

Genus bounds in right-angled Artin groups

Max Forester, Ignat Soroko, and Jing Tao

Abstract

We show that in any right-angled Artin group whose defining graph has chromatic number k , every non-trivial element has stable commutator length at least $1/(6k)$. Secondly, if the defining graph does not contain triangles, then every non-trivial element has stable commutator length at least $1/20$. These results are obtained via an elementary geometric argument based on earlier work of Culler.

1. Introduction

In a topological space X with fundamental group G , a loop $\gamma: S^1 \rightarrow X$ representing an element $g \in G$ may extend to a map of an oriented surface $S \rightarrow X$ with boundary γ . The smallest genus of such a surface is called the *commutator length* (cl) of $g \in G$. The *stable commutator length* (scl) of g is defined to be the limit $\text{scl}(g) = \lim_{n \rightarrow \infty} \text{cl}(g^n)/n$. These quantities have relevance in several areas, particularly low-dimensional topology, bounded cohomology, and dynamics (see [Cal09] and references therein).

Both commutator length and stable commutator length can be very difficult to compute. Understanding the qualitative behavior of scl is a somewhat more tractable problem. For many important classes of groups it has been shown that the spectrum of values of scl has a gap above zero (e.g. [CF10, CFL16, BBF16]). An early result along these lines is due to Culler. In [Cul81] he gave a lower bound for the stable commutator length of elements in a free group F : for every non-trivial $g \in F$,

$$\text{scl}(g) \geq \frac{1}{6}.$$

The purpose of this note is to generalize Culler's argument to the case of right-angled Artin groups (RAAGs) in two different ways. We obtain:

Theorem 1.1. *Let $G = A(\Gamma)$ be a right-angled Artin group whose defining graph Γ has chromatic number k . Then every non-trivial element $g \in G$ satisfies:*

$$\text{scl}(g) \geq \frac{1}{6k}.$$

Theorem 1.2. *Let $G = A(\Gamma)$ be a right-angled Artin group whose defining graph Γ does not contain triangles. Then every non-trivial element $g \in G$ satisfies:*

$$\text{scl}(g) \geq \frac{1}{20}.$$

Note that Theorem 1.2 is not a consequence of Theorem 1.1, as demonstrated by the existence of triangle-free graphs with large chromatic number, such as Mycielski’s graphs [Myc55].

It should be noted that Culler’s result has since been improved: Duncan and Howie [DH91] showed that $\text{scl}(g) \geq 1/2$ for all non-trivial $g \in F$ (see [Che18] for another proof of this result). This is the best possible lower bound for free groups since $\text{scl}([a, b]) = 1/2$ in $\langle a, b \rangle$.

We also note that Heuer, very recently, obtained the same lower bound of $1/2$ for scl in any right-angled Artin group [Heu18]. His method is based on constructing quasimorphisms, as was the previous general lower bound of $1/24$ established in [FFT]. The arguments presented here are quite elementary and geometric in nature, and we believe that they are of independent interest.

Methods

The proofs of Theorems 1.1 and 1.2 find explicit lower bounds for $\text{cl}(g^n)$ in terms of n , by considering a map of a surface S into a Salvetti complex $X(\Gamma)$, with boundary representing g^n . Taking pre-images of the hyperplanes, we obtain one-dimensional submanifolds of S , transversely labeled by generators of the right-angled Artin group. Unlike in the free groups setting, these curves may cross each other, if their labels are generators that commute.

Theorem 1.1 is obtained by showing that the collection of curves includes a suitably large sub-collection of properly embedded arcs that are pairwise disjoint and non-parallel in S .

For Theorem 1.2, we first “tighten” the pattern of curves and reduce the genus of S ; the resulting pattern of curves cuts S into polygons, all having four or more sides. This is where the two-dimensionality assumption is used (otherwise there could be triangles). This combinatorial structure now has non-positive curvature and we can estimate $\chi(S)$ using a Gauss–Bonnet formula.

An easy but crucial step in Culler’s argument is to observe that a cyclically reduced word w in a free group cannot contain overlapping subwords of the form u and u^{-1} . Thus, if w^n contains both u and u^{-1} then $|u| \leq |w|/2$. This fact is used to relate the amount of negative curvature in S to the exponent n .

In the case of a right-angled Artin group G , we must consider pairs of subwords u, u^{-1} appearing in a cyclically reduced word w that represents $g^n \in G$; we need to know that such words always satisfy $|u| \leq |w|/(2n)$. The difficulty is that G has relations and w itself need not be a proper power. We call this property the *non-overlapping property* (by analogy with the free groups case) and prove it for all right-angled Artin groups in Theorem 6.6. This result is needed in both Theorem 1.1 and Theorem 1.2.

The non-overlapping property for right-angled Artin groups appears to be rather non-trivial. For instance, the proof makes use of all of Haglund and Wise’s axioms for special cube complexes from [HW08]. In Section 6 we give the proof, after reviewing some of the terminology and ideas from [FFT] concerning essential characteristic sets in CAT(0) cube complexes.

One might wish to extend the methods of this paper to other non-special group actions on CAT(0) cube complexes. This cannot succeed in complete generality, since there are examples, such as irreducible lattices acting on products of trees [BM00], for which stable commutator length is known to vanish [BM99].

Acknowledgments

This material is based upon work supported by NSF grant DMS-1440140 while the second and third authors were in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester. The second author acknowledges partial support from NSF grants DMS-1107452, DMS-1107263, DMS-1107367 “RNMS: Geometric Structures and Representation Varieties” (the GEAR Network). The second and third authors also acknowledge partial support from NSF grants DMS-1611758 and DMS-1651963. The authors thank the referee for providing helpful comments.

2. Preliminaries

We remind the reader of some relevant definitions.

Definition 2.1 (Commutator length, stable commutator length). Given G , let X be a path connected space with fundamental group G . For any $g \in [G, G]$ the *commutator length* of g is equal to

$$\text{cl}(g) = \min_S \text{genus}(S)$$

where the minimum is taken over all continuous maps of surfaces $f: S \rightarrow X$ such that S is compact, connected, oriented, has one boundary component, and the restriction $f|_{\partial S}: \partial S \rightarrow X$ represents the conjugacy class of g in $\pi_1(X)$. Such a surface will be called (in this paper) an *admissible surface* for g . The condition that $g \in [G, G]$ ensures that admissible surfaces exist.

The *stable commutator length* of $g \in [G, G]$ is defined by the convergent limit

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

See [Cal09] for details on convergence, and for other basic properties and equivalent definitions. If $g^n \in [G, G]$ for some $n \neq 0$ we may define $\text{scl}(g) = \text{scl}(g^n)/n$, which is consistent with the first definition because of the identity $\text{scl}(g^n) = n \text{scl}(g)$. If $g^n \notin [G, G]$ for any $n \neq 0$ then $\text{scl}(g) = \infty$ by convention.

Definition 2.2 (Right-angled Artin group). Let Γ be a finite simplicial graph with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$. The *right-angled Artin group*, or RAAG, associated to Γ is a finitely presented group $A(\Gamma)$ given by the presentation

$$A(\Gamma) = \langle a \in V(\Gamma) \mid [a, b] = 1 \text{ if } (a, b) \in E(\Gamma) \rangle.$$

Definition 2.3 (Salvetti complex). For each $a \in V(\Gamma)$ let S_a^1 be a circle endowed with the structure of a CW complex having a single 0-cell and a single 1-cell. Let $n = \#V(\Gamma)$ be the number of vertices of Γ and let $T = \prod_{a \in V(\Gamma)} S_a^1$ be an n -dimensional torus with the product CW structure. For every complete subgraph $K \subseteq \Gamma$ with $V(K) = \{a_1, \dots, a_k\}$ define a k -dimensional torus T_K as the Cartesian product of CW complexes $T_K = \prod_{i=1}^k S_{a_i}^1$ and observe that T_K can be identified as a combinatorial subcomplex of T . Then the *Salvetti complex associated with $A(\Gamma)$* is

$$X(\Gamma) = \bigcup \{ T_K \subseteq T \mid K \text{ a complete subgraph of } \Gamma \}.$$

Thus $X(\Gamma)$ has a single 0-cell and n 1-cells. Each edge $(a, b) \in E(\Gamma)$ contributes a square 2-cell to $X(\Gamma)$ with the attaching map $aba^{-1}b^{-1}$. In general each complete subgraph $K \subseteq \Gamma$ contributes a k -dimensional cell to $X(\Gamma)$ where $k = \#V(K)$.

Definition 2.4 (Fat Salvetti complex). Given $\varepsilon > 0$, define the *fat Salvetti complex* associated to $A(\Gamma)$ to be the open ε -neighborhood of $X(\Gamma)$ in the torus T , denoted $X_\varepsilon(\Gamma)$. There is a deformation retraction of $X_\varepsilon(\Gamma)$ onto $X(\Gamma)$ for ε sufficiently small. Fix such an ε from now on.

The fat Salvetti complex is convenient for carrying out an easy transversality argument. One could alternatively use an approach to transversality similar to those in [Rou79, Sta87, BF10].

Following [MW02] we now introduce a useful tool for computing the Euler characteristic of a 2-dimensional complex.

Definition 2.5 (Corners, Angled 2-complex). Let X be a locally finite combinatorial 2-complex and v a 0-cell of X . We will refer to the edges of $\text{Link}(v)$ as *corners* of v . X is called an *angled 2-complex* if it has an angle $\angle c \in \mathbf{R}$ associated to each corner c of every 0-cell of X .

Definition 2.6 (Curvature). For every 0-cell v of X , its *curvature* $\kappa(v)$ is defined as

$$\kappa(v) = 2\pi - \pi\chi(\text{Link}(v)) - \sum_{c \in \text{Corners}(v)} \angle c.$$

For every 2-cell f of X , its *curvature* $\kappa(f)$ is defined as

$$\kappa(f) = \sum_{c \in \text{Corners}(f)} \angle c - (P(f) - 2)\pi,$$

where

$$\text{Corners}(f) = \{ \text{edges in } \text{Link}(v) \text{ contained in } f \mid v \text{ is a 0-cell belonging to } f \},$$

and $P(f)$ is the combinatorial length of the boundary of f .

The following theorem was proved in [MW02, Theorem 4.6]:

Combinatorial Gauss–Bonnet Theorem. Let X be a compact angled 2-complex. Then

$$2\pi\chi(X) = \sum_{v \in 0\text{-cells}} \kappa(v) + \sum_{f \in 2\text{-cells}} \kappa(f). \quad \square$$

3. Surfaces with patterns

Mapping a surface to a fat Salvetti complex

From now on we let Γ be a finite simplicial graph. For each $a \in V(\Gamma)$ let e_a be the corresponding oriented edge of $X(\Gamma)$ (so that e_a , considered as a based loop in $X(\Gamma) \subset X_\varepsilon(\Gamma)$, represents the element a of $\pi_1(X_\varepsilon(\Gamma)) = A(\Gamma)$).

Recall that the subcomplex of $X(\Gamma)$ determined by e_a is a circle S_a^1 , and note that there is a retraction $r_a: X_\varepsilon(\Gamma) \rightarrow S_a^1$ which is the restriction of the projection map from the torus T to the factor S_a^1 . We define the *hyperplane dual to a* in $X_\varepsilon(\Gamma)$ to be the pre-image of the midpoint of e_a under this retraction. It is denoted H_a . Note that this is not quite the usual definition of hyperplane, because we are working in the fat Salvetti complex $X_\varepsilon(\Gamma)$. In fact $X_\varepsilon(\Gamma)$ is a manifold and the hyperplanes H_a are codimension-one submanifolds.

Let $f: S \rightarrow X_\varepsilon(\Gamma)$ be an admissible surface for $g \in A(\Gamma)$. The composition $S \rightarrow X_\varepsilon(\Gamma) \hookrightarrow T$ is the product of the component maps $f_a = r_a \circ f$, since $T = \prod_a S_a^1$. Each of these maps $f_a: S \rightarrow S_a^1$ may

be changed by an arbitrarily small homotopy to arrange that the midpoint of e_a is a regular value of f_a . Since $X_\varepsilon(\Gamma)$ is open in T , this can be achieved for all a by homotopies such that the product homotopy is a homotopy of f *inside* $X_\varepsilon(\Gamma)$. Thus, after such a homotopy of f , we may assume that f is simultaneously transverse to all of the hyperplanes H_a . Then, for each a , the pre-image $M_a = f^{-1}(H_a)$ is a compact properly embedded one-dimensional submanifold of S . The submanifolds M_a and M_b may intersect, but they will only do so transversely, in the interior of S , and only if (a, b) is an edge of Γ . If (a, b) is not an edge, then M_a and M_b will be disjoint because $H_a \cap H_b = \emptyset$.

Each component of M_a comes with a transverse orientation (i.e. a choice of normal direction) which we label by the generator a . This is the transverse orientation induced by the orientation of e_a under the map $r_a \circ f$. The opposite transverse orientation is labeled by a^{-1} . Any small arc α in S that crosses M_a in one point will map by f to an arc that crosses H_a in one point. The direction that it crosses in will agree with e_a if and only if α is oriented with the transverse orientation of M_a .

More generally, any path α in S which crosses the submanifolds M_a transversely in distinct points has a corresponding word w_α in the generators of $A(\Gamma)$ and their inverses; the letters of w_α are the labels assigned to the transverse orientations followed by α as it passes through each crossing. If α is a based loop with basepoint p disjoint from the submanifolds M_a , then the word w_α represents the element $[f \circ \alpha]$ in $\pi_1(X_\varepsilon(\Gamma), f(p))$.

In particular, the oriented boundary of S crosses the endpoints of the manifolds M_a transversely, and has an associated cyclic word $w_{\partial S}$. This word represents the conjugacy class of g in $\pi_1(X_\varepsilon(\Gamma))$.

Simplification

We now have a compact surface S together with what we call a *pattern on S* : for each $a \in V(\Gamma)$, a properly embedded submanifold M_a of S with a choice of transverse orientation (labeled a) on each component. The surface with pattern $(S, \{M_a\}_{a \in V(\Gamma)})$ satisfies:

- T1: if M_a and M_b intersect, then (a, b) is an edge of Γ , and the intersections are transverse and occur in the interior of S ;
- T2: the cyclic word $w_{\partial S}$ given by the transverse labels along ∂S represents the conjugacy class of g in $A(\Gamma)$.

At this point we have no further need for the continuous map $f: S \rightarrow X_\varepsilon(\Gamma)$. We will simplify both S and the pattern $\{M_a\}_{a \in V(\Gamma)}$ by applying some moves. These moves will preserve properties T1 and T2. These two properties, along with the additional properties achieved by the moves, are all that will be needed to estimate $\chi(S)$. In order to describe the moves we need some further terminology.

Definition 3.1. A *pattern curve* is a connected component of M_a for any $a \in V(\Gamma)$. A pattern curve is called a *pattern arc* if it is homeomorphic to an interval, and a *pattern loop* if it is homeomorphic to a circle.

An *intersection point* is a point of intersection either between two pattern curves, or between a pattern curve and ∂S .

Now let \mathcal{M} be any union of pattern curves. If we cut $S - \partial S$ along \mathcal{M} , we get pieces, some of which may be open disks. Each such disk has a characteristic map $D^2 \rightarrow S$ giving D^2 the structure of a polygon with some number of sides. Here, a *side* is a maximal connected subset of ∂D^2 mapping

into a single pattern curve of \mathcal{M} or into ∂S . Informally, we say that the original open disk is a polygon with that number of sides (even though distinct sides of D^2 may map to the same arc in S).

Definition 3.2. A *bigon* is a complementary component of \mathcal{M} in $S - \partial S$ which is a polygon with two sides (for some \mathcal{M} , which may be taken to be just one or two pattern curves). A bigon is called *innermost* if every pattern curve that intersects one side also intersects the other side.

A *half-bigon* is a complementary component of \mathcal{M} which is a polygon with three sides, one in ∂S . Call the other two sides *interior* sides. A half-bigon is called *innermost* if every pattern curve that intersects one interior side also intersects the other interior side.

Remark 3.3. If a bigon B is not innermost, then it properly contains a smaller bigon. Similarly, if a half-bigon B is not innermost, then it properly contains either a half-bigon or a bigon. It follows that if there are no innermost bigons or half-bigons, then there are no bigons or half-bigons at all.

Now we are ready to describe the moves for simplifying the pattern on S . Recall that there is no need for us to keep track of continuous maps inducing the patterns, though in principle one could do so; every pattern on S satisfying T1 is induced by a map to $X(\Gamma)$.

The moves are as follows:

1. Remove an innermost bigon (whose sides are not in ∂S). See Figure 1.




Figure 1: Removing an innermost bigon (Move 1)

2. Remove an innermost half-bigon. See Figure 2.




Figure 2: Removing an innermost half-bigon (Move 2)

3. Splice together two endpoints of M_a that land on adjacent cancellable letters of $w_{\partial S}$. See Figure 3.
4. Discard any pattern loop.
5. Perform surgery along a non-separating simple closed curve in S that is disjoint from every pattern curve; that is, replace an annular neighborhood (also chosen disjoint from every pattern curve) with two disks.




Figure 3: Splicing adjacent endpoints (Move 3)

Remark 3.4. The only move that changes the topology of S is move 5. This move increases $\chi(S)$ by 2, but preserves the fact that S is connected and has one boundary component. Such a surface has Euler characteristic at most 1.

Remark 3.5. The only moves that change $w_{\partial S}$ are moves 2 and 3. Move 2 exchanges two adjacent letters of $w_{\partial S}$ which are commuting elements of $A(\Gamma)$. Move 3 removes a subword aa^{-1} or $a^{-1}a$ from $w_{\partial S}$. In both cases, $w_{\partial S}$ still represents g after the move (that is, property T2 is preserved by all moves).

Remark 3.6. In moves 1 and 2, the new pattern curves after the move cross exactly the same pattern curves that they did before the move, by the innermost property of the bigon or half-bigon. Thus property T1 is preserved by all moves (the other cases being obvious).

Now define the *complexity* of $(S, \{M_a\}_{a \in V(\Gamma)})$ to be the sum of three quantities: the total number of pattern curves, the total number of intersection points (including intersections with ∂S), and $1 - \chi(S)$. The complexity is a non-negative integer.

Lemma 3.7. *Each of the moves 1–5 decreases the complexity of $(S, \{M_a\}_{a \in V(\Gamma)})$.*

Proof. Moves 1 and 2 reduce the number of intersection points without changing the other two components of complexity. Move 3 reduces the number of intersection points by 2, and possibly also the number of pattern curves (unless it turns an arc into a loop). It does not change $1 - \chi(S)$. Move 4 reduces the number of pattern curves, and possibly also the number of intersection points. It leaves $1 - \chi(S)$ unchanged. Move 5 reduces $1 - \chi(S)$ by 2 without changing the other two quantities. \square

Starting with $(S, \{M_a\}_{a \in V(\Gamma)})$, one can perform moves, in any order, until the complexity cannot be reduced any further. Since no moves are available, we may conclude several things (in addition to properties T1, T2):

T3: the word $w_{\partial S}$ is cyclically reduced (or move 3 could be performed);

T4: there are no bigons or half-bigons:

An innermost half-bigon, or an innermost bigon not on ∂S , cannot exist (or move 1 or 2 is available). An innermost bigon on ∂S either contains a half-bigon (which contains an innermost half-bigon or bigon, which is a contradiction), or makes move 3 available. The claim now follows from Remark 3.3.

T5: the union $\bigcup_{a \in V(\Gamma)} M_a$ cuts S into disks:

Consider a simple closed curve in $S - \partial S$ that is disjoint from $\mathcal{M} = \bigcup_{a \in V(\Gamma)} M_a$. If it is non-separating then move 5 is available. If it is separating then \mathcal{M} lies entirely on one side, since every pattern curve is a pattern arc meeting ∂S . The other side either is a disk or admits move 5.

T6: if Γ has no triangles then each of these disks is a polygon with at least four sides:

The number of sides cannot be 1 since there are no pattern loops. It cannot be 2 because there are no bigons. If a polygon has three sides and is not a half-bigon, then the sides are on M_a, M_b, M_c with a, b, c forming a triangle in Γ .

Definition 3.8. A surface with pattern $(S, \{M_a\}_{a \in V(\Gamma)})$ satisfying properties T1 and T2, of smallest complexity, is called a *taut surface with pattern for g* . It will then also satisfy T3–T6.

We have just proved:

Proposition 3.9. *Let S_0 be an admissible surface for g in $X_\varepsilon(\Gamma)$. Then there exists a taut surface with pattern $(S, \{M_a\}_{a \in V(\Gamma)})$ for g such that $1 - \chi(S_0) \geq 1 - \chi(S)$. \square*

Bands

Let $(S, \{M_a\}_{a \in V(\Gamma)})$ be a taut surface with pattern for g . We organize the pattern arcs into regions called bands.

Definition 3.10. (Band) A *rectangle* is an embedded disk in S whose boundary is decomposed into four sides, such that two opposite sides are contained in ∂S , and the other two sides are pattern arcs. We also allow a rectangle to be degenerate, consisting of a single pattern arc. The sides in ∂S are called the *boundary sides* and the other two sides are called the *interior sides* of the rectangle. A *band* is a rectangle in S which is maximal with respect to inclusion. Since degenerate rectangles are allowed, every pattern arc is contained in a band.

Remark 3.11. Suppose C_1 and C_2 are pattern arcs that are topologically parallel, meaning that they are homotopic as maps of pairs $(I, \partial I) \rightarrow (S, \partial S)$. Then they must be disjoint and form part of the boundary of a rectangle R , since otherwise there would be a half-bigon.

If C_3 is a third pattern arc that is topologically parallel to the other two, then either it lies in R between C_1 and C_2 , or it cobounds a second rectangle R' with either C_1 or C_2 . In this case, $R \cup R'$ is a rectangle containing all three pattern arcs. More generally, if C_1, \dots, C_k is a maximal family of pairwise parallel pattern arcs, then there is a band bounded by two of them, and all the others lie inside the band. Thus bands correspond to equivalence classes of pattern arcs under the relation of parallelism.

Remark 3.12. Let B be a band, and suppose the pattern arcs inside it are C_1, \dots, C_k , numbered in order along ∂S in one boundary side of the band. Let $a_1^{\varepsilon_1}, \dots, a_k^{\varepsilon_k}$ be their transverse labels along this side. Then the opposite endpoints of C_1, \dots, C_k lie along the other boundary side of B and their transverse labels are $a_k^{-\varepsilon_k}, \dots, a_1^{-\varepsilon_1}$ in order along ∂S , because S is orientable (see Figure 4). Moreover, no other pattern arc can meet ∂S in one of the boundary sides of B , since it would then form a half-bigon with some C_i . Therefore $w_{\partial S}$ contains $u = a_1^{\varepsilon_1} \cdots a_k^{\varepsilon_k}$ and $u^{-1} = a_k^{-\varepsilon_k} \cdots a_1^{-\varepsilon_1}$ as disjoint subwords. In fact $w_{\partial S}$ is partitioned into these subwords, since every letter is the label of one end of a pattern arc, and the set of arcs is partitioned by the bands.

Remark 3.13. We make one last observation in the case where Γ has no triangles, to be used in the proof of Theorem 1.2. By property T6, every polygonal face of S has four or more sides. Let R be a

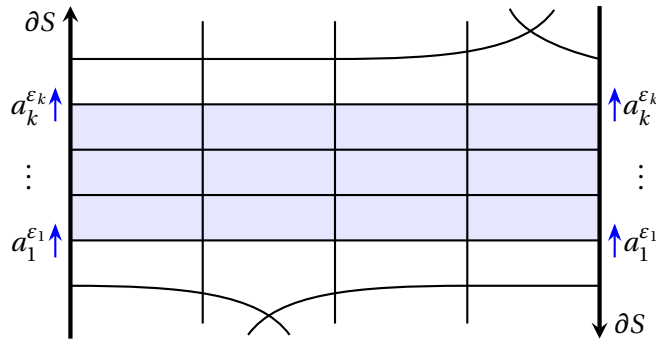


Figure 4: The shaded region is a band, joining two subwords u, u^{-1} of $w_{\partial S}$. In this example Γ has no triangles and so every face has four or more sides.

rectangle and C one of its interior sides. Consider the faces of S that are outside R but have a side in C . If all of these faces have four sides, then the closure of their union is a rectangle R' , and $R \cup R'$ is also a rectangle, properly containing R . Therefore, for any band B , each of its interior sides is adjacent to at least one face with 5 or more sides. See again Figure 4.

Recall now that our goal is to estimate $\text{scl}(g)$. To do this we need to estimate $\text{cl}(g^n)$ in terms of n . The following theorem, proved in Section 6, provides the connection to the exponent n .

Theorem 6.6. *Let w be a cyclically reduced word in the generators of $A(\Gamma)$ representing the conjugacy class of the element g^n in $A(\Gamma)$, and suppose u is a word such that both u and u^{-1} appear as subwords of w (considered as a cyclic word). Then*

$$|u| \leq \frac{|w|}{2n}.$$

Corollary 3.14. *Let $(S, \{M_a\}_{a \in V(\Gamma)})$ be a taut surface with pattern for g^n . Then the total number of bands on S is at least n .*

Proof. Recall from Remark 3.11 that the bands provide a partition of the set of pattern arcs. The subwords u, u^{-1} of $w_{\partial S}$ associated to bands then form a partition of the letters of $w_{\partial S}$. Each band accounts for two subwords, each of length at most $|w_{\partial S}|/(2n)$, so there must be at least n bands. \square

4. Proof of Theorem 1.1

Let $G = A(\Gamma)$ and suppose $g^n \in [G, G]$ for some $n > 0$. Let $S_0 \rightarrow X_\epsilon(\Gamma)$ be an admissible surface for g^n . By Proposition 3.9 there is a taut surface with pattern $(S, \{M_a\}_{a \in V(\Gamma)})$ for g^n such that $1 - \chi(S_0) \geq 1 - \chi(S)$.

Recall that the *chromatic number* of Γ is the smallest number of colors needed to color the vertices of Γ so that no two adjacent vertices have the same color. Let such a coloring with k colors be given. Every pattern arc inherits a color from its transverse label, and whenever two pattern arcs cross, they must have different colors. In particular, the set of pattern arcs having any single color is pairwise disjoint.

For each band in S , assign it the color of one of its pattern arcs. Since there are at least n bands (Corollary 3.14) and only k colors, there must be at least n/k bands of the same color c , for some c

(pigeonhole principle). Taking one pattern arc with color c from each of these bands, we obtain a collection \mathcal{C} of pattern arcs that is pairwise disjoint and non-parallel, of size at least n/k .

The size of any such collection is at most $6\text{genus}(S) - 3$. To see this, enlarge \mathcal{C} to a maximal such collection \mathcal{C}' , which defines an ideal triangulation of $S - \partial S$ (or equivalently, a one-vertex triangulation of $S/\partial S$). Such a triangulation has $6\text{genus}(S) - 3$ edges.

Since

$$\frac{n}{k} \leq 6\text{genus}(S) - 3$$

we have

$$\text{genus}(S_0) \geq \text{genus}(S) \geq \frac{n}{6k} + \frac{1}{2}.$$

The right hand side is a lower bound for $\text{cl}(g^n)$, since S_0 was an arbitrary admissible surface for g^n . Dividing by n and taking $n \rightarrow \infty$, we obtain

$$\text{scl}(g) \geq \frac{1}{6k}.$$

This finishes the proof of Theorem 1.1.

5. Proof of Theorem 1.2

Let $G = A(\Gamma)$, where Γ has no triangles, and suppose $g^n \in [G, G]$ for some $n > 0$. As before, let $S_0 \rightarrow X_\varepsilon(\Gamma)$ be an admissible surface for g^n . By Proposition 3.9 there is a taut surface with pattern $(S, \{M_a\}_{a \in V(\Gamma)})$ for g^n such that $1 - \chi(S_0) \geq 1 - \chi(S)$.

By property T5, S can be given the structure of a combinatorial 2-complex, with 0-cells equal to the intersection points and 1-skeleton equal to $\bigcup_{a \in V(\Gamma)} M_a \cup \partial S$. Each 2-cell, or face, is a polygon as described in Definition 3.1. These faces each have four or more sides by property T6. We further endow S with the structure of an *angled 2-complex* by declaring that every corner has angle $\pi/2$.

Observe that the curvature $\kappa(v)$ of every 0-cell is 0. Indeed, for an interior 0-cell v we have: $\kappa(v) = 2\pi - \pi \cdot 0 - 4 \cdot \frac{\pi}{2} = 0$, and for a 0-cell v on the boundary ∂S we have: $\kappa(v) = 2\pi - \pi \cdot 1 - 2 \cdot \frac{\pi}{2} = 0$. Thus, the combinatorial Gauss–Bonnet formula gives us:

$$2\pi\chi(S) = \sum_{f \in 2\text{-cells}} \kappa(f).$$

For a 2-cell f we have

$$\kappa(f) = \frac{\pi}{2}(\# \text{ of corners of } f) - ((\# \text{ of sides of } f) - 2)\pi = 2\pi - \frac{\pi}{2}(\# \text{ of sides of } f).$$

Therefore:

$$1 - \chi(S) = 1 + \sum_{f \in 2\text{-cells}} \frac{1}{4}((\# \text{ of sides of } f) - 4). \quad (*)$$

Recall that all faces of S have four or more sides. The faces with exactly four sides contribute 0 to the sum in (*). Thus one can sum over only the 2-cells f which have ≥ 5 sides. For simplicity in what follows we will call them *special faces*.

Our goal now is to relate the quantity in (*) to the number of bands on S . We will do this simultaneously for the case of free groups (thus obtaining Culler's bound $\text{scl}(g) \geq 1/6$) and for the case of two-dimensional right-angled Artin groups.

First, we modify the right-hand side of (*) as follows. In the case of free groups the pattern arcs do not intersect each other. Hence the sides of a 2-cell alternate between pattern arcs and arcs in ∂S . Therefore, the number of sides of any 2-cell is always even. Hence the minimal number of sides of a special face is 6. In the right-angled Artin group case, special faces can have 5 or more sides. Thus we have:

$$(\# \text{ of sides of } f) - 4 \geq A \cdot (\# \text{ of sides of } f),$$

where $A = \frac{1}{3}$ for free groups and $A = \frac{1}{5}$ for RAAGs.

Second, recall from Remark 3.13 that each band is adjacent to at least one special face on each of its two sides. It is possible that the two sides of the band are adjacent to the *same* special face, but they will do so in distinct sides of that face. Thus we may count the sides of special faces as follows:

$$\sum_{f \in \text{special faces}} (\# \text{ of sides of } f) \geq (\# \text{ of bands}) \times 2.$$

This inequality can be strict if there is more than one special face on one side of a band, or if there is a special face with one or more sides lying on ∂S . Then these sides will not be accounted for by bands.

In the free group case, each special face has exactly half of its sides lying on ∂S , so bands border exactly half of the total count of the sides of special faces. Thus we have:

$$\sum_{f \in \text{special faces}} (\# \text{ of sides of } f) \geq (\# \text{ of bands}) \times 4.$$

Going back to formula (*), we get:

$$1 - \chi(S) \geq 1 + \frac{B}{4} (\# \text{ of bands}),$$

where $B = \frac{1}{3} \cdot 4$ for free groups and $B = \frac{1}{5} \cdot 2$ for RAAGs.

Corollary 3.14 tells us that $(\# \text{ of bands}) \geq n$. Therefore

$$\text{genus}(S_0) = \frac{1}{2} \cdot (1 - \chi(S_0)) \geq \frac{1}{2} \cdot (1 - \chi(S)) \geq \begin{cases} \frac{1}{2} + \frac{n}{6} & \text{for free groups,} \\ \frac{1}{2} + \frac{n}{20} & \text{for RAAGs.} \end{cases}$$

The right hand side is now a lower bound for $\text{cl}(g^n)$ as before. Dividing by n and taking $n \rightarrow \infty$ yields

$$\text{scl}(g) \geq \begin{cases} \frac{1}{6} & \text{for free groups,} \\ \frac{1}{20} & \text{for RAAGs.} \end{cases}$$

This finishes the proof of Theorem 1.2.

6. The non-overlapping property

The purpose of this section is to prove Theorem 6.6, which was the key ingredient for the estimation of scl in the previous section (see Corollary 3.14). We start with basic definitions and some additional notions from [FFT].

Preliminaries

A CAT(0) cube complex Y is a simply connected polyhedral complex in which the closed cells are standard Euclidean cubes $[0, 1]^n$ of various dimensions, such that any two cubes either have empty intersection or intersect in a single face, and the link of every vertex is a flag complex. The latter condition, called the Gromov Link Condition, guarantees that Y is non-positively curved. The dimension of Y is the dimension of its maximal dimensional cube.

An n -cube $[0, 1]^n$ has n midcubes defined by setting one of the coordinates equal to $1/2$. A *hyperplane* in Y is a closed subspace whose intersection with each cube is either empty or equal to a midcube. Each hyperplane separates Y into two connected components. The closure of either of these two components is called a *half-space*. The hyperplane which bounds a half-space H is denoted by ∂H and the half-space complementary to H is denoted by \bar{H} . The set $\mathcal{H}(Y)$ of half-spaces of Y is partially ordered by inclusion. Two distinct half-spaces H and H' are *nested* if either $H \supset H'$ or $H' \supset H$; they are *tightly-nested* if they are nested and no other half-space is nested between them. We say that H and H' *cross*, denoted $H \pitchfork H'$, if ∂H and $\partial H'$ intersect. When this occurs, there is a square S in Y in which $\partial H \cap S$ and $\partial H' \cap S$ are the two midcubes of S .

Given $H, K \in \mathcal{H}(Y)$, with $H \supset K$, we will call a sequence of half-spaces $\gamma = \{H_0, H_1, \dots, H_n, H_{n+1}\}$ a *chain* of length n from H to K if

$$H = H_0 \supset H_1 \supset \dots \supset H_n \supset H_{n+1} = K.$$

A chain is *taut* if every adjacent pair is tightly-nested. A chain is *longest* if n is largest possible; such chain is necessarily taut. Note that any chain from H to K can be enlarged to be a taut chain. We will call H_m , where $m = n/2$ if n is even or $m = (n + 1)/2$ if n is odd, the *midpoint* of the chain. We make the following observation.

Lemma 6.1. *Suppose $\gamma = \{H, H_1, \dots, H_n, K\}$ and $\gamma' = \{H, H'_1, \dots, H'_n, K\}$ are two longest chains from H to K . Then the midpoints of γ and γ' either cross or coincide.*

Proof. Let H_m and H'_m be the midpoints of γ and γ' respectively. If $H_m \neq H'_m$ and they do not cross, then the only other possibility is that they are nested. If $H_m \supset H'_m$ then the chain

$$\{H, H_1, \dots, H_m, H'_m, \dots, H'_n, K\}$$

from H to K is strictly longer than γ , a contradiction. If $H'_m \supset H_m$ then again one can construct a chain from H to K that is longer than γ . \square

Equip $Y^{(1)}$ with the edge path metric d . Given two vertices x and y , define the following set of half-spaces:

$$[x, y] = \{H \in \mathcal{H}(Y) : x \notin H, y \in H\}.$$

An edge path from x to y is a geodesic if and only if it does not cross any hyperplane twice. Thus $d(x, y) = \#[x, y]$, where $\#[x, y]$ is the cardinality of $[x, y]$. An essential feature of this distance function

is that (Y, d) is a *median space* [Ger97, Rol98]. That is, for every triple of vertices x, y, z , there exists a unique vertex $m = m(x, y, z)$ such that $d(a, b) = d(a, m) + d(m, b)$ for all distinct $a, b \in \{x, y, z\}$. Equivalently, $[a, b]$ is the disjoint union $[a, m] \cup [m, b]$ for all distinct $a, b \in \{x, y, z\}$.

We say that $H \in \mathcal{H}(Y)$ *intersects* an edge e of Y if ∂H intersects e . Every oriented edge $e = (x, y)$ in Y naturally defines a half-space, namely $H = [x, y]$; and conversely, every half-space naturally defines an orientation on all the edges it intersects.

We say that $H \in \mathcal{H}(Y)$ *intersects* a subcomplex $Z \subset Y$ if H intersects an edge of Z . If Z is connected, then whenever $H, K \in \mathcal{H}(Y)$ both intersect Z and are nested, every half-space in between them also intersects Z . A full subcomplex Z of Y is called *convex* if $Z^{(1)}$ is convex in $Y^{(1)}$ with respect to d ; that is, Z contains all edge geodesics between pairs of vertices in Z .

The cube complex of a right-angled Artin group

Let $A(\Gamma)$ be a right-angled Artin group and $X(\Gamma)$ the Salvetti complex associated to $A(\Gamma)$. By equipping each torus with the standard Euclidean metric, we obtain a CAT(0) cube complex structure on the universal cover $Y(\Gamma)$ of $X(\Gamma)$. The dimension of $Y(\Gamma)$ is the dimension of a maximal dimensional torus in $X(\Gamma)$. Equivalently, this is the largest size of a complete subgraph of Γ .

The 2-skeleton of $Y(\Gamma)$ is the Cayley 2-complex of the defining presentation for $A(\Gamma)$. That is, every oriented edge of $Y(\Gamma)$ is labeled by a generator of $A(\Gamma)$ or its inverse, and squares are bounded by edges whose labels correspond to *distinct* vertices of Γ that bound an edge. For any $H \in \mathcal{H}(Y)$, every oriented edge intersected by H has the same label, and therefore every $H \in \mathcal{H}(Y)$ can be labeled accordingly. If $H \in \mathcal{H}(Y)$ has label a , then \bar{H} has label a^{-1} . Note that $A(\Gamma)$ acts on $\mathcal{H}(Y)$ in a label-preserving manner.

Given $H, K \in \mathcal{H}(Y)$, if $H \pitchfork K$, then this occurs in a square, and their labels correspond to distinct adjacent vertices of Γ . If H, K are tightly-nested, then there exists an oriented edge path of length two whose edges are dual to them, and these edges do not bound a square. Thus their labels either are equal, or correspond to distinct vertices of Γ that are not adjacent.

These observations lead easily to the following results. They are a reformulation of the axioms of *special cube complexes* from [HW08], as stated in [FFT, Lemma 7.3].

Proposition 6.2. *The action of $A(\Gamma)$ on $Y(\Gamma)$ satisfies the following properties:*

1. *There do not exist $H \in \mathcal{H}(Y)$, $f \in A(\Gamma)$ such that $f(\bar{H}) = H$.*
2. *There do not exist $H \in \mathcal{H}(Y)$, $f \in A(\Gamma)$ such that $H \pitchfork fH$.*
3. *There do not exist $H \in \mathcal{H}(Y)$, $f \in A(\Gamma)$ such that H and $f(\bar{H})$ are tightly-nested.*
4. *There do not exist a tightly-nested pair $H, K \in \mathcal{H}(Y)$ and $f \in A(\Gamma)$ such that $H \pitchfork f(K)$. □*

As discussed in [FFT, Section 3], every non-trivial $g \in A(\Gamma)$ has is a *combinatorial axis* L , which is a bi-infinite geodesic in the 1-skeleton of $Y(\Gamma)$ such that $g(L) = L$. Let A_g be the set of half-spaces that intersect L . Though combinatorial axes for g are not unique, A_g is independent of the choice of L . There is a decomposition of A_g into two disjoint collections $A_g = A_g^+ \sqcup A_g^-$ such that $H \in A_g^-$ if and only if $\bar{H} \in A_g^+$, and $H \supset gH$ for all $H \in A_g^+$. In other words, every $H \in A_g^+$ contains the positive (attracting) end of L and \bar{H} contains the negative end. Hence, for every pair $H, K \in A_g^+$, either $H \pitchfork K$ or H and K are nested.

Every non-trivial $g \in A(\Gamma)$ has an *essential characteristic set* $Y_g \subset Y(\Gamma)$. It is a convex $\langle g \rangle$ -invariant subcomplex, intersected by all the half-spaces of A_g and no others.

Proposition 6.3 ([FFT, Section 3]). *The essential characteristic set $Y_g \subset Y(\Gamma)$ satisfies the following properties:*

1. *For any $n > 0$, if w is a cyclically reduced word representing the conjugacy class of g^n and x is a vertex of Y_g , then*

$$|w| = d(x, g^n x) = n d(x, gx).$$

Moreover, there is a combinatorial axis $L \subset Y_g$ for g^n , such that every cyclic conjugate of w appears as a word along the labels of L (read in the positive direction). The set of half-spaces crossing L and containing the positive end of L is A_g^+ .

2. *For every $x \in Y_g$ we have $A_g^+ = \bigcup_{n \in \mathbf{Z}} [g^n x, g^{n+1} x]$.* □

Lemma 6.4. *Given $x, y \in Y_g$ with $[x, y] \subset A_g^+$ and $\#[x, y] > \#[x, gx]/2$, if $[fy, fx] \subset A_g^+$ for some $f \in A(\Gamma)$, then there exists $H \in [x, y]$ and $n \in \mathbf{Z}$ such that $g^n H \in [fy, fx]$.*

Proof. In the following, by a *copy* of a half-space $H \in A_g^+$, we will mean $g^k H$ for some $k \in \mathbf{Z}$. Note that since $[fy, fx] \subset A_g^+$, each element of $[fy, fx]$ is a copy of some element in $[x, gx]$. We proceed by considering two cases.

First suppose $[x, y] \subseteq [x, gx]$. In this case, $y = m(x, y, gx)$, so $\#[x, y] + \#[y, gx] = \#[x, gx]$. Since $\#[x, y] > \#[x, gx]/2$, $\#[y, gx] < \#[x, y] = \#[fy, fx]$. Thus, if $[fy, fx]$ contains at most one copy of each element of $[y, gx]$, then it must contain a copy of some element in $[x, y]$ and the statement holds. So suppose $[fy, fx]$ contains at least two copies of some $K \in [y, gx]$. That is, there exist $i < j$ such that $g^i K, g^j K \in [fy, fx]$. By definition, $g^i K = f\bar{H}$ for some $H \in [x, y]$. Consider $g^{i+1} H \in [g^{i+1} x, g^{i+1} y]$. The half-spaces $g^i K$, $g^{i+1} H$, and $g^j K$ have the same labels (up to sign), so no two of them can cross. Hence they are nested in some linear order. Note that $g^{i+1} x \in g^i K - g^{i+1} H$ and $g^j y \in g^{i+1} H - g^j K$, since $i+1 \leq j$. Thus $g^i K \supset g^{i+1} H \supset g^j K$. But this implies $g^{i+1} H \in [fy, fx]$, concluding the argument in this case.

Now suppose that $[x, y] \not\subseteq [x, gx]$. We can also assume that $[x, gx] \not\subseteq [x, y]$, for otherwise we will already be done. Let z be the median of x , y , and gx . Then $[x, gx] = [x, z] \cup [z, gx]$ and $[x, y] = [x, z] \cup [z, y]$, where $[z, gx]$ and $[z, y]$ are both non-empty. Let $K \in [z, y]$ be any element. For any $K' \in [z, gx]$, since K and K' both lie in A_g^+ , they are either nested or they cross. But $gx \in K'$ and $gx \notin K$ and $y \in K$ and $y \notin K'$, so it is impossible for them to be nested. Therefore, K crosses every element of $[z, gx]$. Now consider $f\bar{K} \in [fy, fx]$, which by assumption lies in A_g^+ , and thus $f\bar{K} = g^n H$ for some $H \in [x, gx]$. If $H \in [z, gx]$ then $g^n H$ and $g^n K$ must cross, as H and K do, but this is not possible since $g^n H = f\bar{K}$. Therefore $H \in [x, z] \subset [x, y]$, which concludes the proof. □

Proposition 6.5. *Let $x, y \in Y_g$ be vertices such that $[x, y] \subset A_g^+$ and $\#[x, y] > \#[x, gx]/2$. Then there does not exist $f \in A(\Gamma)$ such that $[fy, fx] \subset A_g^+$.*

Proof. The proof is by contradiction. Suppose such an element f exists. Note that if $[fy, fx] \subset A_g^+$, then $[g^n fy, g^n fx] \subset A_g^+$ for all $n \in \mathbf{Z}$. Thus, using Lemma 6.4 and replacing f by $g^{-n} f$ if necessary, we may assume that there exist $H, K \in [x, y]$ such that $H = f\bar{K}$. It is not possible that $H = K$ or $H \pitchfork K$, by Proposition 6.2(1) and (2), so H and K must be nested. Suppose $H \supset K$. First we claim

that $f\bar{H} = K$. Since both half-spaces are in A_g^+ , they are either transverse, nested, or equal. However, K and $f\bar{H} = f^2K$ cannot be transverse, by Proposition 6.2(2). Thus, if $f\bar{H} \neq K$ then they must be nested. Now consider a longest chain from H to K :

$$\gamma = \{H, H_1, \dots, H_n, K\}.$$

Maximality and nesting are preserved by the action of $A(\Gamma)$, so

$$f\bar{\gamma} = \{f\bar{K}, f\bar{H}_n, \dots, f\bar{H}_1, f\bar{H}\} = \{H, f\bar{H}_n, \dots, f\bar{H}_1, f\bar{H}\}$$

is a longest chain from H to $f\bar{H}$. If $f\bar{H} \supset K$, then

$$\{H, f\bar{H}_n, \dots, f\bar{H}_1, f\bar{H}, K\}$$

is strictly longer than γ , contradicting the choice of γ . Similarly, if $K \supset f\bar{H}$, then

$$\{H, H_1, \dots, H_n, K, f\bar{H}\} = \{f\bar{K}, H_1, \dots, H_n, K, f\bar{H}\}$$

is strictly longer than $f\bar{\gamma}$. This shows that $f\bar{H} = K$. In particular, both γ and $f\bar{\gamma}$ are longest chains from H to K . To proceed with the contradiction, let H_m and H'_m be the midpoints of γ and $f\bar{\gamma}$ respectively. By Lemma 6.1, either $H_m = H'_m$ or they cross. If n is odd, then $H'_m = f\bar{H}_m$. In this case, if $H_m = H'_m$, then this violates (1) of Proposition 6.2. If H_m and H'_m cross, then this violates (2) of Proposition 6.2. If n is even, then $H'_m = f\bar{H}_{m+1}$. In this case, if $H_m = H'_m$, then $f^{-1}(\bar{H}_m) = H_{m+1}$. Since H_m and H_{m+1} are tightly-nested, this violates (3). Finally, if H_m and H'_m cross, then since H_m and H_{m+1} are tightly-nested, this violates (4).

The case $K \supset H$ will yield a similar contradiction. This concludes the proof. \square

Theorem 6.6. *Given $g \in A(\Gamma)$ and $n > 0$, let w be a cyclically reduced word in the generators of $A(\Gamma)$ representing the conjugacy class of the element g^n in $A(\Gamma)$, and suppose u is a word such that both u and u^{-1} appear as subwords of w (considered as a cyclic word). Then*

$$|u| \leq \frac{|w|}{2n}.$$

Proof. Theorem 6.6 follows easily from Proposition 6.5. To see this, let Y_g be the essential characteristic set for g . By Proposition 6.3 there is a path P in Y_g whose labels read u . Let x and y be, respectively, the initial and terminal endpoints of P , and let \bar{P} be the reversal of P (which reads u^{-1}). Since w represents a positive power of g , $[x, y] \subset A_g^+$.

By the same reasoning there is a path P' , from x' to y' , whose labels read u^{-1} , satisfying $[x', y'] \subset A_g^+$. The action of $A(\Gamma)$ is transitive on vertices, so there exists $f \in A(\Gamma)$ such that $fy = x'$. Since \bar{P} and P' both read u^{-1} , we must have $f\bar{P} = P'$, and therefore $[x', y'] = [fy, fx]$. Using Proposition 6.5 we conclude that

$$|u| \leq \frac{d(x, gx)}{2} = \frac{|w|}{2n}. \quad \square$$

References

- [BBF16] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara, *Stable commutator length on mapping class groups*, Ann. Inst. Fourier (Grenoble) **66** (2016), no. 3, 871–898. MR 3494163
- [BF10] Noel Brady and Max Forester, *Density of isoperimetric spectra*, Geom. Topol. **14** (2010), no. 1, 435–472. MR 2578308
- [BM99] M. Burger and N. Monod, *Bounded cohomology of lattices in higher rank Lie groups*, J. Eur. Math. Soc. (JEMS) **1** (1999), no. 2, 199–235. MR 1694584
- [BM00] Marc Burger and Shahar Mozes, *Lattices in product of trees*, Inst. Hautes Études Sci. Publ. Math. (2000), no. 92, 151–194 (2001). MR 1839489
- [Cal09] Danny Calegari, *scl*, MSJ Memoirs, vol. 20, Mathematical Society of Japan, Tokyo, 2009. MR 2527432
- [CF10] Danny Calegari and Koji Fujiwara, *Stable commutator length in word-hyperbolic groups*, Groups Geom. Dyn. **4** (2010), no. 1, 59–90. MR 2566301 (2011a:20109)
- [CFL16] Matt Clay, Max Forester, and Joel Louwsma, *Stable commutator length in Baumslag-Solitar groups and quasimorphisms for tree actions*, Trans. Amer. Math. Soc. **368** (2016), no. 7, 4751–4785. MR 3456160
- [Che18] Lvzhou Chen, *Spectral gap of scl in free products*, Proc. Amer. Math. Soc. **146** (2018), no. 7, 3143–3151. MR 3787373
- [Cul81] Marc Culler, *Using surfaces to solve equations in free groups*, Topology **20** (1981), no. 2, 133–145. MR 605653
- [DH91] Andrew J. Duncan and James Howie, *The genus problem for one-relator products of locally indicable groups*, Math. Z. **208** (1991), no. 2, 225–237. MR 1128707
- [FFT] Talia Fernós, Max Forester, and Jing Tao, *Effective quasimorphisms on right-angled Artin groups*, to appear in Annales de l’Institut Fourier, [arXiv:1602.05637](https://arxiv.org/abs/1602.05637).
- [Ger97] V. N. Gerasimov, *Semi-splittings of groups and actions on cubings*, Algebra, geometry, analysis and mathematical physics (Russian) (Novosibirsk, 1996), Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk, 1997, pp. 91–109, 190. MR 1624115
- [Heu18] Nicolaus Heuer, *Gaps in scl for amalgamated free products and RAAGs*, ArXiv e-prints (2018), 30 pp, [arXiv:1802.01107](https://arxiv.org/abs/1802.01107).
- [HW08] Frédéric Haglund and Daniel T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620. MR 2377497
- [MW02] Jonathan P. McCammond and Daniel T. Wise, *Fans and ladders in small cancellation theory*, Proc. London Math. Soc. (3) **84** (2002), no. 3, 599–644. MR 1888425
- [Myc55] J. Mycielski, *Sur le coloriage des graphs*, Colloq. Math. **3** (1955), 161–162. MR 0069494
- [Rol98] Martin Roller, *Poc sets, median algebras and group actions. An extended study of Dunwoody’s construction and Sageev’s theorem*, Habilitationsschrift, Regensburg, 1998.

- [Rou79] C. P. Rourke, *Presentations and the trivial group*, Topology of low-dimensional manifolds (Proc. Second Sussex Conf., Chelwood Gate, 1977), Lecture Notes in Math., vol. 722, Springer, Berlin, 1979, pp. 134–143. MR 547460
- [Sta87] John R. Stallings, *A graph-theoretic lemma and group-embeddings*, Combinatorial group theory and topology (Alta, Utah, 1984), Ann. of Math. Stud., vol. 111, Princeton Univ. Press, Princeton, NJ, 1987, pp. 145–155. MR 895613

Max Forester: mf@ou.edu

Ignat Soroko: ignatsoroko@lsu.edu

Jing Tao: jing@ou.edu