WORD LENGTH VERSUS LOWER CENTRAL SERIES DEPTH FOR SURFACE GROUPS AND RAAGS

JUSTIN MALESTEIN AND ANDREW PUTMAN

Abstract. For surface groups and right-angled Artin groups, we prove lower bounds on the shortest word in the generators representing a nontrivial element of the k^{th} term of the lower central series.

1. Introduction

Let G be a group and let $\gamma_k(G)$ be its lower central series:

$$\gamma_1(G) = G \quad \text{and} \quad \gamma_{k+1}(G) = [\gamma_k(G), G] \quad \text{for } k \geq 1.$$

If $\gamma_{k+1}(G) = 1$, then G is at most k-step nilpotent. Let S be a finite generating set for G.

Question. *What is the shortest word in $S^{\pm 1}$ representing a nontrivial element in $\gamma_k(G)$? What are the asymptotics of the length of this word as $k \to \infty$?*

The asymptotic question is only interesting for non-nilpotent groups. It is also natural to only consider groups that are residually nilpotent, i.e., such that

$$\bigcap_{k=1}^{\infty} \gamma_k(G) = 1$$

Let G be a non-nilpotent residually nilpotent group with a finite generating set S. Define for $g \in G$ its associated word norm:

$$\|g\|_S = \min\left\{\ell \mid g \text{ can be written as a word of length } \ell \text{ in } S^{\pm 1}\right\}.$$

The *lower central series depth function* is the following function $d_{G,S}\colon \mathbb{N} \to \mathbb{N}$:

$$d_{G,S}(k) = \min\left\{\|g\|_S \mid g \in \gamma_k(G),\, g \neq 1\right\}.$$

Though $d_{G,S}(k)$ depends on the generating set S, its asymptotic behavior as $k \to \infty$ is independent of S. Our goal in this paper is to find bounds on $d_{G,S}(k)$ for several natural classes of groups G.

1.1. **Free groups.** For $n \geq 2$, let F_n be the free group on $S = \{x_1, \ldots, x_n\}$. These are the most fundamental examples of groups that are residually nilpotent but not nilpotent [15], and both lower and upper bounds on $d_{F_n,S}(k)$ have been studied:

- Using the free differential calculus, Fox [9, Lemma 4.2] proved that $d_{F_n,S}(k) \geq \frac{1}{2}k$ for $k \geq 1$. In [17, Theorem 1.2], the authors improved this to $d_{F_n,S}(k) \geq k$.
- In [17, proofs of Theorems 1.3 and 1.5], the authors proved that $d_{F_n,S}(k) \leq \frac{1}{4}(k+1)^2$. Elkasapy–Thom [8, Theorem 2.2] then improved this to a bound that grows like k^c with $c \approx 1.4411$, and later Elkasapy [7, Theorem 1.1] slightly improved this to k^c where $c = \log_\varphi 2$ and φ is the golden ratio.

The growth rate of $d_{F_n,S}(k)$ thus lies between k and $k^{\log_\varphi 2}$, and Elkasapy conjectured that the growth rate is $k^{\log_\varphi 2}$.

Date: January 18, 2024.

JM was supported in part by a Simons Foundation Collaboration Grant 713006. AP was supported in part by NSF Grant DMS-2305183.

Remark 1.1. Kuperberg [14] showed that finding short words that are deep in $\gamma_k(F_n)$ is related to the problem of approximating elements of $\mathrm{SU}(d)$ by elements of a dense subgroup. He showed you can use such short elements of $\gamma_k(F_n)$ in an algorithm that solves the following problem: given a finitely generated dense subgroup Γ of $\mathrm{SU}(d)$, a specified error tolerance, and a specific element $M \in \mathrm{SU}(d)$, construct short words in Γ that approximate M to the given tolerance. See [14] for precise results and more context. □

1.2. **Upper bounds.** Now let G be a non-nilpotent residually nilpotent group with a finite generating set S. If G contains a non-abelian free subgroup, then using the work of Elkasapy–Thom discussed above we can find an upper bound on $d_{G,S}(k)$ that grows[1] like $k^{1.4411}$. However, lower bounds on $d_{G,S}(k)$ do not follow from the analogous results for free groups, so for the rest of this paper we focus on lower bounds.

1.3. **Surface groups.** Let Σ_g be a closed oriented genus $g \geq 2$ surface and let

$$\pi = \pi_1(\Sigma_g) = \langle a_1, b_1, \ldots, a_g, b_g \mid [a_1, b_1] \cdots [a_g, b_g] = 1 \rangle .$$

Here our convention is that $[x, y] = xyx^{-1}y^{-1}$. The surface group π is residually nilpotent but not nilpotent [2, 10], and shares many features with free groups. Since $g \geq 2$, the subgroup of π generated by a_1 and b_1 is a rank 2 free group. As in §1.2 above, this implies a $k^{1.4411}$ upper bound on the growth rate of $d_{\pi,S}(k)$.

However, lower bounds are more problematic. The known lower bounds for free groups use the free differential calculus, and there is no analogue of the free differential calculus for surface groups.[2] The lower bounds for free groups can also be derived using the "Magnus representations" from free groups to units in rings of power series with noncommuting variables, but again it seems hard to construct suitable analogues for surface groups. Nevertheless, we are able to prove the following:

Theorem A. *Let π be a nonabelian surface group with standard generating set $S = \{a_1, b_1, \ldots, a_g, b_g\}$. Then for all $k \geq 1$ we have $d_{\pi,S}(k) \geq \frac{1}{4}k$.*

The $\frac{1}{4}$ in this theorem is probably not optimal. We make the following conjecture:

Conjecture 1.2. *Let π be a nonabelian surface group with standard generating set $S = \{a_1, b_1, \ldots, a_g, b_g\}$. Then $d_{\pi,S}(k) \geq k$ for all $k \geq 1$.*

See §1.6 below for why our proof likely cannot be extended to prove this conjecture.

1.4. **Right-angled Artin groups.** We will derive Theorem A from an analogous result for right-angled Artin groups, which are defined as follows. Let X be a finite graph. The associated *right-angled Artin group* (RAAG) is the group A_X given by the following presentation:

- The generators are the vertex set $V(X)$.
- The relations are $\{[x, y] = 1 \mid x, y \in V(X) \text{ are joined by an edge}\}$.

Example 1.3. The free abelian group $\mathbb{Z}^n$ is the RAAG with X the complete graph on n vertices, and the free group F_n is the RAAG with X a graph with n vertices and no edges. □

[1]Precise upper bounds are more complicated and depend on how the free subgroup is embedded in G.

[2]The free derivatives are derivations $d\colon F_n \to \mathbb{Z}[F_n]$. For a group G, if there exist nontrivial derivations $d\colon G \to \mathbb{Z}[G]$ then $\mathrm{H}^1(G; \mathbb{Z}[G]) \neq 0$. If G has a compact $K(G, 1)$ this implies that G has more than one end [18], so G cannot be a one-ended group like a surface group.

These groups play an important role in many areas of geometric group theory (see, e.g., [3, 22]). Just like free groups and surface groups, they are residually nilpotent [6], and they are only nilpotent if they are free abelian, i.e., if X is a complete graph. The latter fact can be deduced from the basic observation that if Y is a vertex-induced subgraph of X, then the natural map $A_Y \to A_X$ is split injective; indeed, the map $A_X \to A_Y$ that kills the generators which are not vertices of Y is a right inverse for it.

Remark 1.4. More generally, Baudisch [1] proved that any two elements of a RAAG either commute or generate a free subgroup. This implies that any nonabelian subgroup of a RAAG is non-nilpotent. □

Right-angled Artin groups often contain many surface subgroups [4, 5, 13, 20], and we will prove Theorem A by embedding surface groups into RAAGs and studying the lower central series depth function there. The main result we need along these lines is as follows.

Theorem B. *Let X be a finite graph that is not a complete graph, and let $S = V(X)$ be the generating set of A_X. Then for $k \geq 1$ we have $d_{A_X,S}(k) \geq k$.*

Though Theorem B does not seem to previously appear in the literature, it is implicit in the work of Wade (see [21, Lemma 4.7]), and our proof follows his ideas. The key tool is a version of the "Magnus representation" for RAAGs that was introduced by Droms in his thesis [6], generalizing work of Magnus on free groups. The classical Magnus representations are maps from F_n to units in rings of power series with noncommuting variables (see [16, Chapter 5]). They contain much of the same information as the free derivatives.

1.5. **From RAAGs to surface groups.** Let G be a non-nilpotent residually nilpotent group with finite generating set T and let H be the subgroup of G generated by a finite subset $S < G$. Each $s \in S$ can be written as a word in $T^{\pm 1}$, so we can define

$$r = \max\{\|s\|_T \mid s \in S\}.$$

For $h \in H$, we thus have

$$\|h\|_S \geq \frac{1}{r}\|h\|_T.$$

From this, we see that

$$d_{H,S}(k) \geq \frac{1}{r} d_{G,S}(k) \quad \text{for all } k \geq 1.$$

Since all nonabelian surface groups π are subgroups of RAAGs, Theorem B therefore immediately implies a linear lower bound on the lower central series depth function of π. However, the precise constants depend on the embedding into a RAAG, and without further work might depend on the genus g. To get the genus-independent constant $\frac{1}{4}$ from Theorem A, we will have to carefully control the geometry of our embeddings of surface groups into RAAGs and ensure that we can take $r = 4$ in the above.

Remark 1.5. Many other groups can also be embedded in right-angled Artin groups, and the argument above shows that all of them have linear lower bounds on their lower central series depth functions (which are well-defined by Remark 1.4). □

1.6. **Optimal embeddings.** It is natural to wonder if we can improve the $\frac{1}{4}$ in Theorem A by using a more clever embedding into a RAAG. We conjecture that this is not possible:

Conjecture 1.6. *Let π be a nonabelian surface group with standard generating set $S = \{a_1, b_1, \ldots, a_g, b_g\}$, let X be a finite graph, and let $\phi\colon \pi \hookrightarrow A_X$ be an embedding. Then there exists some $s \in S$ such that $\|\phi(s)\|_{V(X)} \geq 4$.*

Remark 1.7. As we will discuss in §3 below, Crisp–Wiest [5] gave an explicit description of all homomorphisms from surface groups to RAAGs in terms of collections of loops on the surface. To prove Conjecture 1.6, what one would have to show is that if $\phi\colon \pi \to A_X$ is a map from a surface group to a RAAG arising from the Crisp–Wiest construction that does not satisfy the conclusion of Conjecture 1.6, then ϕ is not injective. □

1.7. **Sublinearity.** We close by posing the following question:

Question 1.8. *Does there exist a non-nilpotent residually nilpotent group G equipped with a finite generating set S such that $d_{G,S}$ grows sublinearly?*

By Remark 1.5, such a group G cannot be a subgroup of a RAAG.

1.8. **Outline.** We prove Theorem B in §2 and Theorem A in §4. This last section depends on the preliminary §3, which discusses work of Crisp–Wiest parameterizing maps from surface groups to RAAGs.

1.9. **Acknowledgments.** We thank Greg Kuperberg for some useful references.

2. Right-angled Artin groups

Let X be a finite graph with associated right-angled Artin group A_X. In this section, we first discuss some structural results about A_X and then prove Theorem B.

2.1. **Monoid.** In addition to the right-angled Artin group A_X, we will also need the right-angled Artin monoid M_X. This is the associative monoid with the following presentation:

- The generators are the vertices $V(X)$ of X. To distinguish these generators from the corresponding generators of A_X, we will sometimes write them with bold-face letters. In other words, s denotes an element of A_X and $\mathbf{s}$ denotes an element of M_X.
- The relations are $\{\mathbf{xy} = \mathbf{yx} \mid x, y \in V(X) \text{ are joined by an edge}\}$.

There is a monoid homomorphism $M_X \to A_X$ whose image is the set of all elements of A_X that can be represented by "positive words". As we will discuss below, this monoid homomorphism is injective.

2.2. **Normal form.** Let $S = V(X)$ be the generating set for A_X and M_X. Consider a word

$$w = s_1^{e_1} \cdots s_n^{e_n} \quad \text{with } s_1, \ldots, s_n \in S \text{ and } e_1, \ldots, e_n \in \mathbb{Z}.$$

This word represents an element of A_X, and if $e_i \geq 0$ for all $1 \leq i \leq n$ it represents an element of M_X (here for conciseness we are not using our bold-face conventions). We say that w is *fully reduced* if it satisfies the following conditions:

- Each e_i is nonzero.
- For all $1 \leq i < j \leq n$ with $s_i = s_j$, there exists some k with $i < k < j$ such that s_k does not commute[3] with $s_i = s_j$.

Note that this implies in particular that $s_i \neq s_{i+1}$ for all $1 \leq i < n$, so w is reduced as a word in the free group on S. It is clear that every element of A_X and M_X can be represented by a fully reduced word.

This representation is unique in the following sense:

[3]As observed earlier, A_Y embeds in A_X for any vertex-induced subgraph Y, so this is equivalent to s_k being distinct from and not adjacent to $s_i = s_j$.

- Consider fully reduced words

$$w = s_1^{e_1} \cdots s_n^{e_n} \quad \text{and} \quad w' = t_1^{f_1} \cdots t_m^{f_m}$$

representing the same element of A_X or M_X. Then we can obtain w' from w by a sequence of *swaps*, i.e., flipping adjacent terms $s_i^{e_i}$ and $s_{i+1}^{e_{i+1}}$ such that s_i commutes with s_{i+1}.

For A_X, this uniqueness was stated without proof by Servatius [19]. The earliest proof we are aware of is in Green's thesis [11]. Alternate proofs can be found in [5, Proposition 9] and [21, Theorem 4.14]. Using the monoid homomorphism $M_X \to A_X$, the uniqueness for M_X follows[4] from that of A_X. Note that this uniqueness also implies that the monoid homomorphism $M_X \to A_X$ is injective.

The following lemma shows that fully reduced words realize the word norm in A_X:

Lemma 2.1. *Let X be a finite graph. Let $S = V(X)$ be the generating set for A_X. Consider some $w \in A_X$, and represent w by a fully reduced word*

$$w = s_1^{e_1} \cdots s_n^{e_n} \quad \text{with } s_1, \ldots, s_n \in S \text{ and } e_1, \ldots, e_n \in \mathbb{Z}.$$

Then $\|w\|_S = |e_1| + \cdots + |e_n|$.

Proof. Immediate from the uniqueness up to swaps of fully reduced words as well as the fact that taking an arbitrary word and putting it in fully reduced form does not lengthen the word. □

2.3. **Monoid ring.** Let $\mathbb{Z}[M_X]$ be the monoid ring whose elements are formal $\mathbb{Z}$-linear combinations of elements of M_X. Since the relations in M_X are all of the form $\mathbf{xy} = \mathbf{yx}$ for generators $\mathbf{x}$ and $\mathbf{y}$, all words representing an element $\mathbf{m} \in M_X$ have the same length, which we will denote $\ell(\mathbf{m})$. This length function satisfies $\ell(\mathbf{m}_1\mathbf{m}_2) = \ell(\mathbf{m}_1) + \ell(\mathbf{m}_2)$ for $\mathbf{m}_1, \mathbf{m}_2 \in M_X$. For $k \geq 0$, define

$$M_X^{(k)} = \{\mathbf{m} \in M_X \mid \ell(\mathbf{m}) = k\}.$$

The monoid ring $\mathbb{Z}[M_X]$ is a graded ring with $\mathbb{Z}[M_X]_{(k)} = \mathbb{Z}[M_X^{(k)}]$.

2.4. **Partially commuting power series.** Let $I \subset \mathbb{Z}[M_X]$ be the ideal generated by the elements of the generating set $V(X)$. For $k \geq 1$, the ideal I^k consists of $\mathbb{Z}$-linear combinations of $\mathbf{m} \in M_X$ with $\ell(\mathbf{m}) \geq k$. Define

$$\mathcal{P}_X = \varprojlim \mathbb{Z}[M_X]/I^k.$$

Elements of the inverse limit $\mathcal{P}_X$ can be regarded as power series

$$\sum_{k=0}^{\infty} \mathbf{m}_k \quad \text{with } \mathbf{m}_k \in \mathbb{Z}[M_X]_{(k)} \text{ for all } k \geq 0.$$

Each $\mathbf{m}_k$ is a linear combination of products of k generators from $V(X)$, some of which commute and some of which do not. Multiplication works in the usual way:

$$\left(\sum_{k=0}^{\infty} \mathbf{m}_k\right)\left(\sum_{k'=0}^{\infty} \mathbf{m}'_{k'}\right) = \sum_{\ell=0}^{\infty}\left(\sum_{k+k'=\ell} \mathbf{m}_k \mathbf{m}'_{k'}\right).$$

[4]Whether this is a circular argument depends on the proof of uniqueness used for A_X. The geometric proof from [5, Proposition 9] works directly with groups, and does not even implicitly prove anything about monoids.

2.5. **Magnus representation.** We now discuss the Magnus representation of A_X, which was introduced by Droms in his thesis [6], generalizing classical work of Magnus for free groups (see [16, Chapter 5]). See [21] for a survey. The starting point is the observation that for $s \in V(X)$, we have the following identity in $\mathcal{P}_X$:

$$(1+\mathbf{s})(1-\mathbf{s}+\mathbf{s}^2-\mathbf{s}^3+\cdots)=1.$$

In other words, $1+\mathbf{s}$ is a unit in $\mathcal{P}_X$. If generators $s, s' \in V(X)$ commute, then $1+\mathbf{s}$ and $1+\mathbf{s}'$ also commute. It follows that we can define a homomorphism

$$\mu\colon A_X \longrightarrow (\mathcal{P}_X)^{\times}$$

via the formula

$$\mu(s) = 1+\mathbf{s} \quad \text{for } s \in V(X).$$

2.6. **Dimension subgroups and the lower central series.** Recall that $I \subset \mathbb{Z}[M_X]$ is the ideal generated by elements of the generating set $V(X)$. There is a corresponding ideal $\mathcal{I} \subset \mathcal{P}_X$ consisting of all elements with constant term 0. For $k \geq 1$, the k^{th} *dimension subgroup* of A_X, denoted $D_k(A_X)$, is the kernel of the composition

$$A_X \xrightarrow{\mu} \mathcal{P}_X \longrightarrow \mathcal{P}_X/\mathcal{I}^k.$$

In other words, $D_k(A_X)$ consists of elements $w \in A_X$ such that

$$\mu(w) = 1 + (\text{terms of degree at least } k).$$

The most important theorem about $D_k(A_X)$ identifies it with the k^{th} term of the lower central series of A_X:

Theorem 2.2 ([21, Theorem 6.3])**.** *Let X be a finite graph. Then $D_k(A_X) = \gamma_k(A_X)$ for all $k \geq 1$.*

Remark 2.3. In fact, for what follows all we need is the much easier fact that $\gamma_k(A_X) \subset D_k(A_X)$, which appears in Droms's thesis [6]. For this, since $D_1(A_X) = A_X = \gamma_1(A_X)$ it is enough to verify that

$$[D_k(A_X), D_\ell(A_X)] \subset D_{k+\ell}(A_X),$$

which is immediate from the definitions. □

2.7. **Lower bounds for the lower central series of a RAAG.** We close this section by proving Theorem B. As we said in the introduction, the proof closely follows ideas of Wade [21].

Proof of Theorem B. We start by recalling the statement. Let X be a finite graph that is not a complete graph and let $S = V(X)$ be the generating set for A_X. Consider a nontrivial element $w \in A_X$, and let $k = \|w\|_S$ be its word norm in the generating set S. We must prove that $w \notin \gamma_{k+1}(A_X)$. By Theorem 2.2, it is enough to prove that $w \notin D_{k+1}(A_X)$.

Represent w by a fully reduced word:

$$w = s_1^{e_1} \cdots s_n^{e_n} \quad \text{with } s_1,\ldots,s_n \in S \text{ and } e_1,\ldots,e_n \in \mathbb{Z}.$$

By Lemma 2.1, we have

$$\|w\|_S = |e_1| + \cdots + |e_n| \geq n.$$

It is thus enough to prove that $w \notin D_{n+1}(A_X)$. To do this, is enough to prove that a term of degree n appears in $\mu(w) \in \mathcal{P}_X$.

An easy induction shows that for all $1 \leq i \leq n$, we have

$$\mu(s_i^{e_i}) = (1+\mathbf{s}_i)^{e_i} = 1 + e_i\mathbf{s}_i + \mathbf{s}_i^2\mathbf{t}_i \quad \text{for some } \mathbf{t}_i \in \mathcal{P}_X.$$

It follows that

$$\mu(w) = \left(1 + e_1\mathbf{s}_1 + \mathbf{s}_1^2\mathbf{t}_1\right)\left(1 + e_2\mathbf{s}_2 + \mathbf{s}_2^2\mathbf{t}_2\right)\cdots\left(1 + e_n\mathbf{s}_n + \mathbf{s}_n^2\mathbf{t}_n\right). \tag{2.1}$$

Say that some $\mathbf{m} \in M_X$ is *square-free* if it cannot be expressed as a word in the generators $S = V(X)$ for the monoid M_X with two consecutive letters the same generator.[5] It is immediate from the uniqueness up to swaps of fully reduced words that the fully reduced word $\mathbf{s}_1\mathbf{s}_2\cdots\mathbf{s}_n$ represents a square-free element of M_X. When we expand out (2.1), the only square-free term of degree n is

$$e_1e_2\cdots e_n\mathbf{s}_1\mathbf{s}_2\cdots\mathbf{s}_n.$$

It follows that this degree n term survives when we expand out $\mu(w)$, as desired. □

3. Mapping surface groups to RAAGs

Before we can prove Theorem A, we must discuss some work of Crisp–Wiest [5] that parameterizes maps from surface groups to RAAGs. We will not need the most general form of their construction (which they prove can give *any* homomorphism from a surface group to a RAAG), so we will only describe a special case of it. Fix a closed oriented surface Σ and a basepoint $* \in \Sigma$.

3.1. **Crisp–Wiest construction.** A *simple dissection*[6] on Σ is a finite collection $\mathcal{L}$ of oriented simple closed curves on Σ satisfying the following conditions:

- None of the curves contain the basepoint $*$.
- Any two curves in $\mathcal{L}$ intersect transversely.
- There are no triple intersection points between three curves in $\mathcal{L}$.

For a simple dissection $\mathcal{L}$, let $X(\mathcal{L})$ be the graph whose vertices are the curves in $\mathcal{L}$ and where two vertices are joined by an edge if the corresponding curves intersect. Crisp–Wiest [5] proved that the following gives a well-defined homomorphism $\phi\colon \pi_1(\Sigma, *) \to A_{X(\mathcal{L})}$:

- Consider some $x \in \pi_1(\Sigma, *)$. Realize x by an immersed based loop $\mathbf{x}\colon [0,1] \to \Sigma$ that is transverse to all the curves in $\mathcal{L}$ and avoids intersection points between curves of $\mathcal{L}$. If $\mathbf{x}$ is disjoint from all the curves in $\mathcal{L}$, then $\phi(x) = 1$. Otherwise, let
$$0 < t_1 < \cdots < t_n < 1$$
be the collection of all values such that $\mathbf{x}(t_i)$ is contained in some $\gamma_i \in \mathcal{L}$. For $1 \leq i \leq n$, let $e_i = \pm 1$ be the sign of the intersection of $\mathbf{x}$ with the oriented loop γ_i at $x(t_i)$. Then
$$\phi(x) = \gamma_1^{e_1}\cdots\gamma_n^{e_n} \in A_{X(\mathcal{L})}.$$

We will say that ϕ is the map obtained by applying the *Crisp–Wiest construction* to $\mathcal{L}$.

3.2. **Injectivity criterion.** Crisp–Wiest [5] describe an approach for proving that ϕ is injective in certain cases. To describe it, we must introduce some more terminology. For a simple dissection $\mathcal{L}$ on Σ, let

$$G(\mathcal{L}) = \bigcup_{\gamma \in \mathcal{L}} \gamma,$$

which we view as a graph embedded in Σ_g with a vertex for each intersection point between curves in $\mathcal{L}$. We say that $\mathcal{L}$ is a *filling curve system* if each component of $\Sigma \setminus G(\mathcal{L})$ is a disk.

[5]Be warned that it is possible for an element to have one such expression while not being square-free. For instance, if $\mathbf{s}, \mathbf{s}' \in S$ are distinct commuting generators then $\mathbf{s}\mathbf{s}'\mathbf{s}$ is not square-free since $\mathbf{s}\mathbf{s}'\mathbf{s} = \mathbf{s}^2\mathbf{s}'$.

[6]Crisp and Wiest use the term dissection for a collection of curves which satisfy some conditions and have a certain decoration. We add "simple" to indicate that we do not have any decoration.

For a component U of $\Sigma \setminus G(\mathcal{L})$, the boundary of U can be identified with a circuit in the graph $G(\mathcal{L})$. Say that U satisfies the *injectivity criterion* if the following holds for any two distinct edges e and e' in the boundary of U. Let γ and γ' be the oriented curves in $\mathcal{L}$ that contain e and e', respectively. We then require that $\gamma \neq \gamma'$ and that if γ intersects γ', then e and e' are adjacent edges in the boundary of U.

We can now state our injectivity criterion:

Proposition 3.1. *Let Σ be a closed oriented surface equipped with a basepoint $*$ and let $\mathcal{L}$ be a filling simple dissection on Σ. For all components U of $\Sigma \setminus G(\mathcal{L})$, assume that U satisfies the injectivity criterion. Then the map $\phi\colon \pi_1(\Sigma, *) \to A_{X(\mathcal{L})}$ obtained by applying the Crisp–Wiest construction to $\mathcal{L}$ is injective.*

While Proposition 3.1 is not explicitly stated or proved in [5], it is implicit in their work. We present a proof for the convenience of the reader. This requires some preliminaries.

Remark 3.2. The injectivity criterion implies that the dual cubulation to $\mathcal{L}$ we introduce below is a *special cube complex* in the sense of Haglund–Wise [12]. The paper [5] predates [12], and the machinery of [12] is unnecessary for this application. □

3.3. **Salvetti complex.** Let X be a finite graph and let A_X be the corresponding right-angled Artin group. The *Salvetti complex* of A_X, denoted $\mathcal{S}(X)$, is a certain non-positively curved cube complex[7] with $\pi_1(\mathcal{S}(X)) = A_X$. It can be constructed as follows. Enumerate the vertices of X as

$$V(X) = \{v_1, \ldots, v_n\}.$$

Identify S^1 with the the unit circle in $\mathbb{C}$, so $1 \in S^1$ is a basepoint. For a subset $I \subset \{v_1, \ldots, v_n\}$ of cardinality k, let $S_I \cong (S^1)^k$ be

$$S_I = \left\{(z_1, \ldots, z_n) \in (S^1)^n \mid z_i = 1 \text{ for all } i \text{ with } v_i \notin I\right\}.$$

A subset $I \subset \{v_1, \ldots, v_n\}$ is a k-clique of X if the subgraph of X induced by I is a complete subgraph on k vertices. A clique is a set of vertices that forms a k-clique for some k. With these definitions, $\mathcal{S}(X)$ is the union of the S_I as I ranges over cliques in X. The space $\mathcal{S}(X)$ can be given a cube complex structure containing a k-cube for each k-clique in X. In particular, it has a single vertex (i.e., 0-cube) corresponding to the (empty) 0-clique.

3.4. **Dual cubulation.** Now let $\mathcal{L}$ be a filling dissection on Σ_g. We can form a cube complex structure on Σ_g called the *cube complex structure dual to* $\mathcal{L}$ in the following way. We start by defining the 1-skeleton σ of our cube complex structure:

- Put a vertex of σ in the interior of each component of $\Sigma_g \setminus G(\mathcal{L})$. For the component containing the basepoint $*$, the vertex should be $*$.
- For each edge e of $G(\mathcal{L})$, connect the vertices in the components on either side of e by an edge of σ.

A component of $\Sigma_g \setminus G(\mathcal{L})$ looks like the following:

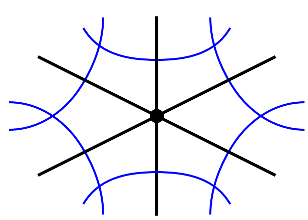

Here $G(\mathcal{L})$ is in blue and σ is in black. Each edge coming out of the vertex of σ shown in the figure terminates in the vertex in the adjacent component.

[7]Here a cube complex is non-positively curved if its universal cover is CAT(0).

Now consider a component C of $\Sigma_g \setminus \sigma$. The component C contains exactly one vertex of $G(\mathcal{L})$, and the boundary of C is composed of four edges of σ as follows:

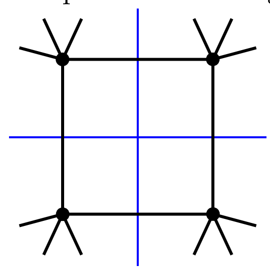

Here again the graph $G(\mathcal{L})$ is blue and σ is black. Complete σ to a cube complex structure by attaching a square to each such C.

3.5. **Proof of Proposition 3.1.** We first recall what we must prove. Let Σ be a closed oriented surface equipped with a basepoint $*$ and let $\mathcal{L}$ be a filling simple dissection on Σ. For all components U of $\Sigma \setminus G(\mathcal{L})$, assume that U satisfies the injectivity criterion. We must prove that the map $\phi\colon \pi_1(\Sigma, *) \to A_{X(\mathcal{L})}$ obtained by applying the Crisp–Wiest construction to $\mathcal{L}$ is injective.

Endow Σ with the cube complex structure dual to $\mathcal{L}$, and let $\mathcal{S}(X(\mathcal{L}))$ be the Salvetti complex of $A_{\mathcal{L}(X)}$. We start by constructing a map of cube complexes $f\colon \Sigma \to \mathcal{S}(X(\mathcal{L}))$ such that

$$f_*\colon \pi_1(\Sigma, *) \to \pi_1(\mathcal{S}(X(\mathcal{L}))) = A_{X(\mathcal{L})}$$

equals ϕ. Define f as follows:

- The map f sends each vertex of Σ to the unique vertex of $\mathcal{S}(X(\mathcal{L}))$.
- For an edge e of Σ that crosses an oriented loop γ of $\mathcal{L}$, the map f takes e isometrically to the loop of $\mathcal{S}(X(\mathcal{L}))$ corresponding to the 1-clique $\{\gamma\}$ of $X(\mathcal{L})$. Orienting e such that the intersection of e with γ is positive, we do this such that $f(e)$ goes around the loop in the direction corresponding to the generator γ of $\pi_1(\mathcal{S}(X(\mathcal{L}))) = A_{X(\mathcal{L})}$.
- For a 2-cube c of Σ centered at an intersection of loops γ_1 and γ_2 of $\mathcal{L}$, the map f sends c isometrically to the 2-cube corresponding to the 2-clique $\{\gamma_1, \gamma_2\}$ of $X(\mathcal{L})$.

With these definitions, it is clear that $f_* = \phi$.

By [5, Theorem 1], the map $f_* = \phi$ will be an injection if for every vertex v of Σ, the map f take the link of v injectively into a full subcomplex of the link of $f(v)$ in $\mathcal{S}(X(\mathcal{L}))$. These links have the following description:

- The vertex v lies in some component U of $\Sigma \setminus G(\mathcal{L})$. The link of v is a cycle whose vertices are precisely the edges of $G(\mathcal{L})$ surrounding U.
- The vertex $f(v)$ is the unique vertex of $\mathcal{S}(X(\mathcal{L}))$. Its link is the following complex:
  - There are two vertices for each generator γ of $A_{X(\mathcal{L})}$ (or alternatively, each $\gamma \in \mathcal{L}$), one corresponding to the positive direction and the other to the negative direction.
  - A collection of vertices forms a simplex if they correspond to distinct generators of $A_{X_{\mathcal{L}}}$ all of which commute.

From this description, we see that the fact that U satisfies the injectivity criterion ensures that f takes the link of v injectively into a full subcomplex of the link of $f(v)$ in $\mathcal{S}(X(\mathcal{L}))$, as desired. $\square$

4. Bounds on surface groups

We now study the lower central series of surface groups and prove Theorem A.

Proof of Theorem A. We start by recalling the statement. For some $g \geq 2$, let Σ_g be a closed oriented genus g surface equipped with a basepoint $*$ and let $S = \{a_1, b_1, \ldots, a_g, b_g\}$ be the standard basis for $\pi = \pi_1(\Sigma_g, *)$. Our goal is to prove that $d_{\pi,S}(k) \geq \frac{1}{4}k$ for all $k \geq 1$. Equivalently, consider some nontrivial $w \in \gamma_k(\pi)$. We must prove that $\|w\|_S \geq \frac{1}{4}k$.

What we will do is find a finite graph X and an injective homomorphism $\phi\colon \pi \to A_X$ such that letting $T = V(X)$ be the generating set for A_X, we have $\|\phi(s)\|_T \leq 4$ for all $s \in S$. We then have $\phi(w) \in \gamma_k(A_X)$, and since ϕ is injective we have $\phi(w) \neq 1$. Since π is nonabelian the graph X is not a complete graph, so we can apply Theorem B to deduce that $\|\phi(w)\|_T \geq k$. Since $\|\phi(s)\|_T \leq 4$ for all $s \in S$, we conclude that

$$\|w\|_S \geq \frac{1}{4}\|\phi(w)\|_T \geq \frac{1}{4}k,$$

as desired.

It remains to construct X and ϕ. We can draw the elements of S as follows, where a_k "encircles" the kth hole from the left:

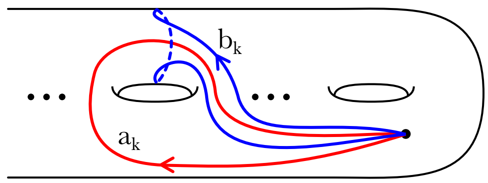

Let

$$\mathcal{L} = \{x_0, \ldots, x_g, y_1, \ldots, y_g, z\}$$

be the following simple dissection on Σ_g:

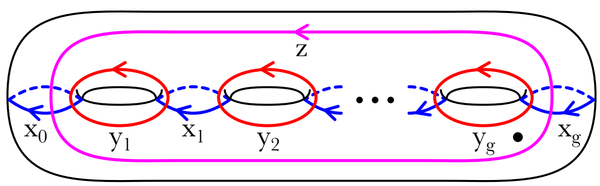

Let $\phi\colon \pi \to A_{X(\mathcal{L})}$ be the homomorphism obtained by applying the Crisp–Wiest construction to $\mathcal{L}$ and let $T = V(X(\mathcal{L}))$ be the generating set for $A_{X(\mathcal{L})}$. There are four components of $\Sigma_g \setminus G(\mathcal{L})$, and by inspection each of them satisfies the injectivity criterion from §3.2. Proposition 3.1 thus implies that ϕ is injective. By construction, the following hold:

$$\phi(a_k) = x_{k-1}x_k^{-1},$$
$$\phi(b_k) = x_k z y_k x_k^{-1}.$$

These formulas imply that $\|\phi(s)\|_T \leq 4$ for all $s \in S$, as desired. □

References

[1] A. Baudisch, Subgroups of semifree groups, Acta Math. Acad. Sci. Hungar. 38 (1981), no.1–4, 19–28. (Cited on page 3.)

[2] G. Baumslag, On generalised free products, Math. Z. 78 (1962), 423–438. (Cited on page 2.)

[3] R. M. Charney, An introduction to right-angled Artin groups, Geom. Dedicata 125 (2007), 141–158. arXiv:math/0610668 (Cited on page 3.)

[4] J. S. Crisp, M. Sageev and M. V. Sapir, Surface subgroups of right-angled Artin groups, Internat. J. Algebra Comput. 18 (2008), no. 3, 443–491. arXiv:0707.1144 (Cited on page 3.)

[5] J. S. Crisp and B. Wiest, Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups, Algebr. Geom. Topol. 4 (2004), 439–472. arXiv:math/0303217 (Cited on pages 3, 4, 5, 7, 8, and 9.)

[6] C. Droms, Graph Groups, PhD thesis, Syracuse University, 1983. (Cited on pages 3 and 6.)

[7] A. I. Elkasapy, A new construction for the shortest non-trivial element in the lower central series, preprint, Oct 2016. arXiv:1610.09725 (Cited on page 1.)
[8] A. I. Elkasapy and A. Thom, On the length of the shortest non-trivial element in the derived and the lower central series, J. Group Theory 18 (2015), no. 5, 793–804. arXiv:1311.0138 (Cited on page 1.)
[9] R. H. Fox, Free differential calculus. I. Derivation in the free group ring, Ann. of Math. (2) 57 (1953), 547–560. (Cited on page 1.)
[10] K. N. Frederick, The Hopfian property for a class of fundamental groups, Comm. Pure Appl. Math. 16 (1963), 1–8. (Cited on page 2.)
[11] E. R. Green, Graph products of groups, PhD thesis, University of Leeds, 1990. (Cited on page 5.)
[12] F. Haglund & D. Wise, Special cube complexes, Geom. Funct. Anal. 17 (2008), no.5, 1551–1620. (Cited on page 8.)
[13] S. Kim, On right-angled Artin groups without surface subgroups, Groups Geom. Dyn. 4 (2010), no. 2, 275–307. arXiv:0811.1946 (Cited on page 3.)
[14] G. Kuperberg, Breaking the cubic barrier in the Solovay-Kitaev algorithm, preprint, June 2023. arXiv:2306.13158 (Cited on page 2.)
[15] W. Magnus, Beziehungen zwischen Gruppen und Idealen in einem speziellen Ring, Math. Ann. 111 (1935), no. 1, 259–280. (Cited on page 1.)
[16] W. Magnus, A. Karrass and D. M. Solitar, *Combinatorial group theory*, second revised edition, Dover Publications, Inc., New York, 1976. (Cited on pages 3 and 6.)
[17] J. Malestein and A. Putman, On the self-intersections of curves deep in the lower central series of a surface group, Geom. Dedicata 149 (2010), 73–84. arXiv:0901.2561 (Cited on page 1.)
[18] G. P. Scott and C. T. C. Wall, Topological methods in group theory, in *Homological group theory (Proc. Sympos., Durham, 1977)*, 137–203, London Math. Soc. Lecture Note Ser., 36, Cambridge Univ. Press, Cambridge. (Cited on page 2.)
[19] H. Servatius, Automorphisms of graph groups, J. Algebra 126 (1989), no. 1, 34–60. (Cited on page 5.)
[20] H. Servatius, C. Droms and B. Servatius, Surface subgroups of graph groups, Proc. Amer. Math. Soc. 106 (1989), no. 3, 573–578. (Cited on page 3.)
[21] R. D. Wade, The lower central series of a right-angled Artin group, Enseign. Math. 61 (2015), no. 3-4, 343–371. arXiv:1109.1722 (Cited on pages 3, 5, and 6.)
[22] D. T. Wise, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012. (Cited on page 3.)

Dept of Mathematics; University of Oklahoma; 601 Elm Ave; Norman, OK 73091
Email address: jmalestein@ou.edu

Dept of Mathematics; University of Notre Dame; 255 Hurley Hall; Notre Dame, IN 46556
Email address: andyp@nd.edu