



ANNALES DE L'INSTITUT FOURIER

Talia FERNÓS, MAX FORESTER & JING TAO

Effective quasimorphisms on right-angled Artin groups

Tome 69, n° 4 (2019), p. 1575-1626.

http://aif.centre-mersenne.org/item/AIF_2019__69_4_1575_0

© Association des Annales de l'institut Fourier, 2019,

Certains droits réservés.



Cet article est mis à disposition selon les termes de la licence
CREATIVE COMMONS ATTRIBUTION – PAS DE MODIFICATION 3.0 FRANCE.
<http://creativecommons.org/licenses/by-nd/3.0/fr/>



EFFECTIVE QUASIMORPHISMS ON RIGHT-ANGLED ARTIN GROUPS

by Talia FERNÓS, Max FORESTER & Jing TAO (*)

ABSTRACT. — We construct families of quasimorphisms on many groups acting on $\text{CAT}(0)$ cube complexes. These quasimorphisms have a uniformly bounded defect of 12, and they “see” all elements that act hyperbolically on the cube complex. We deduce that all such elements have stable commutator length at least $1/24$.

The group actions for which these results apply include the standard actions of right-angled Artin groups on their associated $\text{CAT}(0)$ cube complexes. In particular, every non-trivial element of a right-angled Artin group has stable commutator length at least $1/24$.

These results make use of some new tools that we develop for the study of group actions on $\text{CAT}(0)$ cube complexes: the essential characteristic set and equivariant Euclidean embeddings.

RÉSUMÉ. — Nous construisons des nouvelles familles de quasi-morphismes sur de nombreux groupes agissant sur des complexes cubiques $\text{CAT}(0)$. Ces quasimorphismes ont leur défaut majoré par 12, et sont suffisamment nombreux pour « voir » tous les éléments qui agissent de manière hyperbolique sur le complexe cubique. Nous déduisons que la longueur stable des commutateurs de tous ces éléments est minorée par $1/24$.

Les actions pour lesquelles ces résultats sont vérifiés comprennent l'action standard d'un groupe d'Artin à angles droits sur son complexe cubique associé. En particulier, la longueur stable des commutateurs de tout élément non trivial d'un groupe d'Artin à angles droits est minorée par $1/24$.

Ces résultats reposent sur de nouveaux outils que nous développons pour étudier les actions de groupes sur des complexes cubiques $\text{CAT}(0)$: l'ensemble caractéristique essentiel et les plongements euclidiens équivariants.

Keywords: quasimorphism, stable commutator length, bounded cohomology, $\text{CAT}(0)$ cube complex, right-angled Artin group, median property.

2010 *Mathematics Subject Classification:* 20F65, 20F67, 57M07.

(*) Fernós was partially supported by NSF award DMS-1312928, Forester by NSF award DMS-1105765, and Tao by NSF awards DMS-1311834, DMS-1611758, DMS-1651963. The authors thank the referee for many helpful comments and for pointing out an error in the original proof of Proposition 3.20.

1. Introduction

In this paper, we construct quasimorphisms on groups that admit actions on CAT(0) cube complexes. Our emphasis is on finding quasimorphisms that are both *efficient* and *effective*. By “efficient” we mean that the quasimorphisms have low defect. By “effective” we mean that the quasimorphisms take non-zero values on specified elements of the group. These two qualities, taken together, allow one to establish lower bounds for stable commutator length (scl) in the group.

According to Bavard Duality [1], if φ is a homogeneous quasimorphism of defect at most D and $\varphi(g) \geq 1$, then $\text{scl}(g) \geq 1/2D$. Thus, for the strongest bound on scl, one needs to find effective quasimorphisms with the smallest possible defect.

The quasimorphisms we define have similarities with the “non-overlapping” counting quasimorphisms of Epstein and Fujiwara [15], which in turn are a variation of the Brooks counting quasimorphisms on free groups [6]. If X is a CAT(0) cube complex, there is a notion of a *tightly nested segment* of half-spaces in X . If G acts on X *non-transversely* (see Definition 4.1), then for each tightly nested segment γ there is an associated counting quasimorphism φ_γ . This function counts non-overlapping copies (or G -translates) of γ and $\bar{\gamma}$ inside characteristic subcomplexes of elements of G . Using the *median property* of CAT(0) cube complexes, we show that φ_γ has defect at most 6, and therefore its homogenization $\widehat{\varphi}_\gamma$ has defect at most 12. (Note that this bound is independent of both the length of γ and the dimension of X .)

We now have a large supply of efficient quasimorphisms, but it is by no means clear that any of them are non-trivial. Our main task, given an element $g \in G$, is to find a tightly nested segment γ such that $\widehat{\varphi}_\gamma(g) \geq 1$. This will only be possible for suitable elements g ; for instance, if g is conjugate to g^{-1} , then $\text{scl}(g) = 0$ and every homogeneous quasimorphism vanishes on g .

For our main result we consider cube complexes with group actions that have properties in common with the standard actions of right angled Artin groups on their associated CAT(0) cube complexes. These are called *RAAG-like* actions; see Section 7 and Definition 7.1. Our main theorem is that for such actions, the desired segments γ can be found for every hyperbolic element g . Using Bavard Duality, we obtain:

THEOREM A. — *Let X be a CAT(0) cube complex with a RAAG-like action by G . Then $\text{scl}(g) \geq 1/24$ for every hyperbolic element $g \in G$.*

Since the standard action of a right-angled Artin group on its associated CAT(0) cube complex is RAAG-like, with all non-trivial elements acting hyperbolically, the following corollary is immediate.

COROLLARY B. — *Let G be a right-angled Artin group. Then $\text{scl}(g) \geq 1/24$ for every nontrivial $g \in G$.*

What is perhaps surprising about this result is that there is a *uniform* gap for scl, independent of the dimension of X . Note that in Theorem A we do not assume that X is either finite-dimensional or locally finite; thus Corollary B applies to right-angled Artin groups defined over arbitrary simplicial graphs.

The defining properties of RAAG-like actions arose naturally while working out the arguments in this paper. It turns out, however, that RAAG-like actions are closely related to the *special cube complexes* of Haglund and Wise [19]. That is, if G acts freely on X , then the action is RAAG-like if and only if the quotient complex X/G is special. See Section 7 and Remark 7.4 for the precise correspondence between these notions.

COROLLARY C. — *Let G be the fundamental group of a special cube complex. Then $\text{scl}(g) \geq 1/24$ for every non-trivial $g \in G$.*

This follows from Theorem A since the action of G on the universal cover is RAAG-like, with every non-trivial element acting hyperbolically. Alternatively, it follows from Corollary B and monotonicity, since every such group embeds into a right-angled Artin group.

Related results

There are other gap theorems for stable commutator length in the literature, though in some cases the emphasis is on the existence of a gap, rather than its size. The first such result was Duncan and Howie's theorem [14] that every non-trivial element of a free group has stable commutator length at least $1/2$. In [12] it was shown that in Baumslag–Solitar groups, stable commutator length is either zero or at least $1/12$. A different result in [12] states that if G acts on a tree, then $\text{scl}(g) \geq 1/12$ for every “well-aligned” element $g \in G$. There are also gap theorems for stable commutator length in hyperbolic groups [8, 17] and in mapping class groups (and their finite-index subgroups) [3], where existence of a gap is established. In these cases it is also determined which elements of the group have positive scl. In [8],

the size of the gap in the case of a hyperbolic group is estimated, in terms of the number of generators and the hyperbolicity constant.

In [21, Corollary 6.13], it was shown that every finitely generated right-angled Artin group G embeds into the Torelli subgroup of the mapping class group of a surface. Since scl is positive on the Torelli group [3], monotonicity implies that every non-trivial element of G has positive scl. However, the lower bounds obtained in this way are neither explicit nor uniform. For instance, the genus of the surface needed in [21] grows with the number of generators of G , and this affects the bounds arising in [3] (which go to zero as the genus grows).

There are numerous results on the existence of homogeneous quasimorphisms on groups, where the purpose is to show that the group has non-zero second bounded cohomology. Let $\widetilde{\text{QH}}(G)$ denote the space of homogeneous quasimorphisms on G , modulo homomorphisms. Then $\widetilde{\text{QH}}(G)$ is a subspace of $H_b^2(G; \mathbb{R})$. In [15] it was shown that $\widetilde{\text{QH}}(G)$ is infinite-dimensional for any hyperbolic group G . Recent results in this direction, involving both wider classes of groups and more general coefficient modules, include [20] and [2]. If G is a non-abelian right-angled Artin group then $\widetilde{\text{QH}}(G)$ is infinite-dimensional (via a retraction onto a free subgroup). The quasimorphisms defined in this paper include an infinite collection that is linearly independent in $\widetilde{\text{QH}}(G)$; see Proposition 4.7. These appear to be different from the quasimorphisms one obtains via retractions.

We have mentioned that the *median property* of CAT(0) cube complexes is used to control the defect of our quasimorphisms. The use of medians in this context originated in [10], where they are used to define a bounded cohomology class (the *median class*) which has good functorial properties. This class is defined, and is *non-trivial*, whenever one has a non-elementary group action on a finite-dimensional CAT(0) cube complex. One consequence, among many others, is that $H_b^2(G; M)$ is non-trivial for any such group, for a suitably defined coefficient module M .

Our upper bound of 12 for the defect of the quasimorphisms $\widehat{\phi}_\gamma$ can actually be lowered to 6 in the special case when the CAT(0) cube complex is 1-dimensional; see Remark 4.6. This statement then coincides with Theorem 6.6 of [12], and thus we obtain a new proof of the latter result.

Methods

The fundamental result upon which most of our arguments depend is the existence of *equivariant Euclidean embeddings*, proved in Proposition 5.4.

To state this result, we first note that every element $g \in G$ has a *minimal subcomplex* $M_g \subseteq X$, and if g is hyperbolic then this subcomplex admits a $\langle g \rangle$ -invariant product decomposition $M_g \cong M_g^{\text{ess}} \times X_g^{\text{fix}}$. The action of g on X_g^{fix} is trivial and every edge in M_g^{ess} is on a combinatorial axis for g . We call M_g^{ess} the *essential minimal set* for g . Furthermore, we show that M_g^{ess} is always a finite-dimensional CAT(0) cube complex. However, M_g^{ess} is not always a convex subcomplex of X . We denote by X_g^{ess} its convex hull in X and refer to X_g^{ess} as the *essential characteristic set* for g . The subcomplex X_g^{ess} is in general much more complicated than M_g^{ess} and can have infinite dimension. In Section 3, we give a complete characterization of when X_g^{ess} is finite-dimensional and when X_g^{ess} and M_g^{ess} are the same.

Proposition 5.4 states that under suitable assumptions there is a $\langle g \rangle$ -equivariant embedding of X_g^{ess} into \mathbb{R}^d , where $d = \dim X_g^{\text{ess}}$. That is, there is an embedding of cube complexes $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$ such that the action of $\langle g \rangle$ on X_g^{ess} extends to an action on \mathbb{R}^d (preserving its standard cubing). Furthermore, the embedding induces a bijection between the half-spaces of X_g^{ess} and those of \mathbb{R}^d .

It is well known that any interval in a CAT(0) cube complex admits an embedding into \mathbb{R}^d for some d . This result is proved using Dilworth's theorem on partially ordered sets of finite width; see [5] for details. What is new in our result is the equivariance. In order to prove it, we first state and prove an equivariant version of Dilworth's theorem, Lemma 5.3.

An important aspect of the equivariant Euclidean embedding is that it provides a geometric framework for understanding the fine structure of the set of half-spaces of X_g^{ess} , considered as a partially ordered set. This set becomes identified with the set of half-spaces of \mathbb{R}^d , and the partial ordering from X_g^{ess} is determined by the knowledge of which cubes in \mathbb{R}^d are occupied by X_g^{ess} (cf. Remark 6.1). Tools such as the Quadrant Lemma and the Elbow Lemma (see Section 6) can be used to retrieve information about the partial ordering. These tools become available once X_g^{ess} has been embedded into \mathbb{R}^d .

An outline of the paper

In Section 2 we present background on several topics, including quasimorphisms and stable commutator length, CAT(0) cube complexes, and right-angled Artin groups.

In Section 3 we define the *essential minimal set* and the *essential characteristic set*, and establish their properties. We determine when they agree, and when the latter has finite dimension.

In Section 4 we define *non-transverse* actions. For such actions we also define the quasimorphisms ψ_γ and φ_γ and establish the bounds on defect, using medians. We construct an infinite linearly independent set in $\widetilde{\text{QH}}(A_\Gamma)$, for any non-abelian right-angled Artin group A_Γ .

In Section 5 we prove the equivariant Dilworth theorem, and apply it to prove the existence of equivariant Euclidean embeddings of essential characteristic sets.

In Section 6 we introduce *quadrants* and prove two basic results, the Quadrant Lemma and the Elbow Lemma. These are the primary tools used for studying the essential characteristic set X_g^{ess} once it has been equivariantly embedded into \mathbb{R}^d .

In Section 7 we discuss *RAAG-like* actions on CAT(0) cube complexes.

In Sections 8 and 9 we carry out the rather intricate arguments needed to show that $\widehat{\varphi}_\gamma(g) \geq 1$ for the appropriate choice of γ . Essentially all of the effort in these sections is devoted to showing that X_g^{ess} contains no G -translate of $\bar{\gamma}$.

2. Preliminaries

In this section we establish notation and background for the rest of the paper. We start with the topics of quasimorphisms and stable commutator length. For more detail see [7]. Then we give some background on CAT(0) cube complexes, focusing on the structure of their half spaces and their median structure. More information on these topics can be found in [11, 18, 23, 24, 25]. The section concludes with a brief overview of right-angled Artin groups and properties of their associated CAT(0) cube complexes. These properties lead to the notion of *RAAG-like* actions, to be defined in Section 7.

Notation. — Throughout the paper we use the symbols “ \subset ” and “ \supset ” to denote *strict* inclusion only.

Quasimorphisms and stable commutator length

Let G be any group. A map $\varphi: G \rightarrow \mathbb{R}$ is a *quasimorphism* on G if there is a constant $D \geq 0$ such that for all $g, h \in G$,

$$|\varphi(gh) - \varphi(g) - \varphi(h)| \leq D.$$

The smallest D that satisfies the inequality above is called the *defect* of φ . It is immediate that a quasimorphism is a homomorphism if and only if its defect is 0.

A quasimorphism φ is *homogeneous* if $\varphi(g^n) = n\varphi(g)$ for all $g \in G$ and $n \in \mathbb{Z}$. Given any quasimorphism φ , its *homogenization* $\widehat{\varphi}$ is defined by

$$\widehat{\varphi}(g) = \lim_{n \rightarrow \infty} \frac{\varphi(g^n)}{n}.$$

It is straightforward to check $\widehat{\varphi}$ is a homogeneous quasimorphism. Its defect can be estimated as follows:

LEMMA 2.1. — *If φ is a quasimorphism of defect at most D , then its homogenization has defect at most $2D$.*

Two maps $\varphi, \psi: G \rightarrow \mathbb{R}$ are uniformly close if there exists $D \geq 0$ such that $|\varphi(g) - \psi(g)| \leq D$ for all $g \in G$. It is easy to check that any map uniformly close to a quasimorphism is a quasimorphism. Further, the following statement holds:

LEMMA 2.2. — *If φ is uniformly close to a quasimorphism ψ , then $\widehat{\varphi} = \widehat{\psi}$.*

Proof. — By assumption, there exists $D \geq 0$ such that $|\varphi(g) - \psi(g)| \leq D$ for all $g \in G$. Then

$$\begin{aligned} |\widehat{\varphi}(g) - \widehat{\psi}(g)| &= \left| \lim_{n \rightarrow \infty} \frac{\varphi(g^n)}{n} - \lim_{n \rightarrow \infty} \frac{\psi(g^n)}{n} \right| \\ &= \lim_{n \rightarrow \infty} \frac{|\varphi(g^n) - \psi(g^n)|}{n} \\ &\leq \lim_{n \rightarrow \infty} \frac{D}{n} = 0. \end{aligned} \quad \square$$

Now denote by $[G, G]$ the commutator subgroup of G . Given an element $g \in [G, G]$, the *commutator length* $\text{cl}(g)$ of g is the minimal number of commutators whose product equals g . The commutator length of the identity element is 0. For any $g \in [G, G]$, the *stable commutator length* of g is

$$\text{scl}(g) = \lim_{n \rightarrow \infty} \frac{\text{cl}(g^n)}{n}.$$

Note that $\text{scl}(g^n) = n \text{scl}(g)$ for all $n \in \mathbb{Z}$ and $g \in G$. This formula allows one to define scl for elements that are only virtually in $[G, G]$. By convention, $\text{scl}(g) = \infty$ if no power of g lies in $[G, G]$.

The relationship between stable commutator length and quasimorphisms on G is expressed by Bavard duality [1]. We state the easier direction below:

LEMMA 2.3 (Easy direction of Bavard Duality). — For any $g \in [G, G]$, if φ is a homogeneous quasimorphism on G with defect at most D , then

$$\text{scl}(g) \geq \frac{\varphi(g)}{2D}.$$

CAT(0) cube complexes

A cube of dimension d is an isometric copy of $[0, 1]^d$ with the standard Euclidean metric. A face of a cube is obtained by fixing any number of coordinates to be 0 or 1. This is naturally a cube of the appropriate dimension. A *midcube* is the subset of the cube obtained by fixing one of the coordinates to be $1/2$.

A *cube complex* X is a space obtained from a collection of cubes with some faces identified via isometries. The dimension of X is the dimension of a maximal dimensional cube if it exists; otherwise the dimension of X is infinite. We equip X with the path metric induced by the Euclidean metric on each cube. By Gromov's link condition, X is non-positively curved if and only if the link of every vertex of X is a flag complex. A cube complex X is CAT(0) if and only if it is non-positively curved and simply connected.

Let X be a CAT(0) cube complex. By an *edge path* of length n we will mean a sequence of vertices x_0, \dots, x_n , such that adjacent vertices x_i and x_{i+1} are joined by an edge of X . If $p = x_0, \dots, x_n$ and $q = y_0, \dots, y_m$ are two edge paths with $x_n = y_0$, then their concatenation is the edge path $p \cdot q = x_0, \dots, x_n, y_1, \dots, y_m$.

We will ignore the CAT(0) metric on X and consider the *combinatorial metric* on its vertex set, which measures distance $d(x, y)$ between two vertices x and y as the minimal length of an edge path joining them. An edge path from x to y is a *geodesic* if it has length $d(x, y)$. An infinite sequence of vertices in X is a geodesic if every finite consecutive subsequence is a geodesic.

A *hyperplane* in X is a connected subset whose intersection with each cube of X is either empty or is a midcube. This set always divides X into two disjoint components. The closure of a component is called a *half-space* H of X . The closure of the other component is denoted by \overline{H} . We denote by ∂H the boundary hyperplane of H and note that $\partial H = \partial \overline{H}$.

A subcomplex $C \subseteq X$ is *convex* if every geodesic in X between two of its vertices is contained entirely in C . If $Y \subseteq X$ is any subcomplex, the *convex hull* $C(Y)$ of Y is the smallest convex subcomplex containing Y .

Equivalently, it is the largest subcomplex of X that is contained in the intersection of all half-spaces containing Y .

For any vertices $x, y \in X$, we will denote by $C(x, y)$ the convex hull $C(\{x, y\})$.

A hyperplane ∂H is *dual* to an edge (or vice versa) if ∂H intersects the edge. A half-space H is dual to an edge if ∂H is. A cube C is dual to a hyperplane ∂H if C contains an edge dual to ∂H . The *neighborhood* of ∂H is the union $N(\partial H)$ of all cubes dual to ∂H . By [18, Theorem 2.12], $N(\partial H)$ is convex. Further, there is an involution on $N(\partial H)$ that fixes ∂H pointwise and swaps the endpoints of each edge dual to ∂H (in fact, $N(\partial H) \cong \partial H \times [0, 1]$).

Let $\mathcal{H}(X)$ be the collection of half-spaces of X . This is partially ordered by inclusion. We say two half-spaces are *nested* if they are linearly ordered; they are *tightly nested* if they are nested and there is no third half-space that lies properly between them. The map $\mathcal{H}(X) \rightarrow \mathcal{H}(X)$ sending H to \bar{H} is an order-reversing involution.

Two half-spaces H, H' of X are *transverse*, denoted by $H \pitchfork H'$, if all four intersections

$$H \cap H', \quad H \cap \bar{H}', \quad \bar{H} \cap H, \quad \bar{H} \cap \bar{H}'$$

are non-empty. When this happens, then there is a cube C in X such that $\partial H \cap C$ and $\partial H' \cap C$ are different midcubes of C . More generally, if H_1, \dots, H_n are pairwise transverse, then there is a cube C in X of dimension n such that $\partial H_1 \cap C, \dots, \partial H_n \cap C$ are the n midcubes of C .

Given two vertices $x, y \in X$, the *interval between x and y* is

$$[x, y] = \{H \in \mathcal{H} : y \in H, x \in \bar{H}\}.$$

Two distinct half-spaces $H, H' \in [x, y]$ are always either nested or transverse. The interval $[y, x]$ is exactly the set of half spaces $\{\bar{H} : H \in [x, y]\}$.

An *oriented edge* $e = (x, y)$ is an edge whose vertices x, y have been designated as *initial* and *terminal* respectively. Given an edge path x_0, \dots, x_n , each edge (x_i, x_{i+1}) receives an induced orientation with x_i initial and x_{i+1} terminal. For any oriented edge $e = (x, y)$, the *half space dual to e* is the unique half-space in the interval $[x, y]$; it is dual to e considered as an unoriented edge, and it contains y but not x .

An edge path is a geodesic if and only if it crosses no hyperplane twice. Two geodesics from x to y determine the same set of half-spaces $[x, y]$, and every half-space $H \in [x, y]$ is dual to some edge on every geodesic from x to y . Therefore, the combinatorial distance $d(x, y)$ is the same as the cardinality of $[x, y]$. See [25, Theorem 4.13] for more details.

Ultrafilters

Suppose σ is a function assigning to each hyperplane h in X a half-space H with $\partial H = h$. Then σ is an *ultrafilter* if $\sigma(h)$ and $\sigma(h')$ have non-trivial intersection for every pair of hyperplanes h, h' . An alternative viewpoint is to simply specify the image of σ , as a subset of $\mathcal{H}(X)$ that contains exactly one half-space from each pair $\{H, \overline{H}\}$, such that no two elements are disjoint. For this reason, σ is sometimes called an ultrafilter “on $\mathcal{H}(X)$ ”.

For each vertex v of X there is a *principal ultrafilter* of v , defined by choosing $\sigma(h)$ to be the half-space with boundary h containing v . Neighboring vertices define principal ultrafilters that differ on a single hyperplane (the one that is dual to the edge separating the vertices). Conversely, if two principal ultrafilters differ on a single hyperplane, then the corresponding vertices bound an edge, dual to that hyperplane. Since X is connected, any two principal ultrafilters will differ on finitely many hyperplanes. Indeed, the number of such hyperplanes is precisely the distance between the two vertices.

The principal ultrafilters admit an intrinsic characterization: an ultrafilter σ on $\mathcal{H}(X)$ is principal if and only if it satisfies the descending chain condition: whenever $\{h_i\}$ is a sequence of hyperplanes such that $\sigma(h_i) \supseteq \sigma(h_{i+1})$ for all i , the sequence is eventually constant. It follows that if an ultrafilter differs from a principal one on finitely many hyperplanes, it will also be principal.

Knowledge of the principal ultrafilters on $\mathcal{H}(X)$ completely determines X as a CAT(0) cube complex. The *Sageev construction* is the name for the process of building a cube complex from its partially ordered set of half-spaces. The 1-skeleton of X is determined from principal ultrafilters as already described, and cubes are added whenever their 1-skeleta are present [25].

More generally, let \mathcal{H} be any partially ordered set with an order-reversing free involution $H \mapsto \overline{H}$, such that every interval is finite. The Sageev construction yields a CAT(0) cube complex $X(\mathcal{H})$ whose half-spaces correspond to \mathcal{H} as a partially ordered set with involution [24]. It is often convenient to think of vertices of X as principal ultrafilters, and to identify X with the result of the Sageev construction performed on $\mathcal{H}(X)$.

Medians

Given three vertices $x, y, z \in X$, there is a unique vertex $m = m(x, y, z)$ called the *median* such that $[a, b] = [a, m] \cup [m, b]$ for all pairs $\{a, b\} \subset$

$\{x, y, z\}$. For completeness we sketch the proof, since the standard reference [24] is unpublished.

As an ultrafilter, m is defined by simply assigning to each hyperplane the half-space which contains either two or three of the vertices $\{x, y, z\}$. Two such half-spaces cannot be disjoint, so this rule does indeed define an ultrafilter. This ultrafilter is principal (i.e. it defines a *vertex*) because it differs from the principal ultrafilter of x on finitely many hyperplanes: if H is chosen by m and $x \notin H$, then $y, z \in H$; hence $H \in [x, y] \cap [x, z]$, a finite set. Finally, given $a, b \in \{x, y, z\}$, every half-space containing a and b also contains m , by definition. Thus, no hyperplane can separate m from a and b , and therefore $[a, b] = [a, m] \cup [m, b]$.

A vertex z lies on a geodesic edge path from x to y if and only if $z = m(x, z, y)$. Therefore, $z \in C(x, y)$ if and only if $z = m(x, z, y)$.

Segments

By a *segment* γ of length n we will mean a chain of half-spaces $H_1 \supset H_2 \supset \cdots \supset H_n$ such that H_i and H_{i+1} are tightly nested for all $i = 1, \dots, n-1$. The *inverse* of γ is the segment $\bar{\gamma}: \bar{H}_n \supset \bar{H}_{n-1} \supset \cdots \supset \bar{H}_1$.

Let γ and γ' be segments. We write $\gamma > \gamma'$ if every half-space in γ contains every half-space in γ' . We say that γ and γ' are *nested* if either $\gamma > \gamma'$ or $\gamma' > \gamma$.

DEFINITION 2.4. — *Two segments γ and γ' are said to overlap if either $\gamma \cap \gamma' \neq \emptyset$ or there exist $H \in \gamma$ and $H' \in \gamma'$ with $H \pitchfork H'$. Otherwise, they are non-overlapping.*

LEMMA 2.5. — *Suppose γ_1 and γ_2 are non-overlapping segments that are contained in $[x, y]$. Then γ_1 and γ_2 are nested.*

Proof. — As mentioned above, any two half-spaces in $[x, y]$ are either nested or transverse. Therefore, since γ_1 and γ_2 are non-overlapping, their union is linearly ordered by inclusion. The result follows, since each γ_i is a segment. \square

Right-angled Artin groups

Let Γ be a simplicial graph (i.e. a simplicial complex of dimension at most 1), with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$. The *right-angled Artin*

group A_Γ is defined to be the group with generating set $V(\Gamma)$ and relations $\{[v, w] : \{v, w\} \in E(\Gamma)\}$. That is, two generators commute if and only if they bound an edge in Γ , and there are no other defining relations.

There is a naturally defined non-positively curved cube complex which is a $K(A_\Gamma, 1)$ complex, obtained as a union of tori corresponding to complete subgraphs of Γ (see Davis [13, 11.6], for example). The universal cover X_Γ is a CAT(0) cube complex with a free action by A_Γ . The oriented edges of X_Γ can be labeled by the generators of A_Γ and their inverses in a natural way: each such edge is a lift of a loop representing that generator (or its inverse).

This labeling has the property that two oriented edges are in the same A_Γ -orbit if and only if their labels agree. Also, the oriented edges that are dual to any given half-space will always have the same label, so the label may be assigned to the half-space itself. Half-spaces in the same A_Γ -orbit will have the same label.

The half-space labels lead to several useful observations. Each 2-cell of X_Γ is a square whose boundary is labeled by a commutator $[v, w]$, with $v \neq w$. It follows that no two half-spaces with the same label can be transverse in X_Γ . Since no label equals its inverse, A_Γ acts without inversion on X_Γ . Some additional properties of X_Γ related to the edge-labeling will be discussed in Section 7.

3. Automorphisms and characteristic sets

In this section we discuss automorphisms of CAT(0) cube complexes and their characteristic sets. We define the *essential characteristic set* and the *essential minimal set* of a hyperbolic automorphism, and we determine the structures of these sets. The latter set is always finite-dimensional, whereas the former is a subcomplex which plays an essential role throughout the paper. Toward the end of the section, we characterize when these sets agree (Proposition 3.17) and when the essential characteristic set is finite-dimensional (Corollary 3.19).

Basic notions

Following Haglund [18], an automorphism g of a CAT(0) cube complex X acts with inversion if there is a half-space H such that $g(H) = \overline{H}$. When this occurs, g stabilizes the hyperplane ∂H . For any automorphism g of X ,

the action of g on the cubical subdivision of X is always *without* inversion. Note, however, that for some of our results, we will not be free to perform this modification; see Remark 7.2.

For an automorphism g of X , the *translation distance* of g is $\ell_g = \min_{x \in X} d(x, gx)$, where x ranges over the vertices of X . If g and all of its powers act without inversion, we say that g is *hyperbolic* if $\ell_g > 0$ and *elliptic* otherwise. Haglund showed that when g is hyperbolic, there is an infinite combinatorial geodesic in X that is preserved by g , on which g acts as a translation of magnitude ℓ_g . Any such geodesic will be called a *combinatorial axis* for g . It has a natural orientation, relative to which the translation by g is in the *forward* direction. Note that g and g^{-1} have the same combinatorial axes, but they determine opposite orientations.

Haglund also showed that any two combinatorial axes for g cross the same hyperplanes, in the same directions. That is, the set of half-spaces that are dual to oriented edges in any axis is independent of the choice of axis. We define the *positive half-space axis* of g :

$$A_g^+ = \left\{ H \in \mathcal{H}(X) : \begin{array}{l} H \text{ is dual to a positively oriented edge} \\ \text{in a combinatorial axis for } g \end{array} \right\}.$$

We also define the *negative half-space axis* $A_g^- = \{\overline{H} : H \in A_g^+\}$; note that $A_g^- = A_{g^{-1}}^+$. The *full half-space axis* is $A_g = A_g^+ \sqcup A_g^-$.

If L is a combinatorial axis for g , then for every $H \in A_g^+$, the intersection $L \cap H$ is a ray containing the attracting end of L (since L crosses ∂H exactly once). Note also that $gL \neq L$ for all $L \in A_g$, for otherwise g would fix the unique edge of L dual to H , contradicting hyperbolicity of g .

Remark 3.1. — For any distinct half-spaces $H, H' \in A_g^+$, either $H \supset H'$, $H \subset H'$, or $H' \subset H$. The only other possibilities are that $\overline{H'} \subset H$ or $\overline{H} \subset H'$, but these imply that either $H \cap H'$ or $\overline{H} \cap \overline{H'}$ is empty. But every combinatorial axis for g meets both of these sets in an infinite ray.

Remark 3.2. — For any $H \in A_g^+$ and $n > 0$, if H and $g^n H$ are not transverse, then $H \supset g^n H$. To see this, let L be any oriented combinatorial axis for g . Let $e = (x, y)$ be the oriented edge on L dual to H . Then e lies on a geodesic edge path from x to $g^n x$. In other words, $H \in [x, g^n x]$, and so $g^n x \in H$. Since $x \notin H$, $g^n x \notin g^n H$. It follows that $H \supset g^n H$ (rather than $H \subset g^n H$).

Let G be a group acting on X by automorphisms. We will always assume (here and for the rest of the paper) that all elements of G act without inversion. Under this assumption, Haglund showed that every element $g \in G$ is either elliptic or hyperbolic.

The minimal set, the characteristic set, and their product decompositions

DEFINITION 3.3. — *For any $g \in G$, the minimal set of g is the full subcomplex $M_g \subseteq X$ generated by the vertices of X that realize the translation distance of g .*

Since g and all of its powers act without inversion, there are two types of behavior for M_g . If g is elliptic then M_g is the subcomplex of fixed points of g . If g is hyperbolic then M_g is the smallest full subcomplex containing all combinatorial axes for g . It is non-empty, and every vertex of M_g is on a combinatorial axis, by [18, Corollary 6.2].

Next we define three more sets of half-spaces when $g \in G$ is hyperbolic:

$$S_g = \{H \in \mathcal{H} : H \text{ contains every combinatorial axis of } g\}$$

$$= \{H \in \mathcal{H} : H \text{ contains } M_g\},$$

$$\bar{S}_g = \{H \in \mathcal{H} : H \notin A_g \text{ and } H \text{ contains no combinatorial axis of } g\}$$

$$= \{H \in \mathcal{H} : \bar{H} \in S_g\},$$

$$T_g = \{H \in \mathcal{H} : H \notin A_g \text{ and } \partial H \text{ separates two combinatorial axes of } g\}.$$

Recall that the half-spaces *not* in A_g are exactly those whose boundary hyperplanes do not cross any axis. Thus the aforementioned sets define a partition of $\mathcal{H}(X)$:

$$\mathcal{H}(X) = A_g \sqcup S_g \sqcup \bar{S}_g \sqcup T_g.$$

Remark 3.4. — For any group Γ acting on X , Caprace and Sageev have defined a decomposition of $\mathcal{H}(X)$ into Γ -essential, Γ -half-essential, and Γ -trivial half-spaces [9]. It can be shown that when $\Gamma = \langle g \rangle$ (with g hyperbolic), these three collections of half-spaces coincide with A_g , $(S_g \cup \bar{S}_g)$, and T_g , respectively.

Using this perspective, some of the results below can be derived from results in [9] and [10]. Specifically, Lemma 3.7 is observed in Remark 3.4 of [9], and Lemma 3.8 can be derived from Lemma 2.6 of [10] (see also [16, Remark 2.11]).

For completeness, we include elementary proofs of these results, using the definitions of A_g , S_g , \bar{S}_g , and T_g given above.

LEMMA 3.5. — *Suppose $g \in G$ is hyperbolic. If $H \in A_g$ and $K \in T_g$ then $H \pitchfork K$.*

Proof. — Let L, L' be combinatorial axes of g such that $L \subset K$ and $L' \subset \bar{K}$. Every axis meets both H and \bar{H} . Thus all four intersections $K \cap H, K \cap \bar{H}, \bar{K} \cap H,$ and $\bar{K} \cap \bar{H}$ are non-empty. □

DEFINITION 3.6. — *If $g \in G$ is hyperbolic, the characteristic set of g is the convex hull of M_g , denoted X_g . Equivalently, X_g is the largest subcomplex of X contained in $\bigcap_{H \in S_g} H$.*

The collections of half-spaces A_g and T_g define CAT(0) cube complexes $X_g^{\text{ess}} = X(A_g)$ and $X_g^{\text{ell}} = X(T_g)$ by the Sageev construction, called the *essential characteristic set* and the *elliptic factor* respectively.

LEMMA 3.7. — *Suppose $g \in G$ is hyperbolic. Then there is a $\langle g \rangle$ -equivariant isomorphism of cube complexes $X_g \cong X_g^{\text{ess}} \times X_g^{\text{ell}}$.*

Proof. — First note that since $A_g \pitchfork T_g$, there is an isomorphism $X_g^{\text{ess}} \times X_g^{\text{ell}} \cong X(A_g \cup T_g)$, by [9, Lemma 2.5]. We shall define an embedding $X(A_g \cup T_g) \hookrightarrow X$ and show that its image is X_g .

The map is defined by extending each principal ultrafilter on $A_g \cup T_g$ to an ultrafilter on $\mathcal{H}(X)$ by including every half-space in S_g . These half-spaces have non-trivial intersection with every half-space in $A_g \cup T_g$, and also with each other, so this rule does indeed define an ultrafilter. Moreover, no half-space in S_g is contained in any half-space of $A_g \cup T_g$, and S_g itself contains no descending chains: for each $H \in S_g$, ∂H is separated from M_g by only finitely many hyperplanes. Therefore the descending chain condition is still satisfied. Thus, each vertex of $X(A_g \cup T_g)$ is mapped to a vertex of X . It is clear that adjacent vertices map to adjacent vertices, so the map is an embedding of cube complexes.

Next, the vertices of X_g are exactly the vertices whose principal ultrafilters include all half-spaces of S_g . These are exactly the vertices in the image of our map, so this image is X_g .

Equivariance holds because the $\langle g \rangle$ -actions on $X(A_g)$, $X(T_g)$, and $X = X(\mathcal{H}(X))$ are all simultaneously induced by the action of $\langle g \rangle$ on the half-spaces of X . \square

The next result concerns crossing of half-spaces of A_g . Namely, two such half-spaces cross in X_g^{ess} if and only if they cross in X :

LEMMA 3.8. — *Suppose $g \in G$ is hyperbolic. If $H, H' \in A_g$ and $H \pitchfork H'$ in X , then $H \pitchfork H'$ in X_g^{ess} . That is, there is a square $S \subset X_g^{\text{ess}}$ containing edges e, e' that are dual to H and H' respectively.*

Proof. — Recall that $X_g^{\text{ess}} = X(A_g)$. We may embed X_g^{ess} as a convex subcomplex of X in such a way that the induced map on half-spaces $A_g \rightarrow \mathcal{H}(X)$ is inclusion; this follows from Lemma 3.7, by choosing a vertex $v \in X_g^{\text{ell}}$ and identifying X_g^{ess} with $X_g^{\text{ess}} \times \{v\}$ in X_g .

There is a combinatorial retraction $X \rightarrow X_g^{\text{ess}}$ defined in terms of ultrafilters by restriction: each principal ultrafilter on $\mathcal{H}(X)$ is sent to its

intersection with A_g . The resulting ultrafilter still satisfies the descending chain condition, and therefore defines a vertex in X_g^{ess} . Two adjacent vertices of X will either map to adjacent vertices or to the same vertex. This map extends to cubes, and each cube maps to a cube in X_g^{ess} by a coordinate projection. More specifically, an edge in X is collapsed if and only if its dual half-spaces are not in A_g . It follows that if a square in X is dual to two half-spaces in A_g^+ , then its image in X_g^{ess} is also a square, dual to the same two half-spaces. Thus, if $H, H' \in A_g^+$ are transverse in X , they are transverse in X_g^{ess} . \square

Next we continue to examine the structure of X_g .

LEMMA 3.9. — *Suppose $g \in G$ is hyperbolic. Then g acts as an elliptic automorphism of X_g^{ell} .*

Proof. — If not, any axis L of g acting on $X_g = X_g^{\text{ess}} \times X_g^{\text{ell}}$ would project onto an axis in X_g^{ell} , and L would then cross a hyperplane bounding a half-space in T_g . However, no axis of g crosses such a hyperplane. \square

Accordingly, there is a non-empty subcomplex $X_g^{\text{fix}} \subseteq X_g^{\text{ell}}$ consisting of the fixed points of the $\langle g \rangle$ -action on X_g^{ell} . It is a subcomplex because $\langle g \rangle$ acts without inversion.

LEMMA 3.10. — *Suppose $g \in G$ is hyperbolic. Then there is a $\langle g \rangle$ -invariant subcomplex $M_g^{\text{ess}} \subseteq X_g^{\text{ess}}$ such that $M_g = M_g^{\text{ess}} \times X_g^{\text{fix}}$ under the identification of X_g with $X_g^{\text{ess}} \times X_g^{\text{ell}}$.*

The subcomplex M_g^{ess} is called the *essential minimal set* for g .

Proof. — If x is a vertex of M_g then no half-space of T_g separates x from gx , since x and gx are on a combinatorial axis. Thus the principal ultrafilters at x and at gx agree on half-spaces in T_g . That is, g fixes the second coordinate of x in $X_g^{\text{ess}} \times X_g^{\text{ell}}$. Therefore $M_g \subseteq X_g^{\text{ess}} \times X_g^{\text{fix}}$.

Let M_g^{ess} be the projection of M_g onto the first factor of $X_g^{\text{ess}} \times X_g^{\text{ell}}$, so $M_g \subseteq M_g^{\text{ess}} \times X_g^{\text{fix}}$. Since $\langle g \rangle$ acts trivially on X_g^{fix} , any two vertices of $M_g^{\text{ess}} \times X_g^{\text{fix}}$ with the same first coordinate are moved the same distance by g . It follows that every vertex of $M_g^{\text{ess}} \times X_g^{\text{fix}}$ is moved distance ℓ_g , and hence is in M_g . Since M_g is the full subcomplex spanned by its vertices, we have $M_g = M_g^{\text{ess}} \times X_g^{\text{fix}}$. Further, $\langle g \rangle$ -invariance of M_g^{ess} is clear, because both M_g and its product structure are $\langle g \rangle$ -invariant. \square

LEMMA 3.11. — *Suppose $g \in G$ is hyperbolic. Let e be an edge of M_g which projects to an edge in the factor M_g^{ess} . Then e is on a combinatorial axis of g .*

Proof. — We prove the contrapositive. Let $e = (x, y)$ where x and y are vertices of M_g . If e is not on any combinatorial axis, then y is not on any geodesic from x to gx , so $y \neq m(x, y, gx)$. There must be a half-space containing y but not x or gx . The half-space H dual to e is the only possibility, since $[x, y] = \{H\}$.

Similarly, x is not on any geodesic from y to gy , so there must be a half-space containing x but not y or gy . This can only be \overline{H} , since $[y, x] = \{\overline{H}\}$.

Thus ∂H separates gx from gy , and hence is dual to ge ; therefore $g\partial H = \partial H$. Since g is not an inversion, we have that $gH = H$. Thus $H \notin A_g$ and e does not project to an edge in M_g^{ess} . \square

Relationship between X_g^{ess} and M_g^{ess}

The convex hull of $M_g^{\text{ess}} \subseteq X_g^{\text{ess}}$ is X_g^{ess} . To see this, note that every edge of M_g^{ess} is dual to a half-space of X_g^{ess} , by Lemma 3.11; and conversely, every half-space in A_g^+ is dual to an edge in an axis, and hence to an edge in M_g^{ess} . Hence no half-space of X_g^{ess} contains M_g^{ess} , and therefore $C(M_g^{\text{ess}})$ is the intersection of the empty set of half-spaces of X_g^{ess} .

In this section, we will establish the basic structure of M_g^{ess} and address when M_g^{ess} and X_g^{ess} are the same.

DEFINITION 3.12. — *Let C be any cube in X_g^{ess} and let A be the set of elements in A_g^+ dual to the edges of C . Let x, y be the two vertices of C such that $A = [x, y]$. We will call x the minimal vertex of C and y the maximal vertex of C .*

LEMMA 3.13. — *Let $C \times \{v\}$ be a cube in $M_g^{\text{ess}} \times \{v\} \subseteq M_g^{\text{ess}} \times X_g^{\text{fix}}$ with minimal and maximal vertices x and y . Then $[x, y] \subseteq [x, gx]$. Thus, y lies on some combinatorial axis of g containing x . This axis lies inside $M_g^{\text{ess}} \times \{v\}$.*

Proof. — Let $e = (x, z)$ be any oriented edge in $C \times \{v\}$ with initial vertex x , and let $H \in [x, y]$ be the half-space dual to e . Since $C \subset M_g^{\text{ess}}$, the edge e lies on a combinatorial axis of g . In particular, the vertex z lies on a geodesic edge path from x to gx . Since this geodesic can cross ∂H only once, $gx \in H$. In other words, $H \in [x, gx]$. This is true for every $H \in [x, y]$, so $[x, y] \subseteq [x, gx]$. For the last conclusion, let α be the concatenation of geodesic edge paths from x to y and from y to gx . Then α does not cross any hyperplane twice, since such a hyperplane would separate y from x and gx . The concatenation of α and its g -translates is a combinatorial axis

containing x and y . The axis lies in $M_g^{\text{ess}} \times \{v\}$ by $\langle g \rangle$ -invariance of the product decomposition $M_g^{\text{ess}} \times X_g^{\text{fix}}$. \square

We now show that M_g^{ess} is always finite-dimensional.

LEMMA 3.14. — *Suppose $g \in G$ is hyperbolic. Then M_g^{ess} is finite-dimensional, with dimension bounded by the translation distance ℓ_g of g .*

Proof. — Recall that the distance $d(x, y)$ between two vertices is the same as the cardinality of $[x, y]$. Thus, for any $x \in M_g^{\text{ess}}$, the cardinality of $[x, gx]$ is the same as the translation distance ℓ_g . Let C be any cube in M_g^{ess} , let v be any vertex of X_g^{fix} , and let x and y be the minimal and maximal vertices of $C \times \{v\}$ in $M_g^{\text{ess}} \times X_g^{\text{fix}}$. The dimension of C is the same as the cardinality of $[x, y]$. By Lemma 3.13, we always have $[x, y] \subseteq [x, gx]$, so the dimension of C is bounded by ℓ_g . This is true for all C in M_g^{ess} , whence the result. \square

Our goal now is to relate M_g^{ess} and X_g^{ess} . It turns out that X_g^{ess} may have infinite dimension. An easy example showing that M_g^{ess} and X_g^{ess} can have different dimensions is the glide reflection in \mathbb{R}^2 defined by

$$g(x, y) = (y + 1, x).$$

Then g has translation length 1, so M_g^{ess} is 1-dimensional by Lemma 3.14, but $X_g^{\text{ess}} = \mathbb{R}^2$. See Figure 3.1.

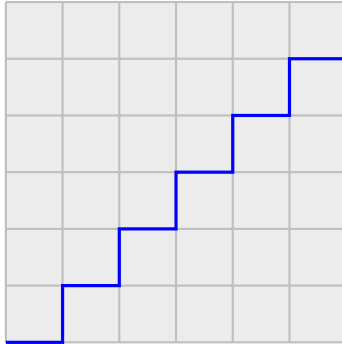


Figure 3.1. A glide reflection g with a unique combinatorial axis. We have $\dim(M_g^{\text{ess}}) = 1$ and $\dim(X_g^{\text{ess}}) = 2$.

This example can be promoted to one in which X_g^{ess} has infinite dimension. Consider $\mathbb{R}^{\mathbb{Z}}$ with its standard integer cubing. Fix the origin $o = (0, 0, \dots)$ and consider the subcomplex $X \subset \mathbb{R}^{\mathbb{Z}}$ generated by the

vertices in $\mathbb{R}^{\mathbb{Z}}$ having at most finitely many non-zero coordinates. Then X is an infinite-dimensional CAT(0) cube complex. Given $x \in X$, let x_i denote its i -th coordinate. Let $g: X \rightarrow X$ be defined by $g(x)_0 = x_1 + 1$ and $g(x)_j = x_{j+1}$ for all other j . Again, g has translation length 1, and M_g^{ess} is 1-dimensional, consisting of a single combinatorial axis with vertices $\{g^n(o)\}$.

Letting $H = [o, go]$, the set of half-spaces $\{H, gH, \dots, g^{d-1}H\}$ are pairwise transverse, and the d -dimensional cube C they cross in is contained in $M_{g^d}^{\text{ess}}$. In particular, since $\ell_{g^d} = d$, we see that $M_{g^d}^{\text{ess}}$ has dimension exactly d . Now X_g^{ess} is infinite-dimensional, since $M_{g^d}^{\text{ess}} \subset X_{g^d}^{\text{ess}} = X_g^{\text{ess}}$ for all $d > 0$. Note that in this example, g has a combinatorial axis in X , but it has no CAT(0) axis; see [4, Example II.8.28].

The above discussion leads to the next definition.

DEFINITION 3.15. — *If $g \in G$ is hyperbolic, we say that $\langle g \rangle$ acts non-transversely on X_g^{ess} if, for every $H \in A_g$, H and gH are not transverse in X_g^{ess} . Note that this occurs if and only if H and gH are not transverse in X , by Lemma 3.8.*

Note that while the inclusion $M_g^{\text{ess}} \hookrightarrow X_g^{\text{ess}}$ induces a bijection on their sets of half-spaces, the partial orderings induced by inclusion on these two sets may be very different; equality holds if and only if $M_g^{\text{ess}} = X_g^{\text{ess}}$. In the following, we will establish some criteria for a square in X_g^{ess} to not be in M_g^{ess} .

LEMMA 3.16. — *Suppose $g \in G$ is hyperbolic. Then:*

- (1) *If H and gH are transverse in X_g^{ess} , then any square in which they cross cannot be contained in M_g^{ess} .*
- (2) *Let e and e' be edges in M_g^{ess} with dual half-spaces $H, H' \in A_g^+$. If e and e' bound a square in X_g^{ess} which is not contained in M_g^{ess} , then either $gH = H'$ or $gH' = H$.*

Proof. — Let S be a square in X_g^{ess} in which H and gH cross. Let x be the minimal vertex of S . If S lies in M_g^{ess} , then by Lemma 3.13, H and gH are in $[x, gx]$, which is a contradiction. This shows (1).

For (2), let S be the square in X_g^{ess} bounded by e and e' . Let x and y be respectively the minimal and maximal vertices of S . Let z be the vertex at which e and e' are incident. We claim that z is distinct from x and y . Since the union of e and e' contains three corners of S , exactly one corner of S , the one opposite z , is not contained in M_g^{ess} . If $z = x$, then y must be the corner not in M_g^{ess} . But in this case, both half-spaces H and H' are contained in $[x, gx]$, by Lemma 3.11, so $y = m(x, y, gx)$. This means that y lies on a

combinatorial axis for g , so y is in M_g^{ess} , a contradiction. Similarly, if $z = y$, then $x = m(g^{-1}y, x, y)$, so $x \in M_g^{\text{ess}}$, again a contradiction. Therefore x, y , and z are distinct and lie in M_g^{ess} . Now, without loss of generality, assume that $[x, z] = H$ and $[z, y] = H'$. Note that $z = m(x, z, gx)$ and $y = m(z, y, gz)$. Since S is not contained in M_g^{ess} , $H' \notin [x, gx]$. On the other hand, $H' \in [z, gz]$. Since $[z, gz] = [z, gx] \cup [gx, gz]$, we must have $H' = [gx, gz] = gH$. \square

PROPOSITION 3.17. — *Suppose $g \in G$ is hyperbolic. Then $M_g^{\text{ess}} = X_g^{\text{ess}}$ if and only if $\langle g \rangle$ acts non-transversely on X_g^{ess} .*

Proof. — If $H \pitchfork gH$ for some $H \in A_g^+$, then, by Lemma 3.16(1), any square in X_g^{ess} in which they cross cannot lie in M_g^{ess} . Therefore, $M_g^{\text{ess}} = X_g^{\text{ess}}$ implies that $\langle g \rangle$ acts non-transversely on X_g^{ess} .

Now suppose $\langle g \rangle$ acts non-transversely on X_g^{ess} . We claim that for any edge $e = (x, y)$ in X_g^{ess} , if $x \in M_g^{\text{ess}}$ then $y \in M_g^{\text{ess}}$. To see this, let $H = [x, y]$. Replacing g by g^{-1} if necessary, we may assume that $H \in A_g^+$. If $y \notin M_g^{\text{ess}}$, then $y \neq m(x, y, gx)$. In particular, $H \notin [x, gx]$. But H must be contained in $[g^n x, g^{n+1} x]$ for some $n \in \mathbb{Z}$, and for this n we have $g^{-n}H \in [x, gx]$. Note that $n > 0$, since $x \notin H$. Because $\langle g \rangle$ acts non-transversely, $g^{-n}H$ and H cannot be transverse, and so $g^{-n}H \supset H$. Since $x \notin g^{-n}H$ and H is the only half space separating x and y , we must have $y \notin g^{-n}H$. But this contradicts the fact that $y \in H$. This finishes the proof of the claim.

To finish the argument, it suffices to observe that the 1-skeleton of X_g^{ess} is connected, and therefore every vertex of X_g^{ess} is in M_g^{ess} . \square

Remark 3.18. — If $\langle g \rangle$ acts non-transversely on X_g^{ess} , then $M_{g^n}^{\text{ess}} = M_g^{\text{ess}}$ for all $n \geq 1$, because $X_{g^n}^{\text{ess}} = X_g^{\text{ess}}$. The latter equality holds because any axis for g is an axis for g^n , and therefore $A_{g^n} = A_g$.

COROLLARY 3.19. — *Suppose $g \in G$ is hyperbolic. The following statements are equivalent.*

- (1) *There is an integer $k > 0$ such that $\langle g^k \rangle$ acts non-transversely on X_g^{ess} .*
- (2) *$X_g^{\text{ess}} = M_{g^k}^{\text{ess}}$ for some $k > 0$.*
- (3) *X_g^{ess} is finite-dimensional.*

Proof. — First we show that (1) \implies (2). Suppose that g^k acts non-transversely on X_g^{ess} . Since $X_g^{\text{ess}} = X_{g^k}^{\text{ess}}$, g^k also acts non-transversely on $X_{g^k}^{\text{ess}}$. By Proposition 3.17, $M_{g^k}^{\text{ess}} = X_{g^k}^{\text{ess}}$.

The implication (2) \implies (3) follows from Lemma 3.14.

Now we show that (3) \implies (1). Suppose X_g^{ess} has dimension d . For every $H \in A_g^+$, we claim that $H \supset g^n H$ for some n satisfying $0 < n \leq d$. Since X_g^{ess} has dimension d , the half-spaces

$$H, gH, g^2H, \dots, g^dH$$

cannot all be pairwise transverse. Thus there exist i, j with $0 \leq i < j \leq d$ such that $g^i H \supset g^j H$, or equivalently, $H \supset g^{j-i} H$. Finally, taking $k = d!$, we have that $n \mid k$ whenever $0 < n \leq d$, and therefore $H \supset g^k H$ for all $H \in A_g^+$. Thus $\langle g^k \rangle$ acts non-transversely on X_g^{ess} . \square

Even though M_g^{ess} is not always a convex subspace of X , it is still always CAT(0), which we show next. The proof will make use of Lemma 3.16(2).

PROPOSITION 3.20. — *Suppose $g \in G$ is hyperbolic. Then the cube complex M_g^{ess} is CAT(0).*

Proof. — Denote by \mathcal{H} the set of half-spaces A_g with partial order induced by M_g^{ess} . That is: $H, H' \in \mathcal{H}$ are incomparable (or transverse) if and only if there is a square in M_g^{ess} in which they cross, and $H \geq H'$ if and only if $H \cap M_g^{\text{ess}} \supseteq H' \cap M_g^{\text{ess}}$. Apply the Sageev construction to \mathcal{H} to obtain a CAT(0) cube complex $X(\mathcal{H})$. There is a natural injective map $f: M_g^{\text{ess}} \rightarrow X(\mathcal{H})$ defined by sending every vertex in M_g^{ess} to its associated principal ultrafilter on \mathcal{H} . This map identifies \mathcal{H} with the set of half-spaces of $X(\mathcal{H})$. Let Y be the image of M_g^{ess} . We now proceed to show that $Y = X(\mathcal{H})$, which will imply that M_g^{ess} is CAT(0).

We claim that for any edge $e = (y, y')$ in $X(\mathcal{H})$, if $y \in Y$ then $y' \in Y$. The result follows, since the 1-skeleton of $X(\mathcal{H})$ is connected. To prove the claim, let $H \in \mathcal{H}$ be the half-space dual to e . Replacing g by g^{-1} if necessary, we may assume that $H \in \mathcal{H}^+$. Let $x \in M_g^{\text{ess}}$ be such that $f(x) = y$. Since H must appear in every axis of g passing through x and $x \notin H$, there exists a geodesic path $x = x_0, \dots, x_{n+1} = x'$ in M_g^{ess} such that $H = [x_n, x_{n+1}]$. If $n = 0$, then $f(x_1) = y'$ and we are done. Now suppose that $n > 0$. Let $H_i = [x_i, x_{i+1}]$ for each i . If $H_i \supset H$ for some $i < n$, then since $x \notin H_i$, $f(x) = y \notin H_i$. But $y' \in H$, and hence $y' \in H_i$, but this is impossible as y and y' are separated by exactly one half-space H . Thus $H_i \pitchfork H$ for $i = 0, \dots, n-1$. We claim now that for each i , there is a square S_i in which H_i and H cross, and x_i is the minimal vertex of S_i . Let $e_i = (x_i, x_{i+1})$ for each i . By assumption, e_{n-1} and e_n are dual to two transverse half-space H_{n-1} and H_n . If they do not bound a square in M_g^{ess} , then by Lemma 3.16, $gH_{n-1} = H_n$ and there is no square in M_g^{ess} in which H_{n-1} and H_n can cross. This contradicts that H_{n-1} and H_n are transverse in M_g^{ess} , so e_{n-1} and e_n must bound a square S_{n-1} in M_g^{ess} .

By the same reasoning, the edge e' parallel to e_n in S_{n-1} and e_{n-2} bound S_{n-2} , in which x_{n-2} is the minimal vertex. Repeating in this way, we find the square S_0 , with minimal vertex x_0 and in which H_0 and H cross. In S_0 there is an edge (x_0, v) dual to H , and $f(v) = y'$. \square

The following proposition will be used in the next section.

PROPOSITION 3.21. — *Suppose $g \in G$ is hyperbolic, and that $\langle g \rangle$ acts non-transversely on X_g^{ess} . Let C be a cube in X_g^{ess} of maximal dimension, A the set of half-spaces in A_g^+ dual to C , and o the minimal vertex of C . Then the following statements hold.*

- (1) *For every pair of half-spaces $H, H' \in A$, either $H \cap gH'$ or $H \supset gH'$.*
- (2) *$K \in [o, go]$ if and only if there exist $H, H' \in A$ such that $H \supseteq K \supset gH'$.*
- (3) *For every $K \in A_g^+$, there exist $r, s \in \mathbb{Z}$ and $H, H' \in A$ such that $g^r H \supset K \supset g^s H'$.*

Proof. — Since $X_g^{\text{ess}} = M_g^{\text{ess}}$ (by Proposition 3.17), there is an axis L for g containing o . It follows that $o \notin gH'$, because $o \notin H'$ (the unique edge e in L dual to H' separates o from ge). Let o^+ be the maximal vertex of C . By Lemma 3.13, o^+ is on a geodesic from o to go . Suppose H and gH' are not transverse. Then $o^+ \notin gH'$, because all half-spaces in $[o, o^+]$ are transverse. Now $o^+ \in H - gH'$, showing that $H \not\subset gH'$. Thus $H \supset gH'$ and (1) holds.

For statement (2), note that $A \subseteq [o, go]$, again by Lemma 3.13. If $H \supseteq K \supset gH'$ for some $H, H' \in A$, then $o \notin K$ and $go \in K$, and therefore $K \in [o, go]$. On the other hand, both C and gC have maximal dimension, so for every $K \in [o, go]$, there exist $H, H' \in A$ such that H and gH' are comparable (or equal) to K . Because $A \subset [o, go]$ and $gA \cap [o, go] = \emptyset$, we must have $H \supseteq K \supset gH'$.

Finally, for (3), we observe that

$$A_g^+ = \bigcup_{n \in \mathbb{Z}} [g^n o, g^{n+1} o].$$

Suppose $K \in [g^n o, g^{n+1} o]$. Applying (2) to $g^{-n}K$, there exist $H, H' \in A$ such that $g^n H \supseteq K \supset g^{n+1} H'$. Since $\langle g \rangle$ acts non-transversely, we also have $g^{n-1}H \supset g^n H$. The conclusion follows. \square

4. Non-transverse actions and efficient quasimorphisms

Here we give a general construction of a large family of quasimorphisms on groups acting on CAT(0) cube complexes. For the construction to succeed (i.e. to achieve bounded defect) we require one assumption.

DEFINITION 4.1. — *Let X be a CAT(0) cube complex with an action by G . The action is non-transverse if it is without inversion and also satisfies: there do not exist $H \in \mathcal{H}(X)$, $g \in G$ with $H \pitchfork gH$.*

This definition agrees with the earlier Definition 3.15 in the case of $\langle g \rangle$ acting on X_g^{ess} . First, such an action is always without inversion. Also, if $H \in A_g$ and H and gH are not transverse, then H and gH are nested by Remark 3.1; hence H and $g^k H$ are not transverse for any k .

Let X be a CAT(0) cube complex with a non-transverse action by G . Let γ be a segment in X , and consider the set $G\gamma = \{g\gamma : g \in G\}$; elements of this set are called *copies* of γ . Define the map $c_\gamma : X^2 \rightarrow \mathbb{Z}$ which assigns to each pair (x, y) the maximal cardinality of a pairwise non-overlapping collection of copies of γ in $[x, y]$.

Define

$$(4.1) \quad \omega_\gamma(x, y) = c_\gamma(x, y) - c_{\bar{\gamma}}(x, y).$$

Observe that $\omega_\gamma(y, x) = -\omega_\gamma(x, y)$ and $\omega_\gamma(gx, gy) = \omega_\gamma(x, y)$ for all $g \in G$.

LEMMA 4.2. — *If the action is non-transverse, then for all $x, y, z \in X$ with $y = m(x, y, z)$, there is a bound*

$$|\omega_\gamma(x, z) - \omega_\gamma(x, y) - \omega_\gamma(y, z)| \leq 2.$$

Proof. — By definition,

$$\begin{aligned} & |\omega_\gamma(x, z) - \omega_\gamma(x, y) - \omega_\gamma(y, z)| \\ &= \left| \left(c_\gamma(x, z) - c_\gamma(x, y) - c_\gamma(y, z) \right) - \left(c_{\bar{\gamma}}(x, z) - c_{\bar{\gamma}}(x, y) - c_{\bar{\gamma}}(y, z) \right) \right|. \end{aligned}$$

It will suffice to show that

$$(4.2) \quad c_\gamma(x, z) \leq c_\gamma(x, y) + c_\gamma(y, z) + 1$$

and

$$(4.3) \quad c_\gamma(x, y) + c_\gamma(y, z) \leq c_\gamma(x, z) + 1,$$

together with analogous statements for $\bar{\gamma}$.

Let $\{g_1\gamma, \dots, g_n\gamma\}$ be a collection of non-overlapping copies of γ in $[x, z]$ of cardinality $n = c_\gamma(x, z)$. By Lemma 2.5 these copies are pairwise nested, and hence up to re-indexing we can assume that

$$(4.4) \quad g_1\gamma > \dots > g_n\gamma.$$

If $g_k\gamma \not\subseteq [x, y] \cup [y, z]$ for some k , then y separates two half-spaces in $g_k\gamma$. It follows from (4.4) that $g_i\gamma \subseteq [x, y]$ for every $i < k$ and $g_i\gamma \subseteq [y, z]$ for every $i > k$. Thus $c_\gamma(x, y) + c_\gamma(y, z) \geq n - 1$, proving (4.2).

Now let $k = c_\gamma(x, y)$ and $\ell = c_\gamma(y, z)$. Let $A = \{g_1\gamma, \dots, g_k\gamma\}$ be a non-overlapping collection of copies of γ in $[x, y]$ and $B = \{g_{k+1}\gamma, \dots, g_{k+\ell}\gamma\}$ a non-overlapping collection of copies in $[y, z]$. As above, by Lemma 2.5, we may re-index A and B to arrange that

$$g_1\gamma > \dots > g_k\gamma \quad \text{and} \quad g_{k+1}\gamma > \dots > g_{k+\ell}\gamma.$$

We claim that $g_i\gamma$ and $g_j\gamma$ (with $i < j$) cannot overlap unless $i = k$ and $j = k + 1$. Discarding $g_k\gamma$, then obtain a non-overlapping collection in $[x, z]$ of cardinality $k + \ell - 1$, proving (4.3).

To prove the claim, suppose that $g_i\gamma \in A$ and $g_j\gamma \in B$ overlap. Then there are half-spaces $H, H' \in \gamma$ such that $g_iH \cap g_jH'$ in X . If $i < k$ then $g_kH' \cap g_jH'$, because $g_iH \supset g_kH'$ and $y \in g_kH' - g_jH'$. However, this contradicts the assumption of a non-transverse action. Hence $i = k$. Similarly, if $j > k + 1$ then $g_iH \cap g_{k+1}H$ because $g_{k+1}H \supset g_jH'$ and $y \in g_iH - g_{k+1}H$. Again, this contradicts non-transversality of the action, and therefore $j = k + 1$. This proves the claim, and equation (4.3). Finally, note that the analogues of (4.2) and (4.3) for $\bar{\gamma}$ are entirely similar. \square

Next define $\delta\omega_\gamma(x, y, z) = \omega_\gamma(x, y) + \omega_\gamma(y, z) + \omega_\gamma(z, x)$.

LEMMA 4.3. — *If the action is non-transverse, then for all $x, y, z \in X$ there is a bound $|\delta\omega_\gamma(x, y, z)| \leq 6$.*

Proof. — Let $m = m(x, y, z)$. By the previous lemma,

$$|\omega_\gamma(a, b) - \omega_\gamma(a, m) - \omega_\gamma(m, b)| \leq 2,$$

where $a, b \in \{x, y, z\}$ are distinct. Then

$$\begin{aligned}
 |\delta\omega_\gamma(x, y, z)| &= |\omega_\gamma(x, y) + \omega_\gamma(y, z) + \omega_\gamma(z, x) \\
 &\quad + \omega_\gamma(x, m) - \omega_\gamma(x, m) + \omega_\gamma(y, m) \\
 &\quad - \omega_\gamma(y, m) + \omega_\gamma(z, m) - \omega_\gamma(z, m)| \\
 &\leq |\omega_\gamma(x, y) - \omega_\gamma(x, m) - \omega_\gamma(m, y)| \\
 &\quad + |\omega_\gamma(y, z) - \omega_\gamma(y, m) - \omega_\gamma(m, z)| \\
 &\quad + |\omega_\gamma(z, x) - \omega_\gamma(z, m) - \omega_\gamma(m, x)| \\
 &\leq 6. \quad \square
 \end{aligned}$$

At this point we are ready to define quasimorphisms associated to γ . We will define two functions, ψ_γ and φ_γ , which produce the same homogeneous quasimorphism $\widehat{\psi}_\gamma = \widehat{\varphi}_\gamma$. The second function φ_γ has the definition we want to use, but ψ_γ is needed to establish the bound on defect.

Fix a base vertex $x_O \in X$ and define $\psi_\gamma: G \rightarrow \mathbb{R}$ by

$$(4.5) \quad \psi_\gamma(g) = \omega_\gamma(x_O, gx_O).$$

Next, for each $g \in G$ choose a vertex $x_g \in X_g$. Define $\varphi_\gamma: G \rightarrow \mathbb{R}$ by

$$(4.6) \quad \varphi_\gamma(g) = \omega_\gamma(x_g, gx_g).$$

LEMMA 4.4. — *If the action is non-transverse, then ψ_γ is a quasimorphism of defect at most 6.*

Proof. — For any $g_1, g_2 \in G$ we have

$$\begin{aligned}
 |\psi_\gamma(g_1g_2) - \psi_\gamma(g_1) - \psi_\gamma(g_2)| &= |\omega_\gamma(x_O, g_1g_2x_O) - \omega_\gamma(x_O, g_1x_O) - \omega_\gamma(x_O, g_2x_O)| \\
 &= |\omega_\gamma(x_O, g_1g_2x_O) + \omega_\gamma(g_1x_O, x_O) + \omega_\gamma(g_2x_O, x_O)| \\
 &= |\omega_\gamma(x_O, g_1g_2x_O) + \omega_\gamma(g_1x_O, x_O) + \omega_\gamma(g_1g_2x_O, g_1x_O)| \\
 &= |\delta\omega_\gamma(x_O, g_1g_2x_O, g_1x_O)| \\
 &\leq 6,
 \end{aligned}$$

by Lemma 4.3. □

LEMMA 4.5. — *If the action is non-transverse, then $\psi_\gamma - \varphi_\gamma$ is uniformly bounded. Hence φ_γ is a quasimorphism, $\widehat{\varphi}_\gamma = \widehat{\psi}_\gamma$, and $\widehat{\varphi}_\gamma$ has defect at most 12.*

Proof. — For any $g \in G$ we have

$$\begin{aligned}
 |\psi_\gamma(g) - \varphi_\gamma(g)| &= |\omega_\gamma(x_O, gx_O) - \omega_\gamma(x_g, gx_g) + \omega_\gamma(gx_O, x_g) - \omega_\gamma(gx_O, x_g)| \\
 &= |\omega_\gamma(x_O, gx_O) + \omega_\gamma(gx_O, x_g) - (\omega_\gamma(gx_O, x_g) + \omega_\gamma(x_g, gx_g)) \\
 &\quad + \omega_\gamma(x_g, x_O) - \omega_\gamma(x_g, x_O)| \\
 &\leq |\omega_\gamma(x_O, gx_O) + \omega_\gamma(gx_O, x_g) + \omega_\gamma(x_g, x_O)| \\
 &\quad + |\omega_\gamma(gx_O, x_g) + \omega_\gamma(x_g, gx_g) + \omega_\gamma(x_g, x_O)| \\
 &= |\delta\omega_\gamma(x_O, gx_O, x_g)| + |\delta\omega_\gamma(gx_O, x_g, gx_g)| \\
 &\leq 12,
 \end{aligned}$$

by Lemma 4.3. This shows that $\psi_\gamma - \varphi_\gamma$ is uniformly bounded. The other conclusions follow immediately from Lemma 2.2 and Lemma 2.1. \square

Note that the equality $\widehat{\varphi}_\gamma = \widehat{\psi}_\gamma$ also implies that this quasimorphism is independent of the choices of basepoints used to define φ_γ and ψ_γ .

Remark 4.6. — The bounds in the preceding lemmas can be improved by a factor of 2 in the special case where X is a 1-dimensional CAT(0) cube complex (that is, a simplicial tree). In this case, half-spaces are never transverse, so two segments overlap if and only if they have non-empty intersection. We obtain an improvement in equation (4.3), which becomes instead

$$(4.3') \quad c_\gamma(x, y) + c_\gamma(y, z) \leq c_\gamma(x, z)$$

since there is no need to discard $g_k\gamma$ from the collection of segments in $[x, z]$. This leads to the bounds

$$|\omega_\gamma(x, z) - \omega_\gamma(x, y) - \omega_\gamma(y, z)| \leq 1$$

in Lemma 4.2, $|\delta\omega_\gamma(x, y, z)| \leq 3$ in Lemma 4.3, and a defect of at most 6 in Lemma 4.5. Thus we have a new proof of Theorem 6.6 of [12], which is the statement that these quasimorphisms have defect at most 6.

At this point, one could enhance Theorem A to say that $\text{scl}(g) \geq 1/12$ when X is a tree, but this already follows from Theorem 6.9 of [12].

Bounded cohomology of right-angled Artin groups

Recall that for any group G , we denote by $\widetilde{\text{QH}}(G)$ the space of homogeneous quasimorphisms on G , modulo homomorphisms. It is a subspace of the second bounded cohomology $H_b^2(G; \mathbb{R})$.

If G is a non-abelian right-angled Artin group, then there is a retraction onto a non-abelian free subgroup H , which induces injections $H_b^2(H; \mathbb{R}) \hookrightarrow H_b^2(G; \mathbb{R})$ and $\widehat{\text{QH}}(H) \hookrightarrow \widehat{\text{QH}}(G)$. Therefore the spaces $H_b^2(G; \mathbb{R})$ and $\widehat{\text{QH}}(G)$ are infinite-dimensional (we thank the referee for suggesting this viewpoint). The next result provides another way of seeing this, using the quasimorphisms $\widehat{\varphi}_\gamma$. We note that the quasimorphisms obtained below appear to be quite different from those arising via the retraction to H , even though their restrictions to H agree.

PROPOSITION 4.7. — *Let $G = A_\Gamma$ be a non-abelian right-angled Artin group, and X the natural cube complex on which G acts. Then there is an infinite family $\{\gamma_i\}$ of segments in $\mathcal{H}(X)$ such that the homogeneous quasimorphisms $\{\widehat{\varphi}_{\gamma_i}\}$ are linearly independent in $\widehat{\text{QH}}(G)$.*

Proof. — Let a, b be standard generators of G which generate a free subgroup $H < G$. We shall show that every “non-overlapping” Brooks quasimorphism on H is the restriction of a quasimorphism $\widehat{\varphi}_\gamma$ for some γ . By [22, Proposition 5.1] there is an infinite linearly independent family of Brooks quasimorphisms in $\widehat{\text{QH}}(H)$, and their extensions will be independent in $\widehat{\text{QH}}(G)$.

If w is a reduced word in a, b , the non-overlapping Brooks quasimorphism $\widehat{B}_w: \langle a, b \rangle \rightarrow \mathbb{R}$ is the homogenization of the quasimorphism $B_w = C_w - C_{\bar{w}}$, where $C_w(g)$ is the maximal number of disjoint subwords of g (considered as a reduced word) which equal w . In the 1-skeleton of X there is an edge path labeled by the word w , starting at a vertex x and ending at y . Because a and b do not commute, no two half-spaces dual to this segment can cross. Thus $[x, y]$ is a *segment*, which we denote by $\gamma(w)$. Modulo the G -action on X , $\gamma(w)$ is uniquely determined by w .

We claim that $B_w(g) = \varphi_{\gamma(w)}(g)$ for every $g \in H$, and therefore \widehat{B}_w is the restriction of $\widehat{\varphi}_{\gamma(w)}$ to H . If an element $g \in H$ is considered as a reduced word, it has a combinatorial axis in X which is labeled by g^∞ . The half-spaces dual to this axis never cross, and so the partial ordering on A_g^+ is a linear ordering. Thus X_g^{ess} is one-dimensional and the axis is an embedded copy of X_g^{ess} in $X_g \subset X$. Let x_g be a vertex on this axis at the beginning of the word g ; this is the basepoint for the definition of $\varphi_{\gamma(w)}(g)$. Now segments in $[x_g, gx_g]$ correspond bijectively with subwords of g via the labelling, and so $B_w(g) = \varphi_{\gamma(w)}(g)$. \square

5. Dilworth's theorem and equivariant embeddings

Let P be a partially ordered set. A *chain* in P is a subset that is linearly ordered. A chain is *maximal* if it is not properly contained in another chain. An *antichain* in P is a subset such that no two elements are comparable to each other. The *width* of P is the maximal cardinality of an antichain (which may be ∞).

LEMMA 5.1 (Dilworth's theorem). — *Let P be a partially ordered set. If P has width $d < \infty$ then there is a partition of P into d chains. Furthermore, there is such a partition such that one of the chains is maximal.*

This first conclusion is the traditional statement of the theorem. The second claim can be proved using Hausdorff's maximal principle.

The partition of P into chains provided by the theorem will be called a *Dilworth partition*.

DEFINITION 5.2. — *Let P be a partially ordered set that admits an order-preserving free action by an infinite cyclic group $\langle g \rangle$. Let A be an antichain in P . We say A is $\langle g \rangle$ -descending if $ga \not\geq a'$ for all $a, a' \in A$. We say that A spans P if for each $p \in P$ there exist $a, a' \in A$ and $r, s \in \mathbb{Z}$ such that $g^r a > p > g^s a'$.*

We further define the subsets

$$[A, gA] = \{p \in P : x \geq p \geq y \text{ for some } x, y \in (A \cup gA)\},$$

which equals

$$\{p \in P : x \geq p \geq y \text{ for some } x \in A, y \in gA\}$$

if A is $\langle g \rangle$ -descending, and

$$[A, gA] = [A, gA] - gA, \quad (A, gA] = [A, gA] - A.$$

LEMMA 5.3 (Equivariant Dilworth theorem). — *Let P be a partially ordered set of width $d < \infty$ with an order-preserving free action by an infinite cyclic group $\langle g \rangle$. Suppose further that there is an antichain A of cardinality d that is both $\langle g \rangle$ -descending and spans P . Then there is a $\langle g \rangle$ -invariant partition of P into d chains whose intersection with $[A, gA]$ is a Dilworth partition which includes a maximal chain in $[A, gA]$.*

Proof. — Apply Lemma 5.1 to the partially ordered set $[A, gA]$ to obtain a partition by chains $[A, gA] = Q_1 \cup \dots \cup Q_d$, with Q_1 maximal in $[A, gA]$. Each Q_i contains exactly one element of A and one of gA , since these

are antichains of cardinality d . We claim that these are the maximal and minimal elements, respectively, of Q_i .

Suppose the unique element a of $A \cap Q_i$ is not maximal in Q_i . If $p \in Q_i$ satisfies $p > a$ then, since $p \in (A, gA]$, we must have $x > p$ for some $x \in A$. Then $x > p > a$, contradicting that A is an antichain. By a similar argument, the unique element of $gA \cap Q_i$ is minimal in Q_i .

Now label the elements of A and define a permutation σ as follows: a_i is the maximal element of Q_i and $ga_{\sigma(i)}$ is the minimal element of Q_i , for $i = 1, \dots, d$. Define the sets

$$P_i = \bigcup_{k \in \mathbb{Z}} g^k Q_{\sigma^k(i)}$$

for each i . Note that for each k , the element $g^k a_{\sigma^k(i)}$ is both the minimum of $g^{k-1} Q_{\sigma^{k-1}(i)}$ and the maximum of $g^k Q_{\sigma^k(i)}$. Hence P_i is a chain, being a concatenation of chains. Since $\langle g \rangle$ acts freely on P , the chains P_i are disjoint. Their union is the set $\bigcup_{k \in \mathbb{Z}} g^k [A, gA]$. It is immediate that $gP_{\sigma(i)} = P_i$, so the partition of $\bigcup_{k \in \mathbb{Z}} g^k [A, gA]$ by the chains P_i is preserved by g . It remains to show that this set is all of P .

Given $p \in P$, let a, a', r, s be given such that $g^r a > p > g^s a'$. First we claim that $s > r$. If not, then $r > s$. Writing $a = a_i$ we have $g^r a_i < g^{r-1} a_{\sigma^{-1}(i)} < \dots < g^s a_{\sigma^{s-r}(i)}$, whence $g^s a' < g^s a_{\sigma^{s-r}(i)}$, a contradiction since $g^s A$ is an antichain.

Next we show that $p \in \bigcup_{k \in \mathbb{Z}} g^k [A, gA]$, by induction on $s - r$. Clearly we may assume that $p \notin \bigcup_{k \in \mathbb{Z}} g^k A$. If $s - r = 1$ then we already have $p \in g^r [A, gA]$. If $s - r > 1$ then consider the (maximal) antichain $g^{r+1} A$. It contains an element $g^{r+1} a''$ which is comparable to p , by maximality. Then either $g^r a > p > g^{r+1} a''$ or $g^{r+1} a'' > p > g^s a'$, and in either case the induction hypothesis yields the conclusion that $p \in \bigcup_{k \in \mathbb{Z}} g^k [A, gA]$. \square

Equivariant Euclidean embeddings

Let \mathbb{R}^d be equipped with its standard integer cubing. Given a coordinate i and an integer n , we define:

$$H_n^i = \{(x_1, \dots, x_d) \in \mathbb{R}^d : x_i \geq n + 1/2\}.$$

Note that H_n^i and H_m^j are transverse in \mathbb{R}^d if and only if $i \neq j$. We also define $\mathcal{H}^i = \{H_n^i : n \in \mathbb{Z}\}$, and set

$$\mathcal{H}^+(\mathbb{R}^d) = \mathcal{H}^1 \sqcup \dots \sqcup \mathcal{H}^d.$$

The set of half-spaces of \mathbb{R}^d is $\mathcal{H}(\mathbb{R}^d) = \mathcal{H}^+(\mathbb{R}^d) \sqcup \mathcal{H}^-(\mathbb{R}^d)$, where $\mathcal{H}^-(\mathbb{R}^d) = \{\bar{H} : H \in \mathcal{H}^+(\mathbb{R}^d)\}$.

PROPOSITION 5.4. — *Let $g \in G$ be hyperbolic and suppose $\langle g \rangle$ acts non-transversely on X_g^{ess} . Let C be a cube in X_g^{ess} of dimension $d = \dim(X_g^{\text{ess}})$ and let A be the set of elements of A_g^+ dual to the edges of C . Then there exist a $\langle g \rangle$ -action on \mathbb{R}^d and a $\langle g \rangle$ -equivariant isometric embedding $\phi: X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$ satisfying the following properties:*

- (1) $\phi(C) = [0, 1]^d \subset \mathbb{R}^d$.
- (2) The induced map $\phi_*: A_g \rightarrow \mathcal{H}(\mathbb{R}^d)$ is a bijection, with $\phi_*(A_g^+) = \mathcal{H}^+(\mathbb{R}^d)$.
- (3) The set $[A, gA] \cap \phi_*^{-1}(\mathcal{H}^1)$ is tightly nested in X_g^{ess} .
- (4) $[A, gA] = [o, go]$, where o is the minimal vertex of C .

By property (2), we can henceforth identify elements of A_g^+ with their corresponding half-spaces in $\mathcal{H}^+(\mathbb{R}^d)$ and we shall denote the corresponding decomposition as $A_g^+ = \mathcal{H}^1 \sqcup \dots \sqcup \mathcal{H}^d$. By property (3), every subsegment of $[A, gA] \cap \mathcal{H}^1$ is tightly nested in A_g^+ . We will call $\gamma = [A, gA] \cap \mathcal{H}^1$ the taut segment of the embedding; $[A, gA] \cap \mathcal{H}^1$ the extended taut segment; and the map ϕ a taut $\langle g \rangle$ -equivariant embedding of X_g^{ess} into \mathbb{R}^d .

Proof of Proposition 5.4. — Let $P = A_g^+$ be partially ordered by inclusion; note that this partial ordering is preserved by $\langle g \rangle$. It has width d since X_g^{ess} has dimension d , and A is an antichain of cardinality d . By Proposition 3.21 (1), A is $\langle g \rangle$ -descending. By Proposition 3.21 (3), A spans P . We also have that $[A, gA] = [o, go]$, by Proposition 3.21 (2), and therefore (4) holds.

Now apply Lemma 5.3 to P to obtain a $\langle g \rangle$ -invariant partition of P into d chains P_1, \dots, P_d . Without loss of generality, we may assume that $P_1 \cap [A, gA]$ is a maximal chain in $[A, gA]$. Note that each chain P_i is bi-infinite, since $\langle g \rangle$ acts freely on it.

For each i , let K_i be the unique element of $P_i \cap A$. There is an order-preserving bijection $P_i \rightarrow \mathcal{H}^i$ induced by sending K_i to H_0^i . The resulting bijection $A_g^+ \rightarrow \mathcal{H}^+(\mathbb{R}^d)$ extends to a bijection $\phi_*: A_g \rightarrow \mathcal{H}(\mathbb{R}^d)$ in an obvious way.

We now define an isometric embedding $\phi: X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$ whose induced map on half-spaces is ϕ_* . For any $x \in \mathbb{R}^d$, denote by x_i its i -th coordinate. Let $v \in X_g^{\text{ess}}$ be any vertex. For each i , let $K \in P_i$ be the largest element such that $v \notin K$. Define $\phi(x)_i = n$, where $\phi_*(K) = H_n^i$. This defines an embedding of the vertices of X_g^{ess} into \mathbb{R}^d . Two vertices v and w in X_g^{ess} bound an oriented edge (v, w) dual to $K \in P_i$ if and only if $\phi(w)_i = \phi(v)_i + 1$

and $\phi(w)_j = \phi(v)_j$ for all $j \neq i$. Therefore ϕ extends to an embedding of the 1-skeleton of X_g^{ess} , and hence extends to all of X_g^{ess} . It is immediate that ϕ induces the same map on half-spaces as ϕ_* , so property (2) holds. By construction, o is mapped to the origin and the vertex of C opposite o is mapped to $(1, \dots, 1)$, so (1) holds.

By $\langle g \rangle$ -invariance of the partition, there is a permutation σ such that $gP_{\sigma(i)} = P_i$. For each i let $n_i = \phi(g(o))_i$. That is, n_i is the shift given by the bijection $g: P_{\sigma(i)} \rightarrow P_i$, relative to the basepoints $K_{\sigma(i)}$ and K_i . Then, for every vertex $v \in X_g^{\text{ess}}$, we have

$$\phi(g(v))_i = \phi(v)_{\sigma(i)} + n_i.$$

This allows us to define an action of $\langle g \rangle$ on \mathbb{R}^d : for every $x \in \mathbb{R}^d$ let $g(x)_i = x_{\sigma(i)} + n_i$. By construction, ϕ is $\langle g \rangle$ -equivariant.

For property (3), note that $[A, gA] \cap \phi_*^{-1}(\mathcal{H}^1) = P_1 \cap [A, gA]$. Suppose $K' \supset K \supset K''$ for some $K \in A_g^+$ and $K', K'' \in P_1 \cap [A, gA]$. There is a unique $i \in \mathbb{Z}$ such that $K \in [g^i A, g^{i+1} A]$. If $i < 0$, then $K \supset H$ for some $H \in A$, which contradicts $K' \supset K$. If $i > 0$, then $gH \supset K$ for some $H \in A$. But $gH \supset K$ contradicts $K \supset K''$, so $K \in [A, gA]$. By maximality, $K \in P_1 \cap [A, gA]$. This shows that $P_1 \cap [A, gA]$ is tightly nested. \square

An example

Let A_Γ be the right-angled Artin group with Γ the pentagon graph:

$$A_\Gamma = \langle a, b, c, d, e \mid [a, b] = [b, c] = [c, d] = [d, e] = [e, a] = 1 \rangle.$$

The element $g = abcde$ is hyperbolic, and part of its essential characteristic set X_g^{ess} is shown in Figure 5.1. The figure also demonstrates the equivariant embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^2$. The action of g on \mathbb{R}^2 (extending the natural action on X_g^{ess}) is by a glide reflection whose axis is a diagonal line through the center of the figure. The A_Γ -invariant labeling of the edges of X_g^{ess} by generators of A_Γ is also shown. For this particular choice of g , the essential characteristic set has the property that the equivariant embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^2$ is unique, up to a change of coordinates in \mathbb{R}^2 by a cubical automorphism. The action on \mathbb{R}^2 is always by a glide reflection, for this g . Other elements have characteristic sets that may embed in more than one way, with g acting on \mathbb{R}^2 either as a translation or a glide reflection (depending on the embedding).




Figure 5.1. The subcomplex X_g^{ess} embedded in \mathbb{R}^2 for the element $g = abcde$ in the pentagon RAAG. The extended action of g on \mathbb{R}^2 is by a glide reflection. The three blue half-spaces (with labels a, c, e) are taken to the three red half-spaces.

The staircase

Our goal in the rest of the paper will be to associate to each hyperbolic element g a segment γ such that $\widehat{\varphi}_\gamma(g) \geq 1$. Bavard Duality then will allow us to conclude that $\text{scl}(g) \geq 1/24$. Here we illustrate one of the difficulties in finding such segments.

Consider \mathbb{R}^2 with its standard integer cubing, and let X be the subcomplex obtained by removing all vertices $(x, y) \in \mathbb{Z}^2$ with $y < x - 1$ (see Figure 5.2). We will refer to X as the *staircase*.

Let $G = \langle g \rangle$, where g is the restriction of the translation $(x, y) \mapsto (x + 2, y + 2)$ to X . Note that $X = X_g = X_g^{\text{ess}}$. Let $x_g = (0, 0)$. Consider the two half-spaces

$$H_1 = \{(x, y) \in X : y \geq 1/2\} \quad \text{and} \quad H_2 = \{(x, y) \in X : x \geq 3/2\}$$

shown in blue on the left hand side of Figure 5.2. The set $\gamma = \{H_1, H_2\}$ is a segment in $[x_g, gx_g]$ (recall that this means γ is tightly nested). For any

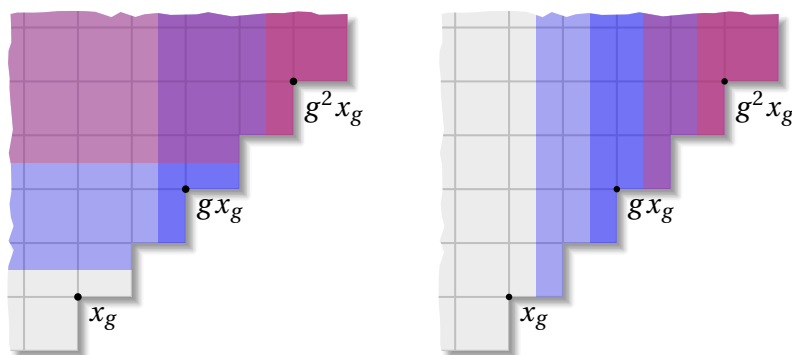


Figure 5.2. Some tightly nested pairs in the staircase: $\{H_1, H_2\}$ in blue, $g\{H_1, H_2\}$ in red.

positive integer n , $g^n H_1$ and H_2 are transverse, so γ and $g^n \gamma$ overlap. It follows that $c_\gamma(g^n) = 1$ for all n , which means that $\widehat{\varphi}_\gamma(g) \leq 0$.

A better choice of segment $\gamma \subset [x_g, gx_g]$ is shown on the right hand side of Figure 5.2. The half-space H_1 has been replaced by $\{(x, y) \in X : x \geq 1/2\}$. In this case, γ and $g\gamma$ do not overlap, and in fact $c_\gamma(g^n) = n$ for all positive n .

This example indicates that from the point of view of an equivariant Euclidean embedding, one should choose a segment γ which lies in a single coordinate direction in \mathbb{R}^d to ensure that $c_\gamma(g^n)$ grows linearly with n . (Keeping $c_\gamma(g^n)$ bounded is a much more serious hurdle to be dealt with in Sections 8 and 9.) It is for this reason that we required one of the chains in the Dilworth partition to be maximal in Lemma 5.3, leading to property (3) in Proposition 5.4. This property ensures that in at least one coordinate direction of \mathbb{R}^d , consecutive half-spaces in \mathbb{R}^d are tightly nested in X_g^{ess} , and therefore define segments in X_g^{ess} .

6. Quadrants

In this section we present two basic tools for working with equivariant Euclidean embeddings: the Quadrant Lemma and the Elbow Lemma. They are useful in determining which cubes in \mathbb{R}^d are occupied by X_g^{ess} . Let x_i and x_j be coordinates of \mathbb{R}^d . We will denote by $p_{ij} : \mathbb{R}^d \rightarrow \mathbb{R}^2$ the projection of \mathbb{R}^d onto the $x_i x_j$ -coordinate plane.

Consider a $\langle g \rangle$ -equivariant embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$, where $d = \dim X_g^{\text{ess}}$. Recall that via this embedding we identify elements of A_g^+ with their corresponding half-spaces \mathcal{H}^+ in $\mathcal{H}(\mathbb{R}^d)$. We will generally suppress the embedding itself and will treat X_g^{ess} as a subcomplex of \mathbb{R}^d .

Remarks 6.1.

(a). — Recall from Lemma 3.8 that if $H, H' \in A_g^+$ then H and H' are transverse in X if and only if they are transverse in X_g^{ess} . When this occurs, they will also be transverse in \mathbb{R}^d , but not conversely.

(b). — Expressing these two half-spaces as H_n^i and H_m^j , the subcomplex $p_{ij}(X_g^{\text{ess}})$ of \mathbb{R}^2 contains the square $[n, n + 1] \times [m, m + 1]$ if and only if H_n^i and H_m^j are transverse in X_g^{ess} . To see this, note that the latter occurs if and only if ∂H_n^i and ∂H_m^j cross in some square in $X_g^{\text{ess}} \subseteq \mathbb{R}^d$. If they cross in the square S , then the image of S in \mathbb{R}^2 is a cube which is dual to both H_n^i and H_m^j , which must be the square $[n, n + 1] \times [m, m + 1]$.

(c). — If $H, H' \in A_g^+$ then H and H' are (tightly) nested in X if and only if they are (tightly) nested in X_g^{ess} . If they are nested in \mathbb{R}^d then they are nested in X_g^{ess} , but not conversely. There is no a priori relation between being tightly nested in X and being tightly nested in \mathbb{R}^d . Half-spaces H and H' may be tightly nested in X and not tightly nested in \mathbb{R}^d , and vice versa.

DEFINITION 6.2. — A quadrant in \mathbb{R}^d is an open set of the form

$$\{(x_1, \dots, x_d) : x_i < n \text{ and } x_j > m\}$$

where $i \neq j$ and $m, n \in \mathbb{Z}$. Often, one of the coordinates x_i or x_j will be designated as the horizontal coordinate. If x_i is horizontal, then the quadrant above is called a northwest quadrant, and if x_j is horizontal, it is called a southeast quadrant.

LEMMA 6.3 (Quadrant Lemma). — Let $H_n^i, H_m^j \in \mathcal{H}^+$ be half-spaces with $i \neq j$ and suppose x_i is horizontal. Then one of the following holds:

- (1) H_n^i and H_m^j are transverse in X_g^{ess} ;
- (2) $H_n^i \supset H_m^j$ in X_g^{ess} and X_g^{ess} is disjoint from the northwest quadrant $\{x_i < n + 1, x_j > m\}$;
- (3) $H_n^i \subset H_m^j$ in X_g^{ess} and X_g^{ess} is disjoint from the southeast quadrant $\{x_i > n, x_j < m + 1\}$.

The quadrant in case 2 or 3 that is disjoint from X_g^{ess} will be called the quadrant generated by H_n^i and H_m^j .

Put another way, if $p_{ij}(X_g^{\text{ess}})$ does not contain the square $[n, n+1] \times [m, m+1]$, then it does not meet the quadrant generated by that square; see Figure 6.1.

Whenever $H \in \mathcal{H}^i$, $K \in \mathcal{H}^j$ are nested in X_g^{ess} with $i \neq j$, denote by $Q(H, K)$ the quadrant generated by this pair of half-spaces. By definition, it is always disjoint from X_g^{ess} .

Proof. — If the first alternative does not hold, then the corresponding half-spaces in X_g^{ess} are nested, by Remark 3.1. That is, one of $H_n^i \cap X_g^{\text{ess}}$, $H_m^j \cap X_g^{\text{ess}}$ contains the other. Suppose $H_n^i \cap X_g^{\text{ess}}$ contains $H_m^j \cap X_g^{\text{ess}}$. If a vertex $v = (v_1, \dots, v_d)$ of X_g^{ess} satisfies $v_j \geq m+1$ then $v \in H_m^j \cap X_g^{\text{ess}}$, so $v \in H_n^i$. Hence $v_i \geq n+1$, showing that $v \notin \{x_i \leq n, x_j \geq m+1\}$. Thus the second alternative holds. Similarly, if $H_m^j \cap X_g^{\text{ess}}$ contains $H_n^i \cap X_g^{\text{ess}}$, then the third alternative holds. \square

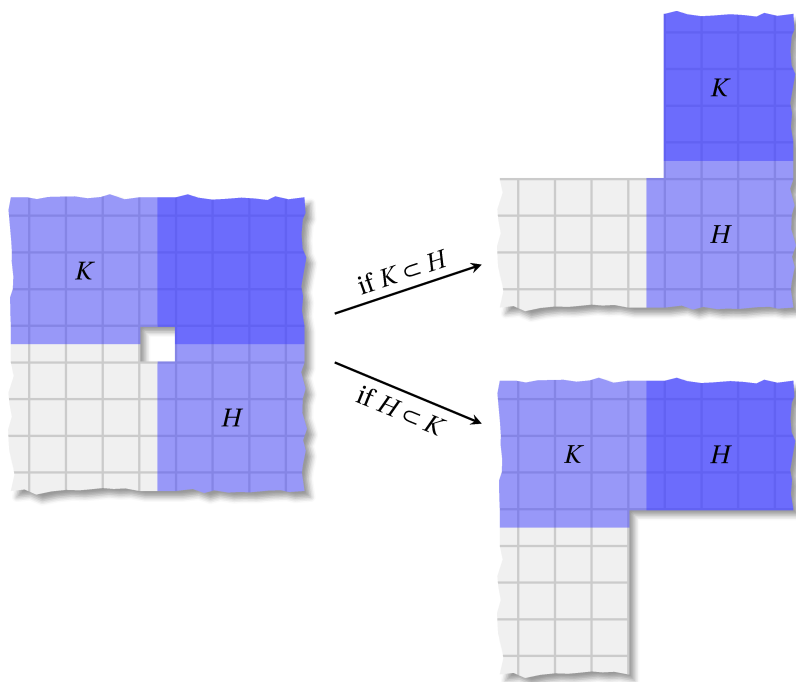


Figure 6.1. The Quadrant Lemma: if X_g^{ess} avoids the interior of a square, it also avoids a northwest or southeast quadrant.

LEMMA 6.4. — Suppose $H \in \mathcal{H}^i$ and $K, K' \in A_g^+ - \mathcal{H}^i$ are such that K, K' are tightly nested and the pairs H, K and H, K' are nested in X_g^{ess} . Let x_i be horizontal. Then the quadrants $Q(H, K)$ and $Q(H, K')$ both face northwest or both face southeast.

Proof. — Suppose without loss of generality that $K \subset K'$. If $Q(H, K)$ faces northwest and $Q(H, K')$ faces southeast, then $K \subset H$ and $H \subset K'$ by the Quadrant Lemma. Now H violates the assumption that K, K' are tightly nested. If $Q(H, K')$ faces northwest and $Q(H, K)$ faces southeast, then $K' \subset H$ and $H \subset K$. Hence $K' \subset K$, a contradiction. \square

LEMMA 6.5 (Elbow Lemma). — Suppose $H_n^i \subset H_m^j$ are tightly nested in X_g^{ess} where $i \neq j$. Then the edges $\{n\} \times [m, m + 1]$ and $[n, n + 1] \times \{m + 1\}$ are contained in $p_{ij}(X_g^{\text{ess}})$.

The two edges form an “elbow” at the corner of the quadrant $Q(H_n^i, H_m^j) = \{x_i > n, x_j < m + 1\}$ (and X_g^{ess} avoids this quadrant, by the Quadrant Lemma).

Proof. — Designate x_i as the horizontal coordinate. We consider the edge $\{n\} \times [m, m + 1]$ (the other case being entirely similar).

If the square $[n - 1, n] \times [m, m + 1]$ is in $p_{ij}(X_g^{\text{ess}})$, then so is the edge $\{n\} \times [m, m + 1]$ and we are done. If not, then the half-spaces H_{n-1}^i and H_m^j are nested in X_g^{ess} . We cannot have $H_{n-1}^i \subset H_m^j$ in X_g^{ess} , because $H_n^i \subset H_{n-1}^i$ and H_n^i, H_m^j are tightly nested. Therefore, $H_m^j \subset H_{n-1}^i$ in X_g^{ess} . By the Quadrant Lemma, the northwest quadrant generated by the square $[n - 1, n] \times [m, m + 1]$ is disjoint from $p_{ij}(X_g^{\text{ess}})$. Similarly, since $H_n^i \subset H_m^j$, the southeast quadrant generated by the square $[n, n + 1] \times [m, m + 1]$ is also disjoint from $p_{ij}(X_g^{\text{ess}})$. The edge $\{n\} \times [m, m + 1]$ now provides the only passage across the strip $\mathbb{R} \times [m, m + 1]$. It must be in $p_{ij}(X_g^{\text{ess}})$, or X_g^{ess} could not contain an axis. \square

Remark 6.6. — The Quadrant Lemma and the Elbow Lemma do not use the fact the embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$ is equivariant. These results hold (with X_g^{ess} replaced with Y) whenever Y is a convex subcomplex of a CAT(0) cube complex X and there is a Euclidean embedding $Y \hookrightarrow \mathbb{R}^d$ that induces a bijection between $\mathcal{H}(Y)$ and $\mathcal{H}(\mathbb{R}^d)$.

7. RAAG-like actions on cube complexes

Recall from Section 2 that every right-angled Artin group A_Γ acts on a CAT(0) cube complex X_Γ , and that the oriented edges of X_Γ admit an

A_Γ -invariant labeling by the generators and their inverses. Also, there is an induced A_Γ -invariant labeling of the half-spaces of X_Γ .

As noted earlier, properties of the half-space labeling lead to many useful observations about X_Γ and its A_Γ -action. The definition below is based on some of these properties of X_Γ .

DEFINITION 7.1. — *Let X be a CAT(0) cube complex with an action by G . The action is RAAG-like if it is without inversion and also satisfies:*

- (1) *there do not exist $H \in \mathcal{H}(X)$, $g \in G$ with $H \pitchfork gH$,*
- (2) *there do not exist tightly nested $H, H' \in \mathcal{H}(X)$, $g \in G$ with $H \pitchfork gH'$,*
- (3) *there do not exist $H \in \mathcal{H}(X)$, $g \in G$ with H and $g\overline{H}$ tightly nested.*

When the G -action on X is understood, we may simply say that X is RAAG-like.

Remark 7.2. — If one has a G -action on X with an inversion, it is customary to perform a cubical subdivision to obtain an action without inversion. We note here that the resulting action will never be RAAG-like, since it will violate property 7.1(3).

LEMMA 7.3. — *For every simplicial graph Γ , the action of A_Γ on X_Γ is RAAG-like.*

Proof. — We have already observed in Section 2 that A_Γ acts without inversion on X_Γ . We have also observed that since boundaries of squares in X_Γ are labeled by commutators $[v, w]$ with $v \neq w$, no two half-spaces in X_Γ with the same label can cross. Property (1) follows immediately.

For (2), suppose H, H' are tightly nested half-spaces in X_Γ . Then there is a vertex $x \in X_\Gamma$ and a pair of edges e, e' both incident to x , such that e is dual to H and e' is dual to H' (modulo orientations). Since H and H' do not cross, the edges e and e' are not in the boundary of a common square; hence their labels do not commute in A_Γ . It follows that no two half-spaces bearing these labels (or their inverses) can cross. In particular, H and gH' cannot cross for any $g \in A_\Gamma$.

For (3), suppose H and $g\overline{H}$ are tightly nested for some $H \in \mathcal{H}(X_\Gamma)$, $g \in A_\Gamma$. Switching H and \overline{H} if necessary, we may assume that $H \subset g\overline{H}$. Since they are tightly nested, there is a pair of (oriented) edges e, e' with common initial vertex x such that e is dual to H and e' is dual to gH . Then e and e' bear the same label v , since the half-space labeling is A_Γ -invariant. However, vertices in X_Γ have exactly one edge incident to them with any

given label (being lifts of the same oriented edge of $K(A_\Gamma, 1)$ at the same initial vertex). This contradiction establishes property (3). \square

Remark 7.4. — The properties of Definition 7.1 correspond precisely to the defining properties of *special cube complexes* due to Haglund and Wise [19], as enumerated in [26]. More specifically, if G acts freely on a CAT(0) cube complex X , then the action is RAAG-like if and only if X/G is special.

The properties correspond as follows. Property (1) means that immersed hyperplanes in X/G are embedded (and hence can simply be called hyperplanes). G acting on X without inversion means that all hyperplanes in X/G are two-sided. Property (2) means that pairs of hyperplanes in X/G do not inter-oscuate. Property (3) means that hyperplanes in X/G do not self-oscuate.

Remark 7.5. — Note that Definition 7.1(1) in particular means that the action of G on X is non-transverse. Therefore, for any hyperbolic element $g \in G$, the action of $\langle g \rangle$ on X_g^{ess} is non-transverse. Hence, by Proposition 3.17, $X_g^{\text{ess}} = M_g^{\text{ess}}$ for all hyperbolic elements $g \in G$.

8. Tightly nested segments in the essential characteristic set

In Section 6, we presented some general tools for studying equivariant Euclidean embeddings of X_g^{ess} . Here we develop more specialized results to be used in proving the main theorem. Generally speaking, these results deal with situations where there is a tightly nested segment $\sigma \subset A_g^+$ in one coordinate direction \mathcal{H}^i , and an element $f \in G$ such that $f\sigma \subset A_g^+$.

For the rest of this section and the next section, we will assume that X is a CAT(0) cube complex with a RAAG-like G -action.

Fix a hyperbolic element $g \in G$ and apply Proposition 5.4 to obtain a taut $\langle g \rangle$ -equivariant embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$. Recall that a cube of maximal dimension $C \subset X_g^{\text{ess}}$ is mapped to $[0, 1]^d \subset \mathbb{R}^d$, and we identify A_g^+ with $\mathcal{H}^+(\mathbb{R}^d) = \mathcal{H}^1 \sqcup \dots \sqcup \mathcal{H}^d$. The set of half-spaces in A_g^+ dual to C is denoted A , and $[A, gA] \cap \mathcal{H}^1 = \{H_0^1, \dots, H_n^1\}$ is a tightly nested segment in A_g^+ . Since $\langle g \rangle$ acts non-transversely on X_g^{ess} , we also have $[A, gA] = [o, go]$, where o denotes the origin in \mathbb{R}^d .

Remark 8.1. — Since the action is assumed to be RAAG-like, property 7.1(1) implies that if $H \in \mathcal{H}^i$ and $hH \in \mathcal{H}^j$ with $i \neq j$ for some

$H \in A_g^+$, $h \in G$, then the quadrant $Q(H, hH)$ exists. Property 7.1(2) implies that if in addition H and $H' \in \mathcal{H}^i$ are tightly nested in X_g^{ess} , then the quadrant $Q(H', hH)$ exists. (Recall that, by definition, $Q(H, K)$ is always disjoint from X_g^{ess} .)

When discussing a quadrant Q of the form $Q(H, hH)$, if $H, H' \in \mathcal{H}^i$ are tightly nested, the quadrant $Q(H', hH)$ faces the same way as Q by Lemma 6.4. It either properly contains Q or is properly contained in Q . If the former occurs, we may refer to $Q(H', hH)$ as an *extended quadrant* for Q .

The first two results below will be used to generate contradictions. The first lemma states that when a certain configuration occurs (involving both a northwest and southeast quadrant) then there is a tightly nested segment that is forced to overlap with a copy of its reverse. The second lemma says that the latter event is impossible. Several of the arguments in the remainder of the paper have the goal of showing that a quadrant faces a particular way (northwest or southeast), with the aim of creating the forbidden configuration, and thereby a contradiction.

LEMMA 8.2. — *Let $\sigma = \{K_0, \dots, K_m\} \subset \mathcal{H}^i$ be tightly nested in X_g^{ess} and suppose that $f\bar{\sigma} \subset A_g^+$ for some $f \in G$. Let x_i be horizontal. Suppose there exist $j \leq j'$ such that $f\bar{K}_j, f\bar{K}_{j'} \notin \mathcal{H}^i$ and $Q(K_j, f\bar{K}_j)$ faces northwest while $Q(K_{j'}, f\bar{K}_{j'})$ faces southeast. Then there is a non-trivial subsegment $\alpha \subset \sigma$ such that $f\bar{\alpha} \subset \mathcal{H}^i$ and $\alpha, f\bar{\alpha}$ overlap.*

Proof. — First note that if $j = j'$ then X_g^{ess} avoids both of the quadrants

$$\{x_i < n + 1, x_j > m\} \quad \text{and} \quad \{x_i > n, x_j < m + 1\}$$

for some $n, m \in \mathbb{Z}$. But then X_g^{ess} avoids the set $\{n < x_i < n + 1\}$ and cannot contain an axis for g . Thus $j < j'$.

For any index k , the quadrant $Q(K_k, f\bar{K}_k)$ is defined if and only if $f\bar{K}_k \notin \mathcal{H}^i$, by Remark 8.1. We may choose j, j' to be an *innermost* pair having the stated properties. Then, for any k between j and j' , we have $f\bar{K}_k \in \mathcal{H}^i$.

Since the pair $K_{j'-1}, K_{j'}$ are tightly nested, there is an extended quadrant $Q(K_{j'-1}, f\bar{K}_{j'})$ which faces southeast (cf. Remark 8.1). There is also an extended northwest quadrant $Q(K_{j+1}, f\bar{K}_j)$, since K_j, K_{j+1} are tightly nested.

If $j' = j + 1$ then $f\bar{K}_j$ and $f\bar{K}_{j'}$ are tightly nested and Lemma 6.4 says that both quadrants $Q(K_j, f\bar{K}_j)$ and $Q(K_j, f\bar{K}_{j'}) = Q(K_{j'-1}, f\bar{K}_{j'})$ face the same way. However, these face northwest and southeast respectively. Therefore, $j' > j + 1$ and the segment $\alpha = \{K_{j+1}, \dots, K_{j'-1}\}$ is non-trivial.

Note that $f\bar{\alpha} \subset \mathcal{H}^i$ by the choice of j, j' . We proceed now to use the Elbow Lemma 6.5 to constrain the location of $f\bar{\alpha}$ along \mathcal{H}^i . In coordinates we have $K_j = H_a^i$ and $K_{j'} = H_b^i$ for some integers $a < b$, and

$$\alpha = \{K_{j+1}, \dots, K_{j'-1}\} = \{H_{a+1}^i, \dots, H_{b-1}^i\}.$$

Write $f\bar{\alpha} = \{H_c^i, \dots, H_{c+|\alpha|-1}^i\}$ for some $c \in \mathbb{Z}$.

Let $f\bar{K}_{j'} = H_{e'}^{i'} \in \mathcal{H}^{i'}$ where $i' \neq i$ and $e \in \mathbb{Z}$. Applying the Elbow Lemma to the tightly nested pair $\{f\bar{K}_{j'-1}, f\bar{K}_{j'}\} = \{H_c^i, H_{e'}^{i'}\}$, we find that the edge $\{c\} \times [e, e+1]$ lies in $p_{ii'}(X_g^{\text{ess}})$. Since X_g^{ess} avoids the quadrant $Q(K_{j'-1}, f\bar{K}_{j'}) = \{x_i > b - 1, x_{i'} < e + 1\}$, we conclude that $c \leq b - 1$. See Figure 8.1.




Figure 8.1. The vertical position of the elbow is aligned with the top of the quadrant $Q(K_{j'-1}, f\bar{K}_{j'})$ as shown. The horizontal position is aligned with the left end of $f\bar{\alpha}$. Since the elbow is outside the quadrant, $f\bar{\alpha}$ cannot be entirely to the right of α .

Now redefine i' and e such that $f\bar{K}_j = H_e^{i'} \in \mathcal{H}^{i'}$ (with $i' \neq i$). Applying the Elbow Lemma to the tightly nested pair $\{f\bar{K}_j, f\bar{K}_{j+1}\} = \{H_e^{i'}, H_{c+|\alpha|-1}^i\}$, we find that $p_{ii'}(X_g^{\text{ess}})$ contains the edge $\{c+|\alpha|\} \times [e, e+1]$. Now X_g^{ess} avoids the quadrant $Q(K_{j+1}, f\bar{K}_j) = \{x_i < a + 2, x_{i'} > e\}$, and therefore $c + |\alpha| \geq a + 2$.

The inequalities $c \leq b - 1$ and $c + |\alpha| \geq a + 2$ say precisely that α and $f\bar{\alpha}$ overlap. \square

The conclusion of the preceding lemma leads directly to a contradiction:

LEMMA 8.3. — *Let $\alpha \subset \mathcal{H}^i$ be tightly nested in X_g^{ess} and suppose that $h\bar{\alpha} \subset \mathcal{H}^i$ for some $h \in G$. Then α and $h\bar{\alpha}$ cannot overlap.*

Proof. — Write $\alpha = \{H_a, \dots, H_{a+k}\}$ and $h\bar{\alpha} = \{H_{b-k}, \dots, H_b\}$ for some $a, b \in \mathbb{Z}$. Then, $hH_{a+j} = \bar{H}_{b-j}$ for each j . The transformation $a + j \mapsto b - j$ either fixes c or exchanges c and $c + 1$, for some $c \in \mathbb{Z}$. If α and $h\bar{\alpha}$ overlap then H_c (and H_{c+1} in the second case) are in $\alpha \cap h\bar{\alpha}$. In the first case h inverts H_c , contrary to the assumption that G acts on X without inversion. In the second case $h\bar{H}_c = H_{c+1}$, violating property 7.1 (3). Thus α and $h\bar{\alpha}$ cannot overlap. \square

The next results perform a technical step that will be used repeatedly in the course of proving the main theorem. They also yield corollaries showing that under certain conditions, quadrants face in particular directions.

LEMMA 8.4 (Southeast quadrant shifting). — *Let $\sigma = \{K_0, \dots, K_m\} \subset \mathcal{H}^i$ be tightly nested in X_g^{ess} and suppose there is an $f \in G$ such that $f\bar{\sigma} \subset A_g^+$ and $f\bar{\sigma} \not\subset \mathcal{H}^i$. Let x_i be horizontal and let k be the smallest index such that $f\bar{K}_k \notin \mathcal{H}^i$. Suppose the quadrant $Q(K_k, f\bar{K}_k)$ faces southeast, so that*

$$Q(K_k, f\bar{K}_k) = \{x_i > a + k, x_j < b + 1\}$$

for some $j \neq i$, $a, b \in \mathbb{Z}$. Then

- (1) X_g^{ess} also avoids the larger quadrant $Q = \{x_i > a - k, x_j < b + 1\}$.
- (2) If $k > 0$ then $f\bar{K}_0 \in \mathcal{H}^i$ and $K_0 \subset f\bar{K}_0$.

Proof. — If $k = 0$ then $Q = Q(K_k, f\bar{K}_k)$ and there is nothing to prove, so assume that $k > 0$. It is implicit from the description of $Q(K_k, f\bar{K}_k)$ that $K_k = H_{a+k}^i$ and $f\bar{K}_k = H_b^j$. Let $\alpha = \{K_0, \dots, K_{k-1}\}$ be the initial segment of σ before K_k and note that $f\bar{\alpha} \subset \mathcal{H}^i$. Writing $f\bar{\alpha} = \{H_{c-k}^i, \dots, H_{c-1}^i\}$ for the appropriate $c \in \mathbb{Z}$, we have $f\bar{K}_{k-1} = H_{c-k}^i$.

Applying the Elbow Lemma to the tightly nested pair $\{f\bar{K}_{k-1}, f\bar{K}_k\} = \{H_{c-k}^i, H_b^j\}$, we find that $p_{ij}(X_g^{\text{ess}})$ contains the edge $e = \{c - k\} \times [b, b + 1]$.

Since e avoids the quadrant $Q(K_k, f\bar{K}_k)$, we must have $c - k \leq a + k$. In fact, by Remark 8.1, e avoids the extended quadrant $Q(K_{k-1}, f\bar{K}_k) = \{x_i > a + k - 1, x_j < b + 1\}$, and so $c - k < a + k$. Thus $f\bar{\alpha}$ cannot be entirely to the right of α in \mathcal{H}^i . By Lemma 8.3 $f\bar{\alpha}$ cannot overlap with α and so it must lie entirely to its left. That is, $c \leq a$. See Figure 8.2.

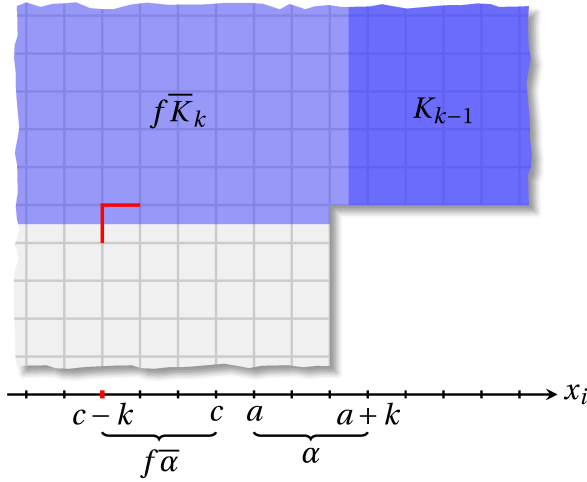


Figure 8.2. The elbow is aligned with the left endpoint of $f\bar{\alpha}$ and the top side of the quadrant $Q(K_{k-1}, f\bar{K}_k)$. Hence $f\bar{\alpha}$ is not to the right of α . Since α and $f\bar{\alpha}$ cannot overlap, we have $c \leq a$. The elbow now generates a larger, “shifted” quadrant.

Now note that the quadrant generated by $f\bar{K}_{k-1}$ and $f\bar{K}_k$ (and avoided by X_g^{ess}) is

$$\begin{aligned} Q(f\bar{K}_{k-1}, f\bar{K}_k) &= \{x_i > c - k, x_j < b + 1\} \\ &\supseteq \{x_i > a - k, x_j < b + 1\}, \end{aligned}$$

proving (1). Finally, note that $K_0 = H_a^i$ and $f\bar{K}_0 = H_{c-1}^i$, and (2) is clear since $c - 1 < a$. □

COROLLARY 8.5. — Let $\sigma = \{K_0, \dots, K_m\} \subset \mathcal{H}^i$ be tightly nested in X_g^{ess} and suppose there is an $f \in G$ such that $f\bar{\sigma} \subset A_g^+$ and $f\bar{\sigma} \not\subset \mathcal{H}^i$. Let x_i be horizontal and let k be the smallest index such that $f\bar{K}_k \notin \mathcal{H}^i$. Suppose there is a vertex v in X_g^{ess} such that $v \in K_0$ and $v \notin f\bar{K}_k$. Then the quadrant $Q(K_k, f\bar{K}_k)$ faces northwest.

Proof. — Set $K_0 = H_a^i$ for some $a \in \mathbb{Z}$, so $K_k = H_{a+k}^i$. We assume $f\bar{K}_k \notin \mathcal{H}^i$, so there exists $j \neq i$ and $b \in \mathbb{Z}$ such that $f\bar{K}_k = H_b^j$. Suppose $Q(K_k, f\bar{K}_k)$ faces southeast; that is:

$$Q(K_k, f\bar{K}_k) = \{x_i > a + k, x_j < b + 1\}.$$

By Lemma 8.4(1), X_g^{ess} also avoids the larger quadrant

$$Q = \{x_i > a - k, x_j < b + 1\}.$$

Since $v \notin f\overline{K}_k$, $v_j \leq b$. But $v \in K_0$, so $v_i \geq a + 1$. So $v \in Q$, which is a contradiction. Therefore, $Q(K_k, f\overline{K}_k)$ faces northwest. \square

The next two results are completely analogous to the previous two, with the same proofs:

LEMMA 8.6 (Northwest quadrant shifting). — Let $\sigma = \{K_0, \dots, K_m\} \subset \mathcal{H}^i$ be tightly nested in X_g^{ess} and suppose there is an $f \in G$ such that $f\overline{\sigma} \subset A_g^+$ and $f\overline{\sigma} \not\subset \mathcal{H}^i$. Let x_i be horizontal and let k be the largest index such that $f\overline{K}_k \not\subset \mathcal{H}^i$. Suppose the quadrant $Q(K_k, f\overline{K}_k)$ faces northwest, so that

$$Q(K_k, f\overline{K}_k) = \{x_i < a - (m - k), x_j > b\}$$

for some $j \neq i$, $a, b \in \mathbb{Z}$. Then

- (1) X_g^{ess} also avoids the larger quadrant

$$Q = \{x_i < a + (m - k), x_j > b\}.$$

- (2) If $k < m$ then $f\overline{K}_m \in \mathcal{H}^i$ and $K_m \supset f\overline{K}_m$.

COROLLARY 8.7. — Let $\sigma = \{K_0, \dots, K_m\} \subset \mathcal{H}^i$ be tightly nested in X_g^{ess} and suppose there is an $f \in G$ such that $f\overline{\sigma} \subset A_g^+$ and $f\overline{\sigma} \not\subset \mathcal{H}^i$. Let x_i be horizontal and let k be the largest index such that $f\overline{K}_k \not\subset \mathcal{H}^i$. Suppose there is a vertex v in X_g^{ess} such that $v \notin K_m$ and $v \in f\overline{K}_k$. Then the quadrant $Q(K_k, f\overline{K}_k)$ faces southeast.

9. Proof of the main theorem

Our goal in this section is to prove Theorem A from the Introduction, which we restate here:

THEOREM 9.1. — Let X be a $CAT(0)$ cube complex with a RAAG-like action by G . Then $\text{scl}(g) \geq 1/24$ for every hyperbolic element $g \in G$.

We continue with the same notation as in the previous section. Fix a taut equivariant embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$. Let C be the cube in X_g^{ess} mapped to $[0, 1]^d$ under the equivariant embedding. Let A be the set of half-spaces in A_g^+ dual to C , so that $[A, gA)$ is a fundamental domain for the action of $\langle g \rangle$ on A_g^+ . We have $[A, gA) = [o, go]$, where $o \in \mathbb{R}^d$ is the origin. Identify $\mathcal{H}(\mathbb{R}^d)$ with A_g . Recall that by property (3) of Proposition 5.4,

the extended taut segment $[A, gA] \cap \mathcal{H}^1$ is tightly nested in X_g^{ess} . Write $[A, gA] \cap \mathcal{H}^1 = \{H_0^1, \dots, H_n^1\}$.

Most of this section is devoted to finding a tightly nested subsegment $\gamma \subseteq [A, gA] \cap \mathcal{H}^1$ such that $\gamma > g\gamma$ and no copy of $\bar{\gamma}$ appears in A_g^+ . Once we find such a γ , then Theorem 9.1 follows immediately; details are laid out in the proof at the end of this section. To arrange that $\gamma > g\gamma$, we may have to pass to a proper subsegment of $[A, gA] \cap \mathcal{H}^1$. On the other hand, if γ is short, it is more likely for A_g^+ to contain a copy of $\bar{\gamma}$. Our approach, therefore, is to use a *maximal g -nested segment*, defined below.

After discussing maximal g -nested segments, we proceed to show that for such a segment γ , no copy of $\bar{\gamma}$ can lie in A_g^+ . First we prove Lemmas 9.4 and 9.5, which are technical statements used to show that quadrants face a particular way, under appropriate conditions. Next comes Proposition 9.6, which states that if $h\bar{\gamma} \subset A_g^+$, then $h\bar{\gamma}$ cannot lie entirely in \mathcal{H}^1 . Hence $h\bar{\gamma}$ contains half-spaces in A_g^+ outside of \mathcal{H}^1 . Such a half-space generates a quadrant, by property 7.1(1). In Propositions 9.7 and 9.8 we prove that the first such quadrant faces northwest, and the last such quadrant faces southeast. In this way, the forbidden configuration of Lemma 8.2 is created, resulting in a contradiction.

Maximal g -nested segments

DEFINITION 9.2. — *A subsegment $\gamma = \{H_\ell^1, \dots, H_r^1\}$ of $[A, gA] \cap \mathcal{H}^1$ is said to be g -nested if $\gamma > g\gamma$ in X_g^{ess} . It is a maximal g -nested segment if it is g -nested and is not properly contained in another g -nested subsegment of $[A, gA] \cap \mathcal{H}^1$.*

Figure 5.1 shows an example where the full segment $\gamma = [A, gA] \cap \mathcal{H}^1$ is not g -nested; this is the segment of blue half-spaces labeled a, c, e . In this example, the subsegment consisting of the pair labeled a, c is a maximal g -nested segment, as is the pair labeled c, e .

Note that for every $H \in A_g^+$ we have $H \supset gH$ in X_g^{ess} by Remark 3.2 and Property 7.1(1). Thus every subsegment of $[A, gA] \cap \mathcal{H}^1$ of length 1 is g -nested. It follows that every $H \in [A, gA] \cap \mathcal{H}^1$ is contained in a maximal g -nested segment.

LEMMA 9.3. — *Let $\gamma = \{H_\ell^1, \dots, H_r^1\}$ be a maximal g -nested subsegment of $[A, gA] \cap \mathcal{H}^1$. Then the following two statements hold:*

- (1) *Either $\ell = 0$ or $H_{\ell-1}^1 \pitchfork g^{-1}H_r^1$ in X_g^{ess} .*
- (2) *Either $r = n$ or $gH_\ell^1 \pitchfork H_{r+1}^1$ in X_g^{ess} .*

Proof. — Suppose $\ell > 0$. If $H_{\ell-1}^1$ and $g^{-1}H_r^1$ are not transverse in X_g^{ess} , then they are nested in X_g^{ess} by Remark 3.1. Since $o \in g^{-1}H_r - H_{\ell-1}$, this means that $g^{-1}H_r^1 \supset H_{\ell-1}^1$ in X_g^{ess} , which is equivalent to $H_r^1 \supset gH_{\ell-1}^1$. Let $\gamma' = \{H_{\ell-1}^1, \dots, H_r^1\}$. We have:

$$H_{\ell-1}^1 \supset \dots \supset H_r^1 \supset gH_{\ell-1}^1 \supset \dots \supset gH_r^1.$$

So γ' is g -nested and γ' properly contains γ , violating maximality of γ . Similarly, if $r < n$ and gH_ℓ^1 and H_{r+1}^1 are not transverse, then the segment $\{H_\ell^1, \dots, H_{r+1}^1\}$ is g -nested and contains γ . \square

We now proceed with the main steps of the proof of Theorem 9.1. The primary goal is to show that a maximal g -nested segment in $[A, gA] \cap \mathcal{H}^1$ never appears in A_g^+ in the reverse direction. The next two lemmas are technical steps that are needed along the way.

LEMMA 9.4. — *Let $\gamma = \{H_\ell^1, \dots, H_r^1\}$ be a maximal g -nested subsegment of $[A, gA] \cap \mathcal{H}^1$. Suppose $h\bar{\gamma} \subset A_g^+$ and $h\bar{\gamma} \not\subset \mathcal{H}^1$, for some $h \in G$. Suppose $\ell > 0$. Let x_1 be horizontal and let j be the smallest integer between ℓ and r such that $h\bar{H}_j^1 \notin \mathcal{H}^1$. Then either the quadrant $Q(H_j^1, h\bar{H}_j^1)$ faces northwest, or there is a vertex v in X_g^{ess} such that $v \notin g^{-1}H_r^1$ and $v \in h\bar{H}_\ell^1$.*

Proof. — By Lemma 9.3, since $\ell > 0$, $H_{\ell-1}^1 \pitchfork g^{-1}H_r^1$ in X_g^{ess} . Therefore, there exists a square S in X_g^{ess} in which they cross. Let v be the unique vertex of S with $v_1 = \ell$ and $v \notin g^{-1}H_r^1$. We now show that $v \in h\bar{H}_\ell^1$ under the assumption that $Q(H_j^1, h\bar{H}_j^1)$ faces southeast.

Let $h\bar{H}_j^1 = H_b^i$ for some $i \neq 1$ and $b \in \mathbb{Z}$. By assumption, X_g^{ess} avoids the quadrant

$$Q(H_j^1, h\bar{H}_j^1) = \{x_1 > j, x_i < b + 1\}.$$

Since $\ell > 0$ and $j \geq \ell$, the half-spaces H_{j-1}^1 and H_j^1 are tightly nested. Thus, by Remark 8.1, X_g^{ess} must further avoid the extended quadrant

$$Q(H_{j-1}^1, h\bar{H}_j^1) = \{x_1 > j - 1, x_i < b + 1\}.$$

If $j = \ell$, then for v to lie outside of $Q(H_{j-1}^1, h\bar{H}_j^1)$, we must have $v_i \geq b + 1$, so $v \in H_b^i = h\bar{H}_\ell^1$.

If $j > \ell$, then applying Lemma 8.4(2) using

$$\{K_0, \dots, K_m\} = \{H_\ell^1, \dots, H_r^1\}, \quad i = 1, \quad k = j, \quad f = h, \quad \sigma = \gamma, \quad a = 0$$

we obtain that $h\bar{H}_\ell^1 \in \mathcal{H}^1$ and $h\bar{H}_\ell^1 \supset H_\ell^1$. In coordinates, this means that $h\bar{H}_\ell^1 = H_c^1$ for some $c < \ell = v_1$. Thus, $v \in H_c^1 = h\bar{H}_\ell^1$. \square

The next lemma is completely analogous to the previous one, with a similar proof.

LEMMA 9.5. — *Let $\gamma = \{H_\ell^1, \dots, H_r^1\}$ be a maximal g -nested subsegment of $[A, gA) \cap \mathcal{H}^1$. Suppose $h\bar{\gamma} \subset A_g^+$ and $h\bar{\gamma} \not\subset \mathcal{H}^1$, for some $h \in G$. Suppose $r < n$. let x_1 be horizontal and let j be the largest integer between ℓ and r such that $h\bar{H}_j^1 \notin \mathcal{H}^1$. Then either the quadrant $Q(H_j^1, h\bar{H}_j^1)$ faces southeast, or there is a vertex v in X_g^{ess} such that $v \in gH_\ell^1$ and $v \notin h\bar{H}_r^1$.*

The next three propositions will form the main body of the argument. The first one shows that if a reverse copy of a maximal g -nested segment appears in A_g^+ , then it cannot lie entirely within \mathcal{H}^1 .

PROPOSITION 9.6. — *Let $\gamma = \{H_\ell^1, \dots, H_r^1\}$ be a maximal g -nested subsegment of $[A, gA) \cap \mathcal{H}^1$. Suppose $h\bar{\gamma} \subset A_g^+$ for some $h \in G$, and that $h\bar{H}_r^1 \in [A, gA)$. Then $h\bar{\gamma} \not\subset \mathcal{H}^1$.*

Proof. — If not then $h\bar{\gamma} \subset \mathcal{H}^1$. Write $h\bar{\gamma} = \{H_a^1, \dots, H_{a+|\gamma|-1}^1\}$, where $a \geq 0$ because $h\bar{H}_r^1 \in [A, gA)$.

Since γ and $h\bar{\gamma}$ cannot overlap (by Lemma 8.3), there are two possibilities for the location of $h\bar{\gamma}$ along \mathcal{H}^1 .

The first case is that $a + |\gamma| \leq \ell$ (i.e. $h\bar{\gamma}$ is to the left of γ). In particular $\ell > 0$ and therefore $H_{\ell-1}^1 \cap g^{-1}H_r^1$ in X_g^{ess} , by Lemma 9.3. Let $g^{-1}H_r^1 = H_b^i$ for some $i \neq 1, b \in \mathbb{Z}$. Note that $b < 0$ because $g^{-1}H_r^1$ contains the origin o . The projection $p_{1i}(X_g^{\text{ess}})$ contains the square $[\ell - 1, \ell] \times [b, b + 1]$, which is dual to both $H_{\ell-1}^1$ and $g^{-1}H_r^1$. Thus there is a vertex $v \in X_g^{\text{ess}}$ such that $v_1 = \ell$ and $v_i = b$.

By property 7.1(1) the half-spaces $g^{-1}H_r^1$ and $h\bar{H}_r^1 = H_a^1$ are not transverse in X_g^{ess} , and hence they generate a quadrant Q disjoint from X_g^{ess} . However, the quadrant $\{x_1 > a, x_i < b + 1\}$ contains v and the quadrant $\{x_1 < a + 1, x_i > b\}$ contains o . These are the two possibilities for Q and thus we have a contradiction (since $v, o \in X_g^{\text{ess}}$).

The second case is that $r < a$ (i.e. $h\bar{\gamma}$ is to the right of γ). Note that $a \leq n$ since $h\bar{H}_r^1 \in [A, gA)$. Hence $r < n$, and $gH_\ell^1 \cap H_{r+1}^1$ in X_g^{ess} by Lemma 9.3. Now redefine $i \neq 1$ and $b \in \mathbb{Z}$ such that $g\gamma \subset \mathcal{H}^i$ and $gH_\ell^1 = H_b^i$. Then $p_{1i}(X_g^{\text{ess}})$ contains the square $[r + 1, r + 2] \times [b, b + 1]$ dual to H_{r+1}^1 and H_b^i . Let $v \in X_g^{\text{ess}}$ be a vertex such that $v_1 = r + 1$ and $v_i = b + 1$.

Let x_1 be horizontal. By property 7.1(1) the half-spaces gH_ℓ^1 and $h\bar{H}_\ell^1 = H_{a+|\gamma|-1}^1$ are not transverse in X_g^{ess} , and hence they generate a quadrant

Q disjoint from X_g^{ess} . The northwest quadrant $\{x_1 < a + |\gamma|, x_i > b\}$ contains v , and therefore cannot be disjoint from X_g^{ess} . Thus $Q(gH_\ell^1, h\overline{H}_\ell^1) = Q(h\overline{H}_\ell^1, gH_\ell^1)$ faces southeast.

Again using property 7.1(1), the half-spaces gH_r^1 and $h\overline{H}_r^1$ are not transverse in X_g^{ess} and generate a quadrant $Q(gH_r^1, h\overline{H}_r^1) = Q(h\overline{H}_r^1, gH_r^1)$ disjoint from X_g^{ess} . If it faces southeast then it must be $\{x_1 > a, x_i < b + |\gamma|\}$. We have $go \notin gH_r^1$ because $o \notin H_r^1$, and $go \in h\overline{H}_r^1$ because $h\overline{H}_r^1 \in [A, gA] = [o, go]$. Therefore go is in this southeast quadrant. Since $go \in X_g^{\text{ess}}$, we conclude that $Q(h\overline{H}_r^1, gH_r^1)$ faces northwest.

Now apply Lemma 8.2 using

$$\{K_0, \dots, K_m\} = \{h\overline{H}_r^1, \dots, h\overline{H}_\ell^1\}, \quad i = 1, \quad f = gh^{-1}, \quad j = 1, \quad j' = m$$

to obtain a contradiction via Lemma 8.3. \square

The next two propositions also deal with a reverse copy of a maximal g -nested segment in A_g^+ . By the previous proposition, there must be a half-space in the segment which lies outside of \mathcal{H}^1 . Such a half-space will generate a quadrant, by property 7.1(1). The two propositions say that the first such quadrant always faces northwest, and the last such quadrant always faces southeast.

PROPOSITION 9.7. — *Let $\gamma = \{H_\ell^1, \dots, H_r^1\}$ be a maximal g -nested subsegment of $[A, gA] \cap \mathcal{H}^1$. Suppose $h\overline{\gamma} \subset A_g^+$, $h\overline{H}_r^1 \in [A, gA]$, and $h\overline{\gamma} \not\subset \mathcal{H}^1$ for some $h \in G$. Let x_1 be horizontal. Let j be the smallest integer between ℓ and r such that $h\overline{H}_j^1 \notin \mathcal{H}^1$. Then the quadrant $Q(H_j^1, h\overline{H}_j^1)$ faces northwest.*

Proof.

Case 1: $\ell = 0$. — In other words, $\gamma = \{H_0^1, \dots, H_r^1\}$. Let v be the vertex of X_g^{ess} with coordinates $v_1 = 1$ and $v_k = 0$ for all $k > 1$. Note that $v \in H_0^1$ and H_0^1 is the only element in $[A, gA]$ with this property. Therefore, since $h\overline{H}_r^1 \in [A, gA]$, if $v \in h\overline{H}_r^1$, then we must have $h\overline{H}_r^1 = H_0^1$. But this contradicts that γ and $h\overline{\gamma}$ cannot overlap by Lemma 8.3, so $v \notin h\overline{H}_r^1$. Since $h\overline{H}_r^1 \supset h\overline{H}_j^1$, $v \notin h\overline{H}_j^1$. Now apply Corollary 8.5 using

$$\{K_0, \dots, K_m\} = \{H_0^1, \dots, H_r^1\}, \quad i = 1, \quad f = h, \quad K_k = H_j^1,$$

to obtain that $Q(H_j^1, h\overline{H}_j^1)$ must face northwest.

Case 2: $\ell > 0$. — We will assume $Q(H_j^1, h\overline{H}_j^1)$ faces southeast and derive a contradiction. By Lemma 9.4, there exists a vertex v in X_g^{ess} such that

$v \notin g^{-1}H_r^1$ and $v \in h\overline{H}_\ell^1$. Note for any j between ℓ and r , $h\overline{H}_\ell^1 \subset h\overline{H}_j^1$, so $v \in h\overline{H}_j^1$.

Let i be the coordinate with $g^{-1}\gamma \subset \mathcal{H}^i$. If $h\overline{H}_j^1 \in \mathcal{H}^i$, then $g^{-1}H_r^1$ and $h\overline{H}_j^1$ are parallel in \mathbb{R}^d , and hence are nested in X_g^{ess} . Since $h\overline{H}_r^1 \in [A, gA]$, we have $o \notin h\overline{H}_r^1$. Then, since $j < r$ we have $h\overline{H}_r^1 \supset h\overline{H}_j^1$, and therefore $o \notin h\overline{H}_j^1$. However, $o \in g^{-1}H_r^1$ and therefore $g^{-1}H_r^1 \supset h\overline{H}_j^1$. This contradicts the existence of v . Therefore, we may assume $h\overline{H}_j^1 \notin \mathcal{H}^i$.

We now forget coordinate x_1 and designate x_i to be the horizontal coordinate. Since $h\overline{\gamma}$ is not entirely contained in \mathcal{H}^i , there is a largest integer j' between ℓ and r such that $h\overline{H}_{j'}^1 \notin \mathcal{H}^i$. Since $v \notin g^{-1}H_r^1$ and $v \in h\overline{H}_{j'}^1$, by Corollary 8.7 using

$$\{K_0, \dots, K_m\} = \{g^{-1}H_\ell^1, \dots, g^{-1}H_r^1\}, \quad f = hg, \quad K_k = g^{-1}H_{j'}^1,$$

we obtain that $Q(g^{-1}H_{j'}^1, h\overline{H}_{j'}^1)$ faces southeast. That is, $h\overline{H}_{j'}^1 \supset g^{-1}H_{j'}^1$, but this is impossible since $o \in g^{-1}H_{j'}^1$ and $o \notin h\overline{H}_{j'}^1$. This yields a contradiction under the assumption that $Q(H_j^1, h\overline{H}_j^1)$ faces southeast, as desired. \square

The next proposition is analogous to the previous one, but the situation is not entirely symmetric because of the assumption throughout that the largest half-space of $h\overline{\gamma}$ lies in $[A, gA]$. For this reason, the next proposition requires an independent proof.

PROPOSITION 9.8. — *Let $\gamma = \{H_\ell^1, \dots, H_r^1\}$ be a maximal g -nested subsegment of $[A, gA] \cap \mathcal{H}^1$. Suppose $h\overline{\gamma} \subset A_g^+$, $h\overline{H}_r^1 \in [A, gA]$, and $h\overline{\gamma} \not\subset \mathcal{H}^1$ for some $h \in G$. Let x_1 be horizontal. Let j be the largest integer between ℓ and r such that $h\overline{H}_j^1 \notin \mathcal{H}^1$. Then the quadrant $Q(H_j^1, h\overline{H}_j^1)$ faces southeast.*

Proof. — In the following, let i and i' be the coordinates with $g\gamma \subset \mathcal{H}^i$ and $h\overline{H}_j^1 \in \mathcal{H}^{i'}$. We will assume that $Q(H_j^1, h\overline{H}_j^1)$ faces northwest, that is, $H_j^1 \supset h\overline{H}_j^1$, and derive a contradiction.

Case 1: $r = n$. — In other words, $\gamma = \{H_\ell^1, \dots, H_n^1\}$.

We first consider the sub-case that $j = n$. Recall that the extended segment

$$[A, gA] \cap \mathcal{H}^1 = \{H_0^1, \dots, H_n^1, H_{n+1}^1\}$$

is tightly nested; in particular, the pair $\{H_n^1, H_{n+1}^1\}$ is tightly nested. Therefore, by Remark 8.1, H_{n+1}^1 and $h\overline{H}_n^1$ must also generate a quadrant that faces northwest; in other words, $H_{n+1}^1 \supset h\overline{H}_n^1$. Let go be the translate of

the origin by g . Since $[A, gA] = [o, go]$ and $H_{n+1}^1 \in gA$, we must have $go \notin H_{n+1}^1$. Thus, $go \notin h\overline{H}_n^1$, but this contradicts the assumption that $h\overline{H}_n^1 = h\overline{H}_r^1 \in [A, gA]$.

Now suppose $j < n$. By Lemma 8.6(2), using

$$\{K_0, \dots, K_m\} = \{H_\ell^1, \dots, H_n^1\}, \quad i = 1, \quad f = h, \quad K_k = H_j^1,$$

we obtain that $h\overline{H}_n^1 \in \mathcal{H}^1$ and $H_n^1 \supset h\overline{H}_n^1$. So we must have that $h\overline{H}_n^1 = H_b^1$, for some $b > n$. But this contradicts the assumption that $h\overline{H}_n^1 \in [A, gA] \cap \mathcal{H}^1 = \{H_0^1, \dots, H_n^1\}$.

Case 2: $r < n$. — In this case, we can apply Lemma 9.5 to the assumption that $Q(H_j^1, h\overline{H}_j^1)$ faces northwest, yielding a vertex v in X_g^{ess} with $v \in gH_\ell^1$ and $v \notin h\overline{H}_r^1$.

If $j = r$ and $i = i'$, then gH_ℓ^1 and $h\overline{H}_r^1$ are parallel in \mathbb{R}^d and hence are nested in X_g^{ess} . Since $go \in h\overline{H}_r^1$ and $go \notin gH_\ell^1$, $h\overline{H}_r^1 \supset gH_\ell^1$. But this contradicts the existence of v .

In all other cases we claim $h\overline{H}_r^1 \notin \mathcal{H}^i$. This is true when $j = r$ and $i \neq i'$, since $h\overline{H}_r^1 = h\overline{H}_j^1 \in \mathcal{H}^{i'}$. In the situation that $j < r$, then by the choice of j , the suffix $\alpha = \{H_{j+1}^1, \dots, H_r^1\}$ has $h\overline{\alpha} \subset \mathcal{H}^1$. But since $r < n$, $gH_\ell^1 \cap H_{r+1}^1$ by Lemma 9.3(2); in particular, since $gH_\ell^1 \in \mathcal{H}^i$, we have $i \neq 1$. This shows that $h\overline{H}_r^1 \notin \mathcal{H}^i$.

We now forget coordinate x_1 and designate x_i to be the horizontal coordinate. Let j' be the smallest integer between ℓ and r such that $h\overline{H}_{j'}^1 \notin \mathcal{H}^i$. Such j' exists since $h\overline{H}_r^1 \notin \mathcal{H}^i$. Our goal now is to use go and v to determine which ways the quadrants $Q(gH_{j'}^1, h\overline{H}_{j'}^1)$ and $Q(gH_r^1, h\overline{H}_r^1)$ face.

Now set

$$\{K_0, \dots, K_m\} = \{gH_\ell^1, \dots, gH_r^1\}, \quad f = hg^{-1}.$$

Since $go \notin gH_r^1$ and $go \in h\overline{H}_r^1$, by Corollary 8.7, where $K_k = gH_r^1$, the quadrant $Q(gH_r^1, h\overline{H}_r^1)$ faces southeast. On the other hand, since $v \in gH_\ell^1$ and $v \notin h\overline{H}_{j'}^1$, by Corollary 8.5, where $K_k = gH_{j'}^1$, the quadrant $Q(gH_{j'}^1, h\overline{H}_{j'}^1)$ faces northwest. Since j' is the smallest index between ℓ and r for which $h\overline{H}_{j'}^1 \notin \mathcal{H}^i$, and r is the largest index, the conclusion of Lemma 8.2 yields a non-trivial subsegment $\alpha \subset g\gamma$ such that $h\overline{\alpha} \subset \mathcal{H}^i$ and α and $h\alpha$ overlap. But this is impossible by Lemma 8.3. This contradiction was obtained under the assumption that $Q(H_j^1, h\overline{H}_j^1)$ faces northwest, concluding the proof. \square

We now tie everything together for the proof of Theorem 9.1.

Proof of Theorem 9.1. — Given a hyperbolic element $g \in G$, fix a taut $\langle g \rangle$ -equivariant embedding $X_g^{\text{ess}} \hookrightarrow \mathbb{R}^d$ using Proposition 5.4, as discussed in the beginning of Section 7. Let $\gamma = \{H_\ell^1, \dots, H_r^1\} \subseteq [A, gA)$ be a maximal g -nested subsegment of $[A, gA) \cap \mathcal{H}^1$.

Suppose $h\bar{\gamma} \in A_g^+$ for some $h \in G$. Replacing $h\bar{\gamma}$ by a $\langle g \rangle$ -translate if necessary, we can assume that $h\bar{H}_r^1 \in [A, gA)$. Declare x_1 to be the horizontal coordinate.

By Proposition 9.6, $h\bar{\gamma}$ cannot be entirely contained in \mathcal{H}^1 . Let j be the smallest index such that $h\bar{H}_j^1 \notin \mathcal{H}^1$. Then, by Proposition 9.7, the quadrant $Q(H_j^1, h\bar{H}_j^1)$ faces northwest. Let j' be the largest index such that $h\bar{H}_{j'}^1 \notin \mathcal{H}^1$. By Proposition 9.8 the quadrant $Q(H_{j'}^1, h\bar{H}_{j'}^1)$ faces southeast. Lemma 8.2 now provides a contradiction, via Lemma 8.3.

Therefore, no copy of $\bar{\gamma}$ appears in A_g^+ . Now consider the counting functions c_γ and $c_{\bar{\gamma}}$ from Section 7. We have $c_{\bar{\gamma}}(o, g^n o) = 0$ for all $n > 0$. Since γ is g -nested we also have $c_\gamma(o, g^n o) = n$ for $n > 0$. Choosing the basepoints $x_{g^n} = o$ for all such n (noting that $X_g \subseteq X_{g^n}$), the resulting homogeneous quasimorphism $\widehat{\varphi}_\gamma$ has value 1 on g . Since $\widehat{\varphi}_\gamma$ has defect at most 12, by Lemma 4.5, Bavard Duality (Lemma 2.3) tells us that $\text{scl}(g) \geq 1/24$. \square

BIBLIOGRAPHY

- [1] C. BAVARD, “Longueur stable des commutateurs”, *Enseign. Math.* **37** (1991), no. 1-2, p. 109-150.
- [2] M. BESTVINA, K. BROMBERG & K. FUJIWARA, “Bounded cohomology with coefficients in uniformly convex Banach spaces”, *Comment. Math. Helv.* **91** (2016), no. 2, p. 203-218.
- [3] ———, “Stable commutator length on mapping class groups”, *Ann. Inst. Fourier* **66** (2016), no. 3, p. 871-898.
- [4] M. R. BRIDSON & A. HAEFLIGER, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften, vol. 319, Springer, 1999, xxii+643 pages.
- [5] J. BRODZKI, S. J. CAMPBELL, E. GUENTNER, G. NIBLO & N. J. WRIGHT, “Property A and CAT(0) cube complexes”, *J. Funct. Anal.* **256** (2009), no. 5, p. 1408-1431.
- [6] R. BROOKS, “Some remarks on bounded cohomology”, in *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, 1978)*, Annals of Mathematics Studies, vol. 97, Princeton University Press, 1981, p. 53-63.
- [7] D. CALEGARI, *scl*, MSJ Memoirs, vol. 20, Mathematical Society of Japan, 2009, xii+209 pages.
- [8] D. CALEGARI & K. FUJIWARA, “Stable commutator length in word-hyperbolic groups”, *Groups Geom. Dyn.* **4** (2010), no. 1, p. 59-90.
- [9] P.-E. CAPRACE & M. SAGEEV, “Rank rigidity for CAT(0) cube complexes”, *Geom. Funct. Anal.* **21** (2011), no. 4, p. 851-891.

- [10] I. CHATTERJI, T. FERNÓS & A. IOZZI, “The median class and superrigidity of actions on CAT(0) cube complexes”, *J. Topol.* **9** (2016), no. 2, p. 349-400, With an appendix by Pierre-Emmanuel Caprace.
- [11] I. CHATTERJI & G. NIBLO, “From wall spaces to CAT(0) cube complexes”, *Int. J. Algebra Comput.* **15** (2005), no. 5-6, p. 875-885.
- [12] M. CLAY, M. FORESTER & J. LOUWSMA, “Stable commutator length in Baumslag–Solitar groups and quasimorphisms for tree actions”, *Trans. Am. Math. Soc.* **368** (2016), no. 7, p. 4751-4785.
- [13] M. W. DAVIS, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, 2008, xvi+584 pages.
- [14] A. J. DUNCAN & J. HOWIE, “The genus problem for one-relator products of locally indicable groups”, *Math. Z.* **208** (1991), no. 2, p. 225-237.
- [15] D. B. A. EPSTEIN & K. FUJIWARA, “The second bounded cohomology of word-hyperbolic groups”, *Topology* **36** (1997), no. 6, p. 1275-1289.
- [16] T. FERNÓS, “The Furstenberg–Poisson boundary and CAT(0) cube complexes”, *Ergodic Theory Dyn. Syst.* **38** (2018), no. 6, p. 2180-2223.
- [17] M. GROMOV, “Volume and bounded cohomology”, *Publ. Math., Inst. Hautes Étud. Sci.* (1982), no. 56, p. 5-99.
- [18] F. HAGLUND, “Isometries of CAT(0) cube complexes are semi-simple”, <http://arxiv.org/abs/0705.3386>, 2007.
- [19] F. HAGLUND & D. T. WISE, “Special cube complexes”, *Geom. Funct. Anal.* **17** (2008), no. 5, p. 1551-1620.
- [20] M. HULL & D. OSIN, “Induced quasicocycles on groups with hyperbolically embedded subgroups”, *Algebr. Geom. Topol.* **13** (2013), no. 5, p. 2635-2665.
- [21] T. KOBERDA, “Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups”, *Geom. Funct. Anal.* **22** (2012), no. 6, p. 1541-1590.
- [22] Y. MITSUMATSU, “Bounded cohomology and l^1 -homology of surfaces”, *Topology* **23** (1984), no. 4, p. 465-471.
- [23] B. NICA, “Cubulating spaces with walls”, *Algebr. Geom. Topol.* **4** (2004), p. 297-309.
- [24] M. ROLLER, “Poc sets, median algebras and group actions. An extended study of Dunwoody’s construction and Sageev’s theorem”, Habilitationsschrift, Regensburg, 1998.
- [25] M. SAGEEV, “Ends of group pairs and non-positively curved cube complexes”, *Proc. Lond. Math. Soc.* **71** (1995), no. 3, p. 585-617.
- [26] D. T. WISE, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, American Mathematical Society, 2012, xiv+141 pages.

Manuscrit reçu le 2 mars 2017,
révisé le 4 décembre 2017,
accepté le 12 juin 2018.

Talia FERNÓS
Department of Mathematics and Statistics
University of North Carolina at Greensboro
317 College Avenue
Greensboro, NC 27412 (USA)
t_fernos@uncg.edu

Max FORESTER
Department of Mathematics
University of Oklahoma
Norman, OK 73019 (USA)
forester@math.ou.edu

Jing TAO
Department of Mathematics
University of Oklahoma
Norman, OK 73019 (USA)
jing@math.ou.edu