1.5 in [CFI16], each action $\Gamma_0 \curvearrowright X_j$ extends to a continuous action of some U_{i_j} on a Γ_0 -invariant median subalgebra³ $A_j \subseteq X_j^{(0)}$. By Lemma 4.6, we actually have $A_j = X_j^{(0)}$.

Observe that each $U_{i_j} \curvearrowright X_j$ is cocompact and essential, since so is the Γ_0 -action. Thus, hyperplane-stabilisers are proper, open subgroups of U_{i_j} and they act cocompactly on the respective hyperplanes by Lemma 2.3. Theorem A in [CDM11] shows that all hyperplane-stabilisers of $U_{i_j} \curvearrowright X_j$ are compact, which means that all hyperplanes of each X_j are compact. Since X_j is hyperplane-essential, it must be a tree. Finally, by Lemma 1.4.7⁴ in [BMZ09], X_j must be U_{i_j} -equivariantly isomorphic to T_{i_j} . We conclude that $k = 2$ and that X is Γ -equivariantly isomorphic to $T_1 \times T_2$. \square

REFERENCES

- [Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.
- [BF95] Mladen Bestvina and Mark Feighn. Stable actions of groups on real trees. *Invent. Math.*, 121(2):287–321, 1995.
- [BF19a] Jonas Beyrer and Elia Fioravanti. Cross ratios and cubulations of hyperbolic groups. *arXiv:1810.08087v2*, 2019.
- [BF19b] Jonas Beyrer and Elia Fioravanti. Cross ratios on CAT (0) cube complexes and marked length-spectrum rigidity. *arXiv preprint arXiv:1903.02447*, 2019.
- [BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

²We immediately have $k = 2$ by quasi-flat rank considerations, but this is not necessary in the proof.

³The stronger statement appearing as Theorem 1.5 in [CFI16] is not true in general. See Section 4 of [Fio17] for a discussion (in particular Theorem 4.4 and Example 4.7).

⁴This appears as Lemma 3.7 in the unpublished version of the article available online.

- [BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.
- [BM00a] Marc Burger and Shahar Mozes. Groups acting on trees: from local to global structure. *Inst. Hautes Études Sci. Publ. Math.*, (92):113–150 (2001), 2000.
- [BM00b] Marc Burger and Shahar Mozes. Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.*, (92):151–194 (2001), 2000.
- [BMZ09] Marc Burger, Shahar Mozes, and Robert J. Zimmer. Linear representations and arithmeticity of lattices in products of trees. In *Essays in geometric group theory*, volume 9 of *Ramanujan Math. Soc. Lect. Notes Ser.*, pages 1–25. Ramanujan Math. Soc., Mysore, 2009.
- [Bou97] Marc Bourdon. Immeubles hyperboliques, dimension conforme et rigidité de Mostow. *Geom. Funct. Anal.*, 7(2):245–268, 1997.
- [Bow98] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [Bow13] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013.
- [BW12] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.
- [Cap19] Pierre-Emmanuel Caprace. Finite and infinite quotients of discrete and indiscrete groups. In *Groups St Andrews 2017 in Birmingham*, volume 455 of *London Math. Soc. Lecture Note Ser.*, pages 16–69. Cambridge Univ. Press, Cambridge, 2019.
- [CDM11] Pierre-Emmanuel Caprace and Tom De Medts. Simple locally compact groups acting on trees and their germs of automorphisms. *Transform. Groups*, 16(2):375–411, 2011.
- [CFI16] Indira Chatterji, Talia Fernós, and Alessandra Iozzi. The median class and superrigidity of actions on CAT(0) cube complexes. *J. Topol.*, 9(2):349–400, 2016. With an appendix by Pierre-Emmanuel Caprace.
- [Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.
- [CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21(4):851–891, 2011.
- [Duf12] Guillaume Dufour. *Cubulations de variétés hyperboliques compactes*. PhD thesis, Paris 11, 2012.
- [Dun79] Martin J. Dunwoody. Accessibility and groups of cohomological dimension one. *Proc. London Math. Soc. (3)*, 38(2):193–215, 1979.
- [Dun85] Martin J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.
- [Fio17] Elia Fioravanti. Superrigidity of actions on finite rank median spaces. *arXiv:1711.07737v1*, 2017.
- [For02] Max Forester. Deformation and rigidity of simplicial group actions on trees. *Geom. Topol.*, 6:219–267, 2002.
- [Gen16] Anthony Genevois. Coning-off CAT(0) cube complexes. *arXiv:1603.06513v1*, 2016.
- [GL07] Vincent Guirardel and Gilbert Levitt. Deformation spaces of trees. *Groups Geom. Dyn.*, 1(2):135–181, 2007.
- [GMRS98] Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.
- [GPR12] Mauricio Gutierrez, Adam Piggott, and Kim Ruane. On the automorphisms of a graph product of abelian groups. *Groups Geom. Dyn.*, 6(1):125–153, 2012.
- [Gur07] Dan Guralnik. Coarse decompositions of boundaries for CAT(0) groups. *arXiv:math/0611006v2*, 2007.

- [Hag07] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. *arXiv:0705.3386v1*, 2007.
- [Hag08] Frédéric Haglund. Finite index subgroups of graph products. *Geom. Dedicata*, 135:167–209, 2008.
- [Hag20] Mark Hagen. Large facing tuples and a strengthened sector lemma. *arXiv preprint arXiv:2005.09536*, 2020.
- [Hea21] Brendan Burns Healy. Acylindrical hyperbolicity of $\text{Out}(W_n)$. *Topology Proc.*, 58:23–36, 2021.
- [HS18] Mark F. Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *arXiv:1609.01313v2*, 2018.
- [HT19] Mark F. Hagen and Nicholas W. M. Touikan. Panel collapse and its applications. *Groups, Geometry, and Dynamics*, 13(4):1285–1334, 2019.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.
- [HW09] G. Christopher Hruska and Daniel T. Wise. Packing subgroups in relatively hyperbolic groups. *Geom. Topol.*, 13(4):1945–1988, 2009.
- [HW14] G. C. Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. *Compos. Math.*, 150(3):453–506, 2014.
- [HW15a] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the general case. *Geom. Funct. Anal.*, 25(1):134–179, 2015.
- [HW15b] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. *Invent. Math.*, 199(2):293–331, 2015.
- [HW16] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *Duke Math. J.*, 165(9):1753–1813, 2016.
- [KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.
- [Lev05] Gilbert Levitt. Automorphisms of hyperbolic groups and graphs of groups. *Geometriae Dedicata*, 114(1):49–70, 2005.
- [LS07] Markus Lohrey and Géraud Sénizergues. When is a graph product of groups virtually-free? *Comm. Algebra*, 35(2):617–621, 2007.
- [Nib02] Graham A. Niblo. The singularity obstruction for group splittings. *Topology Appl.*, 119(1):17–31, 2002.
- [NR03] Graham A. Niblo and Lawrence D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.
- [OW11] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than $1/6$. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.
- [Pet97] Martin Pettet. Virtually free groups with finitely many outer automorphisms. *Transactions of the American Mathematical Society*, 349(11):4565–4587, 1997.
- [Rol98] Martin A. Roller. Poc sets, median algebras and group actions. An extended study of Dunwoody’s construction and Sageev’s theorem. Preprint, University of Southampton, 1998.
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.
- [Sag14] Michah Sageev. CAT(0) cube complexes and groups. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 7–54. Amer. Math. Soc., Providence, RI, 2014.
- [Sha00] Yehuda Shalom. Rigidity of commensurators and irreducible lattices. *Invent. Math.*, 141(1):1–54, 2000.
- [Thu79] William P Thurston. *The geometry and topology of three-manifolds*. Princeton University Princeton, NJ, 1979.

- [Tit70] Jacques Tits. Sur le groupe des automorphismes d'un arbre. In *Essays on topology and related topics (Mémoires dédiés à Georges de Rham)*, pages 188–211. 1970.
- [Wis96] Daniel T. Wise. *Non-positively curved squared complexes: Aperiodic tilings and non-residually finite groups*. ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.)—Princeton University.
- [Wis04] Daniel T. Wise. Cubulating small cancellation groups. *Geom. Funct. Anal.*, 14(1):150–214, 2004.
- [Wis07] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.
- [Wis11] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. *Preprint*, 2011.
- [Wis12a] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.
- [Wis12b] Daniel T. Wise. Recubulating free groups. *Israel J. Math.*, 191(1):337–345, 2012.
- [WW17] Daniel T. Wise and Daniel J. Woodhouse. A cubical flat torus theorem and the bounded packing property. *Israel J. Math.*, 217(1):263–281, 2017.

MAX PLANCK INSTITUTE FOR MATHEMATICS, BONN, GERMANY

Email address: `fioravanti@mpim-bonn.mpg.de`

SCHOOL OF MATHEMATICS, UNIVERSITY OF BRISTOL, BRISTOL, UK

Email address: `markfhagen@posteo.net`